

Computing Approximate Query Answers over Inconsistent Knowledge Bases

Sergio Greco, Cristian Molinaro, Irina Trubitsyna
 University of Calabria, Italy
 {greco,cmolinaro,trubitsyna}@dimes.unical.it

Abstract

Consistent query answering is a principled approach for querying inconsistent knowledge bases. It relies on the notion of a *repair*, that is, a maximal consistent subset of the facts in the knowledge base. One drawback of this approach is that entire facts are deleted to resolve inconsistency, even if they may still contain useful “reliable” information. To overcome this limitation, we propose a new notion of repair allowing values within facts to be updated for restoring consistency. This more fine-grained repair primitive allows us to preserve more information in the knowledge base. We also introduce the notion of a *universal repair*, which is a compact representation of all repairs. Then, we show that consistent query answering in our framework is intractable (coNP-complete). In light of this result, we develop a polynomial time approximation algorithm for computing a sound (but possibly incomplete) set of consistent query answers.

1 Introduction

Reasoning in the presence of inconsistent information is a problem that has attracted much interest in the last decades. Many inconsistency-tolerant semantics for query answering have been proposed, and most of them rely on the notions of *consistent query answer* and *repair*. A consistent answer to a query is a query answer that is entailed by every repair, where a repair is a “maximal” consistent subset of the facts of the knowledge base. Different maximality criteria have been investigated, but all the resulting notions of repair share the same drawback: a fact is either kept or deleted altogether, and deleting entire facts can cause loss of “reliable” information.

Example 1. Consider the knowledge base (D, Σ) where D contains the following facts:

works		
john	cs	nyc
john	math	rome
mary	math	sydney

and Σ is an ontology consisting of the following equality-generating dependency (EGD) σ :

$$\text{works}(E_1, D, C_1) \wedge \text{works}(E_2, D, C_2) \rightarrow C_1 = C_2.$$

As an example, the fact $\text{works}(\text{john}, \text{cs}, \text{nyc})$ states that john is an employee working in the cs department located in nyc. The dependency σ says that every department must be located in a single city. Clearly, the last two facts violate σ , so every repair would discard either of them.

If we pose a query asking for the employees’ name, the only consistent answer is john. However, intuitively, we might consider reliable the information on mary being an employee, as the only uncertainty concerns the math department and its city—roughly speaking, the information in the first column of the works table can be considered “clean”. Dropping entire facts causes loss of information.

To overcome the drawback illustrated above, we propose a notion of repair based on updating values within facts. Update-based repairing allows for rectifying errors in facts without deleting them altogether, thereby preserving consistent values.

Example 2. Consider again the knowledge base of Example 1. Using value updates as the primitive to restore consistency, and assuming that the only uncertain values are math’s cities, we get the following two repairs:

john	cs	nyc	john	cs	nyc
john	math	rome	john	math	sydney
mary	math	rome	mary	math	sydney

If we ask again for the employees’ name, both mary and john are consistent answers.

We consider knowledge bases where ontologies are expressed through (a particular class of) EGDs. Equality-generating dependencies are one of the two major types of data dependencies—the other major type consists of tuple-generating dependencies (TGDs)—and can model several kinds of constraints commonly arising in practice, such as functional dependencies and thus also key dependencies.

We show that consistent query answering in this setting is coNP-complete. Then, we show how to compute a “universal” repair, which compactly represents all repairs and can be computed in polynomial time. A universal repair is a valuable tool to compute approximate query answers. We propose a polynomial time approximation to compute a sound (but possibly incomplete) set of consistent query answers. The basic idea is illustrated in the following example.

Example 3. A universal repair for the two repairs of Example 2 is reported below:

john	cs	nyc
john	math	\perp_1
mary	math	\perp_1

$\perp_1 = \text{rome} \vee \perp_1 = \text{sydney}$

where \perp_1 is a labeled null, and the “global” condition at the bottom restricts the admissible values for \perp_1 , which are rome and sydney. We exploit such a representation for query answering, by combining the condition of the universal repair with provenance information during query evaluation. For instance, if we ask for the departments that are not located in nyc, we get math with condition $\perp_1 \neq \text{nyc}$, which combined with the global condition allows us to conclude that math is a consistent answer (as its city is either rome or sydney, and thus cannot be nyc for sure).

As shown in the following, leveraging universal repairs, our approximation schema yields strictly better results than well-known approximation approaches such as the *intersection of repairs* and the *intersection of closed repairs* semantics [Lembo *et al.*, 2010; Bienvenu *et al.*, 2014].

Below we provide a further example showing that even when knowledge bases are expressed by unary/binary facts only, our approach can preserve more information than the classic notion of repair.

Example 4. Consider the knowledge base (D, Σ) where

$$D = \{\text{emp}(\text{john}, \text{sydney}), \text{stud}(\text{john}, \text{nyc})\} \text{ and} \\ \Sigma = \{\text{emp}(N, C) \wedge \text{stud}(N, C') \rightarrow C = C'\}.$$

Here $\text{emp}(\text{john}, \text{sydney})$ means that john is an employee living in sydney, and $\text{stud}(\text{john}, \text{nyc})$ means that john is a student living in nyc.

The EGD says that every person must be associated with the same city in emp and stud. Clearly, the knowledge base is inconsistent.

Consider now the query asking for the people who are both an employee and a student. Assuming that the only uncertain values are john’s cities, we can still say that john is certainly an employee and a student, even if we are not sure where he lives. This information is preserved by our framework, as john is a consistent answer, but it is completely lost using the classic notion of repair, as there are no consistent answers. In fact, in our framework, there are two repairs obtained from D by changing either sydney into nyc or nyc into sydney—in both of them john is an employee and a student. According to the classical notion of repair, there are two repairs obtained from D by deleting either $\text{emp}(\text{john}, \text{sydney})$ or $\text{stud}(\text{john}, \text{nyc})$ —in the former john is a student but not an employee while in the latter he is an employee but not a student.

Contributions.

We consider the problem of querying possibly inconsistent knowledge bases in the presence of *acyclic* sets of EGDs. In this setting, we propose a new notion of repair based on a more fine-grained repair primitive, which updates values within facts rather than deleting them altogether, enabling us

to preserve more information of the knowledge base. For such a repair strategy, we introduce the notion of a *universal repair*, which is a compact representation of all repairs, can be computed in polynomial time, and can be leveraged to compute consistent query answers.

We then develop a chase-like procedure to compute a universal repair, introducing a formalism that augments instances with provenance information. We show that consistent query answering in our framework is coNP-complete (data complexity). In light of this, we leverage universal repairs and provenance information to develop an approximation algorithm that provides a sound (but possibly incomplete) set of consistent query answers in polynomial time.

Organization.

Section 2 discusses related work. Preliminaries are reported in Section 3. Our inconsistency-tolerant semantics and its complexity are investigated in Section 4. A compact representation of all repairs and its computation are proposed in Section 5. Query evaluation is addressed in Section 6. Our approximation algorithm is introduced in Section 7. Conclusions and directions for future work are reported in Section 8.

2 Related Work

Reasoning in the presence of inconsistent information is a problem that has attracted a great deal of interest in the AI and database communities. Consistent query answering was first proposed in [Arenas *et al.*, 1999]. Query answering under various inconsistency-tolerant semantics for ontologies expressed in DL languages has been studied in [Lembo *et al.*, 2010; 2011; 2015; Bienvenu, 2011; 2012; Rosati, 2011; Bienvenu and Rosati, 2013], and in [Lukasiewicz *et al.*, 2012b; 2015; 2012a] for ontologies expressed by fragments of Datalog+/- . Several notions of maximality for a repair have been considered in [Bienvenu *et al.*, 2014]. Bienvenu and Rosati (2013) have proposed an approach for the approximation of consistent query answers from above and from below. Furfaro *et al.* (2007) proposed an approach based on three-valued logic to compute a sound but possibly incomplete set of consistent query answers.

Different from our proposal, all the approaches above adopt the most common notion of repair, where whole facts are removed. This can cause loss of information, as illustrated in the toy scenario of Example 1—in real-life scenarios, it might well be the case that facts have much more attributes and only a few of them are involved in inconsistencies, leading to significant loss of useful data.

There have also been different proposals adopting a notion of repair that allows values to be updated [Bohannon *et al.*, 2005; Bertossi *et al.*, 2008; Greco and Molinaro, 2008; 2012; Flesca *et al.*, 2010]. Our repair strategy behaves similar to the one of [Bohannon *et al.*, 2005; Greco and Molinaro, 2008; 2012] in that values on the right-hand side of functional dependencies (FDs) are updated. However, those works focus on FDs only. [Bertossi *et al.*, 2008] allow only numerical attributes to be updated, one primary key per relation is allowed, but keys are assumed to be satisfied by the original database: thus, no repairing is possible w.r.t. keys, while we allow it, and we allow much more general constraints. Our

repair strategy can be seen as an instantiation of the value-based family of policies proposed in [Martinez *et al.*, 2014], even though the two approaches differ in how multiple dependencies are handled, and they focus on FDs only. Flesca *et al.* (2010) consider numerical databases and a different class of (aggregate) constraints.

The main difference between this paper and the aforementioned ones is that none of them has investigated the approximate computation of consistent query answers—notice that [Bertossi *et al.*, 2008; Bohannon *et al.*, 2005] have proposed approximation algorithms for computing a repair with minimum distance from the original database, and [Greco and Molinaro, 2008] consider approximate probabilistic query answers. Also, we introduce the notion of a *universal repair*, which compactly represents all repairs and can be used for exact/approximate query answering.

Approximation algorithms for computing sound but possibly incomplete sets of query answers in the presence of nulls have been proposed in [Guagliardo and Libkin, 2016; Libkin, 2015; 2016; Greco *et al.*, 2017; Fiorentino *et al.*, 2018], but no dependencies are considered therein, and thus the database is assumed to be consistent.

3 Preliminaries

We assume the existence of the following pairwise disjoint (countably infinite) sets: a set Const of *constants*, a set Var of *variables*, and a set Null of *labeled nulls*. Nulls are denoted by the symbol \perp subscripted. A *term* is a constant, variable, or null. We also assume a set of *predicates*, disjoint from the aforementioned sets, with each predicate being associated with an *arity*, which is a non-negative integer.

An *atom* A is of the form $p(t_1, \dots, t_n)$, where p is an n -ary predicate and the t_i 's are terms. We write an atom also as $p(\mathbf{t})$, where \mathbf{t} is a sequence of terms. An atom without variables is also called a *fact*. An *instance* is a finite set of facts. A *database* is an instance containing constants only.

A *homomorphism* is a mapping $h : \text{Const} \cup \text{Var} \cup \text{Null} \rightarrow \text{Const} \cup \text{Var} \cup \text{Null}$ that is the identity on Const . Homomorphisms are also applied to atoms and set of atoms in the natural fashion, that is, $h(p(t_1, \dots, t_n)) = p(h(t_1), \dots, h(t_n))$, and $h(S) = \{h(A) \mid A \in S\}$ for any set S of atoms. A *valuation* is a homomorphism ν whose image is Const , that is, $\nu(t) \in \text{Const}$ for every $t \in \text{Const} \cup \text{Var} \cup \text{Null}$.

Conditional instances.

Conditional instances (also known as “conditional tables” [Imielinski and Lipski, 1984; Grahne, 1991]) are instances augmented with conditions restricting the set of admissible values for nulls. Let \mathcal{E} be the set of all expressions, called *conditions*, that can be built using the standard logical connectives $\wedge, \vee, \neg, \Rightarrow$ and expressions of the form $t_i = t_j$, true, and false, where $t_i, t_j \in \text{Const} \cup \text{Null}$. We will also use $t_i \neq t_j$ as a shorthand for $\neg(t_i = t_j)$. We say that a valuation ν *satisfies* a condition ϕ , denoted $\nu \models \phi$, if its assignment of constants to nulls makes ϕ true. Formally, a *conditional instance* (CI) is a pair $\langle I, \Phi \rangle$, where I is an instance and $\Phi \in \mathcal{E}$. The semantics of $C = \langle I, \Phi \rangle$ is given by the set of its *possible worlds*, that is, the set of databases $pw(C) = \{\nu(I) \mid \nu \text{ is a valuation and } \nu \models \Phi\}$.

Equality generating dependencies.

An *equality generating dependency* (EGD) σ is a first-order formula of the form $\forall \mathbf{x} \varphi(\mathbf{x}) \rightarrow x_i = x_j$, where $\varphi(\mathbf{x})$ is a conjunction of atoms (without labeled nulls) whose variables are exactly \mathbf{x} , and x_i and x_j are variables from \mathbf{x} . We call $\varphi(\mathbf{x})$ the *body* of σ , and call $x_i = x_j$ the *head* of σ . We will omit the universal quantification in front of dependencies and assume that all variables are universally quantified. With a slight abuse of notation, we sometimes treat a conjunction as the *set* of its atoms. An instance I *satisfies* σ , denoted $I \models \sigma$, if whenever there exists a homomorphism h s.t. $h(\varphi(\mathbf{x})) \subseteq I$, then $h(x_i) = h(x_j)$. A instance I *satisfies* a set Σ of EGDs, denoted $I \models \Sigma$, if $I \models \sigma$ for every $\sigma \in \Sigma$.

A *knowledge base* (KB) is a pair (D, Σ) , where D is a database and Σ is a finite set of EGDs. It is *consistent* if $D \models \Sigma$, otherwise it is *inconsistent*.

Query language.

The query language we consider is non-recursive safe Datalog with negation. This choice will ease presentation of our approximation algorithm.

A (*positive*) *rule* is of the form $p(\mathbf{x}) \leftarrow \phi(\mathbf{x}, \mathbf{y}), \gamma(\mathbf{z})$, where there are no nulls, $p(\mathbf{x})$ is an atom (whose variables are \mathbf{x}), $\phi(\mathbf{x}, \mathbf{y})$ is a conjunction of atoms (whose variables are \mathbf{x} and \mathbf{y}), $\gamma(\mathbf{z})$ is a conjunction of expressions (whose variables are \mathbf{z}) of the form $t = t'$ or $t \neq t'$ with $t, t' \in \text{Const} \cup \text{Var}$, and $\mathbf{z} \subseteq \mathbf{x} \cup \mathbf{y}$. W.l.o.g. we assume that there cannot be multiple occurrences of the same variable in $\phi(\mathbf{x}, \mathbf{y})$ —so equalities must be made explicit in $\gamma(\mathbf{z})$.

A (*negative*) *rule* is of the form $p(\mathbf{x}) \leftarrow p'(\mathbf{x}), \neg p''(\mathbf{x})$, where $p(\mathbf{x})$, $p'(\mathbf{x})$, and $p''(\mathbf{x})$ are atoms without nulls. W.l.o.g. we assume that there cannot be multiple occurrences of the same variable in $p'(\mathbf{x})$.

In the rules above, $p(\mathbf{x})$ is called the *head* of the rule, and $\phi(\mathbf{x}, \mathbf{y}), \gamma(\mathbf{z})$ (resp. $p'(\mathbf{x}), \neg p''(\mathbf{x})$) is called the *body*.

A *program* P is a finite set of rules. The *predicate graph* of P is a directed graph whose vertices are the predicates appearing in P , and there is a directed edge from a predicate p' to a predicate p if there is a rule in P where p' appears in the body and p appears in the head. As we consider non-recursive Datalog, the predicates of P can be ordered into a sequence $\langle p_1, \dots, p_n \rangle$ according to a topological sorting of the predicate graph. Given a predicate p appearing in P , we use $P[p]$ to denote the set of all rules in P having p in the head.

The *immediate consequence operator* T_P of a program P is defined as follows. For every database D ,

$$T_P(D) = D \cup \{ \nu(p(\mathbf{x})) \mid \begin{array}{l} \text{there exist a valuation } \nu \text{ and} \\ \text{a rule } p(\mathbf{x}) \leftarrow \phi(\mathbf{x}, \mathbf{y}), \gamma(\mathbf{z}) \text{ of } P \text{ s.t.} \\ \nu(\phi(\mathbf{x}, \mathbf{y})) \subseteq D \text{ and } \nu \models \gamma(\mathbf{z}) \} \cup \\ \{ \nu(p(\mathbf{x})) \mid \begin{array}{l} \text{there exist a valuation } \nu \text{ and} \\ \text{a rule } p(\mathbf{x}) \leftarrow p'(\mathbf{x}), \neg p''(\mathbf{x}) \text{ of } P \text{ s.t.} \\ \nu(p'(\mathbf{x})) \in D \text{ and } \nu(p''(\mathbf{x})) \notin D \} \}. \end{array}$$

Given a database D and a program P , we define:

$$\begin{aligned} T_P^0(D) &= D, \\ T_P^i(D) &= T_{P[p_i]}(T_P^{i-1}(D)), \quad \text{for } 1 \leq i \leq n, \\ \text{CN}(P, D) &= T_P^n(D). \end{aligned}$$

A *query* Q is a pair (P, p) , where P is a program and p is a predicate. The result of evaluating Q on a database D is $Q(D) = \{p(\mathbf{t}) \mid p(\mathbf{t}) \in \text{CN}(P, D)\}$.

4 Repairing and Querying Inconsistent KBs

In this section, we present our notion of repair and analyze the computational complexity of two central problems: repair checking and consistent query answering.

We start by defining the EGDs we consider. Let Σ be a set of EGDs. An *argument* of Σ is an expression of the form $p[i]$, where p is an n -ary predicate appearing in Σ and $1 \leq i \leq n$.

Definition 1 (Argument and dependency graphs). *The argument graph of a set Σ of EGDs is a directed graph $G_\Sigma = (V, E)$, where V is the set of all arguments of Σ , and E contains a directed edge from $p[i]$ to $q[j]$ labeled σ iff there is an EGD $\sigma \in \Sigma$ such that:*

- the body of σ contains an atom $p(t_1, \dots, t_n)$ such that either t_i is a constant or t_i is a variable occurring more than once in the body of σ , and
- the body of σ contains an atom $q(u_1, \dots, u_m)$ such that u_j is a variable also appearing in the head of σ .

The dependency graph of Σ is a directed graph $\Gamma_\Sigma = (\Sigma, \Omega)$, where Ω is the following set:

$$\{(\sigma_1, \sigma_2) \mid G_\Sigma = (V, E) \wedge (p[i], q[j], \sigma_1), (q[j], r[k], \sigma_2) \in E\}.$$

We say that Σ is acyclic if its dependency graph is acyclic.

In the following, we write $\sigma_i < \sigma_j$ if $(\sigma_i, \sigma_j) \in \Omega$ or there is σ_k such that $\sigma_i < \sigma_k$ and $\sigma_k < \sigma_j$.

Example 5. Consider the following set of EGDs Σ :

$$\begin{aligned} \sigma_1: & \text{works}(X, Y_1, Z_1) \wedge \text{works}(X, Y_2, Z_2) \rightarrow Y_1 = Y_2 \\ \sigma_2: & \text{works}(X_1, Y, Z_1) \wedge \text{works}(X_2, Y, Z_2) \rightarrow Z_1 = Z_2 \end{aligned}$$

G_Σ has two edges: one from $\text{works}[1]$ to $\text{works}[2]$ labeled σ_1 , and another one from $\text{works}[2]$ to $\text{works}[3]$ labeled σ_2 . Then, Γ_Σ has only the edge (σ_1, σ_2) . Thus, Σ is acyclic.

In the rest of the paper we consider acyclic sets of EGDs—so, from now on, Σ is understood to be acyclic. Even if the tools developed in the paper can be generalized to arbitrary EGDs, we focus on acyclic sets because for them we can easily compute a compact representation of all repairs in polynomial time. This is particularly important for our purpose of developing polynomial time approximation algorithms.

Before introducing our notion of repair, we provide some intuitions. As already mentioned, it is based on updating values within facts. Specifically, we adopt a chase-like procedure that acts as follows: whenever an EGD $\varphi(\mathbf{x}) \rightarrow x_i = x_j$ is not satisfied by a set of facts $\nu(\varphi(\mathbf{x}))$ (for some valuation ν), and thus $\nu(x_i) \neq \nu(x_j)$, then either $\nu(x_i)$ replaces $\nu(x_j)$ or vice versa. A repair is obtained through an exhaustive application of this repair step guaranteeing that the set of changes is minimal. Note that there is a non-deterministic choice to be made when updating values, and this may lead to multiple repairs.

Example 6. Consider the database below and the set of EGDs of Example 5.

john	cs	rome
john	math	rome
mary	math	sydney

By enforcing σ_1 into the first two facts, either *cs* or *math* can be chosen as *john*'s department. If the latter is chosen, then the database D' below is obtained.

Suppose now that σ_2 is enforced into the last two facts of D' . Then, either *rome* or *sydney* can be chosen as *math*'s city. If the former is chosen, then the database D'' below is obtained.

$$D' = \begin{array}{|c|c|c|} \hline \text{john} & \text{math} & \text{rome} \\ \hline \text{john} & \text{math} & \text{rome} \\ \hline \text{mary} & \text{math} & \text{sydney} \\ \hline \end{array} \quad D'' = \begin{array}{|c|c|c|} \hline \text{john} & \text{math} & \text{rome} \\ \hline \text{john} & \text{math} & \text{rome} \\ \hline \text{mary} & \text{math} & \text{rome} \\ \hline \end{array}$$

No further dependency enforcement is applicable at this point and thus D'' is a repair.

Notice that it might be possible to restore consistency in other different ways. For instance, in the first step of the example above, one may modify the employee names. However, we do not consider this option because it is unclear which (different) values should be assigned (any constant in *Const* is a candidate value). For instance, *john* in the first fact might be replaced with *rome*, but this is somewhat arbitrary and indeed does not make much sense. In contrast, our repair strategy chooses candidate values that are in sense “justified” by the content of the database (e.g., in the example above, *john* works for either the *cs* or the *math* department). Moreover, when EGDs are key dependencies, the aforementioned way of restoring consistency may lead to the introduction of entities that are not meaningful. Indeed, our choice has been made by different approaches relying on value updates (e.g., [Bohannon *et al.*, 2005]).

The computation of repairs as informally described above is quite involved in that, when applying a repair step, we have to check that the current database has not been previously obtained (to avoid infinite repair steps) and check minimality of the changes performed. In the following, we show how to deal with such issues and compute repairs efficiently.

When repairing inconsistent databases, conditional instances can be used to keep track of which values must be equal and which constants can be assigned to nulls. For instance, in Example 6 above, after the last dependency enforcement, we have to make sure that the last column contains the same city (which can be either *rome* or *sydney*).

Below we introduce the technical definitions, starting with the notion of a *repair step*.

Definition 2 (Repair step). *Let $\langle I, \Phi \rangle$ be a CI, Σ a set of EGDs, and σ an EGD $\varphi(\mathbf{x}) \rightarrow x_i = x_j \in \Sigma$. Let h be a homomorphism s.t. $h(\varphi(\mathbf{x})) \subseteq h(I)$, $h \models \Phi$, and $h(x_i) \neq h(x_j)$.*

Moreover, let \perp_m be a fresh null and $\langle I', \Phi' \rangle$ the CI obtained from $\langle I, \Phi \rangle$ as follows:

1. *for each fact $p(u_1, \dots, u_n)$ in I , if $\varphi(\mathbf{x})$ contains an atom $p(t_1, \dots, t_n)$ s.t. $h(p(t_1, \dots, t_n)) = h(p(u_1, \dots, u_n))$, then for every $1 \leq k \leq n$, if $t_k \in \{x_i, x_j\}$,*
 - *replace u_k in $p(u_1, \dots, u_n)$ with \perp_m ;*
 - *if $u_k \in \text{Null}$, replace every occurrence of u_k elsewhere with \perp_m ;*
2. *Either $\Phi' = \Phi \wedge (\perp_m = h(x_i))$ or $\Phi' = \Phi \wedge (\perp_m = h(x_j))$.*

We say that $\langle I, \Phi \rangle \xrightarrow{\sigma, h} \langle I', \Phi' \rangle$ is a repair step.

Intuitively, a repair step enforces an EGD that is not satisfied by the knowledge base.

A *repair sequence* of a knowledge base (D, Σ) is a (possibly empty) finite sequence of repair steps $C_i \xrightarrow{\sigma_i, h_i} C_{i+1}$ ($0 \leq i < m$) such that $C_0 = \langle D, \text{true} \rangle$ and $\sigma_i \in \Sigma$ for every i . We also say that the repair sequence is *from* C_0 *to* C_m . We call C_m the *result* of the repair sequence. Also, we say that the repair sequence is *maximal* if there does not exist a repair step of the form $C_m \xrightarrow{\sigma_m, h_m} C_{m+1}$.

It is easy to see that for each repair step $\langle I, \Phi \rangle \xrightarrow{\sigma, h} \langle I', \Phi' \rangle$, if h' is a homomorphism s.t. $h' \models \Phi'$, then it assigns constants to nulls as dictated by Φ' (see item 2 of Definition 2). Thus, all h' satisfying Φ' yield the same database $h'(I')$. For a conditional instance (I, Φ) obtained by means of a repair sequence, we use the notation $\Phi(I)$ to denote the database derived from I by iteratively replacing nulls with constants or other nulls as dictated by Φ .

For ease of presentation, we assume that one can keep of a given fact f in the database during the repair process despite the value changes. Thus, we use $D[f, i]$ to denote the i -th term of a fact f in a database D . Given a repair sequence S from $C_0 = (D, \text{true})$ to $C_m = (I_m, \Phi_m)$, we define the set of changes made by S as $\text{update}(S) = \{(f, i) \mid D[f, i] \neq \Phi_m(I_m)[f, i]\}$, where f is a fact of the database.

Example 7. Consider the database of Example 6 and the set of EGDs of Example 5. By enforcing first σ_1 using the first two facts, then σ_2 using the last two facts, and finally σ_2 using the first two facts, we get, respectively, the CIs C_1 , C_2 and C_3 reported below:

C_1	C_2	C_3																											
<table style="width: 100%; border-collapse: collapse;"> <tr><td>john</td><td>\perp_1</td><td>rome</td></tr> <tr><td>john</td><td>\perp_1</td><td>rome</td></tr> <tr><td>mary</td><td>math</td><td>sydney</td></tr> </table>	john	\perp_1	rome	john	\perp_1	rome	mary	math	sydney	<table style="width: 100%; border-collapse: collapse;"> <tr><td>john</td><td>\perp_1</td><td>rome</td></tr> <tr><td>john</td><td>\perp_1</td><td>\perp_3</td></tr> <tr><td>mary</td><td>math</td><td>\perp_3</td></tr> </table>	john	\perp_1	rome	john	\perp_1	\perp_3	mary	math	\perp_3	<table style="width: 100%; border-collapse: collapse;"> <tr><td>john</td><td>\perp_1</td><td>\perp_4</td></tr> <tr><td>john</td><td>\perp_1</td><td>\perp_4</td></tr> <tr><td>mary</td><td>math</td><td>\perp_4</td></tr> </table>	john	\perp_1	\perp_4	john	\perp_1	\perp_4	mary	math	\perp_4
john	\perp_1	rome																											
john	\perp_1	rome																											
mary	math	sydney																											
john	\perp_1	rome																											
john	\perp_1	\perp_3																											
mary	math	\perp_3																											
john	\perp_1	\perp_4																											
john	\perp_1	\perp_4																											
mary	math	\perp_4																											
$\perp_1 = \text{math}$	$\perp_1 = \text{math} \wedge$ $\perp_3 = \text{sydney}$	$\perp_1 = \text{math} \wedge$ $\perp_3 = \text{sydney} \wedge$ $\perp_4 = \text{rome}$																											

No further repair steps are applicable at this point. Notice that $\perp_3 = \text{sydney}$ can be deleted from the condition of C_3 because it does not play a role anymore. A repair is obtained from C_3 by replacing every occurrence of \perp_1 with math, and every occurrence of \perp_4 with rome.

Definition 3 (Repair). A repair for a knowledge base $K = \langle D, \Sigma \rangle$ is a database $\Phi(J)$ such that there exists a maximal repair sequence S of K from $\langle D, \text{true} \rangle$ to $\langle J, \Phi \rangle$ and $\text{update}(S)$ is minimal w.r.t. \subseteq .

We use $\text{repair}(D, \Sigma)$ to denote the set of all repairs of a knowledge base (D, Σ) . The following simple proposition states that every repair is indeed consistent.

Proposition 1. Let (D, Σ) be a knowledge base. For every $D' \in \text{repair}(D, \Sigma)$, $D' \models \Sigma$.

The consistent answers to a query are defined in the standard way as follows.

Definition 4 (Consistent Query Answers). Let (D, Σ) be a knowledge base and Q a query. The consistent answers to Q over (D, Σ) are

$$Q(D, \Sigma) = \bigcap \{Q(D') \mid D' \in \text{repair}(D, \Sigma)\}.$$

In the context of managing inconsistent knowledge bases, two fundamental problems are *repair checking* and *consistent query answering*, which are defined as follows. Let (D, Σ) be a knowledge base, D' a database, Q a query, and f a fact of constants. Then, the following two problems are defined:

- *repair checking*: decide whether $D' \in \text{repair}(D, \Sigma)$;
- *consistent query answering*: decide whether $f \in Q(D, \Sigma)$.

Below we show that the first problem is in PTIME and the second one is coNP-complete. All the complexity results in this paper are in the data complexity, that is, the query and the EGDs are fixed.

Theorem 2. *Repair checking is in PTIME.*

Proof. (Sketch) Let (D, Σ) be a knowledge base and D' a database. As stated in Theorems 4 and 5 reported in the next section, it is possible to compute in polynomial time a CI $C = \langle I, \Phi \rangle$ such that $\text{pw}(C) = \text{repair}(D, \Sigma)$. After C is computed, for each null \perp_i in I , replace every occurrence of \perp_i in I and Φ with the corresponding constant in D' —if there are different constants corresponding to (different occurrences of) \perp_i , then answer ‘no’. After that, if the updated instance I is equal to D' and the updated condition Φ is true, then answer ‘yes’, otherwise answer ‘no’. Clearly, the entire process can be carried out in polynomial time. \square

Theorem 3. *Consistent query answering is coNP-complete.*

Proof. (Sketch) Membership. Let (D, Σ) be a knowledge base, Q a query, and f a fact of constants. The complementary problem is in NP. In fact, we can guess a database D' and verify that (i) $D' \in \text{repair}(D, \Sigma)$ and (ii) $f \notin Q(D')$. By the definition of repair, the size of D' is linear in the size of D . By Theorem 2, condition (i) can be verified in polynomial time, and obviously condition (ii) too.

Hardness. We reduce 3DNF TAUTOLOGY to our problem. Let ϕ be a 3DNF formula $d_1 \vee \dots \vee d_m$ over variables $X = \{x_1, \dots, x_n\}$. For every literal ℓ in ϕ , we define $\text{var}(\ell) = x_i$ if $\ell = x_i$ or $\ell = \neg x_i$, $\text{tval}(\ell) = \text{t}$ if $\ell = x_i$, and $\text{tval}(\ell) = \text{f}$ if $\ell = \neg x_i$, where t, f, and the x_i 's are constants.

We define D_ϕ as follows. For each disjunct $d_i = (\ell_1 \wedge \ell_2 \wedge \ell_3)$, $1 \leq i \leq m$, D_ϕ contains the following facts:

$$\begin{aligned} & \text{disj}(i, \text{var}(\ell_1), \text{t}, \text{var}(\ell_2), \text{t}, \text{var}(\ell_3), \text{t}), \\ & \text{disj}(i, \text{var}(\ell_1), \text{f}, \text{var}(\ell_2), \text{f}, \text{var}(\ell_3), \text{f}), \\ & \text{disjsat}(i, \text{var}(\ell_1), \text{tval}(\ell_1), \text{var}(\ell_2), \text{tval}(\ell_2), \text{var}(\ell_3), \text{tval}(\ell_3)). \end{aligned}$$

Then, Σ contains the following EGDs:

$$\begin{aligned} & \text{disj}(l, V, A_1, V_2, A_2, V_3, A_3) \wedge \\ & \text{disj}(l', V, A'_1, V'_2, A'_2, V'_3, A'_3) \rightarrow A_1 = A'_1 \\ & \text{disj}(l, V, A_1, V_2, A_2, V_3, A_3) \wedge \\ & \text{disj}(l', V'_1, A'_1, V, A_2, V'_3, A'_3) \rightarrow A_1 = A'_2 \\ & \text{disj}(l, V, A_1, V_2, A_2, V_3, A_3) \wedge \\ & \text{disj}(l', V'_1, A'_1, V'_2, A'_2, V, A_3) \rightarrow A_1 = A'_3 \\ & \text{disj}(l, V_1, A_1, V, A_2, V_3, A_3) \wedge \\ & \text{disj}(l', V'_1, A'_1, V, A'_2, V'_3, A'_3) \rightarrow A_2 = A'_2 \\ & \text{disj}(l, V_1, A_1, V, A_2, V_3, A_3) \wedge \\ & \text{disj}(l', V'_1, A'_1, V'_2, A'_2, V, A_3) \rightarrow A_2 = A'_3 \\ & \text{disj}(l, V_1, A_1, V_2, A_2, V, A_3) \wedge \\ & \text{disj}(l', V'_1, A'_1, V'_2, A'_2, V, A_3) \rightarrow A_3 = A'_3 \end{aligned}$$

It can be easily verified that Σ is acyclic.

Finally, $Q = (P, \text{sat})$, where P contains the following rule:

$$\text{sat}() \leftarrow \text{disjsat}(I, V_1, A_1, V_2, A_2, V_3, A_3), \\ \text{disj}(I, V_1, A_1, V_2, A_2, V_3, A_3).$$

It can be shown that ϕ is a tautology iff $\text{sat}() \in Q(D_\phi, \Sigma)$. \square

Indeed, coNP-hardness holds even for Boolean conjunctive queries. In the next section we introduce the notion of a *universal repair*, which is a compact representation of all repairs for a given knowledge base, and can be computed in polynomial time.

5 Universal Repair

In this section, we show how to compute a conditional instance U , called *universal repair*, that represents all repairs of a knowledge base, that is, the possible worlds of U are exactly the repairs. As shown in the following, such a representation will prove to be a valuable tool for query answering.

Roughly speaking, we generalize the repair step introduced in the previous section in order to keep track of all possible choices, and when each choice should be applied.

Any conjunction of atoms $\varphi(\mathbf{x})$ in the body of an EGD can be rewritten as follows:

- every occurrence of a constant c in $\varphi(\mathbf{x})$ is replaced with a fresh variable z and $(z = c)$ is added to the conjunction;
- for every variable occurring $n > 1$ times in $\varphi(\mathbf{x})$, replace $n-1$ occurrences with $n-1$ fresh variables x_2, \dots, x_n and add $(x = x_2) \wedge \dots \wedge (x = x_n)$ to the conjunction.

As an example, $r(x, y, x) \wedge s(x, a, a)$, where a is a constant, can be rewritten into $r(x, y, x_2) \wedge s(x_3, z_a, z'_a) \wedge (x = x_2) \wedge (x = x_3) \wedge (z_a = a) \wedge (z'_a = a)$. In the following definition, we assume that the body of every EGD has been rewritten as above, and thus it is of the form $\varphi(\mathbf{x}) \wedge \text{eq}(\mathbf{x}')$, where $\varphi(\mathbf{x})$ is a conjunction of atoms (with no constants and no multiple occurrences of the same variable) and $\text{eq}(\mathbf{x}')$ is a conjunction of equalities over variables $\mathbf{x}' \subseteq \mathbf{x}$.

Definition 5 (Universal repair step). *Let $C = \langle I, \Phi \rangle$ be a CI and σ an EGD $\varphi(\mathbf{x}) \wedge \text{eq}(\mathbf{x}') \rightarrow x_i = x_j$. Let h be a homomorphism s.t. $h(\varphi(\mathbf{x})) \subseteq I$ and for which there exists a valuation ν s.t. $\nu \models h(\text{eq}(\mathbf{x}'))$, $\nu \models \Phi$, and $\nu(h(x_i)) \neq \nu(h(x_j))$.*

Moreover, let \perp_ℓ and \perp_m be two fresh nulls, and $\langle I', \Phi' \rangle$ be the CI obtained from $\langle I, \Phi \rangle$ as follows:

1. *for each fact $p(u_1, \dots, u_n)$ in I , if $\varphi(\mathbf{x})$ contains an atom $p(t_1, \dots, t_n)$ s.t. $\nu(h(p(t_1, \dots, t_n))) = \nu(p(u_1, \dots, u_n))$, then for every $1 \leq k \leq n$, if $t_k = x_i$ or $(t_k = x_i)$ appears in $\text{eq}(\mathbf{x}')$ (resp. $t_k = x_j$ or $(t_k = x_j)$ appears in $\text{eq}(\mathbf{x}')$), then replace u_k in $p(u_1, \dots, u_n)$ with \perp_ℓ (resp. \perp_m);*
2. $\Phi' = \Phi \wedge \phi$, where ϕ is defined as follows:

$$(h(\text{eq}(\mathbf{x}')) \Rightarrow ((\perp_m = \perp_\ell) \wedge (\perp_\ell = h(x_i) \vee \perp_m = h(x_j)))) \wedge \\ (\neg(h(\text{eq}(\mathbf{x}')))) \Rightarrow (\perp_\ell = h(x_i) \wedge \perp_m = h(x_j));$$
3. *for each u_k of Item 1 above, if $u_k \in \text{Null}$, replace every occurrence of u_k with \perp_ℓ (resp. \perp_m) in both the current instance and condition Φ' .*

We say that $\langle I, \Phi \rangle \xrightarrow{\sigma, h} \langle I', \Phi' \rangle$ is a universal repair step.

In the previous definition, h is a homomorphism used to map $\varphi(\mathbf{x})$ into I , while ν is a valuation used to check if $h(\varphi(\mathbf{x}))$ may violate σ , that is, $\nu(h(\varphi(\mathbf{x})))$ are admissible facts (cf. $\nu \models \Phi$) that satisfy the join conditions in the body of σ (cf. $\nu \models h(\text{eq}(\mathbf{x}'))$) but do not satisfy the head condition (cf. $\nu(h(x_i)) \neq \nu(h(x_j))$).

A *universal repair sequence* of a knowledge base (D, Σ) is a (possibly empty) finite sequence of universal repair steps

$C_i \xrightarrow{\sigma_i, h_i} C_{i+1}$ ($0 \leq i < m$) such that $C_0 = \langle D, \text{true} \rangle$ and $\sigma_i \in \Sigma$ for every i . We call C_m the *result* of the universal repair sequence. Also, we say that the repair sequence is *maximal* if there does not exist a universal repair step of the form

$C_m \xrightarrow{\sigma_m, h_m} C_{m+1}$. A universal repair sequence is said to be

ordered if for each $C_i \xrightarrow{\sigma_i, h_i} C_{i+1}$ there is no $C_j \xrightarrow{\sigma_j, h_j} C_{j+1}$ such that $j > i$ and $\sigma_j < \sigma_i$ (intuitively, EGDs in Σ must be considered according to a topological sorting of Γ_Σ). If C is the result of an ordered maximal universal repair sequence, then C is called a *universal repair*.

Example 8. Consider the database of Example 6 and the set of EGDs of Example 5. After the application of σ_1 to the first two facts we get the following conditional instance:

john	\perp_1	rome
john	\perp_2	rome
mary	math	sydney

Φ_1

where Φ_1 is as follows:

$$((\text{john} = \text{john}) \Rightarrow ((\perp_1 = \perp_2) \wedge (\perp_1 = \text{cs} \vee \perp_2 = \text{math}))) \wedge \\ ((\text{john} \neq \text{john}) \Rightarrow (\perp_1 = \text{cs} \wedge \perp_2 = \text{math}))$$

By applying σ_2 to the first and third facts, we get:

john	\perp_1	\perp_3
john	\perp_2	rome
mary	math	\perp_4

$\Phi_1 \wedge \Phi_2$

where Φ_2 is as follows:

$$(\perp_1 = \text{math} \Rightarrow ((\perp_3 = \perp_4) \wedge (\perp_3 = \text{rome} \vee \perp_4 = \text{sydney}))) \wedge \\ (\perp_1 \neq \text{math} \Rightarrow (\perp_3 = \text{rome} \wedge \perp_4 = \text{sydney}))$$

By applying σ_2 to the first two facts, we get:

john	\perp_1	\perp_5
john	\perp_2	\perp_6
mary	math	\perp_4

$\Phi_1 \wedge \Phi_2 \wedge \Phi_3$

where Φ_3 is as follows:

$$((\perp_1 = \perp_2) \Rightarrow ((\perp_5 = \perp_6) \wedge (\perp_5 = \perp_3 \vee \perp_6 = \text{rome}))) \wedge \\ ((\perp_1 \neq \perp_2) \Rightarrow (\perp_5 = \perp_3 \wedge \perp_6 = \text{rome}))$$

Then, \perp_3 is replaced with \perp_5 everywhere, that is, in both Φ_2 and Φ_3 , yielding new conditions Φ'_2 and Φ'_3 . The result of this step is the CI consisting of the instance above and the condition $\Phi_1 \wedge \Phi'_2 \wedge \Phi'_3$.

As no further universal repair step can be applied, the conditional instance above is a universal repair.

The following theorems say that all the universal repairs of a knowledge base are equivalent in that they “represent” its repairs and computing one of them can be done in polynomial time.

Theorem 4. *Let U be a universal repair of a knowledge base (D, Σ) . Then, $pw(U) = \text{repair}(D, \Sigma)$.*

Theorem 5. *Computing a universal repair is in PTIME.*

The main idea behind the proof of the previous theorem is to apply universal repair steps according to a topological sorting of the dependency graph of Σ .

6 Query Evaluation

In this section, we show how to evaluate queries over universal repairs and keep track of provenance information. This will be used later on to develop an approximation algorithm. In order to do that, we need to extend conditional instances to allow individual conditions for facts too.

A *conditional fact* is a pair $\langle p(\mathbf{t}), \phi \rangle$, where $p(\mathbf{t})$ is a fact and ϕ is a condition, also called *local condition*.

An *extended conditional instance* (ECI) is a pair $\langle E, \Phi \rangle$, where E is a finite set of conditional facts and Φ is a condition, which we also call *global condition*. Given a valuation ν , we define $\nu(E) = \{\nu(p(\mathbf{t})) \mid \langle p(\mathbf{t}), \phi \rangle \in E \text{ and } \nu \models \phi\}$. Thus, $\nu(E)$ is the database obtained from E by applying ν to every fact and keeping only the facts whose condition is satisfied by ν . The *possible worlds* of an ECI $G = \langle E, \Phi \rangle$ are $pw(G) = \{\nu(E) \mid \nu \text{ is a valuation s.t. } \nu \models \Phi\}$.

We now define how to evaluate queries over ECIs. To this end, we generalize the immediate consequence operator T_P (cf. Section 3) so as to perform a “conditional” evaluation of a query in the presence of conditional facts. We point out that our conditional evaluation is a slight variant of the one for conditional tables [Grahne, 1991]. Below we use $(t_1, \dots, t_n) = (u_1, \dots, u_n)$ as a shorthand for $\bigwedge_{i=1}^n (t_i = u_i)$.

With a slight abuse of notation, given a program P we overload T_P so that when its input is a set E of conditional facts, then $T_P(E) = E \cup E'$, where E' is the following set of conditional facts:

$$\begin{aligned} & \{ \langle h(p(\mathbf{x})), \phi \rangle \mid \text{there exist a homomorphism } h \text{ and} \\ & \quad \text{a rule } p(\mathbf{x}) \leftarrow \bigwedge_{i=1}^n p_i(\mathbf{x}_i, \mathbf{y}_i), \gamma(\mathbf{z}) \text{ of } P \text{ and} \\ & \quad \phi_1, \dots, \phi_n \text{ s.t.} \\ & \quad \langle h(p_i(\mathbf{x}_i, \mathbf{y}_i)), \phi_i \rangle \in E \text{ for every } 1 \leq i \leq n \text{ and} \\ & \quad \phi = h(\gamma(\mathbf{z})) \wedge \bigwedge_{i=1}^n \phi_i \} \cup \\ & \{ \langle h(p(\mathbf{x})), \phi \rangle \mid \text{there exist a homomorphism } h \text{ and} \\ & \quad \text{a rule } p(\mathbf{x}) \leftarrow p'(\mathbf{x}), \neg p''(\mathbf{x}) \text{ of } P \text{ and } \phi' \text{ s.t.} \\ & \quad \langle h(p'(\mathbf{x})), \phi' \rangle \in E \text{ and} \\ & \quad \phi = \phi' \wedge \bigwedge_{\langle p''(\mathbf{t}), \phi'' \rangle \in E} \neg(\phi'' \wedge h(\mathbf{x}) = \mathbf{t}) \}. \end{aligned}$$

Then, $\text{CN}(P, E) = T_P^n(E)$, where T_P^i is defined in exactly the same way as in Section 3.

It can be easily verified that the conditional evaluation above can be carried out in polynomial time. The main reason is that the immediate consequence operator T_P above extends the classical one by just adding conditions of polynomial size to facts, so the time complexity remains polynomial.

Proposition 6. *Given a program P and a finite set E of conditional facts, computing $\text{CN}(P, E)$ is in PTIME.*

Given two conditions ϕ, ϕ' , we write $\phi \models \phi'$ if for every valuation ν , if $\nu \models \phi$, then $\nu \models \phi'$. Below we define the *certain answers* to a query over an extended conditional instance.

Definition 6. *Let $G = \langle E, \Phi \rangle$ be an ECI and $Q = (P, p)$ be a query. The certain answers to Q over G are defined as follows:*

$$\text{cert}(Q, G) = \{p(\mathbf{t}) \mid \langle p(\mathbf{t}), \phi \rangle \in \text{CN}(P, E) \text{ and } \Phi \models \phi \text{ and } \mathbf{t} \text{ has constants only}\}.$$

Thus, to compute the certain answers to Q over G we need to compute $\text{CN}(P, E)$ and check for which conditional facts $\Phi \models \phi$ holds true. The last entailment checking can be carried out by rewriting the formula as a propositional formula and using SAT solvers [Gomes *et al.*, 2008]. Satisfiability solvers are becoming increasingly efficient and effective in solving large satisfiability problems and, despite the worst-case exponential run time of all known algorithms, they are very efficient in solving hard problems with millions of variables and tens of millions of constraints.

The tools developed thus far will be exploited in the next section to devise an approximation algorithm.

7 Approximation Algorithm

In light of Theorem 3, we develop a polynomial time approximation algorithm to compute a sound (but possibly incomplete) set of consistent query answers.

The basic idea of our approximation algorithm is as follows. Given a knowledge base and a query, we first compute a universal repair U of the knowledge base, and then transform U into an ECI U_{true} where each fact is associated with the trivial condition true. After that, the query is “conditionally” evaluated over U_{true} (i.e., as described in the previous section). The result is a set of conditional facts, whose conditions are evaluated together with the condition of U_{true} in order to determine a sound (but possibly incomplete) set of consistent query answers (we will illustrate the entire process in Example 9). For positive queries, a simple approximation can be obtained by deleting facts containing nulls from U_{true} and ignoring its global condition.

First, we need some auxiliary definitions. Given a condition Φ and two nulls $\perp_i, \perp_j \in \text{Null}$, we write $\perp_i \approx_{\Phi} \perp_j$ if $\perp_i = \perp_j$ (or $\perp_j = \perp_i$) appears in Φ . Then, \approx_{Φ}^* is the reflexive and transitive closure of \approx_{Φ} . Given a term $t \in \text{Const} \cup \text{Null}$ and a condition Φ , we denote by $\text{dom}_{\Phi}(t)$ the *domain of t w.r.t. Φ* , defined as follows: $\text{dom}_{\Phi}(c) = \{c\}$ if $c \in \text{Const}$, whereas $\text{dom}_{\Phi}(\perp_i) = \{c \mid \perp_i \approx_{\Phi}^* \perp_k \text{ and } \perp_k = c \text{ appears in } \Phi\}$ if $\perp_i \in \text{Null}$. Intuitively, $\text{dom}_{\Phi}(\perp_i)$ is a set of possible values that \perp_i might assume in some repair.

We assume the strict ordering $\text{false} < \text{unknown} < \text{true}$, and $\neg\text{true} = \text{false}$, $\neg\text{false} = \text{true}$, and $\neg\text{unknown} = \text{unknown}$. Given two conditions φ and Φ , we define $\text{eval}_{\Phi}(\varphi)$ as follows:

$$\begin{aligned} \bullet \text{eval}_{\Phi}((t_i = t_j)) &= \begin{cases} \text{true} & \text{if } t_i = t_j, \\ \text{false} & \text{if } \text{dom}_{\Phi}(t_i) \cap \text{dom}_{\Phi}(t_j) = \emptyset, \\ \text{unknown} & \text{otherwise.} \end{cases} \\ \bullet \text{eval}_{\Phi}((t_i \neq t_j)) &= \neg \text{eval}_{\Phi}((t_i = t_j)) \\ \bullet \text{eval}_{\Phi}((\varphi_1 \wedge \varphi_2)) &= \min\{\text{eval}_{\Phi}(\varphi_1), \text{eval}_{\Phi}(\varphi_2)\}. \end{aligned}$$

- $\text{eval}_\Phi((\varphi_1 \vee \varphi_2)) = \max\{\text{eval}_\Phi(\varphi_1), \text{eval}_\Phi(\varphi_2)\}$.
- $\text{eval}_\Phi((\neg\varphi)) = \neg\text{eval}_\Phi(\varphi)$
- $\text{eval}_\Phi(v) = v$ for $v \in \{\text{true}, \text{unknown}, \text{false}\}$.

The following definition introduces the approximate evaluation of a query over an extended conditional instance.

Definition 7. Let $G = \langle E, \Phi \rangle$ be an ECI and $Q = (P, p)$ be a query. The approximate answers to Q over G are:

$$\text{approx}(Q, G) = \{p(\mathbf{t}) \mid \langle p(\mathbf{t}), \phi \rangle \in \text{CN}(P, E) \text{ and} \\ \text{eval}_\Phi(\phi) = \text{true and} \\ \mathbf{t} \text{ has constants only}\}.$$

A conditional instance $C = \langle I, \Phi \rangle$ can be “lifted to” an ECI by setting all local conditions to true. We denote by C_{true} the ECI $\langle \{p(\mathbf{t}), \text{true}\} \mid p(\mathbf{t}) \in C \rangle, \Phi \rangle$.

The evaluation technique of Definition 7 indeed provides a way to compute a sound (but possibly incomplete) set of consistent query answers, in polynomial time.

Definition 8. Let (D, Σ) be a knowledge base and Q a query. The approximate consistent answers to Q over (D, Σ) are $\tilde{Q}(D, \Sigma) = \text{approx}(Q, U_{\text{true}})$, where U is a universal repair of (D, Σ) .

Example 9. Consider the knowledge base of Example 1 and the query $Q = \langle P, \text{diffCity} \rangle$, asking for the pairs of departments located in different cities, where P consists of the rule:

$$\text{diffCity}(D, D') \leftarrow \text{works}(E, D, C), \text{works}(E', D', C'), C \neq C'.$$

A universal repair, say U , is shown in Example 3. The first step of our approach consists of converting U into an ECI U_{true} by setting all local conditions to true. Thus, U_{true} is as follows:

john	cs	nyc	true
john	math	\perp_1	true
mary	math	\perp_1	true

$$\perp_1 = \text{rome} \vee \perp_1 = \text{sydney}$$

The conditional evaluation of Q over U_{true} gives the following set of conditional facts:

cs	cs	$\text{true} \wedge \text{true} \wedge \text{nyc} \neq \text{nyc}$
cs	math	$\text{true} \wedge \text{true} \wedge \text{nyc} \neq \perp_1$
math	cs	$\text{true} \wedge \text{true} \wedge \perp_1 \neq \text{nyc}$
math	math	$\text{true} \wedge \text{true} \wedge \perp_1 \neq \perp_1$

As $\text{dom}_\Phi(\perp_1) = \{\text{rome}, \text{sydney}\}$, the approximate consistent answers are $\text{diffCity}(\text{cs}, \text{math})$ and $\text{diffCity}(\text{math}, \text{cs})$, which are indeed consistent query answers.

The following theorems state soundness of our approach and its complexity, which is in PTIME (data complexity).

Theorem 7. Let (D, Σ) be a knowledge base and Q a query. Then, $\tilde{Q}(D, \Sigma) \subseteq Q(D, \Sigma)$.

Theorem 8. Let (D, Σ) be a knowledge base and Q a query. Computing $\tilde{Q}(D, \Sigma)$ is in PTIME.

Proof. (Sketch) It follows from Theorem 5, Proposition 6, and because $\text{eval}_\Phi(\varphi)$ can be evaluated in PTIME. \square

We conclude by comparing our approach with two well-known approximation schemes that have been developed in the query answering inconsistency-tolerant area, namely the *intersection of repairs* (IAR) and the *intersection of closed repairs* (ICR) semantics [Lembo *et al.*, 2010; Bienvenu *et al.*, 2014]. The former semantics consists of querying the intersection of all repairs, while the second one consists of querying the intersection of the closure of all repairs. In our setting, since we do not consider tuple-generating dependencies, the two semantics coincide.

As stated in the following theorem, for every positive query (i.e., only positive rules and equalities are allowed), our approximation algorithm provides at least as many answers as IAR (and ICR).

Theorem 9. Let (D, Σ) be a knowledge base and Q a positive query. Then, $Q(D^*) \subseteq \tilde{Q}(D, \Sigma)$, where $D^* = \bigcap \{D' \mid D' \in \text{repair}(D, \Sigma)\}$.

To show a simple case where our approximation algorithm yields strictly more consistent query answers, consider the knowledge base of Example 1, whose repairs are reported in Example 2. Consider also the query asking for the employees’ names. Under the IAR semantics, the intersection keeps only the first fact of the original knowledge base, and thus the other two facts are lost altogether. So the only query answer is john. On the other hand, the universal repair (cf. Example 3) is much more informative: we still have the first fact of the original knowledge base, but we also have the other two facts, where consistent values are unchanged and thus preserved. Our approximation algorithm returns both john and mary.

As a further example, consider Example 6. The intersection of all repairs is empty. On the other hand, the universal repair (cf. Example 8) has valuable information, which again can be profitably exploited for (approximate) query answering.

8 Conclusion

We proposed a framework for query answering over inconsistent KBs based on (i) a notion of repair allowing values within facts to be updated, (ii) a compact representation of all repairs, (iii) an approximation algorithm to compute under-approximations of consistent query answers.

As a direction for future work, we plan to investigate further approximation algorithms based on different conditions’ evaluations. Another issue to be investigated is the generalization of our framework to more general classes of dependencies, such as TGDs or arbitrary EGDs. One way of dealing with TGDs is by inserting new facts. Interestingly, by inserting facts with nulls (in a Chase-like manner) into a universal repair we can get a compact representation of all possible insertions. This also poses the issues of combining different repair primitives, such as fact updates and fact insertions (fact deletions might be considered as well).

References

- [Arenas *et al.*, 1999] Marcelo Arenas, Leopoldo E. Bertossi, and Jan Chomicki. Consistent query answers in inconsistent databases. In *Symposium on Principles of Database Systems (PODS)*, pages 68–79, 1999.

- [Bertossi *et al.*, 2008] Leopoldo E. Bertossi, Loreto Bravo, Enrico Franconi, and Andrei Lopatenko. The complexity and approximation of fixing numerical attributes in databases under integrity constraints. *Information Systems*, 33(4-5):407–434, 2008.
- [Bienvenu and Rosati, 2013] Meghyn Bienvenu and Riccardo Rosati. Tractable approximations of consistent query answering for robust ontology-based data access. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 775–781, 2013.
- [Bienvenu *et al.*, 2014] Meghyn Bienvenu, Camille Bourgaux, and Francois Goasdoué. Querying inconsistent description logic knowledge bases under preferred repair semantics. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 996–1002, 2014.
- [Bienvenu, 2011] Meghyn Bienvenu. First-order expressibility results for queries over inconsistent DL-Lite knowledge bases. In *International Workshop on Description Logics (DL)*, 2011.
- [Bienvenu, 2012] Meghyn Bienvenu. On the complexity of consistent query answering in the presence of simple ontologies. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2012.
- [Bohannon *et al.*, 2005] Philip Bohannon, Michael Flaster, Wenfei Fan, and Rajeev Rastogi. A cost-based model and effective heuristic for repairing constraints by value modification. In *International Conference on Management of Data (SIGMOD)*, pages 143–154, 2005.
- [Fiorentino *et al.*, 2018] Nicola Fiorentino, Sergio Greco, Cristian Molinaro, and Irina Trubitsyna. ACID: A system for computing approximate certain query answers over incomplete databases. In *International Conference on Management of Data (SIGMOD)*, 2018. To appear.
- [Flesca *et al.*, 2010] Sergio Flesca, Filippo Furfaro, and Francesco Parisi. Querying and repairing inconsistent numerical databases. *ACM Transactions on Database Systems*, 35(2):14:1–14:50, 2010.
- [Furfaro *et al.*, 2007] Filippo Furfaro, Sergio Greco, and Cristian Molinaro. A three-valued semantics for querying and repairing inconsistent databases. *Annals of Mathematics and Artificial Intelligence*, 51(2-4):167–193, 2007.
- [Gomes *et al.*, 2008] Carla P. Gomes, Henry A. Kautz, Ashish Sabharwal, and Bart Selman. Satisfiability solvers. In *Handbook of Knowledge Representation*, pages 89–134, 2008.
- [Grahne, 1991] Gösta Grahne. *The Problem of Incomplete Information in Relational Databases*, volume 554 of *Lecture Notes in Computer Science*. Springer, 1991.
- [Greco and Molinaro, 2008] Sergio Greco and Cristian Molinaro. Approximate probabilistic query answering over inconsistent databases. In *International Conference on Conceptual Modeling (ER)*, pages 311–325, 2008.
- [Greco and Molinaro, 2012] Sergio Greco and Cristian Molinaro. Probabilistic query answering over inconsistent databases. *Annals of Mathematics and Artificial Intelligence*, 64(2-3):185–207, 2012.
- [Greco *et al.*, 2017] Sergio Greco, Cristian Molinaro, and Irina Trubitsyna. Computing approximate certain answers over incomplete databases. In *Alberto Mendelzon International Workshop (AMW)*, 2017.
- [Guagliardo and Libkin, 2016] Paolo Guagliardo and Leonid Libkin. Making SQL queries correct on incomplete databases: A feasibility study. In *Symposium on Principles of Database Systems (PODS)*, pages 211–223, 2016.
- [Imielinski and Lipski, 1984] Tomasz Imielinski and Witold Lipski. Incomplete information in relational databases. *Journal of the ACM*, 31(4):761–791, 1984.
- [Lembo *et al.*, 2010] Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi, and Domenico Fabio Savo. Inconsistency-tolerant semantics for description logics. In *International Conference on Web Reasoning and Rule Systems (RR)*, pages 103–117, 2010.
- [Lembo *et al.*, 2011] Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi, and Domenico Fabio Savo. Query rewriting for inconsistent DL-Lite ontologies. In *International Conference on Web Reasoning and Rule Systems (RR)*, pages 155–169, 2011.
- [Lembo *et al.*, 2015] Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi, and Domenico Fabio Savo. Inconsistency-tolerant query answering in ontology-based data access. *Journal of Web Semantics*, 33:3–29, 2015.
- [Libkin, 2015] Leonid Libkin. How to define certain answers. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4282–4288, 2015.
- [Libkin, 2016] Leonid Libkin. Certain answers as objects and knowledge. *Artificial Intelligence*, 232:1–19, 2016.
- [Lukasiewicz *et al.*, 2012a] Thomas Lukasiewicz, Maria Vanina Martinez, and Gerardo I. Simari. Inconsistency handling in Datalog+/- ontologies. In *European Conference on Artificial Intelligence (ECAI)*, pages 558–563, 2012.
- [Lukasiewicz *et al.*, 2012b] Thomas Lukasiewicz, Maria Vanina Martinez, and Gerardo I. Simari. Inconsistency-tolerant query rewriting for linear Datalog+/- . In *Datalog 2.0*, pages 123–134, 2012.
- [Lukasiewicz *et al.*, 2015] Thomas Lukasiewicz, Maria Vanina Martinez, Andreas Pieris, and Gerardo I. Simari. From classical to consistent query answering under existential rules. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 1546–1552, 2015.
- [Martinez *et al.*, 2014] Maria Vanina Martinez, Francesco Parisi, Andrea Pugliese, Gerardo I. Simari, and V. S. Subrahmanian. Policy-based inconsistency management in relational databases. *International Journal of Approximate Reasoning*, 55(2):501–528, 2014.
- [Rosati, 2011] Riccardo Rosati. On the complexity of dealing with inconsistency in description logic ontologies. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1057–1062, 2011.