

Query Rewriting Using Datalog for Duplicate Resolution*

Jaffer Gardezi¹ and Leopoldo Bertossi²

¹ University of Ottawa, SITE,
Ottawa, Canada

² Carleton University, SCS,
Ottawa, Canada

Abstract. Matching Dependencies (MDs) are a recent proposal for declarative entity resolution. They are rules that specify, given the similarities satisfied by values in a database, what values should be considered duplicates, and have to be matched. On the basis of a chase-like procedure for MD enforcement, we can obtain clean (duplicate-free) instances; actually possibly several of them. The clean answers to queries (which we call the resolved answers) are invariant under the resulting class of instances. In this paper, we investigate a query rewriting approach to obtaining the resolved answers (for certain classes of queries and MDs). The rewritten queries are specified in stratified Datalog^{not,s} with aggregation. In addition to the rewriting algorithm, we discuss the semantics of the rewritten queries, and how they could be implemented by means of a DBMS.

1 Introduction

For various reasons, databases may contain different coexisting representations of the same external, real world entity. This can occur, for example, because of errors or because the data comes from different sources using different formats. Those “duplicates” can be entire tuples or values within them. To obtain accurate information, in particular, query answers from the data, those tuples or values should be merged into a single representation.

Identifying and merging duplicates is a process called *entity (or duplicate) resolution* (ER) [15, 9]. Matching dependencies (MDs) are a recent proposal for declarative duplicate resolution [20, 21]. An MD expresses, in the form of a rule, that if the values of certain attributes in a pair of tuples are similar, then the values of other attributes in those tuples should be matched (or merged) into a common value.

For example, the MD $R_1[X_1] \approx R_2[X_2] \rightarrow R_1[Y_1] \doteq R_2[Y_2]$ is a symbolic expression saying that, if an R_1 -tuple and R_2 -tuple have similar values for their attributes X_1, X_2 , then their values for attributes Y_1, Y_2 should be made equal. This is a *dynamic* dependency, in the sense that its satisfaction is checked against a pair of instances: the first one where the antecedent holds, and the second one where the identification of values takes place. This semantics of MDs was sketched in [21].

In this paper we use a refinement of that original semantics that was put forth in [25] (cf. also [26]). It improves wrt the latter in that it disallows changes that are irrelevant to

* Research supported by the NSERC Strategic Network on Business Intelligence (BIN ADC05) and NSERC/IBM CRDPJ/371084-2008.

the duplicate resolution process. Actually, [25] goes on to define the clean versions of the original database instance D_0 that contains duplicates. They are called the *resolved instances* (RIs) of D_0 wrt the given set M of matching dependencies. A resolved instance is obtained as the fixed point of a chase-like procedure that starts from D_0 and iteratively applies or enforces the MDs from M . Each step of this chase generates a new instance by making equal values that are identified as duplicates by the MDs.

In [25] it was shown that resolved instances always exist, and that they have certain desirable properties. For example, the set of allowed changes is just restrictive enough to prevent irrelevant changes, while still guaranteeing existence of resolved instances. The resolved instances that minimize the *overall number* of attribute value changes (associated to a same tuple identifier) wrt the original instance are called *minimally resolved instances* (MRIs). On this basis, given a query Q posed to a database instance D_0 that may contain duplicates, we define the *resolved answers* wrt Σ as the query answers that are true of all the minimally resolved instances [25].

The concept of resolved query answer has similarities to that of *consistent query answer* (CQA) in a database that fails to satisfy a set of integrity constraints [4, 11]. The consistent answers are invariant under the *repairs* of the original instance. However, data cleaning and CQA are different problems. For the former, we want to compute a clean instance, determined by MDs; for the latter, the goal is obtaining semantically correct query answers. MDs are not (static) ICs. In principle, we could see clean instances as repairs, treating MDs similarly to static FDs. However, the existing repair semantics do not capture the matchings as dictated by MDs (cf. [25, 26] for a more detailed discussion).

In this paper, we investigate the *problem of computing the resolved answers*, simply called *resolved answer problem* (RAP). The motivation for addressing this problem is that even in a database instance containing duplicates, much or most of the data may be duplicate-free. One can therefore obtain useful information from the instance without having to perform data cleaning on the instance. This would be convenient if the user does not want, or cannot afford, to go through a data cleaning process. In other situations the user may not have write access to the data being queried, or any access to the data sources, as in virtual data integration systems.

In [27] we identified classes of MDs and conjunctive queries for which RAP can be solved in polynomial time in data complexity. Furthermore, a recursively-defined predicate was introduced for identifying the sets of duplicate values within a database instance. This predicate can be combined with a query, opening the ground for a query rewriting approach to RAP.

In this paper we present a *query rewriting methodology for the RAP problem* (for the identified classes of MDs and queries). It can be used to rewrite the original query Q into a new query Q' , in such a way that the latter, posed as usual to the original instance D_0 , returns the resolved answers to the original query.

More precisely, we show that queries Q (in a restricted but broad class of conjunctive queries) that expect to obtain resolved answers from a given dirty database D can be rewritten into a (non-disjunctive) recursive Datalog^{not,s} query Q' with stratified negation and aggregation. Q' posed to D returns the answers to Q . As expected, such a query can be computed in polynomial time in the size of the initial database. The

recursion arises from the fact that identifying duplicate values requires computing the transitive closure of binary similarity operators. Transitivity is not assumed for similarity operators, and in fact, common similarity relations used in practice, such as those based on the edit distance and related string similarity metrics, are not transitive. Aggregation is needed to enforce the minimality constraint, since this involves finding the frequency of occurrence of values within a set of duplicates.

To the best of our knowledge, this is the first result on query rewriting in the context of MD-based entity resolution. Furthermore, our rewritings into Datalog are non-trivial, in the sense that they are not the result of translations into Datalog of first-order rewritings. Our rewriting uses in an essential manner the elements of the resulting Datalog queries, namely recursion and aggregation. It is worth mentioning that the polynomial-time rewritings for conjunctive queries proposed for consistent query answering have been all been first-order (FO) [4, 18, 24, 31].

On the other hand, the general answer set programs that have been proposed as *repair programs* [5, 7, 29, 19, 17], that specify database repairs and can be used for highly expressive query rewritings, have a higher expressive power and complexity than Datalog programs with stratified negation and aggregation.¹ The attempts in [10] to obtain lower complexity programs for CQA from repair programs for some tractable classes of queries and constraints led back to FO rewritings. Thus classical, i.e. non-disjunctive and stratified, Datalog was missed as an “intermediate” language for CQA.

This paper is organized as follows. In Section 2 we introduce basic concepts and notation on MDs. In Section 3, we define the important concepts used in this paper, in particular, (minimally) resolved instances and resolved answers to queries. Section 4 contains the main results of this paper, that includes a query rewriting algorithm for a special case of the resolved query answer problem. Section 5 concludes the paper and discusses related and future work. Proofs of results can be found in [27].

2 Preliminaries

We consider a relational schema \mathcal{S} that includes an enumerable, possibly infinite domain U , and a finite set \mathcal{R} of database predicates. Elements of U are represented by lower case letters near the beginning of the alphabet. \mathcal{S} determines a first-order (FO) language $L(\mathcal{S})$. An instance D for \mathcal{S} is a finite set of ground atoms of the form $R(\bar{a})$, with $R \in \mathcal{R}$, say of arity n , and $\bar{a} \in U^n$. $R(D)$ denotes the extension of R in D . Every predicate $R \in \mathcal{S}$ has a set of attribute, denoted $attr(R)$. As usual, we sometimes refer to attribute A of R by $R[A]$. We assume that all the attributes of a predicate are different, and that we can identify attributes with *positions* in predicates, e.g. $R[i]$, with $1 \leq i \leq n$. If the i th attribute of predicate R is A , for a tuple $t = (c_1, \dots, c_n) \in R(D)$, $t_R^D[A]$ (usually, simply $t_R[A]$ or $t[A]$ if the instance is understood) denotes the value c_i . For a sequence \bar{A} of attributes in $attr(R)$, $t[\bar{A}]$ denotes the tuple whose entries are the values of the attributes in \bar{A} . Attributes have and may share subdomains of U .

In the rest of this section, we summarize some of the assumptions, definitions, notation, and results from two previous papers, [25] and [27], that we will need.

¹ Under the common assumption that the polynomial hierarchy does not collapse.

We will assume that every relation in an instance has an auxiliary attribute, a surrogate key, holding values that act as *tuple identifiers*. Tuple identifiers are never created, destroyed or changed during the duplicate resolution process. They do not appear in MDs, and are used to identify different versions of the same original tuple that result from the matching process. We usually leave them implicit; and “tuple identifier attributes” are commonly left out when specifying a database schema. However, when explicitly represented, they will be the “first” attribute of the relation. For example, if $R \in \mathcal{R}$ is n -ary, $R(t, c_1, \dots, c_n)$ is a tuple with id t , and is usually written as $R(t, \bar{c})$. We usually use the same symbol for a tuple’s identifier as for the tuple itself. Tuple identifiers are unique over the entire instance. Two instances over the same schema that share the same tuple identifiers are said to be *correlated*, and they can be unambiguously compared tuple by tuple.

As expected, some of the attribute domains, say A , have a built-in binary similarity relation $\approx_A \subseteq \text{Dom}(A) \times \text{Dom}(A)$ that is reflexive and symmetric. Such a relation can be extended to finite lists of attributes (or domains thereof), componentwise. For single attributes or lists thereof, the similarity relation is generically denoted with \approx .

A *matching dependency* (MD) [20] is an expression of the form

$$m : R[\bar{A}] \approx S[\bar{B}] \rightarrow R[\bar{C}] \doteq S[\bar{E}], \quad (1)$$

with $\bar{A} = (A_1, \dots, A_k)$, $\bar{C} = (C_1, \dots, C_{k'})$ lists of (different) attributes from $\text{attr}(R)$; and $\bar{B} = (B_1, \dots, B_k)$, $\bar{E} = (E_1, \dots, E_{k'})$ lists of attributes from $\text{attr}(S)$.²

The set of attributes on the left-hand-side (LHS) of the arrow in m is denoted with $LHS(m)$. Similarly for the right-hand-side (RHS). In relation to (1), the attributes in a *corresponding pair* (A_i, B_i) or (C_i, E_i) are assumed to share a common domain; and in particular, a *similarity relation* \approx_i . In consequence, the condition on the LHS of (1) means that, for a pair of tuples t_1 in R and t_2 in S , $t_1[A_i] \approx_i t_2[B_i]$, $1 \leq i \leq k$. Similarly, the expression on the RHS means $t_1[A_i] \doteq t_2[B_i]$, $1 \leq i \leq k'$. Here, \doteq means that the values should be updated to the same value.

Accordingly, the intended semantics of the MD in (1) is that, for an instance D , if any pair of tuples, $t_1 \in R(D)$ and $t_2 \in S(D)$, satisfy the similarity conditions on the LHS, then for the same tuples (or tuple ids), the attributes on the RHS have to take the same values [21], possibly through updates that may lead to a new version of D .

We assume that all sets M of MDs are in *standard form*, i.e. for no two different MDs $m_1, m_2 \in M$, $LHS(m_1) = LHS(m_2)$. All sets of MDs can be put in this form. MDs in a set M can interact in the sense that a matching enforced by one of them may create new similarities that lead to the enforcement of another MD in M . This intuition is captured through the *MD-graph*.

Definition 1. [26] Let M be a set of MDs in standard form. The *MD-graph* of M , denoted $MDG(M)$, is a directed graph with a vertex m for each $m \in M$, and an edge from m to m' iff $RHS(m) \cap LHS(m') \neq \emptyset$.³ If $MDG(M)$ contains edges, M is called *interacting*. Otherwise, it is called *non-interacting* (NI). \square

² We assume that the MDs are defined in terms of the same schema \mathcal{S} .

³ That is, they share at least one corresponding pair of attributes.

3 Matching Dependencies and Resolved Answers

Updates as prescribed by an MD m are not arbitrary. The updates based on m have to be justified by m , as captured through the notion of *modifiable value* in an instance.

Definition 2. Let D be an instance, M a set of MDs, and \mathcal{P} be a set of pairs (t, G) , where t is a tuple of D and G is an attribute of t . (a) For a tuple $t_R \in R(D)$ and C an attribute of R , the value $t_R^D[C]$ is *modifiable wrt* \mathcal{P} if there exist $S \in \mathcal{R}$, $t_S \in S(D)$, an $m \in M$ of the form $R[\bar{A}] \approx S[\bar{B}] \rightarrow R[\bar{C}] \doteq S[\bar{E}]$, and a corresponding pair (C, E) of (\bar{C}, \bar{E}) in m , such that $(t_S, E) \in \mathcal{P}$ and one of the following holds:

1. $t_R[\bar{A}] \approx t_S[\bar{B}]$, but $t_R[C] \neq t_S[E]$.
2. $t_R[\bar{A}] \approx t_S[\bar{B}]$ and $t_S[E]$ is modifiable wrt $\mathcal{P} \setminus \{(t_S, E)\}$.

(b) The value $t_R^D[C]$ is *modifiable* if it is modifiable wrt $\mathcal{V} \setminus \{(t_R, C)\}$, where \mathcal{V} is the set of all pairs (t, G) with t a tuple of D and G an attribute of t . \square

Definition 2 is recursive. The base case occurs when either case 1 applies (with any \mathcal{P}) or when there is no tuple/attribute pair in \mathcal{P} that can satisfy part (a). Notice that recursion must terminate eventually, since the latter condition must be satisfied when \mathcal{P} is empty, and each recursive call reduces the size of \mathcal{P} .

Example 1. Consider $m : R[A] \approx R[A] \rightarrow R[B] \doteq R[B]$ on schema $R[A, B]$, and the following instance. Assume that the only non-trivial similarities are $a_1 \approx a_2 \approx a_3$ and $b_1 \approx b_2$. Since $a_2 \approx a_3$ and $c_1 \neq c_3$, $t_2[B]$ and $t_3[B]$ are modifiable (base case). With case 2 of Definition 2, since $a_1 \approx a_2$, and $t_2[B]$ is modifiable, we obtain that $t_1[B]$ is modifiable.

$R(D)$	A	B
t_1	a_1	c_1
t_2	a_2	c_1
t_3	a_3	c_3
t_4	b_1	c_3
t_5	b_2	c_3

For $t_5[B]$ to be modifiable, it must be modifiable wrt $\{(t_i, B) \mid 1 \leq i \leq 4\}$, and via t_4 . According to case 2 of Definition 2, this requires $t_4[B]$ to be modifiable wrt $\{(t_i, B) \mid 1 \leq i \leq 3\}$. However, this is not the case since there is no t_i , $1 \leq i \leq 3$, such that $t_4[A] \approx t_i[A]$. Therefore $t_5[B]$ is not modifiable. A symmetric argument shows that $t_4[B]$ is not modifiable. \square

Definition 3. [25] Let D, D' be correlated instances, and M a set of MDs. (D, D') satisfies M , denoted $(D, D') \models M$, iff: 1. For any pair of tuples $t_R \in R(D)$, $t_S \in S(D)$, if there exists an $m \in M$ of the form $R[\bar{A}] \approx S[\bar{B}] \rightarrow R[\bar{C}] \doteq S[\bar{E}]$ and $t_R[\bar{A}] \approx t_S[\bar{B}]$, then for the corresponding tuples (i.e. with same ids) $t'_R \in R(D')$ and $t'_S \in S(D')$, it holds $t'_R[\bar{C}] = t'_S[\bar{E}]$. 2. For any tuple $t_R \in R(D)$ and any attribute G of R , if $t_R[G]$ is *non-modifiable*, then $t'_R[G] = t_R[G]$. \square

Intuitively, D' in Definition 3 is a new version of D that is produced after a single update. Since the update involves matching values (i.e. making them equal), it may produce “duplicate” tuples, i.e. that only differ in their tuple ids. They would possibly be merged into a single tuple in the a data cleaning process. However, we keep the two

versions. In particular, D and D' have the same number of tuples. Keeping or eliminating duplicates will not make any important difference in the sense that, given that tuple ids are never updated, two duplicates will evolve in exactly the same way as subsequent updates are performed. Duplicate tuples will never be subsequently “unmerged”.

This definition of MD satisfaction departs from [21], which requires that updates preserve similarities. Similarity preservation may force undesirable changes [25]. The existence of the updated instance D' for D is guaranteed [25]. Furthermore, wrt [21], our definition does not allow unnecessary changes from D to D' . Definitions 2 and 3 imply that only values of attributes that appear to the right of the arrow in some MD are subject to updates. Hence, they are called *changeable attributes*.⁴

Definition 3 allows us to define a *clean instance* wrt M as the result of a chase-like procedure, each step being satisfaction preserving.

Definition 4. [25] (a) A *resolved instance* (RI) for D wrt M is an instance D' , such that there are instances D_1, D_2, \dots, D_n with: $(D, D_1) \models M$, $(D_1, D_2) \models M, \dots, (D_{n-1}, D_n) \models M$, $(D_n, D') \models M$, and $(D', D') \models M$. We say D' is *stable*. (b) D' is a *minimally resolved instance* (MRI) for D wrt M if it is a resolved instance that minimizes the overall number of attribute value changes wrt D and in relation with the same tuple ids. (c) $MRI(D, M)$ denotes the class of MRIs of D wrt M . \square

Example 2. Consider the MD $R[A] \approx R[A] \rightarrow R[B] \doteq R[B]$ on predicate R , and the instance D . It has several resolved instances, among them, four that

$R(D)$	A	B
t_1	a_1	c_1
t_2	a_1	c_2
t_3	b_1	c_3
t_4	b_1	c_4

$R(D_1)$	A	B
t_1	a_1	c_1
t_2	a_1	c_1
t_3	b_1	c_3
t_4	b_1	c_3

$R(D_2)$	A	B
t_1	a_1	c_1
t_2	a_1	c_1
t_3	b_1	c_1
t_4	b_1	c_1

minimize the number of changes. One of them is D_1 . A resolved instance that is not minimal in this sense is D_2 . \square

In this work, as in [25, 26], we are investigating what we could call “the pure case” of MD-based entity resolution. It adheres to the original semantics outlined in [21], which does not specify how the matchings are to be done, but only which values must be made equal. That is, the MDs have implicit existential quantifiers (for the values in common). The semantics we just introduced formally captures this pure case. We find situations like this in other areas of data management, e.g. with referential integrity constraints, tuple-generating dependencies in general [1], schema mappings in data exchange [8], etc. It can be shown that an RI exists for any instance [25]. It follows immediately that every instance has an MRI. A non-pure case that uses matching functions to realize the matchings as prescribed by MDs is investigated in [13, 14, 6].

Definition 5. [25] Let $\mathcal{Q}(\bar{x})$ be a query expressed in the FO language $L(\mathcal{S})$. A tuple of constants \bar{a} from U is a *resolved answer* to $\mathcal{Q}(\bar{x})$ wrt the set M of MDs, denoted $D \models_M \mathcal{Q}[\bar{a}]$, iff $D' \models \mathcal{Q}[\bar{a}]$, for every $D' \in MRI(D, M)$. $ResAn(D, \mathcal{Q}, M)$ is the set of resolved answers to \mathcal{Q} from D wrt M . \square

⁴ Not to be confused with “modifiability”, that applies to tuples.

Example 3. (example 1 cont.) Since the only MRI for the original instance D is $R(D')$ = $\{\langle t_1, a_1, c_1 \rangle, \langle t_2, a_2, c_1 \rangle, \langle t_3, a_3, c_1 \rangle, \langle t_4, b_1, c_3 \rangle, \langle t_5, b_2, c_3 \rangle\}$, the resolved answers to the query $\mathcal{Q}(x, y) : R(x, y)$ are $\{\langle a_1, c_1 \rangle, \langle a_2, c_1 \rangle, \langle a_3, c_1 \rangle, \langle b_1, c_3 \rangle, \langle b_2, c_3 \rangle\}$. \square

4 Query Rewriting for Resolved Answers

In this section, we present a query rewriting method for retrieving the resolved answers for certain classes of queries and sets of MDs. We provide an intuitive and informal presentation of the rewritten queries. For details and a proof of correctness, see [27].

It has been shown in previous work that the problem of deciding resolved answers (the *resolved answer decision problem*) is generally intractable in data [25, 26, 27]. However, there are tractable cases of this decision problem that are practically relevant [27]. Two of those cases are considered here: that of *non-interacting* (NI) sets of MDs (cf. Definition 1), and that of *hit-set-cyclic* (HSC) sets of MDs, that we now define.

Definition 6. A set M of MDs is *hit-simple-cyclic* iff the following hold: (a) In all MDs in M and in all their corresponding pairs, the two attributes (and predicates) are the same. (b) In all MDs $m \in M$, at most one attribute in $LHS(m)$ is changeable. (c) Each vertex v_1 in $MDG(M)$ is on at least one cycle, or there is a vertex v_2 on a cycle with at least two vertices such that there is an edge from v_1 to v_2 . \square

Example 4. For schema $R[A, C, F, G]$, consider the following set M of MDs:

$$\begin{aligned} m_1 : R[A] \approx R[A] &\rightarrow R[C, F, G] \doteq R[C, F, G], \\ m_2 : R[C] \approx R[C] &\rightarrow R[A, F, G] \doteq R[A, F, G]. \end{aligned}$$

Set M obviously satisfies (a) and (b) of Definition 6. Also, $MDG(M)$ consists of a single cycle through the two vertices, so M satisfies (c). M is then HSC.

Predicate R subject to the given M has two “keys”, $R[A]$ and $R[C]$. Such relations are common in practice. For example, R may be used in a database about people: $R[A]$ could be used for the person’s name, $R[C]$ the address, and $R[F]$ and $R[G]$ for non-distinguishing information, e.g. gender and age. \square

HSC sets have properties similar to those of NI sets wrt the resolved answer problem [27]. For both classes, the value positions identified as duplicates are the same for all MRIs, and they are characterized via equivalence classes of the *tuple-attribute closure*.

Definition 7. [27] Let $M = \{m_i \mid i = 1, \dots, n\}$ be a set of MDs, with $m_i : R_i[\bar{A}_i] \approx_i S_i[\bar{B}_i] \rightarrow R_i[\bar{C}_i] \doteq S_i[\bar{E}_i]$. (a) The *previous set* of m_i , denoted $PS(m_i)$, is the set of all MDs $m_j \in M$ with a path in $MDG(M)$ from m_j to m_i . (b) For an instance D , and tuples ids t_1, t_2 for R, S , resp. (i.e. ids for tuples $t_1 \in R(D), t_2 \in S(D)$): $(t_1, C_i) \approx' (t_2, E_i) \iff t_1[\bar{A}_j] \approx_j t_2[\bar{B}_j]$, where (C_i, E_i) is a corresponding pair of (\bar{C}_i, \bar{E}_i) in m_i and $m_j \in PS(m_i)$. (c) The *tuple-attribute closure* (TA closure) of M wrt D , denoted $TA^{M,D}$, is the reflexive, symmetric, and transitive closure of \approx' . \square

Example 5. (example 4 continued) In this case, $PS(m_1) = PS(m_2) = \{m_1, m_2\}$. Consider the instance D , where the only similarities are: $a_i \approx a_j, b_i \approx b_j, d_i \approx d_j$,

$R(D)$	A	B
t_1	a_1	d_1
t_2	a_2	e_2
t_3	b_1	e_1
t_4	b_2	d_2

$e_i \approx e_j$, with $i, j \in \{1, 2\}$. The relations $(t_{i \bmod 4+1}, A) \approx' (t_{(i+1) \bmod 4+1}, A)$ and $(t_{i \bmod 4+1}, B) \approx' (t_{(i+1) \bmod 4+1}, B)$, $0 \leq i \leq 3$, hold. The TA closure is given by $\{TA(t_i, x, t_j, x) \mid 1 \leq i, j \leq 4, x \in$

$\{A, B\}\}$. Notice that this relation involves just tuple ids and attributes. However, it depends on D through the similarity conditions in (b) in the definition. \square

For the set of MDs as in Definition 7, the TA closure can be specified by Datalog rules. The database predicates in them have a first argument (attribute) to explicitly represent the tuple id. More precisely, for $1 \leq i \leq n$, for each corresponding pair (C, E) of (\bar{C}_i, \bar{E}_i) , and for each $m_j \in PS(m_i)$, we have the rule:⁵

$$(t_1, R_i[C]) \approx' (t_2, S_i[E]) \leftarrow R_i(t_1, \bar{x}), S_i(t_2, \bar{y}), t_1[\bar{A}_j] \approx_j t_2[\bar{B}_j].$$

Additionally, for all attributes A of R_i and ids t of tuples in R_i , we have

$$TA(t, A, t, A) \leftarrow R_i(t, \bar{x}); \quad (2)$$

similarly for S_i . For arbitrary tuple ids t_1, t_2 , and t_3 , and attributes A, B , and C ,

$$TA(t_1, A, t_2, B) \leftarrow TA(t_2, B, t_1, A), \quad (3)$$

$$TA(t_1, A, t_2, B) \leftarrow (t_1, A) \approx' (t_2, B), \quad (4)$$

$$TA(t_1, A, t_3, C) \leftarrow TA(t_1, A, t_2, B), (t_2, B) \approx' (t_3, C). \quad (5)$$

Rules (4) and (5) express that TA is the transitive closure of relation \approx' . Rules (2) and (3), that TA is reflexive and symmetric. A related concept is the *attribute closure*.

Definition 8. [26] Let M be a set of MDs on schema \mathcal{S} . (a) The symmetric binary relation $\dot{=}^r$ relates attributes $R[A], S[B]$ of \mathcal{S} whenever there is an MD m in M where $R[A] \dot{=} S[B]$ appears in $RHS(m)$. (b) The *attribute closure* of M is the reflexive, symmetric, transitive closure of $\dot{=}^r$. (c) $E_{R[A]}$ denotes the equivalence class of attribute $R[A]$ in the attribute closure of M . \square

Example 6. Let M be the set of MDs: $R[A] \approx_1 S[B] \rightarrow R[C] \dot{=} S[D], S[E] \approx_2 T[F] \wedge S[G] \approx T[H] \rightarrow S[D, K] \dot{=} T[J, L], T[F] \approx_3 T[H] \rightarrow T[L, N] \dot{=} T[M, P]$. The equivalence classes of T_{at} are $E_{R[C]} = \{R[C], S[D], T[J]\}$, $E_{S[K]} = \{S[K], T[L], T[M]\}$, and $E_{T[N]} = \{T[N], T[P]\}$. \square

It is easy to show that if $(u_1, A), (u_2, B)$ are in the same equivalence class of tuple-attribute closure, then A and B are in the same equivalence class of attribute closure.

Definition 9. Let M be a set of MDs and D an instance and a a data value. For an equivalence class E of $TA^{M,D}$, the *frequency of a in E* is the quantity $freq^D(a, E) := |\{(t, A) \mid (t, A) \in E, t[A] = a \text{ in } D\}|$. \square

Proposition 1. [27] For M an NI or HSC set of MDs, and D an instance, each MRI for D wrt M is obtained by setting, for each equivalence class E of $TA^{M,D}$, all the values for $t[A]$, with $(t, A) \in E$, to a value a that maximizes $freq^D(a, E)$. \square

⁵ Remember that the first argument in R_i, S_i stands for the tuple id.

Example 7. (example 5 continued) The two equivalence classes of TA closure are $E_1 = \{(t_i, A) \mid 1 \leq i \leq 4\}$ and $E_2 = \{(t_i, B) \mid 1 \leq i \leq 4\}$. All values in the A (B) column of the table have frequency 1 in E_1 (E_2). Thus, there are 16 MRIs, obtained by setting all values in each column to a common value chosen from those in the column. \square

Proposition 1 tells us that the minimally resolved instances for an instance D can be obtained by identifying most frequently occurring values. Thus, resolved query answers from D can be computed by imposing this requirement on the original query. As a consequence, the rewritten queries will become aggregate queries.

In Datalog notation, aggregate queries take the form $P(\bar{a}, \bar{x}, \text{Agg}(\bar{u})) \leftarrow B(\bar{y})$, where P is answer collecting predicate, the body $B(\bar{y})$ represents a conjunction of literals all of whose variables are among those in \bar{y} , \bar{a} is a list of constants, $\bar{x} \cup \bar{u} \subseteq \bar{y}$, and Agg is an aggregate operator such as *Count* or *Sum*. The variables \bar{x} are the “group-by” variables. That is, for each fixed value \bar{b} for \bar{x} , aggregation is performed over all tuples that make $B \frac{\bar{x}}{\bar{b}}$, the instantiation of B on \bar{b} for \bar{x} , true. *Count*(\bar{u}) counts the number of distinct values of \bar{u} , while *Sum*(\bar{u}) sums over all \bar{u} , whether distinct or not.

Our query rewriting methodology for computing resolved answers will be applicable to a certain class of conjunctive queries, the called UJCQ queries defined below. In [27] a counterexample for the general applicability to all conjunctive queries is given.

Definition 10. [27] For a set M of MDs, a conjunctive query Q without built-ins is an *unchangeable join conjunctive query* (UJCQ query) if there are no existentially quantified variables in a join in Q in the position of a changeable attribute. For fixed, M , *UJCQ* denotes this class of queries. \square

In the rest of this paper we assume that we have a fixed and finite set M of MDs that satisfies the hypothesis of Proposition 1. The queries posed to the initial, possibly non-resolved instance belong to *UJCQ*.

The rewritten queries will be in Datalog^{not,s} [1], i.e. Datalog queries with stratified negation and aggregation, and the built-ins \neq and \leq . For simplicity, the rewriting makes use of tuple identifiers only. In the absence of such a surrogate key, whole tuples could be used instead of identifiers.

Given a *UJCQ* query Q , with answer predicate Q :

$$Q(\bar{x}) \leftarrow R_1(\bar{v}_1), R_2(\bar{v}_2), \dots, R_n(\bar{v}_n), \quad (6)$$

the rewritten query Q' is the conjunction of the rewritings Q_i of each of the R_i , to be given in (8) below, i.e.

$$Q'(\bar{x}) \leftarrow Q_1(\bar{v}_1), Q_2(\bar{v}_2), \dots, Q_n(\bar{v}_n). \quad (7)$$

Now, for a fixed atom $R_i(\bar{v}_i)$ in (6), let C be the set of changeable attributes corresponding to a free variable in \bar{v}_i , i.e. also appearing in $Q(\bar{x})$. We denote the list of such variables by \bar{v}_C .

If C is empty, then its rewriting becomes $Q_i(\bar{v}_i) \leftarrow R_i(\bar{v}_i)$. Intuitively, this is because, by Definition 10, only attributes corresponding to free variables can participate in joins, so changes to values of attributes corresponding to bound variables cannot affect satisfaction of the body in (6).

Suppose C is non-empty, and consider $R_i[A] \in C$. From Proposition 1, deciding whether or not all MRIs have the same value v for $R_i[A]$ for a given tuple id t will

involve finding the frequency of v in E for the equivalence class E of the TA closure to which $(t, R_i[A])$ belongs. We introduce aggregation operators to express this count, of values for attributes in $E_{R_i[A]}$ (cf. the remark following Example 6).

We introduce a predicate $C^{R_i[A]}$, with an attribute at the start of its attribute list whose value is the attribute in $E_{R_i[A]}$ over whose values aggregation is performed. For an attribute A and list of variables \bar{v} , we denote with v_A the variable holding the value for A . For each $S[B] \in E_{R_i[A]}$, we have the rule

$$C^{R_i[A]}(S[B], t_1, v_{S[B]}, Count(t_2)) \leftarrow TA(t_1, R_i[A], t_2, S[B]), R_i(t_1, \bar{u}), S(t_2, \bar{v}),$$

in which all predicate arguments are variables except for the attribute labels $S[B]$ and $R_i[A]$, that are constants.

In each tuple in the head predicate of the above expression, the value of the *Count* expression is $|\{t \mid (t, S[B]) \in E, t[S[B]] = v_{S[B]}\}|$, where E is the equivalence class of the TA closure to which $(t_1, R_i[A])$ belongs.

To find the frequency of the value of $v_{S[B]}$ in E , this count must be extended to all attributes in $E_{R_i[A]}$. We introduce the predicate $Total^{R_i[A]}$ for this purpose:

$$Total^{R_i[A]}(t, v, Sum(z)) \leftarrow C^{R_i[A]}(x, t, v, z).$$

Tuples in $Total^{R_i[A]}$ specify in their last argument the frequency of v in the equivalence class of the TA-closure to which $(u, R_i[A])$ belongs.

To compare these aggregate quantities for different values of v , we use the *Compare* predicate:

$$Compare^{R_i[A]}(t, v) \leftarrow Total^{R_i[A]}(t, v, z_1), Total^{R_i[A]}(t, v', z_2), z_1 \leq z_2, v' \neq v.$$

Tuples in $Compare^{R_i[A]}$ consist of a tuple identifier t in R_i and a value v for attribute $R_i[A]$. For such a pair (t, v) there is another value v' whose frequency in the equivalence class of the TA closure to which $(t, R_i[A])$ belongs is at least as large as that of v .

In order for a value to be a ‘‘certain’’ for a given attribute $R_i[A]$ in a given tuple, the tuple and value must not occur as a tuple in $Compare^{R_i[A]}$. Let \bar{v}'_i be \bar{v}_i with all variables in \bar{v}_C replaced with new variables. Then,

$$Q_i(\bar{v}_i) \leftarrow R_i(t, \bar{v}'_i), \bigwedge_{R_i[A] \in C} not\ Compare^{R_i[A]}(t, v_{R_i[A]}), Total^{R_i[A]}(t, v_{R_i[A]}, z). \quad (8)$$

Example 8. Consider the schema $R[ABC], S[CFG], U[HI]$ with non-interacting MDs: $R[A] \approx S[E] \rightarrow R[B] \doteq S[F]$, $S[E] \approx U[H] \rightarrow S[F] \doteq U[I]$, and the *UJCQ* query:

$Q(x, y, z) \leftarrow R(x, y, z), S(u, v, z), U(p, q)$. Since the S and U atoms have no free variables holding the values of changeable attributes, they remain unchanged. Therefore, the rewritten query Q' has the form

$$Q' \leftarrow R'(x, y, z), S(u, v, z), U(p, q), \quad (9)$$

where R' is the rewritten form of R . The only free variable holding the value of a changeable attribute is y . This variable corresponds to attribute $R[B]$, which belongs to the equivalence class $E_{R[B]} = \{R[B], S[F], U[I]\}$. Therefore, we have the rules:

$$\begin{aligned}
C^{R[B]}(R[B], t_1, y, \text{Count}(t_2)) &\leftarrow TA(t_1, R[B], t_2, R[B]), R(t_1, x', y', z'), R(t_2, x, y, z). \\
C^{R[B]}(S[F], t_1, y, \text{Count}(t_2)) &\leftarrow TA(t_1, R[B], t_2, S[F]), R(t_1, x', y', z'), S(t_2, x, y, z). \\
C^{R[B]}(U[I], t_1, y, \text{Count}(t_2)) &\leftarrow TA(t_1, R[B], t_2, U[I]), R(t_1, x', y', z'), U(t_2, x, y). \\
\text{Total}^{R[B]}(t, y, \text{Sum}(u)) &\leftarrow C^{R[B]}(x, t, y, u). \\
\text{Compare}^{R[B]}(t, y) &\leftarrow \text{Total}^{R[B]}(t, y, z_1), \text{Total}^{R[B]}(t, y'', z_2), z_1 \leq z_2, y'' \neq y.
\end{aligned}$$

The rewriting of R becomes

$$R'(x, y, z) \leftarrow R(t, x, y', z), \text{not Compare}^{R[B]}(t, y), \text{Total}^{R[B]}(t, y, w). \quad (10)$$

Thus, the rewriting of the original query is the stratified Datalog program [1] with aggregation consisting of rules (9), (10), plus the five rules preceding (10). \square

In order to obtain the resolved answers to a query on a possibly non-resolved instance D , the resulting Datalog program can be run on D in polynomial time in the size of D . Remarks on implementation and an example are included in an extended version [28].

5 Conclusions

This paper considered a novel approach based on query rewriting to the duplicate resolution problem within the framework of matching dependencies. The transformed queries return the resolved answers to the original query, which are the answers that are true in all minimally resolved instances.

We used minimal resolved instances (MRIs) as our model of a clean database. Another possibility is to use arbitrary, not necessarily minimal, resolved instances (RIs). While MRIs have the advantage of being “closer” to the original instance than RIs, they have the downside of being overly restrictive.

In practice, update values are typically chosen by applying a merging function to the sets of duplicates [9, 13, 14], rather than by imposing a minimal change constraint. RIs are more flexible in that they take into account all ways of choosing the update values that lead to a clean database. We are currently investigating query answering over RIs, identifying tractable cases of the problem that are not tractable for MRIs.

Matching dependencies first appeared in [20], and their semantics is given in [21]. The original semantics was refined in [13, 14], including the use of *matching functions* (MFs) for matching two attribute values. The approach in [13, 14] uses a chase to define clean instances. The MDs are applied one at a time to pairs of tuples, rather than all at once to all tuples as in the present paper. Another important difference is that here we do not use MFs to do a matching, but implicit existential quantifiers for the values in common. When the update values are determined by the matching functions there is no uncertainty arising from different possible choices for update values. Rather, the different clean instances are produced by applying the MDs in different orders. Clean answers are obtained by taking a glb (or lub) over the clean instances wrt a partial ordering that is based on semantic domination of one value by another.

The alternative refinement of the semantics used in this paper was first introduced in [25, 26]. A thorough complexity analysis, as well as the derivation of a query rewriting algorithm for the resolved answer problem was done in [27].

Our work in some ways resembles work on query answering over ontologies [16]. As in our duplicate resolution setting, a chase is applied repeatedly to an initial instance, terminating in a “repaired” instance which is a fixed point of the chase rules. The set of chase rules can include tuple generating dependencies (TGDs) and equality generating dependencies (EGDs). Despite these similarities, our chase differs from those based on EGDs and TGDs in that it does not generate new tuples, but modifies values in existing tuples. Also, despite the fact that MDs are similar to EGDs, issues arise as a result of the non-transitivity of similarity operators that do not occur in the case of EGDs.

In [3], Datalog is used for identifying groups of tuples that could be merged. However, they do not do the merging (a main contribution in our approach) or base their approach on MDs. Actually, that work could be considered as complimentary to ours, in the sense that, in essence, the authors address the problem of identifying similarities. This is the starting point for the actual matchings that we address in this paper.

References

- [1] Abiteboul, S., Hull, R., Vianu, V.: *Foundations of Databases*. Addison-Wesley (1995)
- [2] Afrati, F., Kolaitis, P.: Repair checking in inconsistent databases: Algorithms and complexity. In: Proc. ICDT (2009)
- [3] Arasu, A., Ré, C., Suciu, D.: Large-scale deduplication with constraints using dedupalog. In: Proc. ICDE (2009)
- [4] Arenas, M., Bertossi, L., Chomicki, J.: Consistent query answers in inconsistent databases. In: Proc. PODS (1999)
- [5] Arenas, M., Bertossi, L., Chomicki, J.: Answer sets for consistent query answering in inconsistent databases. *Theory and Practice of Logic Programming* 3(4-5), 393–424 (2003)
- [6] Bahmani, Z., Bertossi, L., Kolahi, S., Lakshmanan, L.: Declarative Entity Resolution via Matching Dependencies and Answer Set Programs. In: Proc. KR 2012 (2012)
- [7] Barceló, P., Bertossi, L., Bravo, L.: Characterizing and Computing Semantically Correct Answers from Databases with Annotated Logic and Answer Sets. In: Bertossi, L., Katona, G.O.H., Schewe, K.-D., Thalheim, B. (eds.) *Semantics in Databases*. LNCS, vol. 2582, pp. 7–33. Springer, Heidelberg (2003)
- [8] Barcelo, P.: Logical foundations of relational data exchange. *SIGMOD Record* 38(1), 49–58 (2009)
- [9] Benjelloun, O., Garcia-Molina, H., Menestrina, D., Su, Q., Euijong Whang, S., Widom, J.: Swoosh: A generic approach to entity resolution. *VLDB Journal* 18(1), 255–276 (2009)
- [10] Bertossi, L.: From database repair programs to consistent query answering in classical logic. In: Proc. AMW. CEUR-WS, vol. 450 (2009)
- [11] Bertossi, L.: *Database Repairing and Consistent Query Answering*. Synthesis Lectures on Data Management. Morgan & Claypool (2011)
- [12] Bertossi, L., Bravo, L., Franconi, E., Lopatenko, A.: The complexity and approximation of fixing numerical attributes in databases under integrity constraints. *Information Systems* 33(4), 407–434 (2008)
- [13] Bertossi, L., Kolahi, S., Lakshmanan, L.: Data cleaning and query answering with matching dependencies and matching functions. In: Proc. ICDT (2011)
- [14] Bertossi, L., Kolahi, S., Lakshmanan, L.: Data cleaning and query answering with matching dependencies and matching functions. *Theory of Computing Systems* (2012), doi:10.1007/s00224-012-9402-7
- [15] Bleiholder, J., Naumann, F.: Data fusion. *ACM Computing Surveys* 41(1), 1–41 (2008)

- [16] Cali, A., Gottlob, G., Lukasiewicz, T., Pieris, A.: A logical toolbox for ontological reasoning. *ACM Sigmod Record* 40(3), 5–14 (2011)
- [17] Caniupan, M., Bertossi, L.: The consistency extractor system: answer set programs for consistent query answering in databases. *Data & Know. Eng.* 69(6), 545–572 (2010)
- [18] Chomicki, J., Marcinkowski, J.: Minimal-change integrity maintenance using tuple deletions. *Information and Computation* 197(1/2), 90–121 (2005)
- [19] Eiter, T., Fink, M., Greco, G., Lembo, D.: Repair localization for query answering from inconsistent databases. *ACM Trans. Database Syst.* 33(2) (2008)
- [20] Fan, W.: Dependencies revisited for improving data quality. In: *Proc. PODS* (2008)
- [21] Fan, W., Jia, X., Li, J., Ma, S.: Reasoning about record matching rules. In: *Proc. VLDB* (2009)
- [22] Flesca, S., Furfaro, F., Parisi, F.: Querying and repairing inconsistent numerical databases. *ACM Trans. Database Syst.* 35(2) (2010)
- [23] Franconi, E., Laureti Palma, A., Leone, N., Perri, S., Scarcello, F.: Census Data Repair: A Challenging Application of Disjunctive Logic Programming. In: Nieuwenhuis, R., Voronkov, A. (eds.) *LPAR 2001. LNCS (LNAI)*, vol. 2250, pp. 561–578. Springer, Heidelberg (2001)
- [24] Fuxman, A., Miller, R.: First-order query rewriting for inconsistent databases. *J. Computer and System Sciences* 73(4), 610–635 (2007)
- [25] Gardezi, J., Bertossi, L., Kiringa, I.: Matching dependencies with arbitrary attribute values: semantics, query answering and integrity constraints. In: *Proc. LID* (2011)
- [26] Gardezi, J., Bertossi, L., Kiringa, I.: Matching dependencies: semantics, query answering and integrity constraints. *Frontiers of Computer Science* 6(3), 278–292 (2012)
- [27] Gardezi, J., Bertossi, L.: Query answering under matching dependencies for data cleaning: Complexity and algorithms, arXiv:1112.5908v1
- [28] Gardezi, J., Bertossi, L.: Query Rewriting using Datalog for Duplicate Resolution (extended version), <http://people.scs.carleton.ca/~bertossi/papers/datalog22Long.pdf>
- [29] Greco, G., Greco, S., Zumpano, E.: A logical framework for querying and repairing inconsistent databases. *IEEE Trans. Knowledge and Data Eng.* 15(6), 1389–1408 (2003)
- [30] Wijsen, J.: Database repairing using updates. *ACM Trans. Database Systems* 30(3), 722–768 (2005)
- [31] Wijsen, J.: On the first-order expressibility of computing certain answers to conjunctive queries over uncertain databases. In: *Proc. PODS* (2010)