

Towards Practical Deletion Repair of Inconsistent DL-programs¹

Thomas Eiter and Michael Fink and Daria Stepanova²

Abstract. Nonmonotonic Description Logic (DL-) programs couple nonmonotonic logic programs with DL-ontologies through queries in a loose way which may lead to inconsistency, i.e., lack of an answer set. Recently defined repair answer sets remedy this but a straightforward computation method lacks practicality. We present a novel evaluation algorithm for deletion repair answer sets based on support sets, which reduces evaluation of *DL-Lite_A* ontology queries to constraint matching. This leads to significant performance gains towards inconsistency management in practice.

1 INTRODUCTION

The need for combining rules with ontologies has led to different approaches (see [12] and references therein). Among them Nonmonotonic Description Logic (DL-) programs [8] embody a loose coupling of rules and a DL-knowledge base (i.e., ontology) which the rules can query through special *DL-atoms*. As the ontology may be enriched before query evaluation, a bidirectional flow of information enables one to solve advanced reasoning tasks on top of ontologies. On the other hand, the information flow may lead to inconsistency, that is, cause a DL-program to lack answer sets (i.e., models).

Example 1. Consider the DL-program Π in Figure 1, which represents information about children of a primary school and their parents in simplistic form [7]. It has an ontology \mathcal{O} consisting of a concept taxonomy (TBox) \mathcal{T} in (1)-(3) and a set \mathcal{A} of assertions (ABox), i.e. facts, about individuals in (4)-(6). The rules \mathcal{P} contain further facts (7), (8) and proper rules: (9) determines fathers from the ontology, upon feeding information to it; (10) checks, informally, against them for local parent information (*ischildof*) that a child has surely not two fathers, unless it is adopted; finally (11)-(12) single out contact persons for children (by default, the parents); for adopted children, fathers in \mathcal{O} are omitted if some other contact exists. Inconsistency arises as *john*, who is not provably adopted, has *pat* as father by the ontology, and by the local information possibly also *alex*.

An inconsistent program may be unusable since it yields no answer set. To cope with this, notions of DL-program repair (under different semantics) have been formalized recently [7]. They give rise to so-called repair answer sets under the supposition of suitable changes to the ABox (data) of the ontology. As for repair computation, however, a natural realization of corresponding algorithms turns out to be too naive and does not scale for practical applications since the number of ABoxes to be checked might be large in general.

Figure 1. DL-program Π over a family ontology

$$\mathcal{O} = \left\{ \begin{array}{ll} (1) \textit{Child} \sqsubseteq \exists \textit{hasParent} & (4) \textit{Male}(\textit{pat}) \\ (2) \textit{Adopted} \sqsubseteq \textit{Child} & (5) \textit{Male}(\textit{john}) \\ (3) \textit{Female} \sqsubseteq \neg \textit{Male} & (6) \textit{hasParent}(\textit{john}, \textit{pat}) \end{array} \right\}$$

$$\mathcal{P} = \left\{ \begin{array}{l} (7) \textit{ischildof}(\textit{john}, \textit{alex}); \quad (8) \textit{boy}(\textit{john}); \\ (9) \textit{hasfather}(X, Y) \leftarrow \textit{DL}[\textit{Male} \uplus \textit{boy}; \textit{Male}](Y) \\ \quad \textit{DL}[\textit{hasParent}](X, Y); \\ (10) \perp \leftarrow \textit{not DL}[\textit{Adopted}](X), Y_1 \neq Y_2, \\ \quad \textit{hasfather}(X, Y_1), \textit{ischildof}(X, Y_2), \\ \quad \textit{not DL}[\textit{Child} \uplus \textit{boy}; \neg \textit{Male}](Y_2); \\ (11) \textit{contact}(X, Y) \leftarrow \textit{DL}[\textit{hasParent}](X, Y), \\ \quad \textit{not omit}(X, Y); \\ (12) \textit{omit}(X, Y) \leftarrow \textit{DL}[\textit{Adopted}](X), Z \neq Y, \\ \quad \textit{hasfather}(X, Y), \textit{contact}(X, Z) \end{array} \right\}$$

We thus propose an alternative approach for repair answer set computation which is based on the notion of support sets. Intuitively, a support set for a ground DL-atom $a = \textit{DL}[\lambda; Q](\mathbf{t})$ is a set of assertions which together with the original TBox is sufficient for the given DL-query $Q(\mathbf{t})$ to be derived.

Our basic method is to precompute small support sets for each DL-atom on a nonground level. Then during the DL-program evaluation, for each candidate interpretation the ground instantiations of the support sets are effectively obtained. These support sets help to prune the answer set search space and they are also used for the ABox repair construction.

This approach is particularly attractive for *DL-Lite_A* ontologies, for which it scales much better than the naive approach [7], as the number of necessary support sets is small and they are easy to compute.

Our contributions on DL-program repair computation are:

- (1) We formally establish support sets as sound and complete structures for DL-atom evaluation avoiding ontology access. Moreover, this characterization is faithfully lifted to the nonground level, enabling scalability of exploitation.
- (2) Nonground support sets for DL-atoms over *DL-Lite_A* ontologies can not only be efficiently computed, but also turn out to be small. We utilize this fact to devise an algorithm for the effective computation of deletion repairs of DL-programs under so-called *flp* semantics and discuss potential generalizations.
- (3) We report experimental results obtained implementing our approach and evaluating its scalability on a set of benchmark scenarios; they provide evidence for the effectiveness of the method.

¹ Supported by the Austrian Science Fund (FWF) project P24090.

² Institute of Information Systems, Vienna University of Technology, Austria. email: {eiter,fink,dasha}@kr.tuwien.ac.at

2 PRELIMINARIES

We start with briefly recalling DL-programs and repair answer sets; see [8, 7] for more details and background.

Syntax. A DL-program is a pair $\Pi = \langle \mathcal{O}, \mathcal{P} \rangle$ of a finite ontology \mathcal{O} and a finite set of rules \mathcal{P} as follows.

- \mathcal{O} is a *DL-knowledge base* (or *ontology*) over a signature $\Sigma_o = \langle \mathbf{I}, \mathbf{C}, \mathbf{R} \rangle$ with a set \mathbf{I} of individuals, a set \mathbf{C} of concept names and a set \mathbf{R} of role names. We assume here throughout that $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$ is a *consistent DL-Lite_A* KB [3] with a *TBox* \mathcal{T} and *ABox* \mathcal{A} , which are sets of axioms capturing taxonomic resp. factual knowledge. Concepts C and roles R obey the following syntax, where $A \in \mathbf{C}$ is an atomic concept and $U \in \mathbf{R}$ is an atomic role:

$$\begin{aligned} C \rightarrow A; \mid \exists R \quad R \rightarrow U \mid U^- \\ B \rightarrow C \mid \neg C \quad S \rightarrow R \mid \neg R \end{aligned}$$

TBox axioms are of the form $C \sqsubseteq B$ and $R \sqsubseteq S$ (inclusion axiom), or (*func* R) (functionality axiom).

An (ABox) *assertion* is a formula $A(c)$ or $U(c, d)$ where $A \in \mathbf{C}$, $U \in \mathbf{R}$, and $c, d \in \mathbf{I}$ (called *positive*) or its negation, i.e., $\neg A(c)$ resp. $\neg U(c, d)$ (*negative*).

Whenever the distinction between concepts and roles is immaterial, P denotes a possibly negated predicate from $\mathbf{C} \cup \mathbf{R}$. We denote by \overline{P} the opposite of P , i.e. $\overline{P} = \neg P$ and $\overline{\overline{P}} = P$.

- \mathcal{P} consists of rules r of the form

$$a_1 \vee \dots \vee a_n \leftarrow b_1, \dots, b_k, \text{not } b_{k+1}, \dots, \text{not } b_m, \quad (1)$$

where each a_i , $0 \leq i \leq n$, is an lp-atom and each b_i , $1 \leq i \leq m$, is either an lp-atom or a *DL-atom*; here

- an *lp-atom* is a first-order atom $p(\mathbf{t})$ with predicate p from a set \mathbf{P} of predicate names disjoint with \mathbf{C} and \mathbf{R} , and constants from the set $\mathcal{C} = \mathbf{I}$;
- a *DL-atom* $a(\mathbf{t})$ is of form $\text{DL}[\lambda; Q](\mathbf{t})$, where (a)

$$\lambda = S_1 op_1 p_1, \dots, S_m op_m p_m, \quad m \geq 0, \quad (2)$$

is such that, for $1 \leq i \leq m$, $S_i \in \mathbf{C} \cup \mathbf{R}$, $op_i \in \{\uplus, \updownarrow\}$ is an *update operator*, and $p_i \in \mathbf{P}$ is an *input predicate* of the same arity as S_i —intuitively, $op_i = \uplus$ (resp., $op_i = \updownarrow$) increases S_i (resp., $\neg S_i$) by the extension of p_i ; (b) $Q(\mathbf{t})$ is a *DL-query*, which is either of the form (i) $C(t)$, where C is a concept and t is a term; (ii) $R(t_1, t_2)$, where R is a role and t_1, t_2 are terms; or (iii) $\neg Q'(\mathbf{t})$ where $Q'(\mathbf{t})$ is from (i)-(ii). We skip (\mathbf{t}) for $\mathbf{t} = \epsilon^3$.

If $n = 0$, the rule is a constraint.

Example 2 (cont'd). In Example 1, the rule (9) contains a DL-atom $\text{DL}[\text{Male} \uplus \text{boy}; \text{Male}](Y)$, which first enriches the concept *Male* in \mathcal{O} by the extension of the predicate *boy* in \mathcal{P} via \uplus , and then queries the concept *Male* over the modified ontology.

Semantics. The semantics of a DL-program $\Pi = \langle \mathcal{O}, \mathcal{P} \rangle$ is defined in terms of its grounding $\text{grd}(\Pi) = \langle \mathcal{O}, \text{grd}(\mathcal{P}) \rangle$ over \mathcal{C} , i.e., $\text{grd}(\mathcal{P})$ contains all ground instances of rules r in \mathcal{P} over \mathcal{C} . In the remainder, by default we assume that Π is ground.

A (Herbrand) *interpretation* of Π is a set $I \subseteq \text{HB}_\Pi$ of ground atoms, where HB_Π is the usual Herbrand base w.r.t. \mathcal{C} and \mathbf{P} ; I satisfies an lp-atom a , if $a \in I$ and a DL-atom a of the form (2) if

$$\mathcal{O} \cup \lambda^I(a) \models Q(\mathbf{c}) \quad (3)$$

³ We disregard here for simplicity the less used constraints-operator and and subsumption queries.

where $\lambda^I(a) = \bigcup_{i=1}^m A_i(I)$, and

- $A_i(I) = \{S_i(\mathbf{t}) \mid p_i(\mathbf{t}) \in I\}$, for $op_i = \uplus$;
- $A_i(I) = \{\neg S_i(\mathbf{t}) \mid p_i(\mathbf{t}) \in I\}$, for $op_i = \updownarrow$.

Satisfaction of a DL-rule r resp. set \mathcal{P} of rules by I is then defined as usual, where I satisfies *not* b_j if I does not satisfy b_j ; I satisfies Π , if it satisfies each $r \in \mathcal{P}$. We denote that I satisfies (is a *model*) of an object o (atom, rule, etc.) by $I \models^{\mathcal{O}} o$.

Example 3 (cont'd). In Example 1, $I = \{\text{ischildof}(\text{john}, \text{alex}), \text{boy}(\text{john})\}$ satisfies $a = \text{DL}[\text{Child} \uplus \text{boy}; \neg \text{Male}](\text{alex})$, as $\mathcal{O} \cup \lambda^I(a) \models \neg \text{Male}(\text{alex})$, while $I \not\models^{\mathcal{O}} \text{DL}[\text{Adopted}](\text{john})$, as the input list is empty and $\mathcal{O} \not\models \text{Adopted}(\text{john})$.

An (*flp*-) *answer set* of $\Pi = \langle \mathcal{T} \cup \mathcal{O}, \mathcal{A} \rangle$ is any interpretation I that is a \subseteq -minimal model of the *flp*-reduct $f\Pi^I$, which maps any set \mathcal{P} of rules and $I \subseteq \text{HB}_\Pi$ to the rule set $f\mathcal{P}^I = \{fr^I \mid r \in \text{grd}(\mathcal{P})\}$, where $fr^I = r$ if the body of r is satisfied, i.e., $I \models^{\mathcal{O}} b_i$, for all b_i , $1 \leq i \leq k$ and $I \not\models^{\mathcal{O}} b_j$, for all $k < j \leq m$; otherwise, fr^I is void.

A DL-program Π is *inconsistent*, if it has no answer set. An interpretation I is an (*flp*-) *deletion repair answer set* of $\Pi = \langle \mathcal{T} \cup \mathcal{A}, \mathcal{P} \rangle$, if it is an *flp*-answer set of some $\Pi' = \langle \mathcal{T} \cup \mathcal{A}', \mathcal{P} \rangle$ where $\mathcal{A}' \subseteq \mathcal{A}$; any such \mathcal{A}' is called a *deletion repair* of Π .

Example 4 (cont'd). As mentioned, Π is inconsistent; if we drop (4) from \mathcal{A} , then $I = \{\text{ischildof}(\text{john}, \text{alex}), \text{boy}(\text{john}), \text{contact}(\text{john}, \text{pat})\}$ is an answer set. Along with the facts (7) and (8) the *flp*-reduct $f\mathcal{P}^I$ contains the ground rule (11), where X and Y are substituted by *john* and *pat* respectively. The interpretation $I_1 = \{\text{ischildof}(\text{john}, \text{alex}), \text{boy}(\text{john})\}$ is a deletion repair answer set with a deletion repair $\mathcal{A}'_1 = \{\text{Male}(\text{pat}), \text{Male}(\text{john})\}$.

In *DL-Lite_A* ontologies, inconsistency arises by few assertions.

Proposition 5 (cf. [3]). *In DL-Lite_A given a TBox \mathcal{T} every \subseteq -minimal ABox \mathcal{A} such that $\mathcal{T} \cup \mathcal{A}$ is inconsistent fulfills $|\mathcal{A}| \leq 2$.*

From this result it is not hard to establish that ≤ 3 distinct constants occur in such an \mathcal{A} .

3 SUPPORT SETS

In this section, we introduce support sets for DL-atoms. Intuitively, a *support set* for a DL-atom $d = \text{DL}[\lambda; Q](\mathbf{t})$ is a portion of its input that, together with ABox assertions, is sufficient to conclude that the query $Q(\mathbf{t})$ will evaluate to true, i.e., that given a subset $I' \subseteq I$ of an interpretation I and a set $\mathcal{A}' \subseteq \mathcal{A}$ of ABox assertions from the ontology, we can conclude that $I \models^{\mathcal{O}} Q(\mathbf{t})$. The evaluation of d w.r.t. I reduces then to test whether some support set $S = I' \cup \mathcal{A}'$ exists; to this end, a sufficient collection of such sets S can be precomputed and stored. Fortunately, for *DL-Lite_A* this is efficiently possible.

To simplify matters and avoid dealing with I' separately, it is convenient to introduce *input assertions* as follows.

Definition 6. Given a DL-atom $d = \text{DL}[\lambda; Q](\mathbf{t})$, we call $P_p(c)$ an *input assertion for d* , if $P \circ p \in \lambda$, $\circ \in \{\uplus, \updownarrow\}$ and $c \in \mathcal{C}$, where P_p is a fresh ontology predicate; \mathcal{A}_d is the set of all such assertions.

For a TBox \mathcal{T} and DL-atom d , let then $\mathcal{T}_d = \mathcal{T} \cup \{P_p \sqsubseteq P \mid P \uplus p \in \lambda\} \cup \{P_p \sqsubseteq \neg P \mid P \updownarrow p \in \lambda\}$, and for an interpretation I , let $\mathcal{O}_d^I = \mathcal{T}_d \cup \mathcal{A} \cup \{P_p(\mathbf{t}) \in \mathcal{A}_d \mid p(\mathbf{t}) \in I\}$. Then we have:

Proposition 7. *For every $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$, DL-atom $d = \text{DL}[\lambda; Q](\mathbf{t})$ and interpretation I , $I \models^{\mathcal{O}} d$ iff $I \models^{\mathcal{O}_d^I} \text{DL}[\epsilon; Q](\mathbf{t})$ iff $\mathcal{O}_d^I \models Q(\mathbf{t})$.*

⁴ Note that additionally considering program facts can be emulated.

Unlike (3), in \mathcal{O}_d^I there is a clear distinction between native assertions and input assertions of d w.r.t. I (via facts P_p and axioms $P_p \sqsubseteq (\neg)P$), mirroring the lp-input of d .

Now, in view of the property that in $DL\text{-Lite}_{\mathcal{A}}$ a single assertion is sufficient to derive a query [3] from a consistent ontology, we obtain support sets as follows.

Definition 8 (Ground Support Sets). Given a ground DL-atom $d = DL[\lambda; Q](\mathbf{t})$, a set S of assertions from $\mathcal{A} \cup \mathcal{A}_d$ is a *support set* for d w.r.t. an ontology $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$, if either

- (i) $S = \{P(\mathbf{c})\}$ and $\mathcal{T}_d \cup S \models Q(\mathbf{t})$, or
- (ii) $S = \{P(\mathbf{c}), P'(\mathbf{d})\}$ such that $\mathcal{T}_d \cup S$ is inconsistent, by $Supp_{\mathcal{O}}(d)$ we denote the set of all such S .

Support sets are linked to interpretations by the following notion.

Definition 9. A support set S of a DL-atom d is *coherent with an interpretation* I , if for each $P_p(\mathbf{c}) \in S$ it holds that $p(\mathbf{c}) \in I$.

Example 10 (cont'd). The set $\{hasParent(john, pat)\}$ is a support set for the DL-atom $DL[; hasParent](john, pat)$ w.r.t. \mathcal{O} , and so is $\{Male(pat)\}$ for the DL-atom $a = DL[Male \uplus boy; Male](pat)$. Moreover, $\{Male_{boy}(pat)\}$ is in $Supp_{\mathcal{O}}(a)$ but incoherent with minimal models of Π .

The evaluation of d w.r.t. I then reduces to the search for coherent support sets.

Proposition 11. Let d be a ground DL-atom, let $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$ be an ontology, and let I be an interpretation. Then, $I \models^{\mathcal{O}} d$ iff some $S \in Supp_{\mathcal{O}}(d)$ exists such that S is coherent with I .

Proof (Sketch). (\Rightarrow) Suppose $d = DL[\lambda; Q](\mathbf{t})$ evaluates w.r.t. \mathcal{O} and I to true, i.e., $\lambda^I(d) \cup \mathcal{O} \models Q(\mathbf{t})$. Towards a contradiction, assume no $S \in Supp_{\mathcal{O}}(d)$ is coherent with I . There are two cases:

(1) $\lambda^I(d) \cup \mathcal{O}$ is consistent. Proposition 5 implies that an assertion $\alpha \in \lambda^I(d) \cup \mathcal{A}$ must exist s.t. $\mathcal{T} \cup \{\alpha\} \models Q(\mathbf{t})$. If $\alpha \in \mathcal{A}$ then $Supp_{\mathcal{O}}(d)$ contains $\{\alpha\}$ by (i) of Definition 8, which trivially is coherent with I and thus contradicts the assumption. If $\alpha \in \lambda^I(d)$, then α is an input assertion of d . For $\alpha_d \in \mathcal{A}_d$, we then obtain that $\{\alpha_d\} \in Supp_{\mathcal{O}}(d)$ according to (i) of Definition 8, again a contradiction due to coherence with I .

(2) $\lambda^I(d) \cup \mathcal{O}$ is inconsistent. From Proposition 5 and consistency of \mathcal{O} , it follows that some $\delta \in \lambda^I(d)$ exists such that either (a) $\mathcal{T} \cup \{\delta\}$ is inconsistent, or (b) some $\gamma \in \mathcal{A} \cup \lambda^I(d)$ exists s.t. $\mathcal{T} \cup \{\delta, \gamma\}$ is inconsistent. In case a), we obtain $\{\delta_d\} \in Supp_{\mathcal{O}}(d)$, for the corresponding input assertion $\delta_d \in \mathcal{A}_d$. By (i) of Definition 8; this is a contradiction, as $\{\delta_d\}$ is coherent with I . In case b), we similarly conclude that either $\{\delta_d, \gamma\} \in Supp_{\mathcal{O}}(d)$ or $\{\delta_d, \gamma_d\} \in Supp_{\mathcal{O}}(d)$, depending on whether $\gamma \in \lambda^I(d)$, according to (ii) of Definition 8. Again this is a support set coherent with I , contradiction.

(\Leftarrow) Suppose some $S \in Supp_{\mathcal{O}}(d)$ is coherent with I . Assume towards a contradiction that $I \not\models^{\mathcal{O}} d$.

Again we consider two cases:

(1) $\mathcal{T}_d \cup S$ is consistent. Then, $\mathcal{T}_d \cup S \models Q(\mathbf{t})$ by item (i) of Definition 8. Since S is coherent with I , we conclude that $\mathcal{O}_d^I \models Q(\mathbf{t})$ which implies $I \models^{\mathcal{O}} d$ by Proposition 7. Contradiction.

(2) $\mathcal{T}_d \cup S$ is inconsistent. Then, due to coherence with I , so is \mathcal{O}_d^I , and trivially $\mathcal{O}_d^I \not\models Q(\mathbf{t})$; again we arrive at a contradiction by concluding that $I \models^{\mathcal{O}} d$ from Proposition 7. \square

As a simple consequence, we get:

Corollary 12. Given a ground DL-atom d and an ontology \mathcal{O} , there exists an interpretation I such that $I \models^{\mathcal{O}} d$ iff $Supp_{\mathcal{O}}(d) \neq \emptyset$.

3.1 Nonground Support Sets

Using support sets, we can completely eliminate the ontology access for the evaluation of DL-atoms. In a naive approach, one precomputes all support sets for all ground DL-atoms with respect to relevant ABoxes, and then uses them during the repair answer set computation. This does not scale in practice, since support sets may be computed that are incoherent with all candidate repair answer sets.

An alternative is to fully interleave the support set computation with the search for repair answer sets. Here we construct coherent ground support sets for each DL-atom and interpretation on the fly. As the input to a DL-atom may change in different interpretations, its support sets must be recomputed, however, since reuse may not be possible; effective optimizations are not immediate.

A better solution is to precompute support sets on a nonground level, that is, schematic support sets, prior to repair computation. Furthermore, in that we may leave the concrete ABox open; the support sets for a DL-atom instance are then easily obtained by syntactic matching. This leads to the following definition.

Definition 13 (Support Set for Nonground DL-atom). Let $d(\mathbf{X}) = DL[\lambda; Q](\mathbf{X})$ be a DL-atom and \mathcal{T} be a TBox. Let $V = \{X, Y, Z\}$ be distinct variables such that $\mathbf{X} \subseteq V$ and let $\mathcal{C} = \{a, b, c\}$.

A *nonground support set* for d w.r.t. \mathcal{T} is a set $S = \{P(\mathbf{Y})\}$ resp. $S = \{P(\mathbf{Y}), P'(\mathbf{Y}')\}$ such that (i) $\mathbf{Y}, \mathbf{Y}' \subseteq V$ and (ii) for each substitution $\theta : V \rightarrow \mathcal{C}$, the instance $S\theta = \{P(\mathbf{Y}\theta)\}$ (resp. $S\theta = \{P(\mathbf{Y}\theta), P'(\mathbf{Y}'\theta)\}$) is a support set for $d(\mathbf{X}\theta)$ w.r.t. $\mathcal{O}_{\mathcal{C}} = \mathcal{T} \cup \mathcal{A}_{\mathcal{C}}$, where $\mathcal{A}_{\mathcal{C}}$ is the set of all possible assertions over \mathcal{C} .

Here $\mathcal{A}_{\mathcal{C}}$ takes care of any possible ABox, by considering the maximal ABox (as $\mathcal{O} \subseteq \mathcal{O}'$ implies $Supp_{\mathcal{O}}(d) \subseteq Supp_{\mathcal{O}'}(d)$); three variables suffice as at most three different constants are involved.

Example 14 (cont'd). Certainly $\{hasParent(X, Y)\}$ is a nonground support set for the DL-atom $DL[; hasParent](X, Y)$, and so are $\{Male(X)\}$ and $\{Male_{boy}(X)\}$ for the DL-atom $a(X) = DL[Male \uplus boy; Male](X)$. Moreover, $\{Male_{boy}(X), Female(X)\}$ is a nonground support set for $a(X)$.

Nonground support sets S are sound, i.e. each instance $S\theta$ that matches with $\mathcal{A} \cup \mathcal{A}_d$ is a support set of the ground DL-atom $d\theta$ w.r.t. $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$; they are also complete, i.e., every support set S of a ground DL-atom d w.r.t. $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$ results as such an instance, and thus can be determined by syntactic matching.

Clearly, support sets as defined above may be subsumed by other support sets (e.g., $\{A(X), R(X, Y)\}$ by $\{A(X)\}$ and removed; for space reasons, we omit further discussion here. In the next section, we discuss how to compute a sufficient set of nonground support sets.

3.2 Determining Nonground Support Sets

Our technique for computing the nonground support sets is based on TBox classification, which is one of the main ontology reasoning tasks. This reasoning service computes complete information about the TBox constraints specified at the conceptual level.

More formally, given a TBox \mathcal{T} over a signature $\Sigma_{\mathcal{O}}$, the TBox classification $Clf(\mathcal{T})$ determines all subsumption relations $P \sqsubseteq (\neg)P'$ between concepts and roles P, P' in $\Sigma_{\mathcal{O}}$ that are entailed by \mathcal{T} . This can be exploited for our goal to compute nonground support sets, more precisely a complete family \mathbf{S} of such sets. Here *completeness* of \mathbf{S} for a (non-ground) DL-atom $d(\mathbf{X})$ w.r.t. \mathcal{O} means that for every

Algorithm 1: *SupRAnsSet*: all deletion repair answer sets

Input: $\Pi = \langle \mathcal{T} \cup \mathcal{A}, \mathcal{P} \rangle$
Output: $f_{lp}RAS(\Pi)$

(a) compute a complete set \mathbf{S} of nongr. supp. sets for the DL-atoms in Π
 (b) **for** $\hat{I} \in AS(\hat{\Pi})$ **do**
 $D_p \leftarrow \{a \mid e_a \in \hat{I}\}; D_n \in \{a \mid ne_a \in \hat{I}\}; \mathbf{S}_{gr}^{\hat{I}} \leftarrow Gr(\mathbf{S}, \hat{I}, \mathcal{A});$
 (c) **if** $\mathbf{S}_{gr}^{\hat{I}}(a) \neq \emptyset$ for $a \in D_p$ and every $S \in \mathbf{S}_{gr}^{\hat{I}}(a)$ for $a \in D_n$ fulfills
 $S \cap \mathcal{A} \neq \emptyset$ **then**
 (d) **for all** $a \in D_p$ **do**
 (e) **if** some $S \in \mathbf{S}_{gr}^{\hat{I}}(a)$ exists s.t. $S \cap \mathcal{A} = \emptyset$ **then** pick next a
 else remove each S from $\mathbf{S}_{gr}^{\hat{I}}(a)$ s.t.
 $S \cap \mathcal{A} \cap \bigcup_{a' \in D_n} \mathbf{S}_{gr}^{\hat{I}}(a') \neq \emptyset$
 (f) **if** $\mathbf{S}_{gr}^{\hat{I}}(a) = \emptyset$ **then** pick next \hat{I}
 end
 (g) $\mathcal{A}' \leftarrow \mathcal{A} \setminus \bigcup_{a' \in D_n} \mathbf{S}_{gr}^{\hat{I}}(a');$
 (h) **if** $f_{lp}FND(\hat{I}, \langle \mathcal{T} \cup \mathcal{A}', \mathcal{P} \rangle)$ **then** output $\hat{I}|_{\Pi}$
end

$\theta: \mathbf{X} \rightarrow \mathbf{I}$ and $S \in Supp_{\mathcal{O}}(d(\mathbf{X}\theta))$, some $S' \in \mathbf{S}$ exists such that $S = S'\theta'$, for some extension θ' of θ to V .

TBox classification is well studied in Description Logics [1]. For example, [9] discusses it for \mathcal{EL} w.r.t. concept hierarchy, and [10] studies it for the OWL 2 QL profile, which is based on *DL-Lite \mathcal{R}* .

Respective algorithms are thus suitable and also easily adapted for the computation of (a complete family of) nonground support sets for a DL-atom $d(\mathbf{X})$ w.r.t. \mathcal{O} . In principle, one can exploit Proposition 7 and resort to \mathcal{T}_a , i.e., compute the classification $Clf(\mathcal{T}_a)$, and determine nonground support sets of $d(\mathbf{X})$ proceeding similar as for computing minimal conflict sets [10]. To determine inconsistent support sets, perfect rewriting [3] can be done over $Pos(\mathcal{T})$, i.e., the TBox obtained from \mathcal{T} by substituting all negated concepts (roles) $\neg C$ ($\neg R, \neg \exists.R, \neg \exists.R^{-}$) with positive replacements \overline{C} ($\overline{R}, \overline{\exists.R}, \overline{\exists.R^{-}}$).

In practice (and as in our implementation), it nonetheless can be worthwhile to compute $Clf(\mathcal{T})$ first, as it is reusable for all DL-atoms. The additional axioms in \mathcal{T}_a , i.e., those of form $P_p \sqsubseteq (\neg)P$ (according to the update operators), are handled when determining the nonground support sets for a particular DL-atom from $Clf(\mathcal{T})$.

Let us now illustrate the computation of nonground support sets.

Example 15 (cont'd). Consider in our running example the DL-atom $a = DL[Male \uplus boy; Male](X)$. For computing a complete family \mathbf{S} of nonground support sets for a w.r.t. \mathcal{O} , we may refer to $\mathcal{T}_a = \mathcal{T} \cup \{Male_{boy} \sqsubseteq Male\}$ and its classification $Clf(\mathcal{T}_a)$. Hence, $S_1 = \{Male(X)\}$ and $S_2 = \{Male_{boy}(X)\}$ are the only unary nonground support sets of a . Further nonground support sets are obtained by computing minimal conflict sets, yielding $\{P(\mathbf{X}), \neg P(\mathbf{X})\}$ for each $P \in \mathbf{C} \cup \mathbf{R}$, as well as $\{Male_{boy}(X), \neg Male(X)\}$, $\{Male(X), Female(X)\}$, and $\{Male_{boy}(X), Female(X)\}$. However, since we are interested in completeness w.r.t. \mathcal{O} and \mathcal{O} is consistent, pairs not involving input assertions can be dropped (as they will not have a match in \mathcal{A}). The remaining two sets are supersets of S_2 , therefore $\mathbf{S} = \{S_1, S_2\}$ is a complete family of support sets for a w.r.t. \mathcal{O} .

4 REPAIR ANSWER SET COMPUTATION

We are now ready to describe our algorithm for computing deletion repair answer sets with the use of support sets. DL-programs are evaluated usually (as in DLVHEX) by means of a rewriting $\hat{\Pi}$ of $grd(\Pi)$, where DL-atoms a are replaced by ordinary atoms e_a (*replacement atoms*), together with a guess on their truth by additional ‘choice’

rules $e_a \vee ne_a$, where ne_a stands for negation of e_a . We denote interpretations of $\hat{\Pi}$ by \hat{I} , and use $\hat{I}|_{\Pi}$ to denote their restriction to the original language of Π .

A basic algorithm for computing repair answer sets was given in [7]. It checks whether candidates \hat{I} (i.e., answer sets of $\hat{\Pi}$) yield repair answer sets by solving a corresponding ontology repair problem (ORP). Intuitively, solutions to an ORP correspond to (potentially) modified ABoxes \mathcal{A}' (not necessarily obtained by deletion, i.e. subsets of \mathcal{A}) such that all DL-atoms evaluate as guessed in \hat{I} (formally \hat{I} is a compatible set of $\langle \mathcal{T} \cup \mathcal{A}', \mathcal{P} \rangle$, cf. [7]); they are computed through multiple calls to the external ontology via a query interface to evaluate DL-atoms under varying ABoxes. A final foundedness check (subset minimality of $\hat{I}|_{\Pi}$ w.r.t. $f\Pi^{\hat{I}|_{\Pi}} = \langle \mathcal{T} \cup \mathcal{A}', f\mathcal{P}^{\hat{I}|_{\Pi}} \rangle$) establishes repair answer sets.

Our new algorithm *SupRAnsSet* (see Algorithm 1), avoids instead multiple interface calls and merely needs to access the ontology once.

Given a (ground) DL-program Π for input, *SupRAnsSet* proceeds as follows. We start (a) by computing a complete family \mathbf{S} of nonground support sets for each DL-atom. Afterwards the replacement program $\hat{\Pi}$ is created and its answer sets are computed one by one. Once an answer set \hat{I} of $\hat{\Pi}$ is found (b), we first determine the sets of DL-atoms D_p (resp. D_n) that are guessed true (resp. false) in \hat{I} . Next, for all ground DL-atoms in $D_p \cup D_n$, the function $Gr(\mathbf{S}, \hat{I}, \mathcal{A})$ instantiates \mathbf{S} to relevant ground support sets, i.e., that are coherent with \hat{I} and match with $\mathcal{A} \cup \mathcal{A}_a$. We then check in (c) for atoms in D_p (resp. D_n) without support (resp. input only support). If either is the case, we skip to (b), the next model candidate, since no repair exists for the current one. Otherwise, in a loop (d) over atoms in D_p —except for those supported input only (e)—we remove support sets S that are conflicting w.r.t. D_n . Intuitively, this is the case if S hinges on an assertion $\alpha \in \mathcal{A}$ that also supports some atom $a' \in D_n$ (hence α needs to be deleted; note that due to consistency of \mathcal{A} , even inconsistent support of a' leaves no choice). If this operation leaves the atom from D_p under consideration without support (check at (f)), then no repair exists and the next model candidate is considered. Otherwise (exiting the loop at (g)), a potential deletion repair \mathcal{A}' is obtained from \mathcal{A} by removing assertions that occur in any support set for some atom $a' \in D_n$. An eventual check (h) for foundedness (minimality) w.r.t. \mathcal{A}' determines whether a deletion repair answer set has been found (output at (i)), or not.

Example 16 (cont'd). Suppose $a = DL[hasParent](john, pat)$, $b = DL[Male \uplus boy; Male](pat)$ and $\{e_a, ne_b\} \subseteq \hat{I}$. Then, $\mathbf{S}_{gr}^{\hat{I}}(a) = \{\{hasParent(john, pat)\}\}$ and we get to the else part of Step (e) where nothing is removed from $\mathbf{S}_{gr}^{\hat{I}}(a)$, since $\mathbf{S}_{gr}^{\hat{I}}(b) = \{\{Male(pat)\}\}$ and $\mathbf{S}_{gr}^{\hat{I}}(a) \cap \mathbf{S}_{gr}^{\hat{I}}(b) = \emptyset$. Hence, at Step (g) we must drop $Male(pat)$ from \mathcal{A} to make \hat{I} a deletion repair answer set.

As can be shown, algorithm *SupRAnsSet* correctly computes the deletion repair answer sets of the input DL-program. For the completeness part, i.e., that all deletion repair answer sets are indeed produced, the following proposition is crucial.

Proposition 17. *Given a DL-program Π , let \hat{I} be an answer set of $\hat{\Pi}$ such that $I = \hat{I}|_{\Pi}$ is an answer set of $\Pi = \langle \mathcal{T} \cup \mathcal{A}, \mathcal{P} \rangle$. If \hat{I} is a compatible set for $\Pi' = \langle \mathcal{T} \cup \mathcal{A}', \mathcal{P} \rangle$ where $\mathcal{A}' \supseteq \mathcal{A}$, then I is an answer set of $\Pi' = \langle \mathcal{T} \cup \mathcal{A}', \mathcal{P} \rangle$.*

Proof (Sketch). Assume that I is an answer set of $\Pi = \langle \mathcal{O}, \mathcal{P} \rangle$, where $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$ and that \hat{I} is a compatible set for $\Pi' = \langle \mathcal{O}', \mathcal{P} \rangle$ where $\mathcal{O}' = \mathcal{T} \cup \mathcal{A}'$ and $\mathcal{A}' \supset \mathcal{A}$. Towards contradiction, suppose I is not an answer set of Π . Hence, $I = \hat{I}|_{\Pi}$ is not a minimal model of $f\Pi^{\hat{I}} = \langle \mathcal{T} \cup \mathcal{A}', f\mathcal{P}^{\hat{I}} \rangle$. That is, some $I' \subset I$ exists such that $I' \models^{\mathcal{O}'} f\mathcal{P}^{\hat{I}}$. We then obtain that also $I' \models^{\mathcal{O}} f\mathcal{P}^{\hat{I}}$; this contradicts that I is an answer set of Π . Indeed, suppose that $I' \not\models^{\mathcal{O}} f\mathcal{P}^{\hat{I}}$. Then

some rule $r \in fP^I$ of form (1) is violated wrt. I' and \mathcal{O} , i.e., (i) $I' \models^{\mathcal{O}} b_i$ for each $1 \leq i \leq k$, (ii) $I' \not\models^{\mathcal{O}} b_j$ for each $k < j \leq m$, and (iii) $I' \not\models^{\mathcal{O}} a_h$ for each $1 \leq h \leq n$. By monotonicity of $I \models^{\mathcal{O}} a$ w.r.t. I and \mathcal{O} , we conclude $I' \models^{\mathcal{O}'} b_i$, $I' \not\models^{\mathcal{O}'} b_j$ (as \hat{I} is a compatible set for both $\hat{\Pi}$ and $\hat{\Pi}'$, and $I \not\models^{\mathcal{O}} b_j$, and $I' \not\models^{\mathcal{O}'} a_h$. But then $I' \not\models^{\mathcal{O}'} fP^I$, which is a contradiction. Hence, I' does not exist and I is an answer set of Π' . \square

Armed with this result, we establish the correctness result.

Theorem 18. *Algorithm SupRAnsSet is sound and complete w.r.t. deletion repair answer sets.*

Proof (Sketch). Soundness. Suppose *SupRAnsSet* outputs $I = \hat{I}|_{\Pi}$. We can get to (i) only if \hat{I} is an answer set of $\hat{\Pi}$; furthermore, by setting $\mathbf{S}_{gr}^{\hat{I}}$ to $Gr(\mathbf{S}, \hat{I}, \mathcal{A})$ in (b) and by the further modifications, it is ensured at (i) that each DL-atom $a \in D_p$ has some coherent support set that matches with \mathcal{A}' (i.e., $Gr(\mathbf{S}, \hat{I}, \mathcal{A}')(a) \neq \emptyset$), while no DL-atom $a' \in D_n$ has such a support set. Thus from Proposition 11, it follows that \hat{I} is a compatible set for $\Pi' = \langle T \cup \mathcal{A}', P \rangle$; hence $I \models \Pi'$. Furthermore, as $fPND(\hat{I}, T \cup \mathcal{A}', P)$ succeeds, I is a minimal model of $f\Pi'^I$. Hence I is an answer set of Π' , and thus a deletion repair answer set of Π .

Completeness. Suppose I is a deletion repair answer set. That is, for some $\mathcal{A}' \subseteq \mathcal{A}$, we have that I is an answer set of $\Pi' = \langle T \cup \mathcal{A}', P \rangle$. This implies that \hat{I} is an answer set of $\hat{\Pi}$ and thus will be considered in (b), with D_p and D_n reflecting the (correct) guess for $I \models^{\mathcal{O}'} a$ for each DL-atom a , where $\mathcal{O}' = T \cup \mathcal{A}'$. From Proposition 11 and completeness of \mathbf{S} , we obtain that each $a \in D_p$ has $Gr(\mathbf{S}, \hat{I}, \mathcal{A}')(a) \neq \emptyset$ and each $a \in D_n$ has $Gr(\mathbf{S}, \hat{I}, \mathcal{A}')(a) = \emptyset$. The initial $\mathbf{S}_{gr}^{\hat{I}}$ is such that $Gr(\mathbf{S}, \hat{I}, \mathcal{A}')(a) \subseteq \mathbf{S}_{gr}^{\hat{I}} = Gr(\mathbf{S}, \hat{I}, \mathcal{A})(a)$ holds for each DL-atom a ; in further steps, the algorithm removes all support sets $S \in Gr(\mathbf{S}, \hat{I}, \mathcal{A})(a)$ for $a \in D_p$ from $\mathbf{S}_{gr}^{\hat{I}}(a)$ such that $S \cap S' \cap \mathcal{A} \neq \emptyset$ for some support set $S' \in Gr(\mathbf{S}, \hat{I}, \mathcal{A})(a')$ and $a' \in D_n$, and removes all assertions in $S' \cap \mathcal{A}$ from \mathcal{A} . Importantly no removed S is in $Gr(\mathbf{S}, \hat{I}, \mathcal{A}')(a)$, since by the assertion that $T \cup \mathcal{A}$ is consistent, $|S' \cap \mathcal{A}| = 1$ must hold. Thus step (g) will be reached, and the variable \mathcal{A}' is assigned an ABox \mathcal{A}'' such that $\mathcal{A}' \subseteq \mathcal{A}'' \subseteq \mathcal{A}$. Since \hat{I} is a compatible set for $\Pi'' = \langle T \cup \mathcal{A}'', P \rangle$ and I is an answer set of Π' , by Proposition 17 I is also an answer set of Π'' , and thus I is a minimal model of $f\Pi''^I = \langle T \cup \mathcal{A}'', fP^I \rangle$. Hence, the test $fPND(\hat{I}, T \cup \mathcal{A}', P)$ in step (h) (where \mathcal{A}' has value \mathcal{A}'') succeeds, and $\hat{I}|_{\Pi}$, i.e., I is output. \square

5 IMPLEMENTATION AND EXPERIMENTS

We implemented the Algorithm *SupRAnsSet* within the DLVHEX evaluation framework,⁵ which allows us to effectively compute deletion repair answer sets for *DL-Lite_A*. In order to profit from existing DLVHEX data structures (e.g. for parsing) and optimization methods (such as nogood learning, etc.), we pursued a declarative ASP approach to realize *SupRAnsSet*. This applies to (a) computing complete nonground support families and to (b) searching candidate deletion repair answer sets and deletion repairs.

As for (a), we use a (stratified) logic program that intuitively mimics perfect rewriting to compute classifications relevant for support set computation on top. The logic program reifies concepts (roles, etc.), as well as positive replacements. Facts express subsumptions in $Pos(T)$, their contra-positives, and the duality of concepts (roles, etc.) and positive replacements of their negation. A simple rule transitively closes the subsumption relation; support sets are naturally expressed (and thus computed), using unary and binary predicates to represent

respective support for a DL-atom, the query and relevant concepts (roles) as from the input signature.

Concerning (b), non-ground rules (see below) are added to $\hat{\Pi}$ for the purpose of filtering candidate deletion repair answer sets as done by *SupRAnsSet*. Therefore, the language of $\hat{\Pi}$ is extended to include support set information. For any nonground support set S that covers a ground DL-atom a , the additional rules are of the form:

$\bar{S}_A \leftarrow r(S), ne_a; \quad sup_a \leftarrow r(S), e_a, not \bar{S}_A; \quad \leftarrow e_a, not sup_a;$
where $r(S)$ is a suitable representation of S , i.e., using predicates $p(\mathbf{X})$ for input assertions $P_p(\mathbf{X})$, resp. $p_P(\mathbf{X})$ ($np_P(\mathbf{X})$) for ABox assertions $P(\mathbf{X})$ ($\neg P(\mathbf{X})$). Furthermore, p_P , np_P , and sup_a are fresh predicates not in \mathcal{P} , and \bar{S}_A refers to $np_P(\mathbf{X})$ (resp. $p_P(\mathbf{X})$) if $S \cap \mathcal{A} \neq \emptyset$ and the assertion is positive (resp. negative), otherwise it is void. With further constraints $\leftarrow p_P(\mathbf{X}), np_P(\mathbf{X})$, the resulting program intuitively prunes candidates \hat{I} , resp. encodes deletion repair candidates, according to *SupRAnsSet*.

Experimental Setup. We have evaluated the approach by comparing it with ordinary answer set computation on two scenarios. They were run on a Linux server with two 12-core AMD 6176 SE CPUs/128GB RAM using DLVHEX 2.3.0; a timeout of 120 secs was set for each run.

Family Benchmark. The first benchmark is derived from our running example. We fixed two ABoxes with \mathcal{A}_{50} and \mathcal{A}_{1000} , of different size, where \mathcal{A}_{50} contains 50 children (7 adopted), 20 female and 32 male adults; and twenty times that many for \mathcal{A}_{1000} . Every child has at most two parents of different sex and the number of children per parent varies from 1 to 3. Rules (11) and (12), not involved in conflicts, have been dropped from \mathcal{P} . Instances are varied in terms of facts over \mathbf{I} included in \mathcal{P} by increasing the probability $p/100$ (p ranges from 5 to 35): for every child c , one additional fact $isChildOf(c, p)$ is added with probability $p/100$ for a random male adult non-parent; as well $boy(c)$ is included with probability $p/100$.

Network Benchmark. In the second scenario a network is described by a fixed ontology \mathcal{O} using a relation *edge*. Some nodes might be broken or blocked. There are overall 70 nodes, (among them 23 broken and rest available). There are 68 edges (among them 5 forbidden). The TBox encodes that if an edge is forbidden, then its endpoint must be blocked, and if a node is known to be broken, then it is automatically blocked, moreover blocked nodes are not available:

$\mathcal{O} = \{\exists.forbid \sqsubseteq Block, \quad Broken \sqsubseteq Block, \quad Block \sqsubseteq \neg Avail\}.$

We consider two DL-programs, P_{guess} and P_{conn} , over \mathcal{O} that are randomly generated as follows. P_{guess} contains for any node n the fact $node(n)$ with probability $p/100$, and the following rules:

- (1) $go(X, Y) \leftarrow open(X), open(Y), DL[; edge](X, Y).$
- (2) $route(X, Z) \leftarrow route(X, Y), route(Y, Z).$
- (3) $route(X, Y) \leftarrow go(X, Y),$
 $\quad not DL[Block \uplus block; forbid](X, Y).$
- (4) $open(X) \vee block(Y) \leftarrow node(X), not DL[; \neg Avail](X).$
- (5) $negIs(X) \leftarrow node(X), route(X, Y), X \neq Y.$
- (6) $\perp \leftarrow node(X), not negIs(X).$

Intuitively, (1), (2) and (3) recursively determine routes over non-blocked (*open*) nodes; where (3) expresses that by default a route is recommended unless it is known to be forbidden. Rule (4) amounts to guessing for each node not known to be unavailable, whether it is blocked or not, i.e. it contains nondeterminism, which makes rule processing challenging. Rules (5) and (6) encode the requirement that no input node is isolated w.r.t. the resulting routes.

The program P_{conn} contains the same random node facts as P_{guess} and additional facts $in(n)$ and $out(n)$, either for each node n with equal probability (i.e., in with $p = 1/2$, out , otherwise). It contains the rules (1) and (2) above, and variants of (3), (4) and (6):

⁵ <http://www.kr.tuwien.ac.at/research/systems/dlvhex>

- (3') $route(X, Y) \leftarrow go(X, Y), not DL[; forbid](X, Y)$.
(4') $open(X) \leftarrow node(X), not DL[; \neg Avail](X)$.
(6') $\perp \leftarrow in(X), out(Y), not route(X, Y)$.

Intuitively, rather than guessing, (4') expresses that each node is open by default unless known to be unavailable, and by (6') the program checks whether each *in*-node is connected to all *out*-nodes (thus ensuring they are available, i.e. not blocked).

p	\mathcal{A}_{50}		\mathcal{A}_{1000}	
	AS	rep	AS	rep
5 (25)	0.12 (0)	0.19 (0)	63.93 (0)	36.70 (0)
10 (25)	0.12 (0)	0.19 (0)	63.77 (0)	37.59 (0)
15 (25)	0.12 (0)	0.20 (0)	63.98 (0)	38.29 (0)
20 (25)	0.12 (0)	0.20 (0)	63.90 (0)	39.17 (0)
25 (25)	0.12 (0)	0.20 (0)	63.82 (0)	39.97 (0)
30 (25)	0.12 (0)	0.21 (0)	63.97 (0)	40.77 (0)
35 (25)	0.12 (0)	0.21 (0)	63.18 (0)	41.41 (0)

Table 1. Family benchmark results

p	P_{conn}		P_{guess}	
	AS	rep	AS	rep
10 (25)	0.19 (0)	0.13 (0)	0.48 (0)	0.16 (0)
20 (25)	0.19 (0)	0.44 (0)	0.48 (0)	0.37 (0)
30 (25)	0.19 (0)	1.28 (0)	0.51 (0)	0.86 (0)
40 (25)	0.19 (0)	2.36 (0)	0.53 (0)	1.31 (0)
50 (25)	0.19 (0)	5.60 (0)	0.57 (0)	3.00 (0)
60 (25)	0.19 (0)	8.83 (0)	0.60 (0)	5.03 (0)
70 (25)	0.19 (0)	15.92 (0)	0.64 (0)	7.27 (0)
80 (25)	0.20 (0)	25.89 (0)	0.69 (0)	11.45 (0)
90 (25)	0.20 (0)	37.33 (0)	0.75 (0)	16.95 (0)
100 (25)	0.20 (0)	55.43 (0)	0.87 (0)	16.38 (0)

Table 2. Network benchmark results

Results. The results that we obtained for these settings are given in Table 1 and Table 2. Instances are grouped by the probability (the value p) used for generating them. For each p , 25 instances were generated and their average running time (in seconds) to compute a first answer set (using DLVHEX) and deletion repair answer set is reported. The naive approach [7] has not been implemented (but expectedly would time out for even the smallest problem instances).

Despite a small overhead for computation of deletion repair answer sets compared to ordinary ones the results for the family problem with \mathcal{A}_{50} scale well. For \mathcal{A}_{1000} the first repair answer set is found quicker than the inconsistency is identified during ordinary answer set computation. Liberal safety [5] exploited in the experiments improves greatly the running time both for repair and ordinary answer set computation modes.

In the Network domain, as expected, running times grow exponentially with the instance size for P_{guess} due to the guessing rule (4). Employing a default in P_{conn} rather than guessing scales linearly, however. For P_{conn} we also considered different splits (rather than $p = 1/2$) between *in* and *out*, although without significant change.

These results indicate the effectiveness of our approach for practical settings.

6 CONCLUSION

Support sets reduce the evaluation of DL-atoms over $DL-Lite_{\mathcal{A}}$ ontologies to constraint matching, providing a base for effective deletion repair answer set computation under *flp*-semantics. Our results, and in particular algorithm *SupRAnsSet*, easily extend to other semantics, e.g., weak repair answer sets [7], and to other notions of repairs. Furthermore, support sets are used in a companion work to compute answer sets of HEX-programs [6]. As external atoms lack an explicit data part (ABox), support sets ought to be more abstract, which makes direct efficient usage less clear; repair is an open issue.

Note that algorithm *SupRAnsSet* constructs in its search all maximal deletion repairs, i.e., \subseteq -maximal repairs $\mathcal{A}' \subseteq \mathcal{A}$ that admit an answer set (however, it also may construct non-maximal repairs). In

this regard it is similar to work on ABox cleaning [11, 13]. There, an inconsistent ontology (with consistent TBox) is repaired by identifying and eliminating minimal conflict sets causing, i.e., explaining, the inconsistency, thus resulting in maximal deletion repairs. However, our setting and work differ fundamentally: (i) the ontology is consistent and inconsistency arises only through the interface of DL-atoms, and (ii) several DL-atom queries have to be considered (entailment or non-entailment) under different potential ABox updates.

More remotely related to our work are approaches for explaining positive and negative answers to conjunctive queries in *DL-Lite* [4, 2], as they apply the perfect reformulation algorithm [3] and then extract explanations in a nontrivial way.

In ongoing work, we consider restricted repairs [7] and an extension of our approach to other DLs like \mathcal{EL} . While no small complete support families exist for \mathcal{EL} in general, this might still apply for practically relevant fragments. Furthermore incomplete support families can be used for optimization (caching) to reduce access to an \mathcal{EL} reasoner.

REFERENCES

- [1] Franz Baader, Diego Calvanese, Deborah L. McGuinness and Daniele Nardi, and Peter F. Patel-Schneider, *The Description Logic Handbook: Theory, Implementation, and Applications*, Cambridge University Press, New York, NY, 2nd edn., 2007.
- [2] Alexander Borgida, Diego Calvanese, and Mariano Rodriguez-Muro, 'Explanation in *DL-Lite*', in *Description Logics*, eds., F. Baader, C. Lutz, and B. Motik, *CEUR Workshop Proc. vol. 353*. CEUR-WS.org, (2008).
- [3] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati, 'Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family', *J. Autom. Reasoning*, **39**(3), 385–429, (2007).
- [4] Diego Calvanese, Magdalena Ortiz, Mantas Simkus, and Giorgio Stefanoni, 'The complexity of explaining negative query answers in *DL-Lite*', in *Proc. 13th International Conf. Principles of Knowledge Representation and Reasoning (KR 2012)*, eds., G. Brewka, T. Eiter, and S. McIlraith, pp. 583–587. AAAI Press, (2012).
- [5] Thomas Eiter, Michael Fink, Thomas Krennwallner, and Christoph Redl, 'Liberal safety for answer set programs with external sources', in *Proc. 27th Conf. Artificial Intelligence (AAAI '13)*, eds., M. desJardins and M. L. Littman, pp. 267–275. AAAI Press, (2013).
- [6] Thomas Eiter, Michael Fink, Christoph Redl, and Daria Stepanova, 'Exploiting support sets for answer set programs with external evaluations', in *Proc. 28th Conf. Artificial Intelligence (AAAI '14)*, eds., C. Brodley and P. Stone. AAAI Press, (2014). To appear.
- [7] Thomas Eiter, Michael Fink, and Daria Stepanova, 'Data repair of inconsistent dl-programs', in *Proc. 23rd International Joint Conference on Artificial Intelligence (IJCAI-13)*, ed., F. Rossi, pp. 869–876. AAAI Press/IJCAI, (2013).
- [8] Thomas Eiter, Giovambattista Ianni, Thomas Lukasiewicz, Roman Schindlauer, and Hans Tompits, 'Combining answer set programming with description logics for the semantic web', *Artificial Intelligence*, **172**(12-13), 1495–1539, (2008).
- [9] Yevgeny Kazakov, Markus Krötzsch, and Frantisek Simancik, 'The incredible ELK - from polynomial procedures to efficient reasoning with \mathcal{EL} ontologies', *J. Autom. Reasoning*, **53**(1), 1–61, (2014).
- [10] Domenico Lembo, Valerio Santarelli, and Domenico Fabio Savo, 'A graph-based approach for classifying OWL 2 QL ontologies', in *Description Logics*, eds., T. Eiter, B. Glimm, Y. Kazakov, and M. Krötzsch, *CEUR Workshop Proc. vol. 1014*, pp. 747–759. CEUR-WS.org, (2013).
- [11] Giulia Masotti, Riccardo Rosati, and Marco Ruzzi, 'Practical abox cleaning in *DL-Lite* (progress report)', in *Description Logics*, eds., R. Rosati, S. Rudolph, and M. Zakharyashev, *CEUR Workshop Proc. vol. 745*. CEUR-WS.org, (2011).
- [12] Boris Motik and Riccardo Rosati, 'Reconciling Description Logics and Rules', *Journal of the ACM*, **57**(5), 1–62, (June 2010).
- [13] Riccardo Rosati, Marco Ruzzi, Mirko Graziosi, and Giulia Masotti, 'Evaluation of techniques for inconsistency handling in OWL 2 QL ontologies', in *International Semantic Web Conference (2)*, eds., P. Cudré-Mauroux et al., *Lecture Notes in Computer Science vol. 7650*, pp. 337–349. Springer, (2012).