

# Explanations for Negative Query Answers under Existential Rules

İsmail İlkan Ceylan<sup>1</sup>, Thomas Lukasiewicz<sup>1</sup>, Enrico Malizia<sup>2</sup>, Cristian Molinaro<sup>3</sup>,  
Andrius Vaicenavičius<sup>1</sup>

<sup>1</sup>University of Oxford, Department of Computer Science, UK

<sup>2</sup>King's College London, Department of Informatics, UK

<sup>3</sup>University of Calabria, DIMES Department, Italy

{ismail.ceylan, thomas.lukasiewicz, andrius.vaicenavicius}@cs.ox.ac.uk,  
enrico.malizia@kcl.ac.uk, cmolinaro@dimes.unical.it

## Abstract

Ontology-mediated query answering is an extensively studied paradigm, where the conceptual knowledge provided by an ontology is leveraged towards more enhanced querying of data sources. A major advantage of ontological reasoning is its interpretability, which allows one to derive explanations for query answers. Indeed, explanations have a long history in knowledge representation, and have also been investigated for ontology languages based on description logics and existential rules. Existing works on existential rules, however, merely focus on understanding why a query is entailed, i.e., explaining positive query answers. In this paper, we continue this line of research and address another important problem, namely, explaining why a query is *not* entailed under existential rules, i.e., explaining negative query answers. We consider various problems related to explaining non-entailments from the abduction literature, and also introduce new problems. For all considered problems, we give a detailed complexity analysis for a wide range of existential rule languages and complexity measures.

## 1 Introduction

Ontology-based query answering aims to enhance querying of data sources with an ontology that encodes domain knowledge. The idea is to view the ontology and the user query as a composite query, called *ontology-mediated query (OMQ)*, and the task of evaluating such queries is then called *ontology-mediated query answering (OMQA)* (Bienvenu et al. 2014). OMQA is an important paradigm in knowledge representation with many application areas. Description logics (DLs) (Baader et al. 2007) and existential rules (Cali, Gottlob, and Kifer 2013; Cali, Gottlob, and Lukasiewicz 2012) are two families of languages, which are commonly used to formulate ontologies.

With the increasing demand for more explainable and interpretable systems, providing explanations for OMQA has recently seen a surge in interest. The most basic problem is that of explaining why a query is entailed, i.e., explaining *positive* query answers. This problem has been studied for ontology languages based on DLs (Borgida, Calvanese, and Rodriguez-Muro 2008; Ceylan et al. 2020) and existential rules (Ceylan et al. 2019). The main idea is to view an explanation as a set of facts from the database, which together with the ontology are sufficient to entail the query.

The complementary problem of explaining why a query is *not* entailed, i.e., explaining *negative* query answers, has only been studied in the context of DLs (Calvanese et al. 2013), where the problem is modeled as an abduction task. Intuitively, a (minimal) explanation for a non-entailment can be seen as a (minimal) set of facts such that, if the database were extended with this set, the query would be entailed, while the knowledge base would remain consistent with the added information. This is a natural choice for explaining non-entailments of monotonic queries, as a non-entailment can be attributed to missing facts.

In this paper, we continue this line of research, and address the problem of explaining negative answers in OMQA based on existential rules (rather than DLs) as underlying ontology languages. While DLs are popular formalisms for modeling ontologies, it is generally agreed that existential rules are well-suited for data-intensive applications, since they allow us to conveniently deal with higher-arity relations, which naturally occur in standard relational databases.

We consider the following computational problems of explaining negative query answers in OMQA with existential rules: (i) deciding whether a given set of facts is a minimal explanation, (ii) deciding whether there exists a minimal explanation at all, (iii) deciding whether a given fact is relevant or necessary, and (iv) deciding whether a given set of facts contains exactly all relevant or exactly all necessary facts. While the problems in (i) to (iii) have been considered before for DLs (Calvanese et al. 2013), they have not been considered yet for existential rules, and the problems in (iv) are considered for the first time in this paper. We give an intuitive example to illustrate the above concepts.

**Example 1.** We encode the functionality of a mechanical system that may experience certain issues as an OMQA problem. The system consists of four parts ( $p_1$ ,  $p_2$ ,  $p_3$ , and  $p_4$ ), and each may cause some issues (among *overheat*, *leakage*, *high\_pressure*, *non\_start*, *burnt\_smell*, *flow\_blockage*, and *pipes\_rattling*) when faulty. For example, (1) if  $p_1$  is faulty, then the issues *overheat* and *leakage* may occur, (2) if  $p_2$  is faulty, then the issues *high\_pressure* and *non\_start* may occur, and (3) if  $p_3$  is faulty, then the issues *leakage* and *burnt\_smell* may occur. This functionality of the system is encoded via facts in a database  $D$  and ontological axioms (i.e., existential rules) in a program  $\Sigma$ .

When issues occur in the system, we would like to under-

stand which parts may be causing the issues, so that we can restrict the search for the causes to a limited number of parts. This is a problem of *fault diagnosis*, which can be modeled via explanations of negative query answers as follows.

In our example, the issues observed in the mechanical system (e.g., *leakage* and *non\_start*) can be encoded as queries. Such queries are not entailed by the database  $D$  alone. However, by adding to  $D$  facts about the potential faultiness of the four parts (forming the set of *hypotheses*  $H$ ), the query can be derived, and hence an explanation built from  $H$  can be interpreted as a possible set of causes for the observed issues.

For example, under the above specification, the observed issues *leakage* and *non\_start* have the *explanations* of (a)  $p_1$  and  $p_2$  being faulty or (b)  $p_2$  and  $p_3$  being faulty. These are also *minimal* explanations, and they are the only minimal ones. Hence,  $p_1$ ,  $p_2$ , and  $p_3$  are all *relevant* facts (which belong to some minimal explanation), while  $p_2$  is the only *necessary* fact (which is in every minimal explanation).

The precise definitions and details for this example will be given later in the paper.  $\triangleleft$

In this paper, we provide a precise picture of the complexity of the following computational tasks around explanations for negative answers in OMQA for a wide range of existential rule languages and under different complexity measures:

- Deciding the existence of a minimal explanation, and deciding whether a given set of facts is an explanation.
- Deciding whether a given fact is relevant (i.e., belongs to at least one minimal explanation), and deciding whether a given set of facts contains exactly all relevant facts.
- Deciding whether a given fact is necessary (i.e., belongs to every minimal explanation), and deciding whether a given set of facts contains exactly all necessary facts.

## 2 Preliminaries

In this section, we recall some basics on existential rules (Cali, Gottlob, and Pieris 2012; Cali, Gottlob, and Kifer 2013; Cali, Gottlob, and Lukasiewicz 2012) and the paradigm of ontology-mediated query answering. We also give some complexity-theoretic background relevant to our study.

### 2.1 General

We consider a relational vocabulary consisting of mutually disjoint, possibly infinite sets  $\mathbf{R}$ ,  $\mathbf{C}$ ,  $\mathbf{N}$ , and  $\mathbf{V}$  of *predicates*, *constants*, *nulls*, and *variables*, respectively. Each predicate is associated with an arity, i.e., a non-negative integer. A *term* is a constant, a null, or a variable. An *atom* is an expression of the form  $P(t_1, \dots, t_n)$ , where  $P$  is an  $n$ -ary predicate, and  $t_1, \dots, t_n$  are terms. A *ground atom* (or *fact*) has only constants as terms. Conjunctions of atoms are often identified with the sets of their atoms.

An *instance*  $I$  is a (possibly infinite) set of atoms containing constants and nulls only. A *database*  $D$  is a finite instance that contains only constants. A *homomorphism* is a mapping  $h: \mathbf{C} \cup \mathbf{N} \cup \mathbf{V} \rightarrow \mathbf{C} \cup \mathbf{N} \cup \mathbf{V}$  that is the identity on  $\mathbf{C}$  and maps  $\mathbf{N}$  to  $\mathbf{C} \cup \mathbf{N}$ . With a slight abuse of notation, homomorphisms are applied also to (sets of) atoms.

A *conjunctive query* (CQ)  $Q(\mathbf{X})$  is a first-order formula of the form  $\exists \mathbf{Y} \phi(\mathbf{X}, \mathbf{Y})$ , where  $\phi(\mathbf{X}, \mathbf{Y})$  is a conjunction of atoms without nulls. The *answer* to  $Q(\mathbf{X})$  over an instance  $I$ , denoted  $Q(I)$ , is the set of all tuples  $\mathbf{t} \in \mathbf{C}^{|\mathbf{X}|}$  for which there is a homomorphism  $h$  such that  $h(\phi(\mathbf{X}, \mathbf{Y})) \subseteq I$  and  $h(\mathbf{X}) = \mathbf{t}$ . A *union of conjunctive queries* (UCQ)  $Q(\mathbf{X})$  has the form  $Q_1(\mathbf{X}) \vee \dots \vee Q_n(\mathbf{X})$ , where each  $Q_i(\mathbf{X})$  is a CQ. The evaluation of  $Q(\mathbf{X})$  over an instance  $I$ , denoted  $Q(I)$ , is defined as the set of tuples  $\bigcup_{1 \leq i \leq n} Q_i(I)$ . A *Boolean UCQ*  $Q$  is a UCQ where all variables are existentially quantified. We say that  $Q$  is *true* over  $I$ , denoted  $I \models Q$ , if  $Q(I) \neq \emptyset$ . In the rest of the paper, we only focus on Boolean UCQs.

### 2.2 Existential Rules

A *tuple-generating dependency* (TGD)  $\sigma$  is a first-order formula of the form  $\forall \mathbf{X} \forall \mathbf{Y} \Phi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \Psi(\mathbf{X}, \mathbf{Z})$ , where  $\Phi(\mathbf{X}, \mathbf{Y})$  is a conjunction of atoms, called the *body* of the TGD and denoted  $body(\sigma)$ , and  $\Psi(\mathbf{X}, \mathbf{Z})$  is a conjunction of atoms, called the *head* of the TGD and denoted  $head(\sigma)$ , all without nulls. Classes of TGDs are also known as *existential rules*, or *Datalog<sup>±</sup> languages* in the literature. An instance  $I$  satisfies  $\sigma$ , written  $I \models \sigma$ , if the following holds: whenever there exists a homomorphism  $h$  such that  $h(\Phi(\mathbf{X}, \mathbf{Y})) \subseteq I$ , then there exists  $h' \supseteq h|_{\mathbf{X}}$ , where  $h|_{\mathbf{X}}$  is the restriction of  $h$  on  $\mathbf{X}$ , such that  $h'(\Psi(\mathbf{X}, \mathbf{Z})) \subseteq I$ .

A *negative constraint* (NC)  $\nu$  is a first-order formula of the form  $\forall \mathbf{X} \Phi(\mathbf{X}) \rightarrow \perp$ , where  $\Phi(\mathbf{X})$  is a conjunction of atoms without nulls, called the *body* of  $\nu$  and denoted  $body(\nu)$ , and  $\perp$  denotes the truth constant *false*. An instance  $I$  satisfies  $\nu$ , written  $I \models \nu$ , if there is no homomorphism  $h$  such that  $h(\Phi(\mathbf{X})) \subseteq I$ .

A *program* (or *ontology*) is a finite set  $\Sigma$  of TGDs and NCs. We denote by  $\Sigma_T$  and  $\Sigma_{NC}$  the subsets of  $\Sigma$  containing the TGDs and NCs of  $\Sigma$ , respectively. An instance  $I$  satisfies  $\Sigma$ , written  $I \models \Sigma$ , if  $I$  satisfies each TGD and NC of  $\Sigma$ . For brevity, we omit the universal quantifiers in front of TGDs and NCs, and use the comma (instead of  $\wedge$ ) for conjoining atoms. Given a class of TGDs  $\mathbb{C}$ , we denote by  $\mathbb{C}_\perp$  the formalism obtained by combining  $\mathbb{C}$  with arbitrary NCs. For a program  $\Sigma$ ,  $\Sigma \in \mathbb{C}$  (resp.,  $\Sigma \in \mathbb{C}_\perp$ ) denotes that all TGDs of  $\Sigma$  belong to  $\mathbb{C}$ .

The Datalog<sup>±</sup> languages that we consider to guarantee decidability are among the most frequently analyzed in the literature, namely, linear (L) (Cali, Gottlob, and Lukasiewicz 2012), guarded (G) (Cali, Gottlob, and Kifer 2013), sticky (S) (Cali, Gottlob, and Pieris 2012), and acyclic TGDs (A), along with the “weak” (proper) generalizations weakly sticky (WS) (Cali, Gottlob, and Pieris 2012) and weakly acyclic TGDs (WA) (Fagin et al. 2005), as well as their “full” (i.e., existential-free) proper restrictions linear full (LF), guarded full (GF), sticky full (SF), and acyclic full TGDs (AF), respectively, and full TGDs (F) in general. We also recall the following further inclusions:  $L \subset G$  and  $F \subset WA \subset WS$ .

### 2.3 Ontology-Mediated Query Answering

The paradigm of ontology-mediated query answering generalizes query answering over databases by incorporating additional background knowledge in terms of an ontology.

$\mathcal{L}$	Data	<i>fp-comb.</i>	<i>ba-comb.</i>	Comb.
L, LF, AF	$\leq AC^0$	NP	NP	PSPACE
S, SF	$\leq AC^0$	NP	NP	EXP
A	$\leq AC^0$	NP	NEXP	NEXP
G	P	NP	EXP	2EXP
F, GF	P	NP	NP	EXP
WS, WA	P	NP	2EXP	2EXP

Table 1: Complexity of OMQA under existential rules.

Formally, an *ontology-mediated query (OMQ)* is a pair  $(Q, \Sigma)$ , where  $Q$  is a query, and  $\Sigma$  is an ontology. Let  $\mathcal{L}$  be an ontology language. If  $\Sigma \in \mathcal{L}$ , then we say that  $(Q, \Sigma)$  is an  $\mathcal{L}$ -OMQ. Consider a database  $D$  and an OMQ  $(Q, \Sigma)$ . The set of *models* of  $(D, \Sigma)$ , denoted  $mods(D, \Sigma)$ , is the set of instances  $\{I \mid I \supseteq D \wedge I \models \Sigma\}$ . We say that  $D$  *entails*  $(Q, \Sigma)$ , denoted  $D \models (Q, \Sigma)$ , if  $I \models Q$  for every  $I \in mods(D, \Sigma)$ . *Ontology-mediated query answering (OMQA)* is the task of deciding whether  $D \models (Q, \Sigma)$  for a given database  $D$  and an OMQ  $(Q, \Sigma)$ . Table 1 summarizes the known complexity results for OMQA in the different classes of TGDs that we consider. We denote by  $OMQA(\mathcal{L})$  the problem of OMQA when restricted over ontologies belonging to  $\mathcal{L}$ . We say that  $(D, \Sigma)$  is *consistent* if  $mods(D, \Sigma) \neq \emptyset$ , otherwise  $(D, \Sigma)$  is *inconsistent*. When  $OMQA(\mathcal{L})$  is restricted to the case where  $(D, \Sigma)$  is consistent, we talk of *consistent-OMQA( $\mathcal{L}$ )*.

A key paradigm in OMQA is the *FO-rewritability* of queries: an OMQ  $(Q, \Sigma)$  is *FO-rewritable*, if there exists a query  $Q_\Sigma$  such that, for all databases  $D$  that are consistent relative to  $\Sigma$ , we have that  $D \models (Q, \Sigma)$  iff  $D \models Q_\Sigma$ . In this case,  $Q_\Sigma$  is called an *FO-rewriting* of  $(Q, \Sigma)$ . A class of programs  $\mathcal{L}$  is *FO-rewritable*, if it admits an FO-rewriting for every query and program in  $\mathcal{L}$ . All languages from Table 1 with  $AC^0$  data complexity are FO-rewritable.

## 2.4 Computational Complexity

In our complexity analysis, we make the standard assumptions (Vardi 1982): the *combined complexity* of query answering is calculated by considering all the components (i.e., the database, the program, and the query) as part of the input. The *bounded-arity combined complexity* (or simply *ba-combined complexity*) assumes that the maximum arity of the predicates in  $\mathbf{R}$  is bounded by a constant integer. The *fixed-program combined complexity* (or simply *fp-combined complexity*) is calculated by considering the ontology as fixed. Finally, the *data complexity* is calculated by considering the database as the input, i.e., everything else is fixed. This paper’s most relevant complexity classes and their relations are as follows:

$$AC^0 \subseteq P \subseteq NP, \text{coNP} \subseteq D^P \subseteq \Sigma_2^P, \Pi_2^P \subseteq D_2^P \subseteq PSPACE \\ \subseteq EXP \subseteq NEXP, \text{coNEXP} \subseteq D^{EXP} \subseteq P^{NEXP} \subseteq 2EXP.$$

## 3 Explanations for Negative Query Answers

In this section, we formally define explanations and minimal explanations for negative query answers in OMQA along with several important computational problems for them.

**Definition 2.** Let  $D$  be a database, let  $(Q, \Sigma)$  be an OMQ, with  $D \not\models (Q, \Sigma)$ , and let  $H$  be a finite set of facts. An *explanation* for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$  is a subset  $E$  of  $H$  such that  $(D \cup E, \Sigma)$  is consistent and  $D \cup E \models (Q, \Sigma)$ . A *minimal explanation* (or *MinEX*) for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$  is an explanation  $E$  for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$  that is inclusion-minimal, i.e., no set  $E' \subsetneq E$  is an explanation for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$ .

Notice that, in order for  $D \not\models (Q, \Sigma)$  to hold,  $(D, \Sigma)$  must be consistent.

**Example 3.** Consider again the mechanical system in Example 1. We can encode the fault diagnosis problem as follows. In the database  $D$ , we store the relationships between the faulty parts and the caused issues. The binary predicate *FaultIssue*( $X, Y$ ) describes that when the part  $X$  is faulty, we may observe the issue  $Y$ . For example, *FaultIssue*( $p_1, \text{overheat}$ ) states that when  $p_1$  is faulty, we may observe an overheat of the system. Consider the following database:

$$D = \{ \text{FaultIssue}(p_1, \text{overheat}), \text{FaultIssue}(p_1, \text{leakage}), \\ \text{FaultIssue}(p_2, \text{high\_pressure}), \text{FaultIssue}(p_2, \text{non\_start}), \\ \text{FaultIssue}(p_3, \text{leakage}), \text{FaultIssue}(p_3, \text{burnt\_smell}), \\ \text{FaultIssue}(p_4, \text{flow\_blockage}) \}.$$

The set  $H$  of the possible causes of the observed issues is:

$$H = \{ \text{Fault}(p_1), \text{Fault}(p_2), \text{Fault}(p_3), \text{Fault}(p_4) \}.$$

OMQs  $(Q, \Sigma)$  then encode the issues that we observe (in  $Q$ ) and our additional knowledge about the functioning of the system (in  $\Sigma$ ). Consider the following program:

$$\Sigma = \{ \text{Fault}(p_1), \text{Fault}(p_4) \rightarrow \perp; \text{Fault}(p_1) \rightarrow \text{Fault}(p_3); \\ \text{Fault}(p_1), \text{Fault}(p_2) \rightarrow \text{Issue}(\text{pipes\_rattling}); \\ \text{FaultIssue}(X, Y), \text{Fault}(X) \rightarrow \text{Issue}(Y) \}.$$

The first rule in  $\Sigma$  (an NC) says that the parts  $p_1$  and  $p_4$  cannot be faulty at the same time. The second one encodes that if  $p_1$  is faulty, then so is  $p_3$ . The third one says that when  $p_1$  and  $p_2$  are both faulty, then we observe rattling pipes. The last one encodes that if  $X$  is a part causing the issue  $Y$ , and  $X$  is faulty, then we observe the issue  $Y$ .

Consider now the following OMQ:

$$(Q = \text{Issue}(\text{leakage}) \wedge \text{Issue}(\text{non\_start}), \Sigma),$$

where  $\Sigma$  is the program defined above. In the query, we state that we observe a leakage and that the system does not start. The OMQ is not entailed by the database  $D$  alone. However, by adding to  $D$  facts taken from  $H$ , the OMQ can be derived, and hence an explanation built from  $H$  can be interpreted as a possible set of causes for the issues observed.  $\triangleleft$

The first computational problem for explanations for negative query answers in OMQA is deciding whether a set of facts is a minimal explanation for a negative answer.

**Problem:** IS-MINEX $\neq(\mathcal{L})$ .

*Input:* A database  $D$ , an  $\mathcal{L}$ -OMQ  $(Q, \Sigma)$ , where  $D \not\models (Q, \Sigma)$ , a finite set of facts  $H$ , and  $E \subseteq H$ .

*Question:* Is  $E$  a MinEX for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$ ?

**Example 4.** Consider again the database  $D$ , the program  $\Sigma$ , and the set  $H$  of Example 3. Then, for the OMQ

$$(Q = \text{Issue}(\text{burnt\_smell}) \wedge \text{Issue}(\text{leakage}), \Sigma),$$

the sets  $E_1 = \{\text{Fault}(p_1)\}$  and  $E_2 = \{\text{Fault}(p_3)\}$  are minimal explanations, while  $E' = \{\text{Fault}(p_1), \text{Fault}(p_3)\}$  is an explanation, but not a minimal one.  $\triangleleft$

Another computational problem is deciding whether there exists a minimal explanation at all for a given OMQ.

**Problem:** MINEX-EXISTS $\neq$ ( $\mathcal{L}$ ).

*Input:* A database  $D$ , an  $\mathcal{L}$ -OMQ  $(Q, \Sigma)$ , where  $D \not\models (Q, \Sigma)$ , and a finite set of facts  $H$ .

*Question:* Is there a MinEX for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$ ?

**Example 5.** Consider again the database  $D$ , the program  $\Sigma$ , and the set  $H$  of Example 3. Then, for the OMQ

$$(Q = \text{Issue}(\text{flow\_blockage}) \wedge \text{Issue}(\text{burnt\_smell}), \Sigma),$$

the set  $E = \{\text{Fault}(p_3), \text{Fault}(p_4)\}$  is an explanation, but not  $E' = \{\text{Fault}(p_1), \text{Fault}(p_4)\}$ , as the simultaneous presence of  $\text{Fault}(p_1)$  and  $\text{Fault}(p_4)$  is not admitted by the NC. On the other hand, for the OMQ

$$(Q = \text{Issue}(\text{flow\_blockage}) \wedge \text{Issue}(\text{overheat}), \Sigma),$$

there is no explanation.  $\triangleleft$

Two other computational problems are recognizing relevant facts and recognizing necessary facts. A fact  $\psi$  is *relevant* (resp., *necessary*) for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$  iff  $\psi$  appears in at least one (resp., in every) MinEX for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$ .

**Problem:** MINEX-REL $\neq$ ( $\mathcal{L}$ ).

*Input:* A database  $D$ , an  $\mathcal{L}$ -OMQ  $(Q, \Sigma)$ , where  $D \not\models (Q, \Sigma)$ , a finite set of facts  $H$ , and a fact  $\psi$ .

*Question:* Is  $\psi$  relevant for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$ ?

**Problem:** MINEX-NEC $\neq$ ( $\mathcal{L}$ ).

*Input:* A database  $D$ , an  $\mathcal{L}$ -OMQ  $(Q, \Sigma)$ , where  $D \not\models (Q, \Sigma)$ , a finite set of facts  $H$ , and a fact  $\psi$ .

*Question:* Is  $\psi$  necessary for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$ ?

**Example 6.** Consider again the database  $D$ , the program  $\Sigma$ , and the set  $H$  of Example 3. Then, for the OMQ

$$(Q = \text{Issue}(\text{non\_start}) \wedge \text{Issue}(\text{leakage}), \Sigma),$$

the sets  $E_1 = \{\text{Fault}(p_1), \text{Fault}(p_2)\}$  and  $E_2 = \{\text{Fault}(p_2), \text{Fault}(p_3)\}$  are all the minimal explanations for the OMQ. Hence, the facts  $\text{Fault}(p_1)$ ,  $\text{Fault}(p_2)$ , and  $\text{Fault}(p_3)$  are relevant, while the only necessary fact is  $\text{Fault}(p_3)$ .  $\triangleleft$

The computational problems introduced so far are those commonly addressed in the context of abductive reasoning and negative answer explanations. In addition, we introduce two novel important problems. The first one asks whether a set  $H'$  of facts contains exactly all the relevant facts, i.e.,  $H'$  is the union of all MinEXs. The problem is particularly interesting, as  $H'$  can be seen as a minimal over-approximation of all MinEXs, i.e., for every MinEX  $E$ , it holds that  $E \subseteq H'$ , and  $H'$  is the smallest set enjoying this property.

**Problem:** MINEX-ALLREL $\neq$ ( $\mathcal{L}$ ).

*Input:* A database  $D$ , an  $\mathcal{L}$ -OMQ  $(Q, \Sigma)$ , where  $D \not\models (Q, \Sigma)$ ,

a finite set of facts  $H$ , and a set  $H' \subseteq H$ .

*Question:* Does  $H'$  contain exactly all the relevant facts for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$ ?

The second novel computational problem that we consider asks whether a set  $H'$  contains exactly all the necessary facts, i.e.,  $H'$  is the intersection of all MinEXs. Interestingly,  $H'$  can be seen as a maximal under-approximation of all MinEXs, i.e., for every MinEX  $E$ , it holds that  $H' \subseteq E$ , and  $H'$  is the biggest set enjoying this property.

**Problem:** MINEX-ALLNEC $\neq$ ( $\mathcal{L}$ ).

*Input:* A database  $D$ , an  $\mathcal{L}$ -OMQ  $(Q, \Sigma)$ , where  $D \not\models (Q, \Sigma)$ , a finite set of facts  $H$ , and a set  $H' \subseteq H$ .

*Question:* Does  $H'$  contain exactly all the necessary facts for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$ ?

**Example 7.** For the OMQ in Example 6

$$(Q = \text{Issue}(\text{non\_start}) \wedge \text{Issue}(\text{leakage}), \Sigma),$$

the sets of all relevant and necessary facts are  $\{\text{Fault}(p_1), \text{Fault}(p_2), \text{Fault}(p_3)\}$  and  $\{\text{Fault}(p_3)\}$ , respectively.  $\triangleleft$

## 4 Positive vs. Negative Explanations

Explaining positive query answers, i.e., why a database  $D$  entails an OMQ  $(Q, \Sigma)$ , where  $\Sigma$  is a set of TGDS, has been recently investigated by Ceylan et al. (2019). Specifically, an *explanation* for  $D \models (Q, \Sigma)$  is a set  $E \subseteq D$  such that  $E \models (Q, \Sigma)$ . A *minimal explanation* for  $D \models (Q, \Sigma)$  is an inclusion-minimal explanation for  $D \models (Q, \Sigma)$ .

In this paper, we deal with “negative” explanations. One may nevertheless wonder whether there are easy reductions mapping one problem into the other and covering the problems that are considered. It is possible, as we show, to leverage a reduction from the positive to the negative setting to derive the lower bounds for the problems IS-MINEX $\neq$  and MINEX-REL $\neq$ . For the remaining problems, however, there is no straightforward correspondence between positive and negative explanation problems, and hence, they required novel proofs.

One important difference between the positive and negative setting is that positive explanations are subsets of the database, while negative explanations are subsets of a set  $H$  of abducible facts and must be added to the entire database. This makes simple reductions from the negative to the positive setting fail. As an example, suppose in our setting that we are given a database  $D$ , an OMQ  $(Q, \Sigma)$ , and a finite set of facts  $H$ . One might try to derive an instance of the positive explanation problem where the database is  $D \cup H$  and the OMQ remains  $(Q, \Sigma)$ . However, a set  $E$  that is a minimal explanation for  $D \cup H \models (Q, \Sigma)$  might not be a minimal explanation for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$  (because  $E$  is not minimal or it is not an explanation at all). As an example, consider the database  $D = \{R(a)\}$ , the OMQ  $(Q, \Sigma)$  with  $Q = \exists X, Y R(X), S(Y)$  and  $\Sigma = \emptyset$ , and the set of abducible facts  $H = \{R(b), S(b)\}$ . Then,  $E = \{R(b), S(b)\}$  is a minimal explanation for  $D \cup H \models (Q, \Sigma)$ , but it is not a minimal explanation for  $D \not\models (Q, \Sigma)$  w.r.t.  $H$  (because it is not minimal).

## 5 Overview of Complexity Results

We give a precise picture of the complexity for all the explanation problems. An overview of our complexity results is given in Tables 2–7. They range from membership in P to 2EXP-completeness; all entries without ‘ $\leq$ ’ are completeness results. For all problems and complexity measures,  $H$  (as well as  $E$  in IS-MINEX $^{\neq}$ ,  $\psi$  in MINEX-REL $^{\neq}$  and MINEX-NEC $^{\neq}$ , and  $H'$  in MINEX-ALLREL $^{\neq}$  and MINEX-ALLNEC $^{\neq}$ ) is always part of the input, that is, it is treated like the database.

MINEX-REL $^{\neq}$  and MINEX-NEC $^{\neq}$  always have complementary complexity classes, except in the *fp*-combined complexity, where MINEX-REL $^{\neq}$  is  $\Sigma_2^p$ -complete, while MINEX-NEC $^{\neq}$  is coNP-complete. Intuitively, the reason is that to decide whether a fact  $\psi$  is *not* necessary, we can guess a subset  $E$  of  $H$  and check whether  $E$  is an explanation that does not include  $\psi$ , and verifying minimality of  $E$  can be avoided. In contrast, deciding relevance requires verifying minimality of a guessed set, which in turn requires verifying *non*-entailment. Such a check is in coNP in the *fp*-combined complexity and yields a membership in  $\text{NP}^{\text{coNP}} = \Sigma_2^p$ . This also explains the difference between MINEX-ALLREL $^{\neq}$  and MINEX-ALLNEC $^{\neq}$  in the *fp*-combined complexity.

The discussion above is also related to why MINEX-EXISTS $^{\neq}$  has a lower complexity than MINEX-REL $^{\neq}$  in the *fp*-combined complexity: a minimal explanation exists iff an explanation (not necessarily minimal) exists; we can then avoid verifying minimality of a guessed subset of  $H$ .

For the linear, full, and sticky languages, as well as for their sublanguages, OMQA is NP-complete both in the *fp*- and in the *ba*-combined complexity. One may expect that, for a fixed problem, the complexity does not change in such cases, as the complexity of OMQA remains the same across them. However, this is not the case for MINEX-EXISTS $^{\neq}$ , MINEX-NEC $^{\neq}$ , and MINEX-ALLNEC $^{\neq}$ , where the complexity goes from NP-, coNP-, and  $D^p$ -completeness in the *fp*-combined complexity to  $\Sigma_2^p$ -,  $\Pi_2^p$ -, and  $D_2^p$ -completeness in the *ba*-combined complexity, respectively. The reason is that for such problems, consistency checking can be done in P in the *fp*-combined complexity, as the program is fixed, and it is not required to check minimality of a guessed set. In contrast, in the *ba*-combined complexity, the program is not fixed anymore, and checking consistency requires a linear number of checks for *non*-entailment, whose cost increases the overall complexity by one level.

We analyze the complexity of IS-MINEX $^{\neq}$  and MINEX-EXISTS $^{\neq}$  in Section 6, of MINEX-REL $^{\neq}$  and MINEX-ALLREL $^{\neq}$  in Section 7, and of MINEX-NEC $^{\neq}$  and MINEX-ALLNEC $^{\neq}$  in Section 8. For each problem, we first discuss membership and then hardness results. We first prove general procedures for the upper bounds that apply to all languages and complexity measures, and we separately show tighter upper bounds, when required. As an example, the data complexity of all problems for FO-rewritable languages is in P. In many cases, this is a tighter upper bound than the one provided by the general procedure, and it is derived by leveraging rewritability into unions of conjunctive queries. Hardness results are stated for the most specific languages that they

$\mathcal{L}$	Data	<i>fp</i> -comb.	<i>ba</i> -comb.	Comb.
$L_{\perp}, LF_{\perp}, AF_{\perp}$	$\leq P$	$D^p$	$D^p$	PSPACE
$S_{\perp}, SF_{\perp}$	$\leq P$	$D^p$	$D^p$	EXP
$A_{\perp}$	$\leq P$	$D^p$	$D^{\text{EXP}}$	$D^{\text{EXP}}$
$G_{\perp}$	P	$D^p$	EXP	2EXP
$F_{\perp}, GF_{\perp}$	P	$D^p$	$D^p$	EXP
$WS_{\perp}, WA_{\perp}$	P	$D^p$	2EXP	2EXP

Table 2: Complexity of IS-MINEX $^{\neq}(\mathcal{L})$ .

apply to (and then hold for all generalizations, of course).

## 6 IS-MINEX $^{\neq}$ and MINEX-EXISTS $^{\neq}$

We now analyze the complexity of recognizing a MinEX and deciding the existence of a MinEX for a given database and an OMQ. We start with IS-MINEX $^{\neq}(\mathcal{L})$ , i.e., deciding whether a given set of facts is a minimal explanation for a negative query answer, and the membership results. The following theorem proves all upper bounds in Table 2.

**Theorem 8.** *For any language  $\mathcal{L}$  considered in this paper, if  $OMQA(\mathcal{L})$  is in the complexity class  $\mathcal{C}$  in the combined (resp., *ba*-combined, *fp*-combined, data) complexity, then IS-MINEX $^{\neq}(\mathcal{L})$  can be decided with a  $\mathcal{C}$  check and a linear number of co- $\mathcal{C}$  checks in the combined (resp., *ba*-combined, *fp*-combined, data) complexity.*

*Proof sketch.* Deciding whether a given set  $E$  of facts is a MinEX requires to carry out essentially three tasks: (1) deciding whether  $(D \cup E, \Sigma)$  is consistent; (2) deciding whether  $D \cup E \models (Q, \Sigma)$ ; and (3) deciding whether  $E$  is inclusion-minimal. Task (1) can be solved by checking that  $D \cup E$  does not entail the body of any negative constraint, which hence consists of a linear number of *non*-entailment tests. Task (2) can be carried out by performing a single entailment test. Task (3) can be decided by checking that all facts in  $E$  are required to entail  $(Q, \Sigma)$ , and this can be done by verifying that removing every single fact  $\psi$  from  $E$  in turn leads to a set of facts not entailing  $(Q, \Sigma)$ . Also this task consists of a linear number of *non*-entailment tests.  $\square$

We now focus on the hardness results for IS-MINEX $^{\neq}(\mathcal{L})$ . Among others, Ceylan et al. (2019) studied the problem of deciding whether a set  $E$  is a minimal explanation for  $D \models (Q, \Sigma)$ , called IS-MINEX $(\mathcal{L})$ . All the lower bounds in Table 2 are obtained from the following result that establishes a connection between “positive” and “negative” explanations.

**Lemma 9.** *For any language  $\mathcal{L}$  considered in this paper, IS-MINEX $(\mathcal{L})$  is reducible in polynomial time to IS-MINEX $^{\neq}(\mathcal{L})$  in the data, *fp*-combined, *ba*-combined, and combined complexity.*

We now focus on the problem MINEX-EXISTS $^{\neq}(\mathcal{L})$  of deciding the existence of (minimal) explanations for negative query answers. We start by providing the membership results. The following theorem proves all the upper bounds in Table 3 but the NP and P ones, for which we need tighter results stated next. Intuitively, to decide whether there exists a minimal explanation for a negative query answer, it suffices

$\mathcal{L}$	Data	$fp$ -comb.	$ba$ -comb.	Comb.
$L_{\perp}, LF_{\perp}, AF_{\perp}$	$\leq P$	NP	$\Sigma_2^P$	PSPACE
$S_{\perp}, SF_{\perp}$	$\leq P$	NP	$\Sigma_2^P$	EXP
$A_{\perp}$	$\leq P$	NP	$P^{NEXP}$	$P^{NEXP}$
$G_{\perp}$	NP	NP	EXP	2EXP
$F_{\perp}, GF_{\perp}$	NP	NP	$\Sigma_2^P$	EXP
$WS_{\perp}, WA_{\perp}$	NP	NP	2EXP	2EXP

Table 3: Complexity of  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$ .

to check whether there is any explanations for the negative query answers, i.e., there is no need to double check the minimality. Hence, we guess a set  $E \subseteq H$  of facts (feasible in NP), and then check via oracle calls that  $E$  is consistent and allows us to entail the OMQ (Tasks (1) and (2) in the proof sketch of Theorem 8).

**Theorem 10.** *For any language  $\mathcal{L}$  considered here, if  $\text{OMQA}(\mathcal{L})$  is in the complexity class  $\mathcal{C}$  in the combined (resp.,  $ba$ -combined,  $fp$ -combined, data) complexity, then  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$  is in  $\text{NP}^{\mathcal{C}}$  in the combined (resp.,  $ba$ -combined,  $fp$ -combined, data) complexity.*

The following theorem provides the NP upper bounds in the  $fp$ -combined complexity in Table 3. The result is obtained from the proof of Theorem 8 with the additional observation that in the  $fp$ -combined setting, for the  $\text{Datalog}^{\pm}$  languages considered, checking whether a set of facts is consistent is feasible in P, because the negative constraints are fixed.

**Theorem 11.**  *$\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$  is in NP in the  $fp$ -combined complexity for all languages  $\mathcal{L}$  considered here.*

We conclude the discussion of the membership results for  $\text{MINEX-EXISTS}^{\neq}$  by showing that the problem is in P in the data complexity for all the FO-rewritable languages considered, namely, L, S, and A (as well as their full restrictions). The proof relies on the fact that, for these languages, any OMQ can be rewritten into an equivalent FO query  $Q$ , which is a union of conjunctive queries. Starting from the values of the tuples in the database and  $H$ , it is possible to build in polynomial time the set  $\mathcal{D}$  containing all the polynomially-many sets of facts that can possibly satisfy  $Q$ . Once  $\mathcal{D}$  is computed, all its elements can be tested in turn, and we verify (in polynomial time) whether any of them is a (non-necessarily minimal) explanation for the negative answer of the OMQ.

**Theorem 12.** *If  $\mathcal{L}$  is FO-rewritable language, then  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$  is in P in the data complexity.*

As for the hardness results for  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$ , the following theorem proves the lower bounds in the data complexity in Table 3 for the non-FO-rewritable languages, namely,  $GF_{\perp}$  and its generalizations. Its proof is a reduction from the satisfiability of 3CNF Boolean formulas  $\phi(X)$ . Intuitively, in a fixed program, there are rules deciding the satisfiability of  $\phi(X)$ . The set  $H$  contains facts associated with truth assignments for the variables  $X$ . A candidate explanation  $E$  encodes a truth assignment  $\sigma_E$  for the variables  $X$ , and  $E$  is actually an explanation iff  $\sigma_E$  satisfies  $\phi(X)$ .

**Theorem 13.**  *$\text{MINEX-EXISTS}^{\neq}(GF_{\perp})$  is NP-hard in the data complexity.*

The hardness results in Table 3 in the  $fp$ -combined,  $ba$ -combined, and combined complexity can mainly be obtained from the following lemma showing a link between the complexity of consistent- $\text{OMQA}(\mathcal{L})$  and of  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$ . The  $\Sigma_2^P$ -hardness and the  $P^{NEXP}$ -hardness results in Table 3 require tighter theorems, which will be stated next.

**Lemma 14.** *For any language  $\mathcal{L}$  considered here, consistent- $\text{OMQA}(\mathcal{L})$  is reducible in polynomial time to  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$  in the data,  $fp$ -combined,  $ba$ -combined, and combined complexity.*

The following theorem proves the  $\Sigma_2^P$ -hardness results (in the  $ba$ -combined complexity) in Table 3. The proof uses a reduction from the  $\Sigma_2^P$ -complete problem  $\text{QBF}_{2,\forall}^{DNF}$ : decide the validity of a quantified Boolean formula  $\Phi = \exists X \forall Y \phi(X, Y)$ , where  $\phi(X, Y)$  is in 3DNF. We need to encode the validity of the quantified formula. In  $H$ , there are facts associated with the various truth assignments for the variables  $X$ . The test for the unsatisfiability of  $\phi(X, Y)$ , when a truth assignment for  $X$  is encoded in a candidate explanation, is captured with a negative constraint. A candidate explanation  $E$  encoding a truth assignment  $\sigma_E$  for  $X$  is actually an explanation iff  $\forall Y \phi(X/\sigma_E, Y)$  is valid. So, if  $\forall X \exists Y \neg \phi(X, Y)$  holds, then there are no explanations.

**Theorem 15.**  *$\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$  is  $\Sigma_2^P$ -hard in the  $ba$ -combined complexity for all languages  $\mathcal{L}$  considered here.*

The following theorem proves the two  $P^{NEXP}$ -hardness results in Table 3 for A in the  $ba$ -combined and combined complexity. The result can be shown via a reduction from the problem  $ETP$  defined by Eiter, Lukasiewicz, and Predoiu (2016): given a triple  $(m, TP_1, TP_2)$ , where  $m$  is a number in unary, and  $TP_1$  and  $TP_2$  are two tiling problems for the exponential square  $2^m \times 2^m$ , decide whether there exists an initial tiling conditions  $w$  of length  $m$  such that  $TP_1$  has a solution with  $w$  and  $TP_2$  has no solution with  $w$ . The idea of the reduction is to have the candidate explanations taken from  $H$  to encode the possible initial tiling conditions. Then, there are rules that allow us to derive the query if  $TP_1$  has a solution with the initial condition  $w$  encoded in the explanation, and  $TP_2$  has no solution with  $w$ .

**Theorem 16.**  *$\text{MINEX-EXISTS}^{\neq}(A_{\perp})$  is  $P^{NEXP}$ -hard in the  $ba$ -combined and combined complexity.*

## 7 $\text{MINEX-REL}^{\neq}$ and $\text{MINEX-ALLREL}^{\neq}$

In this section, we deal with the problems regarding relevant facts, i.e., facts appearing in at least one minimal explanation. In particular, we analyze the problem of deciding whether a fact is relevant and the problem of deciding whether a set of facts is the set of all and only the relevant facts. We start by looking at the problem  $\text{MINEX-REL}^{\neq}(\mathcal{L})$  of deciding whether a fact is relevant or not, and first provide the membership results. The following theorem proves all the upper bounds in Table 4, except for those in the data complexity for FO-languages, for which we need a tighter statement. Intuitively, to decide whether  $\psi$  is relevant, it suffices to guess a set  $E$  of facts containing  $\psi$  (feasible in NP), and then, via an oracle call, check that  $E$  is a minimal explanation.

$\mathcal{L}$	Data	$fp\text{-comb.}$	$ba\text{-comb.}$	Comb.
$L_{\perp}, LF_{\perp}, AF_{\perp}$	$\leq P$	$\Sigma_2^P$	$\Sigma_2^P$	PSPACE
$S_{\perp}, SF_{\perp}$	$\leq P$	$\Sigma_2^P$	$\Sigma_2^P$	EXP
$A_{\perp}$	$\leq P$	$\Sigma_2^P$	$P^{NEXP}$	$P^{NEXP}$
$G_{\perp}$	NP	$\Sigma_2^P$	EXP	2EXP
$F_{\perp}, GF_{\perp}$	NP	$\Sigma_2^P$	$\Sigma_2^P$	EXP
$WS_{\perp}, WA_{\perp}$	NP	$\Sigma_2^P$	2EXP	2EXP

Table 4: Complexity of  $\text{MINEX-REL}^{\neq}(\mathcal{L})$ .

$\mathcal{L}$	Data	$fp\text{-comb.}$	$ba\text{-comb.}$	Comb.
$L_{\perp}, LF_{\perp}, AF_{\perp}$	$\leq P$	$D_2^P$	$D_2^P$	PSPACE
$S_{\perp}, SF_{\perp}$	$\leq P$	$D_2^P$	$D_2^P$	EXP
$A_{\perp}$	$\leq P$	$D_2^P$	$P^{NEXP}$	$P^{NEXP}$
$G_{\perp}$	$D^P$	$D_2^P$	EXP	2EXP
$F_{\perp}, GF_{\perp}$	$D^P$	$D_2^P$	$D_2^P$	EXP
$WS_{\perp}, WA_{\perp}$	$D^P$	$D_2^P$	2EXP	2EXP

Table 5: Complexity of  $\text{MINEX-ALLREL}^{\neq}(\mathcal{L})$ .

**Theorem 17.** For any language  $\mathcal{L}$  considered here, if  $\text{IS-MINEX}^{\neq}(\mathcal{L})$  is in the complexity class  $\mathcal{C}$  in the combined (resp.,  $ba$ -combined,  $fp$ -combined, data) complexity, then  $\text{MINEX-REL}^{\neq}(\mathcal{L})$  is in  $\text{NP}^{\mathcal{C}}$  in the combined (resp.,  $ba$ -combined,  $fp$ -combined, data) complexity.

For the FO-rewritable languages, the following theorem proves the P-membership results in the data complexity shown in Table 4 that are not obtained via Theorem 17. To prove this result, we exploit, also in this case, the rewritability of the OMQ into an FO query  $Q$  and the possibility to compute in polynomial time the set  $\mathcal{D}$  containing all the polynomially many sets of facts that can possibly satisfy  $Q$ . With  $\mathcal{D}$  at hand, its elements can be considered in turn to verify whether there exists a minimal explanation containing the fact to be tested for relevance.

**Theorem 18.** For any FO-rewritable language  $\mathcal{L}$  considered here,  $\text{MINEX-REL}^{\neq}(\mathcal{L})$  is in P in the data complexity.

We now consider the hardness results. To study the hardness of  $\text{MINEX-REL}^{\neq}(\mathcal{L})$ , we can resort to its “positive” variant. Ceylan et al. (2019) studied the problem of deciding whether a fact  $\psi$  belongs to some minimal explanation for  $D \models (Q, \Sigma)$ , called  $\text{MINEX-REL}(\mathcal{L})$ . It is possible to show that  $\text{MINEX-REL}(\mathcal{L})$  reduces to  $\text{MINEX-REL}^{\neq}(\mathcal{L})$ , thereby obtaining all the hardness results in Table 4.

**Theorem 19.** For any language  $\mathcal{L}$  considered here,  $\text{MINEX-REL}(\mathcal{L})$  is reducible in polynomial time to  $\text{MINEX-REL}^{\neq}(\mathcal{L})$  in the data,  $fp$ -combined,  $ba$ -combined, and combined complexity.

We now analyze the problem  $\text{MINEX-ALLREL}^{\neq}(\mathcal{L})$  of deciding whether a set contains all and only the relevant facts. We focus on the membership results first. The following theorem proves all the upper bounds in Table 5. Intuitively, to check that  $H'$  is the set of all and only the relevant facts, it suffices to check that all facts in  $H'$  are relevant, and all facts outside  $H'$  are not relevant.

**Theorem 20.** For any language  $\mathcal{L}$  considered here, if  $\text{MINEX-REL}^{\neq}(\mathcal{L})$  is in the complexity class  $\mathcal{C}$  in the combined (resp.,  $ba$ -combined,  $fp$ -combined, data) complexity, then  $\text{MINEX-ALLREL}^{\neq}(\mathcal{L})$  can be decided with a  $\mathcal{C}$  check and a  $co\text{-}\mathcal{C}$  in the combined (resp.,  $ba$ -combined,  $fp$ -combined, data) complexity.

As for hardness, the following theorem proves all the hardness results in Table 5 in the data complexity for the non-FO-rewritable languages via a reduction from the classical  $D^P$ -complete problem SAT-UNSAT: given two Boolean formulas in 3CNF  $\phi(X)$  and  $\psi(X)$ , decide whether  $\phi$  is satisfiable, and  $\psi$  is not satisfiable. In a fixed program, there are the rules to check the satisfiability of the formulas. In  $H$ , there are facts to encode possible truth assignments for the Boolean variables  $X$ , plus some additional facts to recognize the satisfaction of  $\phi$  and  $\psi$ , for the specific truth assignments in the candidate explanations. The idea is to have all facts in  $H$  to be relevant, apart from the one associated with the satisfaction of  $\psi$ .

**Theorem 21.**  $\text{MINEX-ALLREL}^{\neq}(\text{GF}_{\perp})$  is  $D^P$ -hard in the data complexity.

The following theorem proves all the  $D_2^P$  lower bounds in Table 5 in the  $fp$ -combined complexity. It also gives hardness results matching the upper bounds for  $L_{\perp}$ ,  $F_{\perp}$ ,  $S_{\perp}$ , and their specializations in the  $ba$ -combined complexity. We show a reduction from the prototypical  $D_2^P$ -complete problem  $\text{QBF}_{2,\forall,\neg}^{CNF} \wedge \text{QBF}_{2,\exists}^{CNF}$ : decide the validity of two quantified Boolean formulas  $\Phi = \exists X \forall Y \neg \phi(X, Y)$  and  $\Psi = \forall X \exists Y \psi(X, Y)$  (notice that by an argument similar to the one in the proof of Theorem 3.9 in (Lukasiewicz and Malizia 2017),  $X$  and  $Y$  can be assumed w.l.o.g. to be the same in  $\phi(X, Y)$  and  $\psi(X, Y)$ ). The tests of satisfaction of  $\phi(X, Y)$  and  $\psi(X, Y)$  are in the query. In  $H$ , there are facts representing truth assignments for variables  $X$ , plus additional facts to recognize the validity of  $\Phi$  and  $\Psi$ : we check for the relevance or not of these additional facts.

**Theorem 22.** For any language  $\mathcal{L}$  considered here,  $\text{MINEX-ALLREL}^{\neq}(\mathcal{L})$  is  $D_2^P$ -hard in the  $fp$ -combined complexity.

The remaining hardness results in the  $ba$ -combined complexity as well as all the hardness results in the combined complexity are established via the following theorem, which shows that  $\text{MINEX-EXISTS}^{\neq}$  can be reduced to the complement of  $\text{MINEX-ALLREL}^{\neq}$ .

**Theorem 23.** For any language  $\mathcal{L}$  considered here,  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$  is reducible in polynomial time to the complement of  $\text{MINEX-ALLREL}^{\neq}(\mathcal{L})$  in the data,  $fp$ -combined,  $ba$ -combined, and combined complexity.

## 8 $\text{MINEX-NEC}^{\neq}$ and $\text{MINEX-ALLNEC}^{\neq}$

In this section, we study the problems regarding necessary facts, i.e., facts appearing in all the minimal explanations. In particular, we analyze the problem of deciding whether a fact is necessary and the problem of deciding whether a set of facts is the set of all and only the necessary facts.

We now focus on the problem  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$  of deciding whether a fact is necessary. As for the membership

results, the following theorem proves all upper bounds in Table 6 but the coNP and P ones, for which we need tighter results, as stated next. Intuitively, we can *disprove* that a fact  $\psi$  is necessary by checking that there is a (non-necessarily minimal) explanation excluding  $\psi$ . Hence, we can answer by guessing a set  $E \subseteq H \setminus \{\psi\}$  of facts (feasible in NP), and then check via oracle calls that  $E$  is consistent and entails the OMQ (Tasks (1) and (2) in the proof sketch of Theorem 8).

**Theorem 24.** *For any language  $\mathcal{L}$  considered here, if  $OMQA(\mathcal{L})$  is in the complexity class  $\mathcal{C}$  in the combined (resp., *ba-combined*, *fp-combined*, *data*) complexity, then  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$  is in  $co\text{-}(\text{NP}^{\mathcal{C}})$  in the combined (resp., *ba-combined*, *fp-combined*, *data*) complexity.*

The following theorem provides the coNP upper bounds in the *fp-combined* complexity in Table 6. The result is obtained from the proof of Theorem 24 with the additional observation that in the *fp-combined* setting, for the  $\text{Datalog}^{\pm}$  languages considered, checking whether a set of facts is consistent is feasible in P, because the negative constraints are fixed.

**Theorem 25.** *For any language  $\mathcal{L}$  considered here,  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$  is in coNP in the *fp-combined* complexity.*

We conclude the discussion of the membership results for  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$  by showing that the problem is in P in the data complexity for all the FO-rewritable languages considered, namely, L, S, and A (as well as their full restrictions). For the FO-rewritable languages, the following theorem proves the P membership results in the data complexity shown in Table 6 that are not obtained via Theorem 24. To prove this result, we exploit the rewritability of an OMQ into an FO query  $Q$  and the possibility to compute in polynomial time the set  $\mathcal{D}$  containing all the polynomially many sets of facts that can possibly satisfy  $Q$ . Once  $\mathcal{D}$  has been computed, its elements can be considered in turn to verify whether there exists a non-necessarily minimal explanation excluding the fact to be tested for necessity. If such an explanation is found, we answer no, otherwise we answer yes.

**Theorem 26.** *For any FO-rewritable language  $\mathcal{L}$  considered here,  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$  is in P in the data complexity.*

Regarding the hardness results of  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$ , they can be proven by showing a reduction from the complement of  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$  to  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$ .

**Theorem 27.** *For any language  $\mathcal{L}$  considered here, the complement of  $\text{MINEX-EXISTS}^{\neq}(\mathcal{L})$  is reducible in polynomial time to  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$  in the data, *fp-combined*, *ba-combined*, and combined complexity.*

We now study the problem  $\text{MINEX-ALLNEC}^{\neq}(\mathcal{L})$  of deciding whether a set contains all and only the necessary facts. We start by looking at the membership results. The following theorem proves all upper bounds in Table 7. Intuitively, to check that  $H'$  is the set of all and only the necessary facts, it suffices to check that all facts in  $H'$  are necessary and all facts outside  $H'$  are not necessary.

**Theorem 28.** *For any language  $\mathcal{L}$  considered here, if  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$  is in the complexity class  $\mathcal{C}$  in the combined (resp., *ba-combined*, *fp-combined*, *data*) complexity, then  $\text{MINEX-ALLNEC}^{\neq}(\mathcal{L})$  can be decided by a check in*

$\mathcal{L}$	Data	<i>fp-comb.</i>	<i>ba-comb.</i>	Comb.
$L_{\perp}, LF_{\perp}, AF_{\perp}$	$\leq P$	coNP	$\Pi_2^P$	PSPACE
$S_{\perp}, SF_{\perp}$	$\leq P$	coNP	$\Pi_2^P$	EXP
$A_{\perp}$	$\leq P$	coNP	$P^{\text{NEXP}}$	$P^{\text{NEXP}}$
$G_{\perp}$	coNP	coNP	EXP	2EXP
$F_{\perp}, GF_{\perp}$	coNP	coNP	$\Pi_2^P$	EXP
$WS_{\perp}, WA_{\perp}$	coNP	coNP	2EXP	2EXP

Table 6: Complexity of  $\text{MINEX-NEC}^{\neq}(\mathcal{L})$ .

$\mathcal{L}$	Data	<i>fp-comb.</i>	<i>ba-comb.</i>	Comb.
$L_{\perp}, LF_{\perp}, AF_{\perp}$	$\leq P$	$D^P$	$D_2^P$	PSPACE
$S_{\perp}, SF_{\perp}$	$\leq P$	$D^P$	$D_2^P$	EXP
$A_{\perp}$	$\leq P$	$D^P$	$P^{\text{NEXP}}$	$P^{\text{NEXP}}$
$G_{\perp}$	$D^P$	$D^P$	EXP	2EXP
$F_{\perp}, GF_{\perp}$	$D^P$	$D^P$	$D_2^P$	EXP
$WS_{\perp}, WA_{\perp}$	$D^P$	$D^P$	2EXP	2EXP

Table 7: Complexity of  $\text{MINEX-ALLNEC}^{\neq}(\mathcal{L})$ .

$\mathcal{C}$  and a check in  $co\text{-}\mathcal{C}$  in the combined (resp., *ba-combined*, *fp-combined*, *data*) complexity.

As for the hardness results, the following theorem proves all the hardness results in Table 7 in the data complexity for the languages considered here that are not FO-rewritable. The proof is via a reduction from SAT-UNSAT. In  $H$ , there are facts associated with the possible truth assignments for the variables  $X$  of the formulas, plus additional facts allowing to recognize that the formulas  $\phi$  and  $\psi$  are unsatisfiable. In this case, the idea is to have all facts in  $H$  to be non-necessary, apart from the one associated with the unsatisfaction of  $\psi$ .

**Theorem 29.**  *$\text{MINEX-ALLNEC}^{\neq}(\text{GF}_{\perp})$  is  $D^P$ -hard in the data complexity.*

The next theorem proves all the hardness results in Table 7 in *fp-combined* complexity. Also here, we provide a reduction from SAT-UNSAT. The set  $H$  includes a couple of facts allowing to recognize the unsatisfiability of  $\phi$  and  $\psi$ . We check that the fact associated with the unsatisfiability of  $\psi$  is necessary.

**Theorem 30.** *For any language  $\mathcal{L}$  considered here,  $\text{MINEX-ALLNEC}^{\neq}(\mathcal{L})$  is  $D^P$ -hard in the *fp-combined* complexity.*

The following theorem proves the  $D_2^P$ -hardness in the *ba-combined* complexity for all the languages we consider. The theorem then provides matching lower bounds for the  $D_2^P$ -completeness results in Table 7, while for the remaining languages we need tighter lower bounds shown next.

The proof uses a reduction from the  $D_2^P$ -complete problem  $\text{QBF}_{2,\forall,\neg}^{\text{CNF}} \wedge \text{QBF}_{2,\exists}^{\text{CNF}}$ : given two quantified Boolean formulas  $\Phi = \exists X \forall Y \neg \phi(X, Y)$  and  $\Psi = \forall Z \exists K \psi(Z, K)$ , decide whether  $\Phi$  and  $\Psi$  are both valid. In  $H$ , there are facts associated with the various valid truth assignments for the variables in  $X$  and  $Z$ . The test for the satisfiability of  $\phi(X, Y)$  and  $\psi(Z, K)$ , when a truth assignment for  $X$  and  $Z$  is encoded in a candidate explanation, is captured with two negative constraints. Additionally,  $H$  contains facts encoding a special truth assignment (using a special truth value) for the variables



in  $X$  and  $Z$  that allow us to bypass the satisfiability tests for  $\phi(X, Y)$  and  $\psi(Z, K)$ . Then,  $H'$  contains the facts encoding the special truth assignment for  $Z$ . All facts encoding valid truth assignments are not necessary. Thus,  $H'$  contain exactly all necessary facts iff all facts encoding the special truth assignment for  $Z$  are necessary and all facts encoding the special truth assignment for  $X$  are not necessary. Then,  $\Phi$  is valid iff we can build an explanation not using any of the facts encoding the special truth assignment for  $X$ ;  $\Psi$  is valid iff every explanation that we can build must contain all facts encoding the special truth assignment for  $Z$ .

**Theorem 31.** *For any language  $\mathcal{L}$  here considered, MINEX-ALLNEC $^{\neq}$ ( $\mathcal{L}$ ) is D $_2^P$ -hard in the *ba*-combined complexity.*

The remaining hardness results in the *ba*-combined complexity as well as all the hardness results in the combined complexity are established via the following theorem, which shows that MINEX-EXISTS $^{\neq}$  can be reduced to the complement of MINEX-ALLNEC $^{\neq}$ .

**Theorem 32.** *For any language  $\mathcal{L}$  considered here, MINEX-EXISTS $^{\neq}$ ( $\mathcal{L}$ ) is reducible in polynomial time to the complement of MINEX-ALLNEC $^{\neq}$ ( $\mathcal{L}$ ) in the data, *fp*-combined, *ba*-combined, and combined complexity.*

## 9 Related Work

The literature on explanations is very rich. Abstracting away from subtle differences, the study of explanations can be classified in accordance to the (i) underlying *logical formalism*, (ii) *reasoning task* to be explained (e.g., query answering and consistency checking), and (iii) *type of consequences* to be explained (e.g., entailments vs. non-entailments). We focus on existential rules as the underlying logical formalism, on OMQA as the reasoning task, and on non-entailments.

Our study may be seen as an instance of abductive reasoning, explaining observations in terms of sets of hypotheses that may have led to those observations. Abductive reasoning has been studied for several formalisms, such as propositional logic (Eiter and Gottlob 1995), logic programs (Eiter, Gottlob, and Leone 1997b), default theories (Eiter, Gottlob, and Leone 1997a), probabilistic temporal logic (Molinaro, Sliva, and Subrahmanian 2014), and DLs (Du et al. 2011; Du, Wang, and Shen 2014; Del-Pinto and Schmidt 2019; Wang et al. 2015; Klarman, Endriss, and Schlobach 2011; Calvanese et al. 2013). The study of ABox abduction by Klarman, Endriss, and Schlobach (2011) is related to our approach, as they also view a set of assertions (i.e., facts) as explanations. The closest work to ours is the one by Calvanese et al. (2013): for a given query that is *not* entailed from the knowledge base, the idea is to find a set of assertions such that, when they are added to the ABox, the entailment holds (avoiding inconsistencies). They consider arbitrary, inclusion-minimal, and cardinality-minimal explanations. Differently from their work, we focus on existential rules, and also consider novel explanation problems. One important difference is in the way abducibles are defined, which in turn affects how explanations are defined: in our case, explanations are subsets of a given set  $H$  of abducible *facts*; in contrast, they start with a set of abducible *predicates*, called *signature*, and explanations are sets of facts whose predicate belongs to the signature.

This formulation is suitable when abducible predicates are known, but relevant constants may not be known. The same problem formalization has been adopted in (Du et al. 2011; Du, Wang, and Shen 2014). The latter work also identified a class of TBoxes guaranteeing the existence of a finite number of minimal explanations (as there can be infinitely many minimal explanations). Furthermore, they introduced a special kind of minimal explanations, called *representative* explanations, from which all minimal explanations can be retrieved. Different works have focused on the efficient computation of explanations, e.g., see (Du et al. 2011; Du, Wang, and Shen 2014; Wang et al. 2015).

Explanations for positive entailments have been studied in DLs extensively; see, e.g., (McGuinness and Borgida 1995; Borgida, Franconi, and Horrocks 2000), where explanations are in the form of proofs. Most of the subsequent works focused on explanations, which are minimal subsets of the given theory (Schlobach and Cornet 2003). This is known as *axiom pinpointing* (Kalyanpur et al. 2007; Baader and Suntisrivaraporn 2008; Peñaloza and Sertkaya 2017), and such explanations are called *justifications* (Horridge, Parsia, and Sattler 2008). These works focus on DLs, and on explaining classical reasoning tasks. Explanations for (positive) OMQA has been first investigated for the *DL-Lite* family of languages (Borgida, Calvanese, and Rodriguez-Muro 2008). Recently, a thorough investigation is given by Ceylan et al. (2019) for existential rules and by Ceylan et al. (2020) for DLs. These works address the problem of explaining why a query is entailed, as opposed to our approach.

Explaining query answers has also been studied in the context of *inconsistency-tolerant query answering*. The problem has been addressed in (Du, Wang, and Shen 2015; Bienvenu, Bourgaux, and Goasdoué 2019) for DL languages and in (Hecham et al. 2017; Lukasiewicz, Malizia, and Molinaro 2020) for existential rule languages. The main difference between this paper and the aforementioned ones is that we provide explanations under standard semantics and for *consistent* knowledge bases.

## 10 Summary and Outlook

We have addressed the problem of explaining why a query is not entailed in OMQA under existential rules. (Minimal) Explanations for a non-entailment can be seen as a (minimal) set of facts such that, if the database were extended with this set, the query would be entailed, while the knowledge base would remain consistent with the added information. We have conducted a detailed complexity analysis for various explanation problems, for a wide range of existential rule languages under different complexity measures.

As for future work, it would be interesting to study our framework under different minimality criteria (e.g., by considering cardinality-minimal explanations) or to express other forms of preferences among explanations. Another interesting direction for future work is to study the problems of computing all explanations or relevant/necessary facts.

## Acknowledgments

This work was supported by the Alan Turing Institute under the UK EPSRC grant EP/N510129/1, the AXA Research Fund, and by the EPSRC grants EP/R013667/1, EP/L012138/1, and EP/M025268/1.

## References

- Baader, F., and Suntisrivaraporn, B. 2008. Debugging SNOMED CT using axiom pinpointing in the description logic  $\mathcal{EL}^+$ . In *Proc. KR-MED*, 1–7.
- Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F., eds. 2007. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2nd edition.
- Bienvenu, M.; Cate, B. T.; Lutz, C.; and Wolter, F. 2014. Ontology-based data access: A study through disjunctive Datalog, CSP, and MMSNP. *ACM Trans. Database Syst.* 39(4):33:1–33:44.
- Bienvenu, M.; Bourgaux, C.; and Goasdoué, F. 2019. Computing and explaining query answers over inconsistent *DL-Lite* knowledge bases. *J. Artif. Intell. Res.* 64:563–644.
- Borgida, A.; Calvanese, D.; and Rodriguez-Muro, M. 2008. Explanation in the *DL-Lite* family of description logics. In *Proc. OTM*, 1440–1457.
- Borgida, A.; Franconi, E.; and Horrocks, I. 2000. Explaining *ALC* subsumption. In *Proc. ECAI*, 209–213.
- Calì, A.; Gottlob, G.; and Kifer, M. 2013. Taming the infinite chase: Query answering under expressive relational constraints. *J. Artif. Intell. Res.* 48:115–174.
- Calì, A.; Gottlob, G.; and Lukasiewicz, T. 2012. A general Datalog-based framework for tractable query answering over ontologies. *J. Web Sem.* 14:57–83.
- Calì, A.; Gottlob, G.; and Pieris, A. 2012. Towards more expressive ontology languages: The query answering problem. *Artif. Intell.* 193:87–128.
- Calvanese, D.; Ortiz, M.; Simkus, M.; and Stefanoni, G. 2013. Reasoning about explanations for negative query answers in *DL-Lite*. *J. Artif. Intell. Res.* 48:635–669.
- Ceylan, İ. İ.; Lukasiewicz, T.; Malizia, E.; and Vaiceniavičius, A. 2019. Explanations for query answers under existential rules. In *Proc. IJCAI*, 1639–1646.
- Ceylan, İ. İ.; Lukasiewicz, T.; Malizia, E.; and Vaiceniavičius, A. 2020. Explanations for ontology-mediated query answering in description logics. In *Proc. ECAI*.
- Del-Pinto, W., and Schmidt, R. A. 2019. ABox abduction via forgetting in  $\mathcal{ALH}$ . In *Proc. AAAI*, 2768–2775.
- Du, J.; Qi, G.; Shen, Y.; and Pan, J. Z. 2011. Towards practical ABox abduction in large OWL DL ontologies. In *Proc. AAAI*, 1160–1165.
- Du, J.; Wang, K.; and Shen, Y. 2014. A tractable approach to abox abduction over description logic ontologies. In *Proc. AAAI*, 1034–1040.
- Du, J.; Wang, K.; and Shen, Y. 2015. Towards tractable and practical ABox abduction over inconsistent description logic ontologies. In *Proc. AAAI*, 1489–1495.
- Eiter, T., and Gottlob, G. 1995. The complexity of logic-based abduction. *J. ACM* 42(1):3–42.
- Eiter, T.; Gottlob, G.; and Leone, N. 1997a. Abduction from logic programs: Semantics and complexity. *Theor. Comput. Sci.* 189(1-2):129–177.
- Eiter, T.; Gottlob, G.; and Leone, N. 1997b. Semantics and complexity of abduction from default theories. *Artif. Intell.* 90(1-2):177–223.
- Eiter, T.; Lukasiewicz, T.; and Predoiu, L. 2016. Generalized consistent query answering under existential rules. In *Proc. KR*, 359–368.
- Fagin, R.; Kolaitis, P. G.; Miller, R. J.; and Popa, L. 2005. Data exchange: Semantics and query answering. *Theor. Comput. Sci.* 336(1):89–124.
- Hecham, A.; Arioua, A.; Stapleton, G.; and Croitoru, M. 2017. An empirical evaluation of argumentation in explaining inconsistency-tolerant query answering. In *Proc. DL*.
- Horridge, M.; Parsia, B.; and Sattler, U. 2008. Laconic and precise justifications in OWL. In *Proc. ISWC*, 323–338.
- Kalyanpur, A.; Parsia, B.; Horridge, M.; and Sirin, E. 2007. Finding all justifications of OWL DL entailments. In *Proc. ISWC/ASWC*, 267–280.
- Klarman, S.; Endriss, U.; and Schlobach, S. 2011. ABox abduction in the description logic *ALC*. *J. Autom. Reasoning* 46(1):43–80.
- Lukasiewicz, T., and Malizia, E. 2017. A novel characterization of the complexity class  $\Theta_k^P$  based on counting and comparison. *Theor. Comput. Sci.* 694:21–33.
- Lukasiewicz, T.; Malizia, E.; and Molinaro, C. 2020. Explanations for inconsistency-tolerant query answering under existential rules. In *Proc. AAAI*, 2909–2916.
- McGuinness, D. L., and Borgida, A. 1995. Explaining subsumption in description logics. In *Proc. IJCAI*, 816–821.
- Molinaro, C.; Sliva, A.; and Subrahmanian, V. S. 2014. Super-solutions: Succinctly representing solutions in abductive annotated probabilistic temporal logic. *ACM Trans. Comput. Log.* 15(3):18:1–18:35.
- Peñaloza, R., and Sertkaya, B. 2017. Understanding the complexity of axiom pinpointing in lightweight description logics. *Artif. Intell.* 250:80–104.
- Schlobach, S., and Cornet, R. 2003. Non-standard reasoning services for the debugging of description logic terminologies. In *Proc. IJCAI*, 355–360.
- Vardi, M. Y. 1982. The complexity of relational query languages. In *Proc. STOC*, 137–146.
- Wang, Z.; Chitsaz, M.; Wang, K.; and Du, J. 2015. Towards scalable and complete query explanation with OWL 2 EL ontologies. In *Proc. CIKM*, 743–752.