

Explanation in *DL-Lite*

Alexander Borgida¹, Diego Calvanese², Mariano Rodriguez-Muro²

¹ Dept. of Computer Science
Rutgers University, USA

² Faculty of Computer Science
Free University of Bozen-Bolzano, Italy

Abstract. The paper addresses the problem of explaining some reasoning tasks associated with the *DL-Lite* Description Logic. Because of the simplicity of the language, standard concept level reasoning is quite easy, and the only contribution is an alternate, more accessible syntax, plus a novel focus on *brevity of proofs*. Of greater interest is the explanation of reasoning in finite models, which is motivated by the use of *DL-Lite* for database access. The fame of *DL-Lite* rests on its ability to answer efficiently conjunctive queries over KBs, and the paper makes three contributions in this regard: (1) a method for explaining why a value b was returned by a query; (2) a method for finding minimal explanations for why a conjunctive query is *unsatisfiable*; (3) the beginnings of a theory for explaining why a value b was **not** returned by a query.

1 Introduction

It is by now widely accepted that end-users of information systems which do more than simple fact retrieval require some sort of facility for *explaining* the reason for their answers. For example, in the area of deductive databases there has been work on explaining answers returned by Datalog-query processors [1]. In Description Logics (DLs), starting from [2], there have been papers studying the explanation of deductions such as concept subsumption [3] and knowledge base inconsistency [4]. More generally, the work on the Inference Web [5] has produced a substrate on which general explanation facilities for reasoners can be built.

It is worth recalling that in much prior research on explanation, as here, the focus has been on finding inference rules and presentations whose *form is acceptable to naive end-users*, and on finding explanation proofs from the traces of an optimized implementation; simply translating/normalizing into FOL variants and presenting resolution proofs are not acceptable.

The present research considers the problem of explaining reasoning and query answering for *DL-Lite*, a DL that was introduced to capture most of the features of conceptual modeling languages such as ER and UML, while at the same time maintaining efficient query processing (low data complexity), as in standard databases (DBs) [6]. As for any DL, there are standard judgments such as concept subsumption, TBox and ABox consistency, which require relatively straightforward explanations. Because of its use for DB conceptual modeling, and the fact that DBs are almost always considered to represent *finite structures* in which the conceptual models are interpreted, one important novel feature of the above reasoning tasks is the possibility of requiring finite models; since *DL-Lite* lacks the finite-model property, this enables additional inferences.

A second distinguishing feature of *DL-Lite* is the emphasis on *conjunctive query (CQ) processing*. Although DB query answering can be viewed formally as truth evaluation in a finite (portion of a) model, this is not always acceptable as an explanation, e.g., when the model has millions of tuples. Therefore the focus here will be on *minimizing* explanations. In addition, *DL-Lite* also uses inference in query answering, by modifying the query, so we also need to explain these inferences in a transparent manner.

After providing formal background on *DL-Lite* and explanations in Section 2, the remainder of the paper considers the problem of explaining various kinds of answers in *DL-Lite*: TBox and ABox reasoning, including the finite model case (Section 3); why a value is returned as one of the answers to a CQ posed to a *DL-Lite* ABox (Section 4); why a particular value was *not* returned in a CQ answer (Section 5). Although the last problem is only formalized, not resolved, a special-purpose solution is given for the case when this occurs because the query is itself unsatisfiable.

2 Background

For reasons of space, we deal only with the *DL-Lite_F* variant of the *DL-Lite* family [6]. However, the present work can be extended with little effort also to the other variants of *DL-Lite* with role hierarchies and attributes, so we continue to use the term *DL-Lite*.

Starting from atomic concepts A and atomic roles P , in *DL-Lite* one can construct so called *basic concepts* B of the form A , $\exists P$, and $\exists P^-$. In addition, we make use of the negation $\neg B$ of a basic concept B . The semantics is the standard one, i.e., given an interpretation $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$, a concept C denotes a subset $C^{\mathcal{I}}$ of the interpretation domain $\Delta^{\mathcal{I}}$, and a role P denotes a binary relation $P^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, satisfying $(\exists P)^{\mathcal{I}} = \{x \mid \exists y.(x, y) \in P^{\mathcal{I}}\}$, $(\exists P^-)^{\mathcal{I}} = \{y \mid \exists x.(x, y) \in P^{\mathcal{I}}\}$, and $(\neg B)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus B^{\mathcal{I}}$.

A knowledge base (KB) $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ consists of a *TBox* \mathcal{T} of *terminological axioms* and an *ABox* \mathcal{A} of *facts* about individuals. Axioms in \mathcal{T} have the form $B_1 \sqsubseteq B_2$, $B_1 \sqsubseteq \neg B_2$, and (funct P) and assert, respectively, subsumptions and disjointness between two basic concepts B_1 and B_2 , and the functionality of role P . An interpretation \mathcal{I} satisfies $B \sqsubseteq C$ if $B^{\mathcal{I}} \subseteq C^{\mathcal{I}}$, and it satisfies (funct P) if $\forall x, y, z.(x, y) \in P^{\mathcal{I}} \wedge (x, z) \in P^{\mathcal{I}} \rightarrow y = z$. Each individual d in \mathcal{A} is interpreted as a domain element $d^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ under the unique name assumption. Assertions in \mathcal{A} have the form $A(d)$ and $P(d_1, d_2)$, and they are satisfied in an interpretation \mathcal{I} if $d^{\mathcal{I}} \in A^{\mathcal{I}}$ and $(d_1^{\mathcal{I}}, d_2^{\mathcal{I}}) \in P^{\mathcal{I}}$, respectively.

An interpretation that satisfies all assertions in \mathcal{K} is called a *model* of \mathcal{K} . The reasoning tasks of KB satisfiability, subsumption, classification, and concept consistency are defined in the standard way, by referring to the set of models of a KB. *Finite model reasoning* restricts the attention only to models whose domain is finite.

We are interested also in answering conjunctive queries (CQs) over *DL-Lite* KBs. A CQ over a *DL-Lite* KB \mathcal{K} has the form $Q(\mathbf{x}) \leftarrow \text{conj}(\mathbf{x}, \mathbf{y})$, where $\text{conj}(\mathbf{x}, \mathbf{y})$ is a conjunction of atoms whose predicates are atomic concepts and roles of \mathcal{K} , and whose variables are among \mathbf{x} and \mathbf{y} . $Q(\mathbf{x})$ is interpreted in \mathcal{I} as the set $Q(\mathbf{x})^{\mathcal{I}}$ of tuples \mathbf{d} of elements of $\Delta^{\mathcal{I}}$ such that, when we assign \mathbf{d} to \mathbf{x} , the formula $\exists \mathbf{y}.\text{conj}(\mathbf{x}, \mathbf{y})$ evaluates to true in \mathcal{I} . The set of *certain answers* to $Q(\mathbf{x})$ over \mathcal{K} is the set of tuples \mathbf{c} of ABox individuals such that $\mathbf{c}^{\mathcal{I}} \in Q(\mathbf{x})^{\mathcal{I}}$, for every model \mathcal{I} of \mathcal{K} . $Q(\mathbf{x})$ is *consistent* wrt \mathcal{K} , if there is a model \mathcal{I} of \mathcal{K} such that $Q(\mathbf{x})^{\mathcal{I}} \neq \emptyset$. $Q_1(\mathbf{x})$ is *contained in* $Q_2(\mathbf{x})$ wrt \mathcal{K} if $Q_1(\mathbf{x})^{\mathcal{I}} \subseteq Q_2(\mathbf{x})^{\mathcal{I}}$, for every model \mathcal{I} of \mathcal{K} . When the ABox \mathcal{A} is treated as a “closed DB” (i.e., a single interpretation), $Q(\mathbf{x})^{\mathcal{A}}$ denotes the evaluation of $Q(\mathbf{x})$ in \mathcal{A} .

As far as explanations are concerned, it is almost universally accepted that they are formal *proofs*, constructed from premises using rules of inference. For better explanations, proofs should be as short as possible, and inference rules as self-evident as possible. In addition to the proof, there is also the matter of its *presentation*. There is a decided preference (e.g., [2, 5]) for tree-shaped proofs, produced by rules of inference with a single conclusion, and zero or more antecedents. Such proofs support, through follow-up questions, the interactive and gradual unfolding of relevant parts under user control. Although [5] suggests a specific XML-based syntax for inference rule schemas, we will use the more concise notation used in Programming Languages, and first applied to DLs in [7], but which can easily be translated to the rules in [5].

To illustrate the notation, we use the inference rule **Isa-trans** in Figure 1, expressing the transitivity of the \sqsubseteq relationship. Here, the rule schema is named **Isa-trans**, and it expresses that, if $B_1 \sqsubseteq B_2$ and $B_2 \sqsubseteq B_3$ can be deduced from \mathcal{T} , then so can $B_1 \sqsubseteq B_3$; the side-condition of the rule requires B_1 , B_2 , and B_3 to be concept expressions.

3 Explanations of Standard Inferences

3.1 Modified Syntax of *DL-Lite*

A number of notions in *DL-Lite* (and their notation), such as existential constraints, inverses of roles, and complements of concepts are mathematically too sophisticated for users familiar only with notations like UML diagrams. Hence, we propose to eliminate the notations $\exists P$ and $\exists P^-$, and replace them by the more familiar notions of “the current domain” and “the current range of role P ”, written as $\text{domain}(P)$ and $\text{range}(P)$ ¹. The only remaining use of role inverses is in axioms of the form (funct P^-), as in (funct *makes*⁻). Normally, we would suggest replacing *makes*⁻ by *madeBy*, and $\text{domain}(\textit{makes})$ (resp., $\text{range}(\textit{makes})$) by $\text{range}(\textit{madeBy})$ (resp., $\text{domain}(\textit{madeBy})$). This leaves the case of having both (funct P) and (funct P^-) in the TBox (e.g., (funct *wife*), (funct *wife*⁻)). For lack of space, we leave the treatment of this case to an extended version of the paper. As for eliminating \neg , we also replace each axiom of the form $B_1 \sqsubseteq \neg B_2$, expressing disjointness of two basic concepts, by the axiom (disjoint $B_1 B_2$), having the same semantics.

As a result of the above simplifications, subsumption axioms will now only relate atomic concepts and/or current domains/ranges of roles. And in addition to subsumption, we have axioms for disjointness of concepts, functionality of roles, and declarations of names for role inverses.

3.2 TBox Reasoning

Subsumption reasoning in *DL-Lite* is a particularly simple form of structural subsumption, in part because there are no nested concepts. Therefore one does not need any of the complications suggested in [2], such as atomic concepts etc., and for normalized concepts, the standard \sqsubseteq -inference rules for givens, reflexivity and transitivity suffice. Moreover, a *DL-Lite* concept can only be unsatisfiable either due to subsumption by disjoint concepts, or because it is subsumed by the current domain (resp., range) of a role whose current range (resp., domain) is unsatisfiable. Hence, we get (see Figure 1)

¹ The use of the word “*current*” is meant to emphasize the distinction from OWL “domain”, which describe the *potential* set of objects to which a property may apply.

$\frac{\mathcal{T} \vdash B_1 \sqsubseteq B_2}{\mathcal{T} \vdash B_2 \sqsubseteq B_3}$	$\frac{\mathcal{T} \vdash B_1 \sqsubseteq B_2}{\mathcal{T} \vdash B_1 \sqsubseteq B_3}$	B_1, B_2, B_3 concepts	Given $\frac{\mathcal{T}, \alpha \vdash \alpha}{\mathcal{T}, \alpha \vdash \alpha}$ α any axiom
IsA-refl $\frac{\mathcal{T} \vdash B \sqsubseteq B}{\mathcal{T} \vdash B \sqsubseteq B}$	B a concept		
Inc-disj $\frac{\mathcal{T} \vdash B \sqsubseteq B_1 \quad \mathcal{T} \vdash B \sqsubseteq B_2}{\mathcal{T}, (\text{disjoint } B_1 B_2) \vdash B \sqsubseteq \perp}$	C, B_1, B_2 concepts		
Inc-dom $\frac{\mathcal{T} \vdash \text{range}(R) \sqsubseteq \perp}{\mathcal{T} \vdash \text{domain}(R) \sqsubseteq \perp}$	R a role		
Inc-rng $\frac{\mathcal{T} \vdash \text{domain}(R) \sqsubseteq \perp}{\mathcal{T} \vdash \text{range}(R) \sqsubseteq \perp}$	R a role		
Func-comp $\frac{\mathcal{T}, (\text{func } R_1), \dots, (\text{func } R_k) \vdash (\text{func } R_1 \circ \dots \circ R_k)}{\mathcal{T}, (\text{func } R_1), \dots, (\text{func } R_k) \vdash (\text{func } R_1 \circ \dots \circ R_k)}$			
Rng-comp $\frac{\mathcal{T} \vdash \text{domain}(R_k) \sqsubseteq \text{range}(R_1 \circ \dots \circ R_{k-1})}{\mathcal{T} \vdash B_2 \sqsubseteq \text{range}(R_k)}$	$k \geq 2$		
Dom-comp $\frac{\mathcal{T} \vdash \text{domain}(R_1) \sqsubseteq B_1}{\mathcal{T} \vdash \text{domain}(R_1 \circ \dots \circ R_k) \sqsubseteq B_1}$	B_1, B_2 concepts;		
Inherit $\frac{\mathcal{T} \vdash B_1 \sqsubseteq B_2 \quad \langle \mathcal{T}, \mathcal{A} \rangle \vdash \text{Holds}(B_1(e))}{\langle \mathcal{T}, \mathcal{A} \rangle \vdash \text{Holds}(B_2(e))}$	e an individual		

Fig. 1. Inference rules

Proposition 1 *For DL-Lite, the following is a sound and complete set of rules of inference: {Given, IsA-refl, IsA-trans, Nothing, Inc-disj, Inc-dom, Inc-rng}.*

This means that all reasoning about \sqsubseteq reduces to simple classification of atomic concepts and expressions denoting the current domains/ranges of roles, i.e., computing the so-called Hasse diagram induced by the \sqsubseteq axioms in the TBox, plus detecting inconsistency. Once this is done, explaining $B_1 \sqsubseteq B_2$ for satisfiable concepts involves only finding the shortest path between them.

To explain unsatisfiability of a concept B , we have to find two \sqsubseteq -paths from B to two concepts, say B_1, B_2 , asserted to be disjoint, whose total length is minimal². To find this, one can use a breadth-first-search strategy, where rule **Isa-trans** is applied in parallel waves, and each concept inferred to subsume B is decorated with a superscript that is one higher than that of the antecedent concept, thereby capturing the wave when it was added. (B is assigned superscript 0.) Then, if a disjointness is detected the first time between ${}^n B_1$ and ${}^j B_2$ (with $n \geq j$, say), then the length of the explanation for this is $n + j$. Unfortunately, this may not be the shortest explanation, because the total length of such an explanation could be as high as $2n$; on the other hand, there might be another disjointness between, say, ${}^1 B_3$ and ${}^{n+1} B_4$ added in the $(n + 1)$ st pass say, which would only have total explanation length $n + 2$. To find the shortest explanation one must therefore continue beyond pass n up to pass $2n - 1$.

3.3 ABox Reasoning

In *DL-Lite*, one infers new facts about existing individuals by applying inclusion axioms, and recognizing that $R(a, b)$ entails $\text{domain}(R)(a)$ and $\text{range}(R)(b)$.³

To detect inconsistencies in a KB, one simply looks for objects belonging to concepts which can be deduced to be subsumed by \perp , or for objects violating functionality constraints. The one nontrivial aspect is that, as usual, we prefer shorter explanations. In the case of inconsistent ABoxes, one wants *the shortest derivation of a conflict* from the

² We would have to consider also paths from B to some $\text{domain}(R)$ (resp., $\text{range}(R)$) for which $\text{range}(R)$ (resp., $\text{domain}(R)$) is unsatisfiable, but we simplify here for lack of space.

³ ABox reasoning can also be formalized with inference rules, using a judgment $\text{Holds}(\beta)$ for facts β , as in rule **Inherit**, but lack of space prevents us from giving these axioms.

original ABox – one with fewest rule applications. To find this, one can use a strategy similar to the one described above for finding the shortest proof of inconsistency, using rule **Inherit** instead of rule **Isa-trans**.

We note that while the above looks for evidence of knowledge base inconsistency, this is not the same problem as diagnosing errors in the knowledge base. Pinpointing [4], and related orthogonal techniques are much more likely to be useful for this task.

3.4 Reasoning in Finite Models

Example 1. Consider the following TBox

$$\begin{array}{ll} (\text{func } tutors) & Student \sqsubseteq \text{range}(tutors) \\ \text{domain}(tutors) \sqsubseteq TA & TA \sqsubseteq Student \end{array}$$

Since, *tutors* is a function, there can be at most as many values in its range as in its domain. Since the current range of *tutors* contains all *Students*, there can be at most as many *Students* as values in the domain of *tutors*. And since $\text{domain}(tutors)$ is contained in the set of *TAs*, there can be at most as many *Students* as *TAs*. If, in addition, now one has that $TA \sqsubseteq Student$, this implies that there are at most as many *TAs* as *Students*, and therefore the number of *TAs* and *Students* is the same. In an infinite model, this leads to no new conclusions, even if one recalls that *TA* is a subset of *Student*. However, in a finite model, these two facts imply that the extensions of *TA* and *Student* must be identical, which means that a new subsumption has been inferred: $Student \sqsubseteq TA$.

Clearly, the above pattern can be generalized by replacing *tutors* with the composition of an arbitrary set of roles $R_1 \circ R_2 \circ \dots \circ R_k$, obtaining rule **Same-cardinality**:

$$\frac{\begin{array}{ll} \mathcal{T} \vdash (\text{func } R_1 \circ \dots \circ R_k) & \mathcal{T} \vdash B_2 \sqsubseteq \text{range}(R_1 \circ \dots \circ R_k) \\ \mathcal{T} \vdash \text{domain}(R_1 \circ \dots \circ R_k) \sqsubseteq B_1 & \mathcal{T} \vdash B_1 \sqsubseteq B_2 \end{array}}{\mathcal{T} \vdash B_2 \sqsubseteq B_1} \quad \begin{array}{l} B_1, B_2 \text{ concepts;} \\ R_1 \text{ a role} \end{array}$$

The remaining question is how one can deduce properties of a composition of roles, given only *DL-Lite* axioms. First, rule **Func-comp** captures that if all roles are functions, then their composition will be a function. And since the current domain of a composition is contained in the current domain of the first role, we also have rule **Dom-comp**. However, $B_2 \sqsubseteq \text{range}(R_1 \circ \dots \circ R_k)$ does not follow from $B_2 \sqsubseteq \text{range}(R_k)$ alone, because the *current* range of the composition may be smaller, if not all values in $\text{domain}(R_k)$ are reached by $R_1 \circ \dots \circ R_{k-1}$. So one also needs the entire current domain of R_k to be contained in the current range of $R_1 \circ \dots \circ R_{k-1}$, leading to the rule **Rng-comp**. It follows from results in [8] that these are *all* the possible additional subsumption inferences needed for the finite model case.

Proposition 2 *The only additional inference rules needed for deducing subsumptions in the case of DL-Lite TBoxes with finite domains are Same-cardinality, Func-comp, Dom-comp, Rng-comp.*

As far as explanations are concerned, this is a prime example where the user will need separate explanations for the rules of inference themselves.

4 Explaining positive answers for queries over an ABox

Consider first the simple issue of answering CQs over a regular DB. To explain why $Q(b)$ is true in a DB requires showing why the DB, treated as an interpretation, makes the body of the query evaluate to true. For CQs, the only interesting aspect is exhibiting the values used for existentially quantified variables. For example, if MIMI is an answer to query $Q_0(x) \leftarrow Student(x), supervisedBy(x, y), teaches(y, z)$, one would locate some “witness” values ANNA and ENG101 for variables y and z , and then explain that $Student(MIMI)$, $supervisedBy(MIMI, ANNA)$, and $teaches(MIMI, ENG101)$ are given in the DB.

In the case of *DL-Lite*, the difficulty is that the ABox is not a closed DB, but instead must be “completed” according to the axioms. For example, if we have $Student(MIMI)$ and $Student \sqsubseteq Person$, then we must also add $Person(MIMI)$; and if we have $Professor(TOM)$ and $Professor \sqsubseteq \text{domain}(teaches)$, then one can conclude that there is some hypothetical individual, say $@c$, representing what TOM teaches, and that $teaches(TOM, @c)$ holds. Such hypothetical individuals may also get additional properties. Unfortunately, the result can be an infinite DB since the axioms may contain cyclic dependencies; e.g., $Nat \sqsubseteq \text{domain}(succ)$, $\text{range}(succ) \sqsubseteq Nat$.

From theoretical results [6], we do however know that from the TBox \mathcal{T} and the ABox \mathcal{A} one can derive a (generally infinite) canonical model $can(\mathcal{T}, \mathcal{A})$, representing all possible DBs extending \mathcal{A} , as far as CQ evaluation is concerned. Essentially, $can(\mathcal{T}, \mathcal{A})$ extends the facts in the ABox as illustrated in the example above. The crucial property of $can(\mathcal{T}, \mathcal{A})$ is that each CQ Q can be answered by evaluating it, as in regular DBs, over only a finite “small” portion of $can(\mathcal{T}, \mathcal{A})$ (whose size depends on Q).

We exploit this fact to explain why $Q(b)$ holds, by essentially constructing the *finite part* of $can(\mathcal{T}, \mathcal{A})$ that is needed to explain the truth of the query body for $Q(b)$. In order to generate the relevant part of $can(\mathcal{T}, \mathcal{A})$, we resort to a (suitably adapted version of) the algorithm for query answering in *DL-Lite*. Query answering in *DL-Lite* [6] is performed by first rewriting the original query $Q(x)$ into a set $\mathcal{S} = \{Q_0(x), \dots, Q_k(x)\}$ of alternate queries, then evaluating these over the original ABox (treated as a closed DB), and finally returning the union of the results.

Example 2. Consider the following TBox \mathcal{T}

- | | |
|---|--|
| (1) $PhD \sqsubseteq Student$ | (3) $\text{range}(supervisedBy) \sqsubseteq Professor$ |
| (2) $PhD \sqsubseteq \text{domain}(supervisedBy)$ | (4) $Professor \sqsubseteq \text{domain}(teaches)$ |

and the query $Q_0(x) \leftarrow Student(x), supervisedBy(x, y), teaches(y, z)$. The *DL-Lite* rewriting algorithm would then rewrite Q_0 into the following set of queries:

- $$\begin{aligned}
 Q_0(x) &\leftarrow Student(x), supervisedBy(x, y), teaches(y, z) \\
 Q_1(x) &\leftarrow PhD(x), supervisedBy(x, y), teaches(y, z) \\
 Q_2(x) &\leftarrow PhD(x), supervisedBy(x, y), Professor(y) \\
 Q_3(x) &\leftarrow PhD(x), supervisedBy(x, y), supervisedBy(w, y) \\
 Q_4(x) &\leftarrow PhD(x), supervisedBy(x, y) \\
 Q_5(x) &\leftarrow PhD(x), PhD(x)
 \end{aligned}$$

We recall briefly, using the above query as an example, the basic steps of the rewriting algorithm that are necessary to understand its use in our explanation setting; for the full

details we refer to [6]. Essentially, the algorithm makes use of replacement and unification steps. A replacement can be applied to an atom when the corresponding predicate appears on the right hand side of an inclusion axiom; e.g., query $Q_1(x)$ is obtained from $Q_0(x)$ by replacing $Student(x)$ with $PhD(x)$, due to axiom (1). Similarly, $Q_2(x)$ is obtained from $Q_1(x)$ by making use of axiom (4).⁴ On the other hand, a unification step collapses two atoms with the same predicate when the corresponding variables can be unified; e.g., query $Q_4(x)$ is obtained from $Q_3(x)$ by unifying the two atoms $supervisedBy(x, y)$ and $supervisedBy(w, y)$.

The key observation is that the replacement rewriting steps correspond to the “inverses” of the additions made to the canonical model of the knowledge base. Hence, we can use them to guide the explanation of why a given individual is in the certain answers to the original query. Suppose that an individual b is part of the certain answers to $Q(x)$, and we want to explain why⁵. In our example, suppose the ABox contains $PhD(\text{BOB})$, and no other facts about BOB, and we want to explain why BOB is in the answer to $Q_0(x)$. To do so, we proceed as follows.

Step 1. We compute the set \mathcal{S} of rewritings of $Q(x)$, building a data structure that tracks how (changed) atoms in one query are derived from a predecessor query.

In our example, queries $Q_0(x)$ to $Q_5(x)$ will be part of the set \mathcal{S} of rewritings, though if there were other axioms, there may be many more, irrelevant to BOB.

Step 2. Since $Q(b)$ was true, we select from \mathcal{S} a rewriting $Q_k(x)$ that produces b as an answer when directly evaluated over the ABox (viewed as a closed DB)⁶. This means that there is an assignment θ of ABox individuals to the variables in $Q_k(x)$ such that $\theta(x) = b$, and for each atom β in $Q_k(x)$, $\theta(\beta)$ is an ABox fact.

In our example, we would select $Q_5(x)$, with $\theta(x) = \text{BOB}$, where $PhD(\text{BOB})$ in the ABox would satisfy the two identical atoms of $Q_5(x)$.

Step 3. Traversing backwards the sequence of rewritings from $Q_k(x)$ to $Q(x)$, we extend the substitution θ to the variables in only the intervening queries, by keeping track of unifications, or assigning to such variables newly introduced *skolem constants* (corresponding to objects introduced by the inclusion axioms). More precisely, when going backwards from Q_i to Q_j (the query from which Q_i was generated), if no variable has been eliminated when rewriting Q_j to Q_i , then θ need not be extended. When a variable z has been eliminated because it has been unified with a variable y , then we set $\theta(z) = \theta(y)$. While when a variable z has been eliminated in Q_i by replacing an atom $A(y)$ for $P(y, z)$ (resp., $P(z, y)$), due to inclusion axiom $A \sqsubseteq \text{domain}(P)$ (resp., $A \sqsubseteq \text{range}(P)$), then we set $\theta(z) = @c$, where $@c$ is a fresh skolem constant. In this way, when one reaches the original query $Q(x)$, θ will have assigned to each variable appearing in it either an ABox individual or a skolem constant. In analogy to the case of standard DBs, one then initially shows to the user $\theta(Q(x))$, i.e., the set of atoms of the original query to which θ has been applied.

⁴ Technical note: a replacement where a variable (z , in the example) is removed from the resulting query is allowed only when such a variable does not occur anywhere else in the query.

⁵ Everything below can be generalized for the case when the query returns a tuple of values.

⁶ By completeness of the query answering algorithm based on query rewriting [6], such a rewriting $Q_k(x)$ will always exist.

In our example, we would get $\theta(x) = \text{BOB}$ for $Q_5(x)$, and the following new assignments: $\theta(y) = @1$ for $Q_4(x)$, $\theta(w) = \theta(x) = \text{BOB}$ for $Q_3(x)$, and $\theta(z) = @2$ for $Q_1(x)$. The resulting initial explanation shown to the user would be the sequence of ground atoms matching the query conjuncts:

$$Q_0(\text{BOB}) \leftarrow \text{Student}(\text{BOB}), \text{supervisedBy}(\text{BOB}, @1), \text{teaches}(@1, @2)$$

Step 4. At this point, the user is allowed to ask follow-up explanations for any of the atoms $f = \theta(\beta)$ appearing above. The data structure set up in Step 1 allows one to find either that f is given in the ABox, or the first replacement (not unification!) which replaced β by β' say. The explanation for f is then $\theta(\beta')$ together with the inclusion axiom supporting it. In our example, if the user asks for the explanation of $\text{Student}(\text{BOB})$, the answer is “*PhD*(BOB), because *PhDs* are *Students*.” obtained from the Q_0 to Q_1 rewriting, based on axiom (1). And the explanation of $\text{teaches}(@1, @2)$, is “*Professor*(BOB), because *Professors* are in the domain of *teaches*” obtained from the Q_1 to Q_2 rewriting, based on axiom (4).

A critical point in the above explanation procedure is the selection of the rewriting $Q_k(x)$ in Step 2. First, we need an efficient way of selecting only the indexes k of those rewritings that actually contribute to $Q(b)$. This can be done by associating to each of the queries $Q_i(x)$ in \mathcal{S} a distinct tag, and returning such tags as part of the answer. If $\mathcal{S} = \{Q_1(x), \dots, Q_n(x)\}$, then the modified rewriting is written in SQL as:

```
SELECT x, 1 AS tag FROM ... WHERE Q1(x) UNION
SELECT x, 2 AS tag FROM ... WHERE Q2(x) UNION ...
```

If more than one query in \mathcal{S} returns b as an answer, it might be best to explain the one whose total number of atoms *plus* number of applied transformations is minimal. (This assumes that users find explanations of ground atom look-ups as easy to understand as axiom applications. If not, one can weight the latter more heavily.)

5 Failed Answers to Conjunctive Queries — First Steps

In this section we begin to consider explaining why an individual b was *not* returned by the query $Q(b)$; such explanations are needed because, presumably, the user was expecting $Q(b)$ to be satisfied in the current KB, but in fact it was not.

There are actually two alternative reasons for this: the query itself is unsatisfiable (so it can never return any answers), or just that b was not one of the values returned.

5.1 Conjunctive query unsatisfiability in *DL-Lite*

The reason to treat this case separately is because it is hard to see why one would want to pose a query that would never be able to return an answer.

In the case of ordinary CQs over regular DBs, every query is satisfiable, so this problem does not arise. However, in *DL-Lite*, this no longer holds because of the presence of integrity constraints in the TBox. For example, a *DL-Lite* query would be unsatisfiable if it had a sub-query $B_1(x), B_2(x)$ where B_1 and B_2 are disjoint concepts.

In general, one way to determine whether a query $Q(x)$ is unsatisfiable wrt a TBox \mathcal{T} is to check if it is contained wrt \mathcal{T} in the query $Q_\perp(x) \leftarrow \perp(x)$. Following [9], we can show that, in general, containment of a query $Q_1(x)$ in a query $Q_2(x)$ wrt \mathcal{T} can

be checked by adapting the standard technique based on reducing query containment to query evaluation, i.e.: (i) The atoms in the body of $Q_1(x)$ are viewed as a collection of facts in an ABox \mathcal{A}_{Q_1} , except that the query variables (e.g., x, y) are replaced by constant (e.g., c_x, c_y) that do *not* have the unique name assumption. (ii) Functionality is propagated in \mathcal{A}_{Q_1} , i.e., if $P(a, b)$ and $P(a, b')$ are encountered, $(\text{funct } P) \in \mathcal{T}$, and either of b or b' (say b) comes from a variable name, then b is unified with b' . Let \mathcal{A}'_{Q_1} be the resulting ABox. (iii) $Q_2(x)$ is evaluated over $\langle \mathcal{T}, \mathcal{A}'_{Q_1} \rangle$. Then, $Q_1(x)$ is contained in $Q_2(x)$ wrt \mathcal{T} iff c_x is returned as answer in step (iii). Now, when $Q_2(x)$ is $\perp(x)$, the only way it can return c_x is when \mathcal{A}'_{Q_1} is inconsistent. The explanation of such inconsistencies was discussed in Section 3, though it needs to be enriched with the explanation of the propagation of functionality.

5.2 Explaining failed answers in ordinary databases

Considering the case of CQs over a regular DB (a KB with empty TBox) we are interested in the reason why some value b was *not* returned as part of the answer to query $Q(x)$, e.g.,

$$Q_1(x) \leftarrow \text{Italian}(x), \text{friendOf}(x, y), \text{Woman}(y), \text{drives}(y, z), \text{Ferrari}(z)$$

This corresponds to explaining why the body of the query $Q(b)$ evaluates to false in the interpretation consisting of the DB, e.g., why $\neg \exists y, z. \text{Italian}(b) \wedge \text{friendOf}(b, y) \wedge \dots \wedge \text{Ferrari}(z)$ evaluates to false. Unfortunately, attempting a direct, value-level explanation of this would likely be impractical because y and z are actually universally quantified, and one would need to exhibit all combinations of values for them in the DB. Note however that one kind of failed answer that can be explained trivially is the non-membership of a (set of) individuals in a table: there is nothing to do (except maybe let the user see the values that are in the table).

To the best of our knowledge, the problem of explaining negated queries was addressed only in [10], which considers Datalog queries with negation. Unfortunately, the solution proposed there takes a rule such as $p(x) \leftarrow t(x, y), s(y, z)$, and explains “Why not $p(2)$?” with the statement “because $t(2, 4)$ but not $s(4, Z)$ ”. But this deals with only *one* possible value of y (namely 4), while one actually needs to explain why *none* of the possible values work. Therefore, we devote the rest of this section to uncovering intuitions about appropriate explanations, formalizing these, and then sketching some special cases in which the intuitions can be realized.

So consider query $Q_1(x)$ above. Suppose that TED was not returned as one of the answers, i.e., $Q_1(\text{TED})^A$ was false. First, note that the simplest explanation for this would be the fact that TED is not Italian *or* has no friends at all, i.e., TED is not in the current domain of *friendOf*. (A simple way to handle the second case would be to expand the query so that every time it contains an atom like *friendOf*(x, y), an additional conjunct of the form $\text{domain}(\text{friendOf})(x)$ is added.) Second, an explanation like “TED had no female friends who drive a Ferrari” would be considered inappropriate by natural-language speakers if, in fact, it was the case that TED had no friends at all, or even if TED had no female friends, because the first sentence *presupposes* the existence of objects of the other kind. Third, even if the explanation “TED was Italian but had no friends” broke no presuppositions, it would be considered inferior to one which

simply stated that “TED had no friends”, since the second one is more concise. Fourth, after an explanation such as “TED had no women friends who drive” (with no failed presuppositions) one should allow explanations for several follow-up questions, especially “Who are all the women friends of TED?”, with possible follow up questions of why any are present in this set; or, for any object h in the domain of *drives*, “Why isn’t h one of TED’s women friends?”. We generalize, suggesting the following definition:

Definition 1. *The explanation of why $Q(b)^A$ failed for a satisfiable CQ $Q(x)$ is either a single unsatisfied atom β in $Q(b)$ (a trivial base case to explain) or starts with finding a sub-query $subq_Q(b, y)$ of $Q(b)$ plus an additional atom $atom_Q(y)$ in $Q(b)$ such that: (i) $S := subq_Q(b, y)^A$ is non-empty (to avoid failed presuppositions). (ii) None of the elements d in S satisfy $atom_Q$, i.e., are in the table for $atom_Q$ (to ensure that $subq_Q(b, y) \wedge atom_Q(y)$ is also a failing sub-query of $Q(b)$). (iii) $subq_Q(b, y)$ is minimal, in the sense that no atoms can be removed from it while preserving (i) and (ii).*

Since for any $Q(b)$, there may be more than one such pair of formulas ($subq_Q, atom_Q$), there are additional possibilities for minimizing explanations, including: (i) choose the $subq_Q$ with smallest size, assuming that users rarely want to consider the exact values in S ; (ii) minimize the size of set S , so there are fewer individuals to be displayed when follow-up questions arise. Let us focus here on the first criterion only.

Given a CQ Q , one can first construct a directed graph G_Q that has as vertices the variables and constants of Q , and edges corresponding to the atoms in Q . Unfortunately, in general, it is reasonable to expect that minimizing the size of $subq_Q$ will be a combinatorially difficult problem (in the length of the query) because it involves, in principle, considering all subgraphs of G_Q . (A formal proof of the intractability of the problem is deferred to another paper.)

One special case where the problem is tractable occurs for what we call “path queries”: ones whose role-induced edges form a single chain in the graph, starting from the constant b (whose absence is to be explained), and where every node has a small, bounded number of associated concepts⁷. Query $Q_1(x)$ is an example of such a path query. In this case, one can start from the node for b , and consider successively longer paths, performing for each the relational join of the tables named by the edges. The first path $\pi \circ T$ such that π does not return the empty set, but $\pi \bowtie T$ does, provides a minimal size $subq_Q$, equal to π .

This approach can be generalized to “tree queries” by performing a breadth-first search through the paths of the query graph that start from the root. And in fact one can try to find spanning trees whose join returns the empty set even for arbitrary graphs (though finding the minimal one is likely to remain a hard problem). The case where no sub-query $subq_Q$ is a tree remains an even more open problem because we have not yet found a strategy for presenting $subq_Q$, and follow-up queries, to the user.

5.3 Generalizing to failed answers in *DL-Lite* ABoxes

Even assuming some solution to the problem of explaining non-answers to a CQ in a DB, generalizing this to the case of *DL-Lite* is a significant task. In particular, if, as usual in *DL-Lite*, the original query $Q(x)$ is expanded to the set of queries $\mathcal{S} =$

⁷ This number is bounded in any case by the cardinality of the TBox alphabet.

$\{Q_1(x), Q_2(x), \dots\}$, then in order to explain why b is not in the certain answers to $Q(x)$, it is not sufficient to explain why each $Q_i(b)$ fails over the ABox, for $i = 1, 2, \dots$; one must also be prepared to explain why there are no additional queries in the set \mathcal{S} — a highly non-trivial task akin to the proof in [6].

6 Conclusions

We have tackled the problem of explaining *DL-Lite* reasoning, which we view as (i) requiring inference rules for building proofs, and (ii) finding short proofs. For standard concept level reasoning, essentially we rely on an alternate more accessible syntax. Of greater novelty and complexity is the explanation of reasoning in finite models. Since *DL-Lite* is intended to support efficient CQ evaluation over a DL KB, we address for the first time explanation for this, providing three contributions: (i) a theoretically sound technique for explaining why a value was returned by a query; (ii) explanation of why a CQ is *unsatisfiable*, which rests on an algorithm for finding minimal length proofs of the unsatisfiability of an ABox; (iii) explanation of why a value was **not** returned by a CQ, a problem that, surprisingly, has not been addressed even for CQs over ordinary DBs. For lack of space we do not provide inference rules for all aspects, which would make clear how to implement our explanations following [5].

Acknowledgments This research has been partially supported by the EU FP6 FET project TONES (Thinking ONtologiES) under contract number FP6-7603, and by the U.S. DHS under ONR grant N00014-7-1-0150.

References

1. Shmueli, O., Tsur, S.: Logical diagnosis of LDL programs. In: Proc. of ICLP'90. (1990)
2. McGuinness, D.L., Borgida, A.: Explaining subsumption in description logics. In: Proc. of IJCAI'95. (1995) 816–821
3. Deng, X., Haarslev, V., Shiri, N.: A framework for explaining reasoning in description logics. In: Working Notes of the AAAI Fall Symposium on “Explanation-aware Computing”. (2005)
4. Schlobach, S.: Debugging and semantic clarification by pinpointing. In: Proc. of ESWC 2005. (2005) 226–240
5. McGuinness, D.L., Pinheiro da Silva, P.: Explaining answers from the Semantic Web: the Inference Web approach. *J. of Web Semantics* **1**(4) (2004) 397–413
6. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *J. of Automated Reasoning* **39**(3) (2007) 385–429
7. Borgida, A.: From type systems to knowledge representation: Natural semantics specifications for description logics. *J. of Intelligent and Cooperative Information Systems* **1**(1) (1992)
8. Rosati, R.: Finite model reasoning in *DL-Lite*. In: Proc. of ESWC 2008. (2008)
9. Rosati, R.: On the decidability and finite controllability of query processing in databases with incomplete information. In: Proc. of PODS 2006. (2006) 356–365
10. Specht, G.: Generating explanation trees even for negations in deductive database systems. In: Proc. of the 5th Workshop on Logic Programming Environments (LPE'93). (1993) 8–13