

From Causes for Database Queries to Repairs and Model-Based Diagnosis and Back

Babak Salimi¹ and Leopoldo Bertossi²

1 Carleton University, School of Computer Science, Ottawa, Canada
bsalimi@scs.carleton.ca

2 Carleton University, School of Computer Science, Ottawa, Canada
bertossi@scs.carleton.ca

Abstract

In this work we establish and investigate connections between causality for query answers in databases, database repairs wrt. denial constraints, and consistency-based diagnosis. The first two are relatively new problems in databases, and the third one is an established subject in knowledge representation. We show how to obtain database repairs from causes and the other way around. Causality problems are formulated as diagnosis problems, and the diagnoses provide causes and their responsibilities. The vast body of research on database repairs can be applied to the newer problem of determining actual causes for query answers and their responsibilities. These connections, which are interesting *per se*, allow us, after a transition -inspired by consistency-based diagnosis- to computational problems on hitting sets and vertex covers in hypergraphs, to obtain several new algorithmic and complexity results for database causality.

1998 ACM Subject Classification H.2.3 [Database Management] Languages, H.2.4 [Database Management] Systems, I.2.3 [Artificial Intelligence] Deduction and Theorem Proving, I.2.4 [Artificial Intelligence] Knowledge Representation Formalisms and Methods, F.1.3 [Computation of Abstract Devices] Complexity Measures and Classes

Keywords and phrases causality, diagnosis, repairs, consistent query answering, integrity constraints

Digital Object Identifier 10.4230/LIPIcs.ICDT.2015.342

1 Introduction

When querying a database, a user may not always obtain the expected results, and the system could provide some explanations. They could be useful to further understand the data or check if the query is the intended one. Actually, the notion of explanation for a query result was introduced in [41], on the basis of the deeper concept of *actual causation*.

A tuple t is an *actual cause* for an answer \bar{a} to a conjunctive query Q from a relational database instance D if there is a *contingent* set of tuples Γ , such that, after removing Γ from D , \bar{a} is still an answer, but after further removing t from $D \setminus \Gamma$, \bar{a} is not an answer anymore. Here, Γ is a set of tuples that has to accompany \bar{a} for it to be a cause. Actual causes and contingent tuples are restricted to be among a pre-specified set of *endogenous tuples*, which are admissible, possible candidates for causes, as opposed to *exogenous tuples*, which may also be present in the database. In rest of this paper, whenever we simply say “cause”, we mean “actual cause”.

In applications involving large data sets, it is crucial to rank potential causes by their *responsibilities* [42, 41], which reflect the relative (quantitative) degrees of their causality for a query result. The responsibility measure for a cause is based on its *contingency sets*: the smallest (one of) its contingency sets, the strongest it is as a cause.



© Babak Salimi and Leopoldo Bertossi;
licensed under Creative Commons License CC-BY
18th International Conference on Database Theory (ICDT'15).

Editors: Marcelo Arenas and Martín Ugarte; pp. 342–362



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Actual causation, as used in [41], can be traced back to [29, 30], which provides a model-based account of causation on the basis of *counterfactual dependence*. Responsibility was introduced in [18], to capture the intuitive notion of *degree of causation*.

Apart from the explicit use of causality, research on explanations for query results has focused mainly, and rather implicitly, on provenance [12, 13, 14, 21, 36, 34, 53]. A close connection between causality and provenance has been established in [41]. However, causality is a more refined notion that identifies causes for query results on the basis of user-defined criteria, and ranks causes according to their responsibilities [42].

Consistency-based diagnosis [47], a form of model-based diagnosis [52, sec. 10.3], is an area of knowledge representation. The problem here is, given the *specification* of a system in some logical formalism and a usually unexpected *observation* about the system, to obtain *explanations* for the observation, in the form of a diagnosis for the unintended behavior.

In a different direction, a database instance, D , that is expected to satisfy certain integrity constraints may fail to do so. In this case, a *repair* of D is a database D' that does satisfy the integrity constraints and *minimally departs* from D . Different forms of minimality can be applied and investigated. A *consistent answer* to a query from D and wrt. the integrity constraints is a query answer that is obtained from all possible repairs, i.e. is invariant or certain under the class of repairs. These notions were introduced in [2] (surveys of the area can be found in [7, 9]). Although not in the context of repairs, consistency-based diagnosis has been applied to consistency restoration of a database wrt. integrity constraints [27].

These three forms of reasoning, namely inferring causes from databases, consistency-based diagnosis, and consistent query answering (and repairs) are all *non-monotonic* [49]. For example, a (most responsible) cause for a query result may not be such anymore after the database is updated. Furthermore, they all reflect some sort of *uncertainty* about the information at hand. In this work we establish natural, precise, useful, and deeper connections between these three reasoning tasks.

More precisely, we unveil a strong connection between computing causes and their responsibilities for conjunctive query answers, on one hand, and computing repairs in databases wrt. *denial constraints*, on the other. These computational problems can be reduced to each other. In order to obtain repairs wrt. a *set* of denial constraints from causes, we investigate causes for queries that are *unions of conjunctive queries*, and develop algorithms to compute causes and responsibilities.

We show that inferring and computing actual causes and their responsibilities in a database setting become diagnosis reasoning problems and tasks. Actually, a causality-based explanation for a conjunctive query answer can be viewed as a diagnosis, where in essence the first-order logical reconstruction of the relational database provides the system description [48], and the observation is the query answer. We also establish a bidirectional connection between diagnosis and repairs.

Being the causality problems the main focus of this work, we take advantage of algorithms and complexity results both for consistency-based diagnosis; and database repairs and consistent query answering [9]. In this way, we obtain new complexity results for the main problems of causality, namely computing actual causes, determining their responsibilities, and obtaining most responsible causes; and also for their decision versions. In particular, we obtain fixed-parameter tractable algorithms for some of them. More precisely, our main results are as follows:¹ (the complexity results are all in data complexity)

¹ A few of the results included here appear in [49].

1. For a boolean conjunctive query and its associated denial constraint (the former being its violation view), we establish a precise connection (characterization and computational reductions) between actual causes for the query (being true) and the subset- and cardinality-repairs of the instance wrt. the denial constraint. We obtain causes from repairs.
2. We obtain repairs from causes, for which we extend the treatment of causality to unions of conjunctive queries (to represent multiple denial constraints). We characterize an actual cause's responsibility in terms of cardinality-repairs. We provide algorithms to compute causes and their (minimal) contingency sets for unions of conjunctive queries. The causes can be computed in PTIME.
3. We establish a precise connection between consistency-based diagnosis for a boolean conjunctive query being unexpectedly true according to a system description, and causes for the query being true. In particular, we show how to compute actual causes, their contingency sets, and responsibilities using the diagnosis characterization. Hitting-set-based algorithmic approaches to diagnosis inspire our algorithmic/complexity approaches to causality.
4. We reformulate the causality problems as hitting set problems and vertex cover problems on hypergraphs, which allows us to apply results and techniques for the latter to causality.
5. (a) Checking minimal contingency sets can be done in PTIME. (b) The responsibility (decision) problem for conjunctive queries becomes NP -complete. (c) However, it is fixed-parameter tractable when the parameter is the inverse of the responsibility bound. (d) The functional problem of computing the causes' responsibilities is $FP^{NP(\log(n))}$ -complete, and deciding most responsible causes is $P^{NP(\log(n))}$ -complete.
6. The structure of the resulting hitting-set problem allows us to obtain efficient parameterized algorithms and good approximation algorithms for computing causes and minimal contingency sets.
7. On the basis of the causality/repair connection, and the dichotomy result for causality [41], we obtain a dichotomy result for the complexity of deciding the existence of repairs of a certain size wrt. single, self-join-free denial constraints.
8. We discuss extensions and open issues that deserve investigation.

The paper is structured as follows. Section 2 introduces technical preliminaries for relational databases, causality in databases, database repairs and consistent query answering, consistency-based diagnosis, and relevant complexity classes. Section 3 characterizes actual causes and responsibilities in terms of database repairs. Section 3 characterizes repairs and consistent query answering in terms of causes and contingency sets for queries that are unions of conjunctive queries; and presents an algorithm for computing both of the latter. Section 5 formulates causality problems as consistency-based diagnosis problems, and the latter as repair problems. Section 6 shows complexity and algorithmic results; in particular a fixed-parameter tractability result for causes' responsibilities. Finally, Section 7 discusses several relevant issues, connections and open problems around causality in databases. Proofs of results without an implicit proof in this paper can be found in [50].

2 Preliminaries

We consider relational database schemas of the form $\mathcal{S} = (U, \mathcal{P})$, where U is the possibly infinite database domain of *constants* and \mathcal{P} is a finite set of *database predicates*² of fixed

² As opposed to built-in predicates (e.g. \neq) that we assume do not appear, unless explicitly stated otherwise.

arities. A database instance D compatible with \mathcal{S} can be seen as a finite set of ground atomic formulas (in databases aka. atoms or tuples), of the form $P(c_1, \dots, c_n)$, where $P \in \mathcal{P}$ has arity n , and $c_1, \dots, c_n \in U$. A *conjunctive query* (CQ) is a formula $\mathcal{Q}(\bar{x})$ of the first-order (FO) logic language, $\mathcal{L}(\mathcal{S})$, associated to \mathcal{S} of the form $\exists \bar{y}(P_1(\bar{t}_1) \wedge \dots \wedge P_m(\bar{t}_m))$, where the $P_i(\bar{t}_i)$ are atomic formulas, i.e. $P_i \in \mathcal{P}$, and the \bar{t}_i are sequences of terms, i.e. variables or constants. The \bar{x} in $\mathcal{Q}(\bar{x})$ shows all the free variables in the formula, i.e. those not appearing in \bar{y} . If \bar{x} is non-empty, the query is *open*. If \bar{x} is empty, the query is *boolean* (a BCQ), i.e. the query is a sentence, in which case, it is true or false in a database, denoted by $D \models \mathcal{Q}$ and $D \not\models \mathcal{Q}$, respectively. A sequence \bar{c} of constants is an answer to an open query $\mathcal{Q}(\bar{x})$ if $D \models \mathcal{Q}[\bar{c}]$, i.e. the query becomes true in D when the variables are replaced by the corresponding constants in \bar{c} .

An *integrity constraint* is a sentence of language $\mathcal{L}(\mathcal{S})$, and then, may be true or false in an instance for schema \mathcal{S} . Given a set IC of integrity constraints, a database instance D is *consistent* if $D \models IC$; otherwise it is said to be *inconsistent*. In this work we assume that sets of integrity constraints are always finite and logically consistent. A particular class of integrity constraints is formed by *denial constraints* (DCs), which are sentences κ of the form: $\forall \bar{x} \neg(A_1(\bar{x}_1) \wedge \dots \wedge A_n(\bar{x}_n))$, where $\bar{x} = \bigcup \bar{x}_i$ and each $A_i(\bar{x}_i)$ is a database atom, i.e. predicate $A \in \mathcal{P}$. (The atoms may contain constants.) Denial constraints are exactly the negations of BCQs.

Causality and Responsibility. Assume that the database instance is split in two, i.e. $D = D^n \cup D^x$, where D^n and D^x denote the sets of *endogenous* and *exogenous* tuples, respectively. A tuple $t \in D^n$ is called a *counterfactual cause* for a BCQ \mathcal{Q} , if $D \models \mathcal{Q}$ and $D \setminus \{t\} \not\models \mathcal{Q}$. A tuple $t \in D^n$ is an *actual cause* for \mathcal{Q} if there exists $\Gamma \subseteq D^n$, called a *contingency set*, such that t is a counterfactual cause for \mathcal{Q} in $D \setminus \Gamma$ [41]. We will concentrate mostly on CQs. However, the definition of actual causes and contingency sets can be applied without a change to monotone queries in general [41].

The *responsibility* of an actual cause t for \mathcal{Q} , denoted by $\rho_D(t)$, is the numerical value $\frac{1}{|\Gamma|+1}$, where $|\Gamma|$ is the size of the smallest contingency set for t . We can extend responsibility to all the other tuples in D^n by setting their value to 0. Those tuples are not actual causes for \mathcal{Q} .

► **Example 1.** Consider $D = D^n = \{R(a_4, a_3), R(a_2, a_1), R(a_3, a_3), S(a_4), S(a_2), S(a_3)\}$, and the query $\mathcal{Q} : \exists x \exists y (S(x) \wedge R(x, y) \wedge S(y))$. It holds: $D \models \mathcal{Q}$.

Tuple $S(a_3)$ is a counterfactual cause for \mathcal{Q} . If $S(a_3)$ is removed from D , \mathcal{Q} is not true anymore. Therefore, the responsibility of $S(a_3)$ is 1. Besides, $R(a_4, a_3)$ is an actual cause for \mathcal{Q} with contingency set $\{R(a_3, a_3)\}$. If $R(a_3, a_3)$ is removed from D , \mathcal{Q} is still true, but further removing $R(a_4, a_3)$ makes \mathcal{Q} false. The responsibility of $R(a_4, a_3)$ is $\frac{1}{2}$, because its smallest contingency sets have size 1. Likewise, $R(a_3, a_3)$ and $S(a_4)$ are actual causes for \mathcal{Q} with responsibility $\frac{1}{2}$.

For the same \mathcal{Q} , but with $D = \{S(a_3), S(a_4), R(a_4, a_3)\}$, and the partition $D^n = \{S(a_4), S(a_3)\}$ and $D^x = \{R(a_4, a_3)\}$, it turns out that both $S(a_3)$ and $S(a_4)$ are counterfactual causes for \mathcal{Q} . ◀

Notation: $\mathcal{CS}(D^n, D^x, \mathcal{Q})$ denotes the set of actual causes for BCQ \mathcal{Q} (being true) from instance $D = D^n \cup D^x$. When $D^n = D$ and $D^x = \emptyset$, we sometimes simply write: $\mathcal{CS}(D, \mathcal{Q})$.

Database Repairs. Given a set IC of integrity constraints, a *subset repair* (simply, S-repair) of a possibly inconsistent instance D for schema \mathcal{S} is an instance D' for \mathcal{S} that satisfies

IC and makes $\Delta(D, D') = (D \setminus D') \cup (D' \setminus D)$ minimal under set inclusion. $Srep(D, IC)$ denotes the set of S-repairs of D wrt. IC [2]. Similarly, D' is a *cardinality repair* (simply C-repair) of D if D' satisfies IC and minimizes $|\Delta(D, D')|$. $Crep(D, IC)$ denotes the class of C-repairs of D wrt. IC . C-repairs are S-repairs of minimum cardinality.

For DCs, S-repairs and C-repairs are obtained from the original instance by deleting an S-minimal, resp. C-minimal, set of tuples.³ More generally, different *repair semantics* may be considered to restore consistency wrt. general integrity constraints. They depend on the kind of allowed updates on the database (i.e. tuple insertions/deletions, changes of attribute values), and the minimality conditions on repairs (e.g. subset-minimality, cardinality-minimality, etc.). Given D and IC , a repair semantics determines a class of intended or preferred repairs [9, sec. 2.5].

Given a repair semantics, RS , \bar{c} is a *consistent answer* to an open query $Q(\bar{x})$ if $D' \models Q[\bar{c}]$ for every RS -repair D' . A BCQ is *consistently true* if it is true in all RS -repairs. If \bar{c} is a consistent answer to $Q(\bar{x})$ wrt. S-repairs, we say it is an *S-consistent answer*. Similarly for *C-consistent answers*. Consistent query answering for DCs under S-repairs was investigated in detail [17]. C-repairs and consistent query answering were investigated in detail in [39]. (Cf. [9] for more references.)

Consistency-Based Diagnosis. Consistency-based diagnosis, a form of model-based diagnosis [52, sec. 10.4], considers problems $\mathcal{M} = (SD, COMPS, OBS)$, where SD is the description in logic of the intended properties of a system under the *explicit* assumption that all the *components* in $COMPS$, are working normally. OBS is a FO sentence that represents the observations. If the system does not behave as expected (as shown by the observations), then the logical theory obtained from $SD \cup OBS$ plus the explicit assumption, say $\bigwedge_{c \in COMPS} \neg Ab(c)$, that the components are indeed behaving normally, becomes inconsistent. Ab is an *abnormality* predicate.⁴

The inconsistency is captured via the *minimal conflict sets*, i.e. those minimal subsets $COMPS'$ of $COMPS$, such that $SD \cup OBS \cup \{\bigwedge_{c \in COMPS'} \neg Ab(c)\}$ is inconsistent. As expected, different notions of minimality can be used at this point.

A *minimal diagnosis* for \mathcal{M} is a minimal subset Δ of $COMPS$, such that $SD \cup OBS \cup \{\neg Ab(c) \mid c \in COMPS \setminus \Delta\} \cup \{Ab(c) \mid c \in \Delta\}$ is consistent. That is, consistency is restored by flipping the normality assumption to abnormality for a minimal set of components, and those are the ones considered to be (jointly) faulty. The notion of minimality commonly used is S-minimality, i.e. a diagnosis that does not have a proper subset that is a diagnosis. We will use this kind of minimality in relation to diagnosis. Diagnosis can be obtained from conflict sets [47].

Complexity Classes. We recall some complexity classes [46] used in this paper. FP is the class of functional problems associated to decision problem in the class $PTIME$, i.e. that are solvable in polynomial time. P^{NP} (or Δ_2^P) is the class of decision problems solvable in polynomial time by a machine that makes calls to an NP oracle. For $P^{NP(\log(n))}$ the number of calls is logarithmic. It is not known if $P^{NP(\log(n))}$ is strictly contained in P^{NP} . $FP^{NP(\log(n))}$ is similarly defined.

³ We will usually say that a set is S-minimal in a class of sets \mathcal{C} if it minimal under set inclusion in \mathcal{C} . Similarly, a set is C-minimal if it is minimal in cardinality within \mathcal{C} .

⁴ Here, and as usual, the atom $Ab(c)$ expresses that component c is (behaving) abnormal(ly).

3 Actual Causes From Database Repairs

Let $D = D^n \cup D^x$ be an instance for schema \mathcal{S} , and $\mathcal{Q}: \exists \bar{x}(P_1(\bar{x}_1) \wedge \dots \wedge P_m(\bar{x}_m))$ a BCQ. \mathcal{Q} may be unexpectedly true, i.e. $D \models \mathcal{Q}$. Now, $\neg \mathcal{Q}$ is logically equivalent to the DC $\kappa(\mathcal{Q}): \forall \bar{x} \neg(P_1(\bar{x}_1) \wedge \dots \wedge P_m(\bar{x}_m))$. The requirement that $\neg \mathcal{Q}$ holds can be captured by imposing $\kappa(\mathcal{Q})$ on D . Due to $D \models \mathcal{Q}$, it holds $D \not\models \kappa(\mathcal{Q})$. So, D is inconsistent wrt. $\kappa(\mathcal{Q})$, and could be repaired.

Repairs for (violations of) DCs are obtained by tuple deletions. Intuitively, a tuple that participates in a violation of $\kappa(\mathcal{Q})$ in D is an actual cause for \mathcal{Q} . S-minimal sets of tuples like this are expected to correspond to S-repairs for D and $\kappa(\mathcal{Q})$. More precisely, given an instance $D = D^n \cup D^x$, a BCQ \mathcal{Q} , and a tuple $t \in D^n$, we consider the class containing the sets of differences between D and those S- or C-repairs that do not contain t , and are obtained by removing a subset of D^n :

$$\mathcal{DF}^s(D, D^n, \kappa(\mathcal{Q}), t) = \{D \setminus D' \mid D' \in \text{Srep}(D, \kappa(\mathcal{Q})), t \in (D \setminus D') \subseteq D^n\}, \quad (1)$$

$$\mathcal{DF}^c(D, D^n, \kappa(\mathcal{Q}), t) = \{D \setminus D' \mid D' \in \text{Crep}(D, \kappa(\mathcal{Q})), t \in (D \setminus D') \subseteq D^n\}. \quad (2)$$

It holds $\mathcal{DF}^c(D, D^n, \kappa(\mathcal{Q}), t) \subseteq \mathcal{DF}^s(D, D^n, \kappa(\mathcal{Q}), t)$. Now, any $s \in \mathcal{DF}^s(D, D^n, \kappa(\mathcal{Q}), t)$ can be written as $s = s' \cup \{t\}$. From the S-minimality of S-repairs, $D \setminus (s' \cup \{t\}) \models \kappa(\mathcal{Q})$, but $D \setminus s' \models \neg \kappa(\mathcal{Q})$, i.e. $D \setminus (s' \cup \{t\}) \not\models \mathcal{Q}$, but $D \setminus s' \models \mathcal{Q}$. So, t is an actual cause for \mathcal{Q} with contingency set s' .

► **Proposition 2.** *Given $D = D^n \cup D^x$, and a BCQ \mathcal{Q} , $t \in D^n$ is an actual cause for \mathcal{Q} iff $\mathcal{DF}^s(D, D^n, \kappa(\mathcal{Q}), t) \neq \emptyset$.* ◀

► **Proposition 3.** *Given $D = D^n \cup D^x$, a BCQ \mathcal{Q} , and $t \in D^n$:*

- (a) *If $\mathcal{DF}^s(D, D^n, \kappa(\mathcal{Q}), t) = \emptyset$, then $\rho(t) = 0$.*
- (b) *Otherwise, $\rho(t) = \frac{1}{|s|}$, where $s \in \mathcal{DF}^s(D, D^n, \kappa(\mathcal{Q}), t)$ and there is no $s' \in \mathcal{DF}^s(D, D^n, \kappa(\mathcal{Q}), t)$ such that $|s'| < |s|$.* ◀

► **Corollary 4.** *Given $D = D^n \cup D^x$, and a BCQ \mathcal{Q} , $t \in D^n$ is a most responsible actual cause for \mathcal{Q} iff $\mathcal{DF}^c(D, D^n, \kappa(\mathcal{Q}), t) \neq \emptyset$.* ◀

► **Example 5.** (ex. 1 cont.) Consider the same instance D and query \mathcal{Q} . In this case, the DC $\kappa(\mathcal{Q})$ is, in Datalog notation, a negative rule: $\leftarrow S(x), R(x, y), S(y)$.

Here, $\text{Srep}(D, \kappa(\mathcal{Q})) = \{D_1, D_2, D_3\}$ and $\text{Crep}(D, \kappa(\mathcal{Q})) = \{D_1\}$, with $D_1 = \{R(a_4, a_3), R(a_2, a_1), R(a_3, a_3), S(a_4), S(a_2)\}$, $D_2 = \{R(a_2, a_1), S(a_4), S(a_2), S(a_3)\}$, $D_3 = \{R(a_4, a_3), R(a_2, a_1), S(a_2), S(a_3)\}$.

For tuple $R(a_4, a_3)$, $\mathcal{DF}^s(D, D, \kappa(\mathcal{Q}), R(a_4, a_3)) = \{D \setminus D_2\} = \{\{R(a_4, a_3), R(a_3, a_3)\}\}$, which, by Propositions 2 and 3, confirms that $R(a_4, a_3)$ is an actual cause, with responsibility $\frac{1}{2}$. For tuple $S(a_3)$, $\mathcal{DF}^s(D, D, \kappa(\mathcal{Q}), S(a_3)) = \{D \setminus D_1\} = \{S(a_3)\}$. So, $S(a_3)$ is an actual cause with responsibility 1. Similarly, $R(a_3, a_3)$ is an actual cause with responsibility $\frac{1}{2}$, because $\mathcal{DF}^s(D, D, \kappa(\mathcal{Q}), R(a_3, a_3)) = \{D \setminus D_2, D \setminus D_3\} = \{\{R(a_4, a_3), R(a_3, a_3)\}, \{R(a_3, a_3), S(a_4)\}\}$.

It holds $\mathcal{DF}^s(D, D, \kappa(\mathcal{Q}), S(a_2)) = \mathcal{DF}^s(D, D, \kappa(\mathcal{Q}), R(a_2, a_1)) = \emptyset$, because all repairs contain $S(a_2), R(a_2, a_1)$. This means they do not participate in the violation of $\kappa(\mathcal{Q})$ or contribute to make \mathcal{Q} true. So, they are not actual causes for \mathcal{Q} , confirming the result in Example 1.

$\mathcal{DF}^c(D, D, \kappa(\mathcal{Q}), S(a_3)) = \{S(a_3)\}$. From Corollary 4, $S(a_3)$ is the most responsible cause. ◀

► **Remark 6.** *The results in this section can be easily extended to unions of BCQs without built-ins, i.e. essentially FO monotone queries without built-ins. This can be done by associating a DC to each disjunct of the query, and considering the corresponding problems for database repairs wrt. several DCs (cf. Section 4.1).* ◀

4 Database Repairs From Actual Causes

We now characterize repairs for inconsistent databases wrt. a set of DCs in terms of actual causes, and reduce their computation to computation of causes. Consider an instance D for schema \mathcal{S} , and a set of DCs Σ on \mathcal{S} . For each $\kappa \in \Sigma$, of the form $\kappa: \leftarrow A_1(\bar{x}_1), \dots, A_n(\bar{x}_n)$, consider its associated *violation view* defined by a BCQ, namely $V^\kappa: \exists \bar{x}(A_1(\bar{x}_1) \wedge \dots \wedge A_n(\bar{x}_n))$. Next, consider the query obtained as the union of the individual violation views: $V^\Sigma := \bigvee_{\kappa \in \Sigma} V^\kappa$, a *union of BCQs* (UBCQs). Clearly, D violates (is inconsistent wrt.) Σ iff $D \models V^\Sigma$. It is easy to verify that D is consistent wrt. Σ iff $\mathcal{CS}(D, \emptyset, V^\Sigma) = \emptyset$, i.e. there are no actual causes for V^Σ to be true when all tuples are endogenous.

Now, let us collect all *S-minimal contingency sets* associated with an actual cause t for V^Σ :

$$\begin{aligned} \mathcal{CT}(D, D^n, V^\Sigma, t) &:= \{s \subseteq D^n \mid D \setminus s \models V^\Sigma, D \setminus (s \cup \{t\}) \not\models V^\Sigma, \text{ and} \\ &\quad \forall s'' \subsetneq s, D \setminus (s'' \cup \{t\}) \models V^\Sigma\}. \end{aligned} \quad (3)$$

Notice that for $s \in \mathcal{CT}(D, D^n, V^\Sigma, t)$, $t \notin s$. If $t \in \mathcal{CS}(D, \emptyset, V^\Sigma)$ and $s \in \mathcal{CT}(D, D^n, V^\Sigma, t)$, from the definition of actual cause and the S-minimality of s , it holds that $s'' = s \cup \{t\}$ is an S-minimal subset of D with $D \setminus s'' \not\models V^\Sigma$. So, $D \setminus s''$ is an S-repair for D . Then, the following holds.

► **Proposition 7.** For an instance D and a set DCs Σ , $D' \subseteq D$ is an S-repair for D wrt. Σ iff, for every $t \in D \setminus D'$: $t \in \mathcal{CS}(D, \emptyset, V^\Sigma)$ and $D \setminus (D' \cup \{t\}) \in \mathcal{CT}(D, D, V^\Sigma, t)$. ◀

To establish a connection between most responsible actual causes and C-repairs, collect the most responsible actual causes for V^Σ :

$$\mathcal{MRC}(D, V^\Sigma) := \{t \in D \mid t \in \mathcal{CS}(D, \emptyset, V^\Sigma), \nexists t' \in \mathcal{CS}(D, \emptyset, V^\Sigma) \text{ with } \rho(t') > \rho(t)\}.$$

► **Proposition 8.** For instance D and set of DCs Σ , $D' \subseteq D$ is a C-repair for D wrt. Σ iff, for every $t \in D \setminus D'$: $t \in \mathcal{MRC}(D, V^\Sigma)$ and $D \setminus (D' \cup \{t\}) \in \mathcal{CT}(D, D, V^\Sigma, t)$. ◀

Actual causes for V^Σ , with their contingency sets, account for the violation of some $\kappa \in \Sigma$. Removing those tuples from D should remove the inconsistency. From Propositions 7 and 8 we obtain:

► **Corollary 9.** Given an instance D and a set DCs Σ , the instance obtained from D by removing an actual cause, resp. a most responsible actual cause, for V^Σ together with any of its S-minimal, resp. C-minimal, contingency sets forms an S-repair, resp. a C-repair, for D wrt. Σ . ◀

► **Example 10.** Consider $D = \{P(a), P(e), Q(a, b), R(a, c)\}$ and $\Sigma = \{\kappa_1, \kappa_2\}$, with $\kappa_1: \leftarrow P(x), Q(x, y)$ and $\kappa_2: \leftarrow P(x), R(x, y)$. The violation views are $V^{\kappa_1}: \exists xy(P(x) \wedge Q(x, y))$ and $V^{\kappa_2}: \exists xy(P(x) \wedge R(x, y))$. For $V^\Sigma := V^{\kappa_1} \vee V^{\kappa_2}$, $D \models V^\Sigma$. D is inconsistent wrt. Σ .

With all tuples endogenous, $\mathcal{CS}(D, \emptyset, V^\Sigma) = \{P(a), Q(a, b), R(a, c)\}$. Its elements are associated with sets of S-minimal contingency sets: $\mathcal{CT}(D, D, V^\Sigma, Q(a, b)) = \{\{R(a, c)\}\}$,

$\mathcal{CT}(D, D, V^\Sigma, R(a, c)) = \{\{Q(a, b)\}\}$, $\mathcal{CT}(D, D, V^\Sigma, P(a)) = \{\emptyset\}$. From Corollary 9, and $\mathcal{CT}(D, D, V^\Sigma, R(a, c))$, $D_1 = D \setminus (\{R(a, c)\} \cup \{Q(a, b)\}) = \{P(a), P(e)\}$ is an S-repair. So is $D_2 = D \setminus (\{P(a)\} \cup \emptyset) = \{P(e), Q(a, b), R(a, c)\}$. These are the only S-repairs.

Furthermore, $\mathcal{MRC}(D, V^\Sigma) = \{P(a)\}$. From Corollary 9, D_2 is also a C-repair for D . ◀

An actual cause t with any of its S-minimal contingency sets determines a unique S-repair. The last example shows that, with different combinations of a cause and one of its contingency sets, we may obtain the same repair (e.g. for the first two \mathcal{CT} s). So, we may have more minimal contingency sets than minimal repairs. However, we may still have exponentially many minimal contingency sets, so as we may have exponentially many minimal repairs.

► **Example 11.** Consider $D = \{R(1, 0), R(1, 1), \dots, R(n, 0), R(n, 1), S(1), S(0)\}$ and the DC $\kappa: \leftarrow R(x, y), R(x, z), S(y), S(z)$. D is inconsistent wrt. κ . There are exponentially many S-repairs of D : $D' = D \setminus \{S(0)\}$, $D'' = D \setminus \{S(1)\}$, $D_1 = D \setminus \{R(1, 0), \dots, R(n, 0)\}$, ..., $D_{2^n} = D \setminus \{R(1, 1), \dots, R(n, 1)\}$. The C-repairs are only D' and D'' .

For the BCQ V^κ associated to κ , $D \models V^\kappa$, and $S(1)$ and $S(0)$ are actual causes for V^κ (counterfactual causes with responsibility 1). All tuples in R are actual causes, each with exponentially many S-minimal contingency sets. For example, $R(1, 0)$ has the S-minimal contingency set $\{R(2, 0), \dots, R(n, 0)\}$, among exponentially many others (any set built with just one element from each of the pairs $\{R(2, 0), R(2, 1)\}$, ..., $\{R(n, 0), R(n, 1)\}$ is one). ◀

The characterization results obtained so far extend those in [49] for single DCs.

4.1 Causes for unions of conjunctive queries

If we want to compute repairs wrt. sets of DCs from causes for UBCQs using, say Corollary 9, we first need an algorithm for computing the actual causes and their (minimal) contingency sets for UBCQs. These algorithms could be used as a first stage for the computation of S-repairs and C-repairs wrt. sets of DCs. However, these algorithms (cf. Section 4.2) are also interesting *per se*.

The PTIME algorithm for computing actual causes in [41] is for single conjunctive queries, but does not compute the actual causes' contingency sets. Actually, doing the latter increases the complexity, because *deciding responsibility*⁵ of actual causes is NP-hard [41] (which would be tractable if we could efficiently compute all (minimal) contingency sets).⁶ In principle, an algorithm for responsibilities can be used to compute C-minimal contingency sets, by iterating over all candidates, but Example 11 shows that there can be exponentially many of them.

We first concentrate on the problem of computing actual causes for UBCQs, without their contingency sets, which requires some notation.

► **Definition 12.** Given $\mathcal{Q} = C_1 \vee \dots \vee C_k$, with each C_i a BCQ, and an instance D :

- (a) $\mathfrak{S}(D)$ is the collection of all S-minimal subsets of D that satisfy a disjunct C_i of \mathcal{Q} ;
- (b) $\mathfrak{S}^n(D)$ consists of the S-minimal subsets s of D^n for which there exists a $s' \in \mathfrak{S}(D)$ with $s \subseteq s'$ and $s \setminus s' \subseteq D^x$. ◀

$\mathfrak{S}^n(D)$ contains all S-minimal sets of endogenous tuples that simultaneously (and possibly accompanied by exogenous tuples) make the query true. It is easy to see that $\mathfrak{S}(D)$ and

⁵ For a precise formulation, see Definition 31.

⁶ Actually, [41] presents a PTIME algorithm for computing responsibilities for a restricted class of CQs.

$\mathfrak{S}^n(D)$ can be computed in polynomial time in the size of D . Now, generalizing a result for CQs in [41], actual causes for a UBCQs can be computed in PTIME in the size of D without computing contingency sets.

► **Proposition 13.** *Given $D = D^x \cup D^n$ and a UBCQ \mathcal{Q} .*

- (a) *t is an actual cause for \mathcal{Q} iff there is $s \in \mathfrak{S}^n(D)$ with $t \in s$.*
- (b) *The decision problem (about membership of) $CPD := \{(D^x, D^n, t) \mid t \in D^n, \text{ and } t \in CS(D^n, D^x, \mathcal{Q})\}$ belongs to PTIME.* ◀

► **Example 14.** (ex. 10 cont.) Consider the query $\mathcal{Q}: \exists xy(P(x) \wedge Q(x, y)) \vee \exists xy(P(x) \wedge R(x, y))$, and assume that for D , $D^n = \{P(a), R(a, c)\}$ and $D^x = \{P(e), Q(a, b)\}$. It holds $\mathfrak{S}(D) = \{\{P(a), Q(a, b)\}, \{P(a), R(a, c)\}\}$. Since $\{P(a)\} \subseteq \{P(a), R(a, c)\}$, $\mathfrak{S}^n(D) = \{\{P(a)\}\}$. So, $P(a)$ is the only actual cause for \mathcal{Q} . ◀

4.2 Contingency sets for unions of conjunctive queries

It is possible to develop a (naive) algorithm that accepts as input an instance $D = D^n \cup D^x$, and a UBCQ \mathcal{Q} , and returns $CS(D, D^n, \mathcal{Q})$, and also, for each $t \in CS(D, D^n, \mathcal{Q})$, its (set of) S-minimal contingency sets $\mathcal{CT}(D, D^n, \mathcal{Q}, t)$. The basis for the algorithm is a correspondence between the actual causes for \mathcal{Q} with their contingency sets and a *hitting-set problem*.⁷

More precisely, for a fixed UBCQ \mathcal{Q} , consider the *hitting-set framework* $\mathfrak{H}^n(D) = \langle D^n, \mathfrak{S}^n(D) \rangle$, with $\mathfrak{S}^n(D)$ as in Definition 12. Different decision problems can be imposed on it. The S-minimal *hitting sets* (HSs) for $\mathfrak{H}^n(D)$ correspond to actual causes with their S-minimal contingencies for \mathcal{Q} . Most responsible causes for \mathcal{Q} are in correspondence with minimum hitting sets for $\mathfrak{H}^n(D)$. Notice that these hitting sets are all subsets of D^n .

► **Proposition 15.** *For $D = D^x \cup D^n$ and a UBCQ \mathcal{Q} , it holds:*

- (a) *t is an actual cause for \mathcal{Q} with S-minimal contingency set Γ iff $\Gamma \cup \{t\}$ is an S-minimal HS for $\mathfrak{H}^n(D)$.*
- (b) *t is a most responsible actual cause for \mathcal{Q} with C-minimal contingency set Γ iff $\Gamma \cup \{t\}$ is a minimum HS for $\mathfrak{H}^n(D)$.* ◀

► **Example 16.** (ex. 10 and 14 cont.) D and \mathcal{Q} are as before, but we now all tuples are endogenous. Here, $\mathfrak{S}(D) = \mathfrak{S}^n(D) = \{\{P(a), Q(a, b)\}, \{P(a), R(a, c)\}\}$. $\mathfrak{H}^n(D)$ has two S-minimal HSs: $H_1 = \{P(a)\}$ and $H_2 = \{Q(a, b), R(a, c)\}$. Each of them implicitly contains an actual cause (any of its elements) with an S-minimal contingency set (what's left after removing the actual cause). H_1 is also the C-minimal hitting set, and contains the most responsible actual cause, $P(a)$. ◀

► **Remark 17.** *For $\mathfrak{H}^n(D) = \langle D^n, \mathfrak{S}^n(D) \rangle$, $\mathfrak{S}^n(D)$ can be computed in PTIME, and its elements are bounded in size by $|\mathcal{Q}|$, which is the maximum number of atoms in one of \mathcal{Q} 's disjuncts. This is a special kind of hitting-set problems. For example, deciding if there is a hitting set of size at most k as been called the d -hitting-set problem [44], and d is the bound on the size of the sets in the set class. In our case, d would be $|\mathcal{Q}|$.* ◀

⁷ If \mathcal{C} is a collection of non-empty subsets of a set S , a subset $S' \subseteq S$ is a *hitting set* for \mathcal{C} if, for every $C \in \mathcal{C}$, $C \cap S' \neq \emptyset$. S' is an S-minimal HS if no proper subset of it is also an HS. S is a minimum HS if it has minimum cardinality.

4.3 Causality, repairs and consistent answers

Corollary 9 and Proposition 15 can be used to compute repairs. If the classes of S- and C-minimal HSs for $\mathfrak{H}^n(D)$ (with $D^n = D$) are available, computing S- and C-repairs will be in PTIME in the sizes of those classes. However, it is well known that computing minimal HSs is a complex problem. Actually, as Example 11 implicitly shows, we can have exponentially many of them in $|D|$; so as exponentially many minimal repairs for a D wrt. a denial constraint.⁸ So, the complexity of contingency sets computation is in line with the complexities of computing hitting sets and repairs.

The computation of causes, contingency sets, and most responsible causes via minimal/minimum HS computation can then be used to compute repairs and decide about repair questions. Since the HS problems in our case are of the d -hitting set kind, good algorithms and approximations for the latter (cf. Section 6.1) could be used in the context of repairs (all this via Corollary 9 and Proposition 15).

Consider an instance D (with all tuples endogenous) and a set Σ of DCs. For the disjunctive violation view V^Σ , the following result is obtained from Propositions 7 and 8, and Corollary 9.

► **Corollary 18.** *For an instance D and set DCs Σ , it holds:*

- (a) *For every $t \in \mathcal{CS}(D, V^\Sigma)$, there is an S-repair that does not contain t .*
- (b) *For every $t \in \mathcal{MRC}(D, V^\Sigma)$, there is a C-repair that does not contain t .*
- (c) *For every $D' \in \text{Srep}(D, \Sigma)$ and $D'' \in \text{Crep}(D, \Sigma)$, $D \setminus D' \subseteq \mathcal{CS}(D, V^\Sigma)$ and $D \setminus D'' \subseteq \mathcal{MRC}(D, V^\Sigma)$.* ◀

For a projection-free, and a possibly non-boolean CQ \mathcal{Q} , we are interested in its consistent answers from D wrt. Σ . For example, for $\mathcal{Q}(x, y, z) : R(x, y) \wedge S(y, z)$, the S-consistent (C-consistent) answers would be of the form (a, b, c) , where $R(a, b)$ and $S(b, c)$ belong to all S-repairs (C-repairs) of D . From Corollary 18, (a, b, c) is an S-consistent, resp. C-consistent, answer iff $R(a, b)$ and $S(b, c)$ belong to D , but they are not actual causes, resp. most responsible actual causes, for V^Σ .

► **Proposition 19.** *For an instance D , a set of DCs Σ , and a projection-free CQ $\mathcal{Q}(\bar{x}) : P_1(\bar{x}_1) \wedge \dots \wedge P_k(\bar{x}_k)$:*

- (a) *\bar{c} is an S-consistent answer iff, for each i , $P_i(\bar{c}_i) \in (D \setminus \mathcal{CS}(D, V^\Sigma))$.*
- (b) *\bar{c} is a C-consistent answer iff, for each i , $P_i(\bar{c}_i) \in (D \setminus \mathcal{MRC}(D, V^\Sigma))$.* ◀

► **Example 20.** (ex. 10 cont.) Consider $\mathcal{Q}(x) : P(x)$. We had $\mathcal{CS}(D, V^\Sigma) = \{P(a), Q(a, b), R(a, c)\}$, $\mathcal{MRC}(D, V^\Sigma) = \{P(a)\}$. Then, a is both an S- and a C-consistent answer. ◀

Notice that Proposition 19 can easily be extended to conjunction of ground atomic queries. Actually, from it we obtain the following result that will be useful later on.

► **Corollary 21.** *Given D , a set of DCs Σ , the ground atomic query $\mathcal{Q} : P(c)$ is C-consistently true if $P(c) \in D$ and it is not a most responsible cause for V^Σ .* ◀

► **Example 22.** For $D = \{P(a, b), R(b, c), R(a, d)\}$ and the DC $\kappa : \leftarrow P(x, y), R(y, z)$: $\mathcal{CS}(D, V^\kappa) = \mathcal{MRC}(D, V^\kappa) = \{P(a, b), R(b, c)\}$. From Proposition 19, the ground atomic query $\mathcal{Q} : R(a, d)$ is both S- and C-consistently true in D wrt. κ , because, $D \setminus \mathcal{CS}(D, V^\kappa) = D \setminus \mathcal{MRC}(D, V^\kappa) = \{R(a, d)\}$. ◀

⁸ An example of this kind for FDs is given in [4]. However, FDs form a special class of DCs that involve equality. Consequently, their violation views involve inequality.

The CQs considered in Proposition 19 and its Corollary 21 are not the particularly interesting, but will use those results to obtain relevant results for causality later on, e.g. Theorem 41.

5 Diagnosis: Query Answer Causality and Repairs

Let $D = D^n \cup D^x$ be an instance for schema \mathcal{S} , and $\mathcal{Q} : \exists \bar{x}(P_1(\bar{x}_1) \wedge \cdots \wedge P_m(\bar{x}_m))$, a BCQ. Assume \mathcal{Q} is, possibly unexpectedly, true in D . So, for the associated DC $\kappa(\mathcal{Q}) : \forall \bar{x} \neg(P_1(\bar{x}_1) \wedge \cdots \wedge P_m(\bar{x}_m))$, $D \not\models \kappa(\mathcal{Q})$. \mathcal{Q} is our *observation*, for which we want to find explanations, using a consistency-based diagnosis approach.

For each predicate $P \in \mathcal{P}$, we introduce predicate Ab_P , with the same arity as P . A tuple in its extension is *abnormal* for P . The “system description”, SD , includes, among other elements, the original database, expressed in logical terms, and the DC being true “under normal conditions”. More precisely, we consider the following *diagnosis problem*, $\mathcal{M} = (SD, D^n, \mathcal{Q})$, associated to \mathcal{Q} . The FO system description, SD , contains the following elements:

- (a) $Th(D)$, which is Reiter’s logical reconstruction of D as a FO theory [48] (cf. Example 23).
- (b) Sentence $\kappa(\mathcal{Q})^{Ab}$, which is $\kappa(\mathcal{Q})$ rewritten as follows:

$$\kappa(\mathcal{Q})^{Ab} : \forall \bar{x} \neg(P_1(\bar{x}_1) \wedge \neg Ab_{P_1}(\bar{x}_1) \wedge \cdots \wedge P_m(\bar{x}_m) \wedge \neg Ab_{P_m}(\bar{x}_m)). \quad (4)$$

This formula can be refined by applying the abnormality predicate, Ab , to endogenous tuples only. For this we need to use additional auxiliary predicates End_P , with the same arity of $P \in \mathcal{S}$, which contain the endogenous tuples in P ’s extension (see Example 23).

- (c) The inclusion dependencies: $\forall \bar{x}(Ab_P(\bar{x}) \rightarrow P(\bar{x}))$, $\forall \bar{x}(End_P(\bar{x}) \rightarrow P(\bar{x}))$, and $\forall \bar{x}(Ab_P(\bar{x}) \rightarrow End_P(\bar{x}))$, for each $P \in \mathcal{P}$.

The last entry, \mathcal{Q} , in \mathcal{M} is the *observation*, which together with SD will produce an inconsistent theory, because we make the initial and explicit assumption that all the abnormality predicates are empty (equivalently, that all tuples are normal), i.e. we consider, for each predicate P , the sentence⁹

$$\forall \bar{x}(Ab_P(\bar{x}) \rightarrow \mathbf{false}), \quad (5)$$

where, **false** is a propositional atom that is always false. Actually, the second entry in \mathcal{M} tells us how we can restore consistency, namely by (minimally) changing the abnormality condition on tuples in D^n . In other words, the rules (5) are subject to qualifications: some endogenous tuples may be abnormal. Each diagnosis shows an S-minimal set of endogenous tuples that are abnormal.

► **Example 23.** (ex. 1 cont.) For the instance $D = \{S(a_3), S(a_4), R(a_4, a_3)\}$, with $D^n = \{S(a_4), S(a_3)\}$, consider the diagnostic problem $\mathcal{M} = (SD, \{S(a_4), S(a_3)\}, \mathcal{Q})$, with SD containing:

- (a) Predicate completion axioms: $\forall xy(R(x, y) \leftrightarrow x = a_4 \wedge y = a_3)$, $\forall x(S(x) \leftrightarrow x = a_3 \vee x = a_4)$, $\forall xy(End_R(x, y) \leftrightarrow \mathbf{false})$, $\forall x(End_S(x) \leftrightarrow x = a_3 \vee x = a_4)$.
Unique names assumption: $a_4 \neq a_3$.
- (b) $\kappa(\mathcal{Q})^{Ab} : \forall xy \neg(S(x) \wedge End_S(x) \wedge \neg Ab_S(x) \wedge R(x, y) \wedge End_R(x, y) \wedge \neg Ab_R(x, y) \wedge S(y) \wedge \neg Ab_S(y))$.

⁹ Notice that these can also be seen as DCs, since they can be written as $\forall \bar{x} \neg Ab_P(\bar{x})$.

- (c) $\forall xy(Ab_R(x, y) \rightarrow R(x, y)), \forall x(Ab_S(x) \rightarrow S(x)), \forall xy(End_R(x, y) \rightarrow R(x, y)),$
 $\forall x(End_S(x) \rightarrow S(x)), \forall xy(Ab_R(x, y) \rightarrow End_R(x, y)), \forall x(Ab_S(x) \rightarrow End_S(x)).$
 The normality assumptions for tuples are: $\forall xy(Ab_R(x, y) \rightarrow \mathbf{false}), \forall x(Ab_S(x) \rightarrow \mathbf{false}).$ ◀

Now, the observation is \mathcal{Q} (being true), obtained by evaluating query \mathcal{Q} on (theory of) D . In this case, $D \not\models \kappa(\mathcal{Q})$. Since all the abnormality predicates are assumed to be empty, $\kappa(\mathcal{Q})$ is equivalent to $\kappa(\mathcal{Q})^{Ab}$, which also becomes false wrt D . As a consequence, $SD \cup \{(5)\} \cup \{\mathcal{Q}\}$ is an inconsistent FO theory. A diagnosis is a set of endogenous tuples that, by becoming abnormal, restore consistency.

► **Definition 24.**

- (a) A *diagnosis* for \mathcal{M} is a $\Delta \subseteq D^n$, such that $SD \cup \{Ab_P(\bar{c}) \mid P(\bar{c}) \in \Delta\} \cup \{\neg Ab_P(\bar{c}) \mid P(\bar{c}) \in D \setminus \Delta\} \cup \{\mathcal{Q}\}$ is consistent.
 (b) $\mathcal{D}(\mathcal{M}, t)$ denotes the set of S-minimal diagnoses for \mathcal{M} that contain a tuple $t \in D^n$. (c) $\mathcal{MCD}(\mathcal{M}, t)$ denotes the set of C-minimal diagnoses in $\mathcal{D}(\mathcal{M}, t)$. ◀

By definition, $\mathcal{MCD}(\mathcal{M}, t) \subseteq \mathcal{D}(\mathcal{M}, t)$. Diagnoses for \mathcal{M} and actual causes for \mathcal{Q} are related.

- **Proposition 25.** Consider $D = D^n \cup D^x$, a BCQ \mathcal{Q} , and the diagnosis problem \mathcal{M} associated to \mathcal{Q} . Tuple $t \in D^n$ is an actual cause for \mathcal{Q} iff $\mathcal{D}(\mathcal{M}, t) \neq \emptyset$. ◀

The responsibility of an actual cause t is determined by the cardinality of the diagnoses in $\mathcal{MCD}(\mathcal{M}, t)$.

- **Proposition 26.** For $D = D^n \cup D^x$, a BCQ \mathcal{Q} , the associated diagnosis problem \mathcal{M} , and a tuple $t \in D^n$, it holds:
 (a) $\rho_D(t) = 0$ iff $\mathcal{MCD}(\mathcal{M}, t) = \emptyset$.
 (b) Otherwise, $\rho_D(t) = \frac{1}{|s|}$, where $s \in \mathcal{MCD}(\mathcal{M}, t)$. ◀

- **Example 27.** (ex. 23 cont.) \mathcal{M} has two diagnosis: $\Delta_1 = \{S(a_3)\}$ and $\Delta_4 = \{S(a_4)\}$. Here, $\mathcal{D}(\mathcal{M}, S(a_3)) = \mathcal{MCD}(\mathcal{M}, S(a_3)) = \{\{S(a_3)\}\}$ and $\mathcal{D}(\mathcal{M}, S(a_4)) = \mathcal{MCD}(\mathcal{M}, S(a_4)) = \{\{S(a_4)\}\}$. From Propositions 25 and 26, $S(a_3)$ and $S(a_4)$ are actual cases, with responsibility 1. ◀

In consistency-based diagnosis, minimal diagnoses can be obtained as S-minimal HSs of the collection of S-minimal *conflict sets* (cf. Section 2) [47]. In our case, conflict sets are S-minimal sets of endogenous tuples that, if not abnormal (only endogenous ones can be abnormal), and together, and possibly in combination with exogenous tuples, make (4) false. It is easy to verify that the conflict sets of \mathcal{M} coincide with the sets in $\mathfrak{S}(D^n)$ (cf. Definition 12 and Remark 17). As a consequence, conflict sets for \mathcal{M} can be computed in PTIME, the HSs for \mathcal{M} contain actual causes for \mathcal{Q} , and the HS problem for the diagnosis problems is of the d -hitting-set kind. The connection between consistency-based diagnosis and causality allows us, in principle, to apply techniques for the former, e.g. [25, 43], to the latter.

- **Example 28.** (ex. 23 cont.) The diagnosis problem $\mathcal{M} = (SD, \{S(a_4), S(a_3)\}, \mathcal{Q})$ gives rise to the HS framework $\mathfrak{H}^n(D) = \langle \{S(a_4), S(a_3)\}, \{\{(S(a_3), S(a_4))\}\} \rangle$, with $\{S(a_3), S(a_4)\}$ corresponding to the conflict set $c = \{S(a_4), S(a_3)\}$. $\mathfrak{H}^n(D)$ has two minimum HSs: $\{S(a_3)\}$ and $\{S(a_4)\}$, which are the S-minimal diagnosis for \mathcal{M} . Then, the two tuples are actual causes for \mathcal{Q} (cf. Proposition 25). From Proposition 26, $\rho_D(S(a_3)) = \rho_D(S(a_4)) = 1$. ◀

The solutions to the diagnosis problem can be used for computing repairs.

► **Proposition 29.** Consider a database instance D with only endogenous tuples, a set of DCs of the form $\kappa: \forall \bar{x} \neg (P_1(\bar{x}_1) \wedge \dots \wedge P_m(\bar{x}_m))$, and their associated “abnormality” integrity constraints¹⁰ in (4) (in this case we do not need End_P atoms). Each S -minimal diagnosis Δ gives rise to an S -repair of D , namely $D_\Delta = D \setminus \{P(\bar{c}) \in D \mid Ab_P(\bar{c}) \in \Delta\}$; and every S -repair can be obtained in this way. Similarly, for C -repairs using C -minimal diagnoses. ◀

► **Example 30.** (ex. 27 cont.) The instance $D = \{S(a_3), S(a_4), R(a_4, a_3)\}$ has three (both S - and C -) repairs wrt. the DC $\kappa: \forall xy \neg (S(x) \wedge R(x, y) \wedge S(y))$, namely $D_1 = \{S(a_3)\}$, $D_2 = \{S(a_4)\}$, and $D_3 = \{R(a_4, a_3)\}$. They can be obtained as $D_{\Delta_1}, D_{\Delta_2}, D_{\Delta_3}$ from the only (S - and C -) diagnoses, $\Delta_1 = \{S(a_3)\}$, $\Delta_2 = \{S(a_4)\}$, $\Delta_3 = \{R(a_4, a_3)\}$, resp. ◀

The kind of diagnosis problem we introduced above can be formulated as a *preferred-repair problem* [9, sec. 2.5] (see [51] for a general approach to prioritized repairs). For this, it is good enough to materialize tables for the auxiliary predicates Ab_P and End_P , and consider the DCs of the form (4) (with the End_P atoms if not all tuples are endogenous), plus the DCs (5). The initial extensions for the Ab_P predicates are empty. If D is inconsistent wrt. this set of DCs, the S -repairs that are obtained by only *inserting* endogenous tuples into the extensions of the Ab_P predicates correspond to S -minimal diagnosis, and each S -minimal diagnosis can be obtained in this way.

6 Complexity Results

There are three main computational problems in database causality. For a BCQ \mathcal{Q} and database D , they are:

- (a) The *causality problem* (CP) that is about computing the actual causes for \mathcal{Q} .
- (b) The *responsibility problem* (RP) that is about computing the responsibility $\rho_D(t)$ of a given actual cause t . Since a tuple that is not an actual cause has responsibility 0, the latter problem subsumes the former.
- (c) Computing the *most responsible actual causes* (MRC).

These problems have corresponding *decision versions*. Both CP and its decision version, CPD, are solvable in polynomial time [41], which can be extended to UBCQs (cf. Proposition 13). We consider the decision version of the second problem.

► **Definition 31.** For a BCQ \mathcal{Q} , the *responsibility decision problem* (RPD) is (deciding about membership of) $\mathcal{RPD}(\mathcal{Q}) = \{(D^x, D^n, t, v) \mid t \in D^n, v \in \{0\} \cup \{\frac{1}{k} \mid k \in \mathbb{N}^+\}, D := D^x \cup D^n \models \mathcal{Q} \text{ and } \rho_D(t) > v\}$. ◀

The complexity analysis of RPD in [41] is restricted to conjunctive queries without self-joins, for which a dichotomy result holds: depending on the syntactic structure of a query, RPD is either in PTIME or is NP-hard. Here, we generalize the complexity analysis for RPD to general CQs.

We will also investigate the decision version, MRCD, of MRC, i.e. about deciding most responsible actual causes. This is a natural problem, because actual causes with the highest responsibility tend to provide most interesting explanations for query answers [41, 42].

► **Definition 32.** For a BCQ \mathcal{Q} , the *most responsible cause decision problem* is $\mathcal{MRCD}(\mathcal{Q}) = \{(D^x, D^n, t) \mid t \in D^n \text{ and } 0 < \rho_D(t) \text{ is a maximum for } D := D^x \cup D^n\}$. ◀

¹⁰Notice that these are not denial constraints.

We start by analyzing a more basic decision problem: *S-minimal contingency checking* (MCCD).

► **Definition 33.** For a BCQ \mathcal{Q} , $\text{MCCD}(\mathcal{Q}) := \{(D^x, D^n, t, \Gamma) \mid \Gamma \in \mathcal{CT}(D^n \cup D^x, D^n, \mathcal{Q}, t)\}$ (cf. (3)). ◀

Due to the results in Sections 3 and 4, it is clear that there is a close connection between MCCD and the *S-repair checking* problem in consistent query answering [9, chap. 5], about deciding if instance D' is an S-repair of instance D wrt. a set of integrity constraints. Actually, the following result is obtained from the membership of the S-repair checking problem of LOGSPACE for DCs [1, prop. 5].

► **Proposition 34.** For a BCQ \mathcal{Q} , $\text{MCCD}(\mathcal{Q}) \in \text{PTIME}$. ◀

We could also consider the decision problem defined as in Definition 33, but with C-minimal Γ . We will not use results about this problem in the following. Furthermore, its connection with the C-repair checking problem is less direct. As one can see from Section 3, C-minimal contingency sets correspond to a repair semantics somewhere between the S-minimal and C-minimal repair semantics (a subclass of *Srep*, but a superclass of *Crep*): It is about an S-minimal repair with minimum cardinality that does not contain a particular tuple.

Now we establish that RPD is NP-complete for CQs in general. The NP-hardness is shown in [41]. Membership of NP is obtained using Proposition 34.

► **Theorem 35.**

(a) For every BCQ \mathcal{Q} , $\text{RPD}(\mathcal{Q}) \in \text{NP}$.

(b) [41] There are CQs \mathcal{Q} for which $\text{RPD}(\mathcal{Q})$ is NP-hard. ◀

In order to better understand the complexity of the problem, RP, of computing responsibility, we will investigate the *functional*, non-decision version of the problem.

The main source of complexity when computing responsibilities is related to the hitting-set problem associated to $\mathfrak{H}^n(D) = \langle D^n, \mathfrak{S}^n(D) \rangle$ in Remark 17. In this case, it is about computing the cardinality of a minimum hitting set that contains a given vertex (tuple) t . That this is a kind of *d-hitting-set problem* [44] will be useful in Section 6.1.

Our responsibility problem can also be seen as a *vertex cover problem* on the hypergraph $\mathfrak{G}^n(D) = \langle D^n, \mathfrak{E}^n(D) \rangle$ associated to $\mathfrak{H}^n(D) = \langle D^n, \mathfrak{S}^n(D) \rangle$. In it, the set of hyperedges $\mathfrak{E}^n(D)$ coincides with the collection $\mathfrak{S}^n(D)$. Determining the responsibility of a tuple t becomes the problem on hypergraphs of determining the size of a minimum vertex cover (VC)¹¹ that contains vertex t (among all VCs that contain the vertex). Again, in this problem the hyperedges are bounded by $|\mathcal{Q}|$.¹²

► **Example 36.** For $\mathcal{Q} : \exists xy(P(x) \wedge R(x, y) \wedge P(y))$, and $D = D^n = \{P(a), P(c), R(a, c), R(a, a)\}$, $\mathfrak{S}(D) = \mathfrak{S}^n(D) = \{\{P(a), R(a, a)\}, \{P(a), P(c), R(a, c)\}\}$. D is the set of vertices of hypergraph $\mathfrak{G}^n(D)$, and its hyperedges are $\{P(a), R(a, a)\}$, $\{P(a), P(c), R(a, c)\}$. The following are the minimal VCs: $vc_1 = \{P(a)\}$, $vc_2 = \{P(c), R(a, a)\}$, $vc_3 = \{R(a, a), R(a, c)\}$. Then, $P(a)$ is an actual cause with responsibility 1. The other tuples are actual causes with responsibility $\frac{1}{2}$. ◀

¹¹ A set of vertices is a VC for a hypergraph if it intersects every hyperedge. Obviously, when we talk of *minimum* VC, we are referring to minimal in cardinality.

¹² We recall that repairs of databases wrt. DCs can be characterized as maximal independent sets of *conflict hypergraphs* (conflict graphs in the case of FDs) whose vertices are the database tuples, and hyper-edges connect tuples that together violate a DC [4, 17].

To simplify the presentation, we will formulate and address our computational problems as problems for graphs (instead of hypergraphs). However, our results still hold for hypergraphs [39]. Actually, the following *representation lemma* holds.

► **Lemma 37.** *There is a fixed database schema \mathcal{S} and a BCQ $\mathcal{Q} \in L(\mathcal{S})$, without built-ins, such that, for every graph $G = (V, E)$ and $v \in V$, there is an instance D for \mathcal{S} and a tuple $t \in D$, such that the size of a minimum VC of G containing v equals the responsibility of t as an actual cause for \mathcal{Q} .* ◀

Having represented our responsibility problem as a graph-theoretic problem, we first consider the following *membership minimal VC problem* (MMVC): Given a graph $G = (V, E)$, a vertex $v \in V$, determine the size of a minimum VC of G that contains v .

► **Lemma 38.** *Given a graph G and a vertex v in it, there is a graph G' extending G that can be constructed in polynomial time in $|G|$, such that the size of a minimum VC for G that contains v and the size of a minimum VC for G' coincide.* ◀

From this lemma and the $FP^{NP(\log(n))}$ -completeness of determining the size of a maximum clique in a graph [35], we obtain:

► **Proposition 39.** *MMVC problem for graphs is $FP^{NP(\log(n))}$ -complete.* ◀

From Lemma 37 and Proposition 39 we obtain the complexity result for RP. Membership can also be obtained from Theorem 35.

► **Theorem 40.**

- (a) *For every BCQ without built-ins, \mathcal{Q} , computing the responsibility of a tuple as a cause for \mathcal{Q} is in $FP^{NP(\log(n))}$.*
- (b) *There is a database schema and a BCQ \mathcal{Q} , without built-ins, such that computing the responsibility of a tuple as a cause for \mathcal{Q} is $FP^{NP(\log(n))}$ -complete.* ◀

Now we address the most responsible causes problem, MRCD. We use the connection with consistent query answering of Section 4.3, namely Corollary 21, and the $P^{NP(\log(n))}$ -completeness of consistent query answering under the C-repair semantics for queries that are conjunctions of ground atoms and a particular DC [39, theo. 4].

► **Theorem 41.**

- (a) *For every BCQ without built-ins, $MRCD(\mathcal{Q}) \in P^{NP(\log(n))}$.*
- (b) *There is a database schema and a BCQ \mathcal{Q} , without built-ins, for which $MRCD(\mathcal{Q})$ is $P^{NP(\log(n))}$ -complete.* ◀

From Proposition 15 and the $FP^{NP(\log(n))}$ -completeness of determining the size of C-repairs for DCs [39, theo. 3], we obtain the following for the computation of the highest responsibility value.

► **Proposition 42.**

- (a) *For every BCQ without built-ins, computing the responsibility of the most responsible causes is in $FP^{NP(\log(n))}$.*
- (b) *There is a database schema and a BCQ \mathcal{Q} , without built-ins, for which computing the responsibility of the most responsible causes is $FP^{NP(\log(n))}$ -complete.* ◀

6.1 FPT of responsibility

We need to cope with the intractability of computing most responsible causes. The area of *fixed parameter tractability* (FPT) [26] provides tools to attack this problem. In this regard, we recall that a decision problem with inputs of the form (I, p) , where p is a distinguished parameter of the input, is fixed parameter tractable (or belongs to the class FPT), if it can be solved in time $O(f(|p|) \cdot |I|^c)$, where c and the hidden constant do not depend on $|p|$ or $|I|$, and f does not depend on $|I|$.

In our case, the *parameterized version of the decision problem* $\mathcal{RPD}(\mathcal{Q})$ (cf. Definition 31) is denoted with $\mathcal{RPD}^p(\mathcal{Q})$, and the distinguished parameter is k , such that $v = \frac{1}{k}$. That $\mathcal{RPD}^p(\mathcal{Q})$ belongs to FPT can be obtained from its formulation as a d -hitting-set problem (d being the fixed upper bound on the size of the sets in the set class); in this case about deciding if there is a HS that contains the given tuple t that has cardinality smaller than k . This problem belongs to FPT.

► **Theorem 43.** *For every BCQ \mathcal{Q} , $\mathcal{RPD}^p(\mathcal{Q})$ belongs to FPT, where the parameter is the inverse of the responsibility bound.* ◀

The proof of this result is interesting *per se*, and we sketch it here. First, there is a PTIME parameterized algorithm for the d -hitting-set problem about deciding if there is a HS of size at most k that runs in time $O(e^k + n)$, with n the size of the underlying set and $e = d - 1 + o(d^{-1})$ [44]. In our case, $n = |D|$, and $d = |\mathcal{Q}|$ (cf. also [24]).

Now, to decide if the responsibility of a given tuple t is greater than $v = \frac{1}{k}$, we consider the associated hypergraph $\mathfrak{G}^n(D)$, and we decide if it has a VC that contains t and whose size is less than k . In order to answer this, we use Lemma 38, and build the extended hypergraph \mathfrak{G}' . The size of a minimum VC for \mathfrak{G}' gives the size of the minimum VC of $\mathfrak{G}^n(D)$ that contains t . If $\mathfrak{G}^n(D)$ has a VC that contains t of size less than k , then \mathfrak{G}' has a VC of size less than k . If \mathfrak{G}' has a VC of size less than k , its minimum size for a VC is less than k . Since this minimum is the same as the size of a minimum VC for $\mathfrak{G}^n(D)$ that contains t , $\mathfrak{G}^n(D)$ has a VC of size less than k that contains t . As a consequence, it is good enough to decide if \mathfrak{G}' has a VC of size less than k . For this, we use the HS formulation of this hypergraph problem, and the already mentioned FPT algorithm.

This result and the corresponding algorithm show that the higher the required responsibility degree, the lower the computational effort needed to compute the actual causes with at least that level of responsibility. In other terms, parameterized algorithms are effective for computing actual causes with high responsibility or most responsible causes. In general, parameterized algorithms are very effective when the parameter is relatively small [26].

Now, in order to compute most responsible causes, we could apply, for each actual cause t , the just presented FPT algorithm on the hypergraph $\mathfrak{G}^n(D)$, starting with $k = 1$, i.e. asking if there is VC of size less than 1 that contains t . If the algorithm returns a positive result, then t is a counterfactual cause, and has responsibility 1. Otherwise, the algorithm will be launched with $k = 2, 3, \dots, |D^n|$, until a positive result is returned. (The procedure can be improved through binary search on $k = 1, 2, 3, \dots, m$, with m possibly much smaller than $|D|$.)

The complexity results and algorithms provided in this section can be extend to UBCQs. This is due to Remark 6 and the construction of $\mathfrak{G}^n(D)$, which the results in this section build upon.

For the d -hitting-set problem there are also efficient parameterized approximation algorithms [11]. They could be used to approximate the responsibility problem. Furthermore, approximation algorithms developed for the minimum VC problem on bounded hypergraphs

[31, 45] should be applicable to approximate most responsible causes for query answers. Via the causality/repair connection (cf. Section 4.3), it should be possible to develop approximation algorithms to compute S-repairs of particular sizes, C-repairs, and consistent query answers wrt. DCs.

6.2 The causality dichotomy's reflection on repairs

In [41] the class of *linear* CQs is introduced. For them, computing tuple responsibilities is tractable. Roughly speaking, a BCQ is linear if its atoms can be ordered in a way that every variable appears in a continuous sequence of atoms, e.g. $Q_1: \exists xvyu(A(x) \wedge S_1(x, v) \wedge S_2(v, y) \wedge R(y, u) \wedge S_3(y, z))$ is linear, but not $Q_2: \exists xyz(A(x) \wedge B(y) \wedge C(z) \wedge W(x, y, z))$, for which RPD is *NP*-hard [41]. The class of BCQs for which computing responsibility (more precisely, our \mathcal{RPD} decision problem) is tractable can be extended to *weakly linear*.¹³ Now, the dichotomy result in [41] says that for a BCQ Q without self-joins, RDP is tractable when Q is weakly-linear, but *NP*-hard, otherwise. Due to the causality/repair connection of Section 4, we can obtain the following results for database repairs.

► Theorem 44.

- (a) For single weakly-linear DCs, C-repair checking and deciding if the size of a C-repair is larger than a bound are both tractable.¹⁴
- (b) For single, self-join free DCs κ , and the problem $\text{RepSize}(\kappa)$ of deciding if there is a repair D' for a given input instance D and a tuple $t \in D$ with $|D'| \geq m$ and $t \notin D'$,¹⁵ the following dichotomy holds: (b1) If κ is weakly-linear, $\text{RepSize}(\kappa)$ is tractable. (b2) Otherwise, it is *NP*-complete. ◀

This dichotomy result for repairs shows that interesting results in one of the areas (causality, in this case) have counterparts in some of the others. The form the reincarnation of the known result takes in the new area (repairs, in this case) is interesting *per se*.

Notice that both problems in (a) in Theorem 44 may be intractable even for single DCs [39]. More specifically, C-repair checking can be *coNP*-hard for single DCs [39, 1]. Actually, the single DC used in [39, lemma 4] is of the form $\kappa: \leftarrow V(x), V(y), E(x, y, z)$, whose associated BCQ is not weakly-linear. As a matter of fact, this BCQ is a *NP*-hard for RDP [41].

7 Discussion and Conclusions

In this research we have unveiled and formalized some first interesting relationships between causality in databases, database repairs, and consistency-based diagnosis. These connections allow us to apply results and techniques developed for each of them to the others. This is particularly beneficial for causality in databases, where still a limited number of results and techniques have been obtained or developed.

The connections we established here inspired complexity results for causality, e.g. Theorems 40 and 41, and were used to prove them. We appealed to several non-trivial results (and the proofs thereof) about repairs/CQA obtained in [39]. It is also the case that the well-established hitting-set approach to diagnosis inspired a similar approach to causal

¹³ Computing sizes of minimum contingency sets is reduced to the max-flow/min-cut problem in a network.

¹⁴ A DC κ is weakly-linear if the corresponding BCQ V^κ is weakly-linear. In this way any adjective that applies to BCQs can be applied to DCs.

¹⁵ More precisely, D' is a subset of D that satisfies κ . Here, $0 \leq m \leq n = |D|$.

responsibility, which in its turn allowed us to obtain results about its fixed-parameter tractability. It is also the case that diagnostic reasoning, as a form of non-monotonic reasoning, can provide a solid foundation for causality in databases and query answer explanation, in general [15, 16].

Our work creates a theoretical basis for deeper and mathematically more complex investigations. In particular, it also opens interesting research directions, some of which are briefly discussed below.

Preferred causes for queries. In Section 3 we characterized causes and most responsible causes in terms of S-repairs and C-repairs, resp. This could be generalized by using the notion of *preferred repair* [51]. These are repairs whose minimization correspond to a *priority relationship*, \preceq , between instances. Let assume it defines a corresponding class of preferred repairs, $\preceq Rep$. Inspired by (1), we can define, for a BCQ Q : $\mathcal{DF}^{\preceq}(D, D^n, \kappa(Q), t) := \{D \setminus D' \mid D' \in \preceq Rep(D, \kappa(Q)), t \in (D \setminus D') \subseteq D^n\}$, and, $t \in D^n$ is a \preceq -cause iff $\mathcal{DF}^{\preceq}(D, D^n, \kappa(Q), t) \neq \emptyset$. In this way, a whole class of preferences on causes can be introduced, which is natural problem [42].¹⁶

Endogenous repairs. The partition of a database into endogenous and exogenous tuples may also be of interest in the context of repairs. Considering that we should have more control on endogenous tuples than on exogenous ones, which may come from external sources, it makes sense to consider *endogenous repairs*. They are obtained by updates (of any kind) on endogenous tuples. For example, in the case of DCs, endogenous repairs would be obtained by deleting endogenous tuples only. If there are no repairs based on endogenous tuples, a preference condition could be imposed on repairs [54, 51], privileging those that change exogenous the least. (Of course, it could also be the other way around, that is we may feel more inclined to change exogenous tuples than our endogenous ones.)

As a further extension, it could be possible to assume that combinations of (only) exogenous tuples never violate the integrity constraints, which could be checked at upload time. In this sense, there would be a part of the database that is considered to be consistent, while the other is subject to possible repairs. (For slightly related research, see [28].)

Objections to causality. Causality as introduced by Halpern and Pearl in [29, 30], aka. HP-causality, is the basis for the notion of causality in [41]. HP-causality has been the object of some criticism [32], which is justified in some (more complex, non-relational) settings, specially due to the presence of different kinds of *logical variables* (or lack thereof). In our context the objections do not apply: variables just say that a certain tuple belongs to the instance (or not); and for relational databases the closed-world assumption applies. In [32], the definition of HP-causality is slightly modified. In our setting, this modified definition does not change actual causes or their properties.

ASP specification of causes. S-repairs can be specified by means of *answer set programs* (ASPs) [3, 6], and C-repairs too, with the use of weak program constraints [3]. This should allow for the introduction of ASPs in the context of causality, for specification and reasoning. There are also ASP-based specifications of diagnosis [23] that could be brought into a more complete picture.

¹⁶In [40] the possibility of introducing weights in the partition is considered, in this way imposing a form of preference on causes.

Causes and functional dependencies, and beyond. Functional dependencies are DCs with conjunctive violation views with inequality, and are still monotonic. There is much research on repairs and consistent query answering for functional dependencies, and more complex integrity constraints [9]. In causality, mostly CQs without built-ins have been considered. The repair connection could be exploited to obtain results for causality and CQs with inequality, and also other classes of queries.

View updates and abduction. Abduction [19, 22] is another form of model-based diagnosis, and is related to the subjects investigated in this work. The *view update problem*, about updating a database through views, is a classical problem in databases that has been treated through abduction [33, 20]. User knowledge imposed through view updates creates or reflects *uncertainty* about the base data, because alternative base instances may give an account of the intended view updates. The view update problem, specially in its particular form of *deletion propagation*, has been recently related in [37, 38] to causality as introduced in [41]. (Notice only tuple deletions are used with violation views and repairs associated to DCs.)

Database repairs are also related to the view update problem. Actually, *answer set programs* (ASP) for database repairs [6] implicitly repair the database by updating intentional, annotated predicates. Even more, in [8], in order to protect sensitive information, databases are explicitly and virtually “repaired” through secrecy views that specify the information that has to be kept secret. These are prioritized repairs that have been specified via ASPs. Abduction has been explicitly applied to database repairs [5]. The deep interrelations between causality, abductive reasoning, view updates and repairs are the objects of our ongoing research efforts [10].

Acknowledgments. Research funded by NSERC Discovery, and the NSERC Strategic Network on Business Intelligence (BIN). Conversations with Alexandra Meliou during Leo Bertossi’s visit to U. of Washington in 2011 are much appreciated. He is also grateful to Dan Suciu and Wolfgang Gatterbauer for their hospitality. L. Bertossi is grateful to Benny Kimelfeld for stimulating conversations. Part of the research was developed by L. Bertossi at *LogicBlox* and *The Center for Semantic Web Research* (Chile). Their support is much appreciated.

References

- 1 Afrati, F. and Kolaitis, P. Repair Checking in Inconsistent Databases: Algorithms and Complexity. *Proc. ICDT 2009* pp. 31-41.
- 2 Arenas, M., Bertossi, L. and Chomicki, J. Consistent Query Answers in Inconsistent Databases. *Proc. ACM PODS*, 1999, pp. 68-79.
- 3 Arenas, M., Bertossi, L. and Chomicki, J. Answer Sets for Consistent Query Answers. *Theory and Practice of Logic Programming*, 2003, 3(4&5):393-424.
- 4 Arenas, M., Bertossi, L., Chomicki, J., He, X., Raghavan, V. and Spinrad, J. Scalar Aggregation in Inconsistent Databases. *Theoretical Computer Science*, 2003, 296:405-434.
- 5 Arieli, O., Denecker, M., Van Nuffelen, B. and Bruynooghe, M. Coherent Integration of Databases by Abductive Logic Programming. *J. Artif. Intell. Res.*, 2004, 21:245-286.
- 6 Barcelo, P., Bertossi, L. and Bravo, L. Characterizing and Computing Semantically Correct Answers from Databases with Annotated Logic and Answer Sets. In *Semantics of Databases*, Springer LNCS 2582, 2003, pp. 1-27.
- 7 Bertossi, L. Consistent Query Answering in Databases. *ACM SIGMOD Record*, 2006, 35(2):68-76.

- 8 Bertossi, L. and Li, L. Achieving Data Privacy through Secrecy Views and Null-Based Virtual Updates. *IEEE Transaction on Knowledge and Data Engineering*, 2013, 25(5):987-1000.
- 9 Bertossi, L. *Database Repairing and Consistent Query Answering*. Morgan & Claypool, Synthesis Lectures on Data Management, 2011.
- 10 Bertossi, L. and Salimi, B. Unifying Causality, Diagnosis, Repairs and View-Updates in Databases. Presented at the First International Workshop on Big Uncertain Data (BUDA 2014). Posted at: arXiv:1405.4228 [cs.DB].
- 11 Brankovic, L., and H. Fernau, H. Parameterized Approximation Algorithms for Hitting Set. In *Approximation and Online Algorithms*, 2012, Springer LNCS 7164, pp. 63-76.
- 12 Buneman, P., Khanna, S. and Tan, W. C. Why and Where: A Characterization of Data Provenance. *Proc. ICDT*, 2001, pp. 316-330.
- 13 Buneman, P. and Tan, W. C. Provenance in Databases. *Proc. ACM SIGMOD*, 2007, pp. 1171-1173.
- 14 Cheney, J., Chiticariu, L. and Tan, W. C. Provenance in Databases: Why, How, And Where. *Foundations and Trends in Databases*, 2009, 1(4): 379-474.
- 15 Cheney, J., Chong, S., Foster, N., Seltzer, M. I. and Vansummeren, S. Provenance: A Future History. *OOPSLA Companion (Onward!)*, 2009, pp. 957-964.
- 16 Cheney, J. Is Provenance Logical? *Proc. LID*, 2011, pp. 2-6.
- 17 Chomicki, J. and Marcinkowski, J. Minimal-Change Integrity Maintenance Using Tuple Deletions. *Information and Computation*, 2005, 197(1-2):90-121.
- 18 Chockler, H. and Halpern, J. Y. Responsibility and Blame: A Structural-Model Approach. *J. Artif. Intell. Res.*, 2004, 22:93-115.
- 19 Console, L. and Torasso, P. A Spectrum of Logical Definitions of Model-Based Diagnosis. *Computational Intelligence*, 1991, 7:133-141.
- 20 Console, L., Sapino M. L. and Theseider-Dupre, D. The Role of Abduction in Database View Updating. *J. Intell. Inf. Syst.*, 1995, 4(3): 261-280.
- 21 Cui, Y., Widom, J. and Wiener, J. L. Tracing the Lineage of View Data in a Warehousing Environment. *ACM Trans. Database Syst.*, 2000, 25(2):179-227.
- 22 Eiter, T., Gottlob, G. and Leone, N. Abduction from Logic Programs: Semantics and Complexity. *Theor. Comput. Sci.*, 1997, 189(1-2):129-177.
- 23 Eiter, Th., Faber, W., Leone, N. and Pfeifer, G. The Diagnosis Frontend of the DLV System. *AI Commun.*, 1999, 12(1-2):99-111.
- 24 Fernau, H. Parameterized Algorithmics for d -Hitting Set. *Int. J. Comput. Math.*, 2010, 87(14):3157-3174.
- 25 Feldman, A., Provan G., and Gemund A.V. Approximate model-based diagnosis using greedy stochastic search. *Journal of Artificial Intelligence Research (JAIR)*, 2010, 87(14):3157-3174.
- 26 Flum, J. and Grohe, M. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science, Springer Verlag, 2006.
- 27 Gertz, M. Diagnosis and Repair of Constraint Violations in Database Systems. PhD Thesis, Universität Hannover, 1996.
- 28 Greco, S., Pijcke, F. and Wijsen, J. Certain Query Answering in Partially Consistent Databases. *PVLDB*, 2014, 7(5):353-364.
- 29 Halpern, J., and Pearl, J. Causes and Explanations: A Structural-Model Approach: Part 1 *Proc. UAI*, 2001, pp. 194-202.
- 30 Halpern, J., and Pearl, J. Causes and Explanations: A Structural-Model Approach: Part 1. *British J. Philosophy of Science*, 2005, 56:843-887.

- 31 Halperin, E. Improved Approximation Algorithms for the Vertex Cover Problem in Graphs and Hyper-Graphs. *Proc. ACM-SIAM Symposium on Discrete Algorithms*, 2000, pp. 329-337.
- 32 Halpern, J. Appropriate Causal Models and Stability of Causation. *Proc. KR'14*, 2014.
- 33 Kakas A. C. and Mancarella, P. Database Updates through Abduction. *Proc. VLDB*, 1990, pp. 650-661.
- 34 Karvounarakis, G. and Green, T. J. Semiring-Annotated Data: Queries and Provenance? *SIGMOD Record*, 2012, 41(3):5-14.
- 35 Krentel, M. The Complexity of Optimization Problems. *J. Computer and Systems*, 1988, 36:490-509.
- 36 Karvounarakis, G. Ives, Z. G. and Tannen, V. Querying Data Provenance. *Proc. ACM SIGMOD*, 2010, pp. 951-962.
- 37 Kimelfeld, B. A Dichotomy in the Complexity of Deletion Propagation with Functional Dependencies. *Proc. ACM PODS*, 2012.
- 38 Kimelfeld, B., Vondrak, J. and Williams, R. Maximizing Conjunctive Views in Deletion Propagation. *ACM Trans. Database Syst.*, 2012, 37(4):24.
- 39 Lopatenko, A. and Bertossi, L. Complexity of Consistent Query Answering in Databases under Cardinality-Based and Incremental Repair Semantics. *Proc. ICDT*, 2007, Springer LNCS 4353, pp. 179-193. Extended version posted at: arXiv:cs/0604002 [cs.DB].
- 40 Meliou, A., Gatterbauer, W. and Suciu, D. Bringing Provenance to its Full Potential Using Causal Reasoning. *Proc. TaPP*, 2011.
- 41 Meliou, A., Gatterbauer, W. Moore, K. F. and Suciu, D. The Complexity of Causality and Responsibility for Query Answers and Non-Answers. *Proc. VLDB*, 2010, pp. 34-41.
- 42 Meliou, A., Gatterbauer, W., Halpern, J. Y., Koch, C., Moore K. F. and Suciu, D. Causality in Databases. *IEEE Data Eng. Bull*, 2010, 33(3):59-67.
- 43 Mozetic, I, Holzbaur, C. Controlling the Complexity in Model-Based Diagnosis *Annals of Mathematics and Artificial Intelligence*, 1994, 11(1-4): 297-314.
- 44 Niedermeier, R. and Rossmanith, P. An efficient fixed-parameter algorithm for 3-hitting set. In *J. Discrete Algorithms*, 2003 1(1):89-102.
- 45 Okun, M. On Approximation of the Vertex Cover Problem in Hypergraphs. In *Discrete Optimization*, 2005,2(1):101-111.
- 46 Papadimitriou, Ch. *Computational Complexity*. Addison-Wesley, 1994.
- 47 Reiter, R. A Theory of Diagnosis from First Principles. *Artificial Intelligence*, 1987, 32(1):57-95.
- 48 Reiter, R. Towards a Logical Reconstruction of Relational Database Theory. In *On Conceptual Modelling*, Springer, 1984, pp. 191-233.
- 49 Salimi, B. and Bertossi, L. Causality in Databases: The Diagnosis and Repair Connections. Presented at *The 15th International Workshop on Non-Monotonic Reasoning (NMR 2014)*. Posted at: arXiv:1404.6857 [cs.DB].
- 50 Salimi, B. and Bertossi, L. From Causes for Database Queries to Repairs and Model-Based Diagnosis and Back. Extended version of this paper. Posted at: arXiv:1412.4311 [cs.DB].
- 51 Staworko, S., Chomicki, J. and Marcinkowski, J. Prioritized Repairing and Consistent Query Answering in Relational Databases. *Ann. Math. Artif. Intell.*, 2012, 64(2-3):209-246.
- 52 Struss, P. Model-based Problem Solving. In *Handbook of Knowledge Representation*, chap. 10. Elsevier, 2008.
- 53 Tannen, V. Provenance Propagation in Complex Queries. In *Buneman Festschrift*, 2013, Springer LNCS 8000, pp. 483-493.
- 54 Yakout, M., Elmagarmid, A., Neville, J., Ouzzani, M. and Ilyas, I. Guided Data Repair. *PVLDB*, 2011, 4(5):279-289.