

Explaining Repaired Data with CFDs

Joeri Rammelaere
University of Antwerp
Antwerp, Belgium

joeri.rammelaere@uantwerpen.be

Floris Geerts
University of Antwerp
Antwerp, Belgium

floris.geerts@uantwerpen.be

ABSTRACT

Many popular data cleaning approaches are rule-based: Constraints are formulated in a logical framework, and data is considered dirty if constraints are violated. These constraints are often discovered from data, but to ascertain their validity, user verification is necessary. Since the full set of discovered constraints is typically too large for manual inspection, recent research integrates user feedback into the discovery process. We propose a different approach that employs user interaction *only at the start* of the algorithm: a user manually cleans a small set of dirty tuples, and we infer the constraint underlying those repairs, called an *explanation*. We make use of conditional functional dependencies (CFDs) as the constraint formalism. We introduce XPLODE, an on-demand algorithm which discovers the best explanation for a given repair. Guided by this explanation, data can then be cleaned using state-of-the-art CFD-based cleaning algorithms. Experiments on synthetic and real-world datasets show that the best explanation can typically be inferred using a limited number of modifications. Moreover, XPLODE is substantially faster than discovering all CFDs that hold on a dataset, and is robust to noise in the modifications.

PVLDB Reference Format:

Joeri Rammelaere and Floris Geerts. Explaining Repaired Data with CFDs. *PVLDB*, 11(11): xxxx-yyyy, 2018.
DOI: 10.14778/3236187.3236193

1. INTRODUCTION

Dirty data is a significant problem for any sizeable organization. The amounts of data being processed have skyrocketed, and manual cleaning has long become infeasible. Research on data cleaning has provided methods for the automatic detection and repairing of various kinds of dirtiness. These methods are often rule-based: Constraints to which a database should adhere are encoded in some declarative constraint formalism, and repair algorithms modify the data such that no constraints are violated (see [22, 13] for an overview of rule-based approaches to data cleaning).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were invited to present their results at The 44th International Conference on Very Large Data Bases, August 2018, Rio de Janeiro, Brazil.

Proceedings of the VLDB Endowment, Vol. 11, No. 11
Copyright 2018 VLDB Endowment 2150-8097/18/07.
DOI: 10.14778/3236187.3236193

To use these repair methods, the challenge is then to obtain the correct constraints for repairing. Constraints can be discovered automatically from data [7, 15, 21, 30, 28, 33, 37, 18, 10], but such methods often discover an excessive amount of constraints. Many of these are spurious or invalid, and hence not suitable for actually repairing the data. To determine which constraints are *semantically valid* and *useful for repairing*, semantic knowledge of the data is required, which typically only a human user can provide.

We thus wish to harness the user’s knowledge, while limiting the amount of interaction required. We therefore consider user interaction in the form of *manual repairs*. It is natural to assume that users know how certain errors *should be repaired*, based on expertise. Nevertheless, the *formal* constraints underlying these repair actions may be unknown to the user or hard to formulate precisely. Indeed, verifying candidate constraints may involve the inspection of the entire dataset to ensure that their violation sets coincide precisely with the errors in the data. The initial user effort can possibly be guided by high-precision error detection and data profiling methods [32, 2, 31]: even if such methods don’t detect all errors, and cannot repair them, they can bootstrap the repairing process. We develop a method that generates constraints that are *consistent* with the repairs made by a user. We refer to such constraints as *explanations*. Intuitively, when such an explanation is used for *repairing*, the user provided repairs are left *intact*. As such, we maximally take the user input into account. The actual repairs, based on the explanations, can subsequently be performed using any state-of-the-art repair algorithm, such as those presented in [4, 12, 16, 17, 23, 25], among others.

In this paper we will focus on the problem of finding *the best possible explanation* for a given (partial) repair. Hence, our method should be seen as a single, core component in a *larger interactive data quality process* in which a *full* repair is gradually constructed – by interweaving manual repairs by a user and automatic repairs based on explanations. As more repaired tuples become available, the corresponding “best” explanations become more likely to be the *correct ones for repairing*, and better assistance can be offered to the user.

Methodology. As constraint formalism we use the class of *conditional functional dependencies* (CFDs) [14], widely used in data cleaning. CFDs are easy to interpret, yet more expressive than regular functional dependencies (FDs). Our proposed method, XPLODE (for **eXplanation on demand**), extracts a single constraint, the *best explanation*, from a dirty dataset and an associated repair. We introduce a scoring function to assess the quality of an explanation, and

Table 1: Running example: A customers dataset.

TID	CC	AC	PN	NM	STR	CT	ZIP
1	01	908	1111111	Mike	Tree Ave.	LA MH	07974
2	01	908	1111111	Rick	Tree Ave.	GLA MH	07974
3	01	212	2222222	Joe	5th Ave	NYC	01202
4	01	908	2222222	Jim	Elm Str.	MH	07974
5	44	131	3333333	Ben	High St.	EDI	EH41DT
6	44	131	4444444	Ian	High St.	EDI	EH41DT
7	44	908	4444444	Ian	Port PI	MH	W1B1JH
8	44 01	131	2222222	Sean	3rd Str.	UN	01202

prove that the outcome of XPLODE is equivalent to a “naive” method that first discovers all CFDs, and then uses post-processing to find the explanation with maximal score.

Our algorithm is “on-demand” in that it dynamically generates candidate explanations, hereby attempting to quickly discover the best possible explanation of the partial repair. A similar on-demand exploration was used in [19] to discover constant patterns that fit a single given CFD. We generalize the theoretical foundation of their algorithm, extending it to a larger class of evaluation functions, called loose anti-monotonic [5]. This class includes useful functions such as a minimum, maximum, or average, and our scoring function.

Underlying XPLODE is a CFD discovery process that employs the user’s modifications to navigate the search space of constraints, and quickly finds the best explanation. What differentiates XPLODE from other constraint discovery methods that employ user responses to prune the search space [20, 34] is that we *only require modifications* as user input, and consider *general CFDs*. By contrast, in [20] a constant CFD discovery process is bootstrapped by only a *single* modification, followed by a number of questions to the user about the validity of the *constraints*. Similarly, in [34] FDs capturing errors are found by post-processing a set of approximate FDs. The post-processing again involves questions to the user about the validity of data and *constraints*. Our method only requires *positive* feedback in the form of correct modifications as opposed to a user frequently having to *invalidate* a candidate constraint. Hence, our method spends little time on constraints that are not suitable for repairing. Moreover, by considering all user feedback together, our method becomes more robust to small mistakes, instead of requiring every individual user interaction to be fully correct.

Motivating Example. For illustration, we borrow the running example from [15]. Table 1 shows a dirty version D_{dirty} and a clean version D_{rep} of the data. In the clean version, the three crossed out values are replaced by those next to it. In other words, a user is to repair D_{dirty} by changing the cities (CT) in $t_1[CT]$ from “LA” to “MH” and in $t_2[CT]$ from “GLA” to “MH”, and the country code (CC) in $t_8[CC]$ from 44 to 01. We assume that a user is faced with only the dirty data, and unable to exactly formalize the CFD that is supposed to hold. Nevertheless, based on experience or real-world knowledge, a user may be able to clean some tuples. We then wish to derive a constraint underlying the modifications made by the user. Since the obtained repair is assumed to be partial, a CFD discovery algorithm can only suggest a useful CFD by discovering approximate CFDs, i.e., CFDs that only partially hold. However, there are too many approximate CFDs for a user to inspect. Instead, we discover CFDs based on corrections made by the user.

For example, suppose that the user corrects the “error” in t_1 by restoring $t_1[CT]$ back to its correct value of “MH”.

Suppose that an algorithm is at hand that only discovers CFDs that somehow “explain” this correction. Intuitively, this corresponds to the CFDs “becoming cleaner” in the repaired version. Two such FDs are $([ZIP, AC] \rightarrow CT)$ and $([AC, CC] \rightarrow CT)$. Indeed, if $t_1[CC] = \text{“MH”}$, then both FDs can be made to hold by removing a single tuple (t_2). Before the correction, two deletions were required. That is, these FDs have become cleaner. When considering CFDs as well, $([CC, PN] \rightarrow CT, (01, 1111111, MH))$ and $(NM \rightarrow CT, (Mike, MH))$, and many other CFDs become candidate explanations. Previous approaches would now ask multiple questions to the user, as to which of the candidate CFDs are (in)valid, until a semantically valid CFD is obtained. In this paper, we propose to let the user continue with repairing the dirty instance. For example, the user may decide to correct $t_8[CC]$ back to 01. One can verify (as we did experimentally) that, for certain thresholds on support and confidence, as will be defined later, only the FD $\varphi = ([AC, CC] \rightarrow CT)$ can be related to the two modifications made by the user. This FD, saying that country code and area code uniquely determine city, is known to be semantically valid on this data. Moreover, φ can now be used to automatically correct $t_2[CT]$ to “MH”, using any CFD-based repair algorithm.

Summary of Contributions.

1. We formally define what it means for a CFD to *explain* a set of modifications. To differentiate between different explanations, we define a scoring function based on the number of explained modifications. (Section 3)
2. We design an algorithm, XPLODE, that discovers the best explanation, i.e., the explanation of highest score. Moreover, XPLODE is an *on-demand* algorithm that avoids a full exploration of the search space, when possible. To this aim, XPLODE leverages an upper bound of the scoring function which is *loose anti-monotonic*. We discuss how XPLODE can be modified to discover *multiple explanations*. (Section 4)
3. To further increase the efficiency of XPLODE, we introduce an approximate scoring function that can be computed in time linear in the number of changes made by the user. By contrast, the actual scoring function has an inherent exponential dependency on the number of changes. (Section 5)
4. We experimentally show that our method can discover the correct CFD for repairing from only a small number of modifications, saving considerable user effort compared to manual validation of constraints ranked by baselines such as confidence. Moreover, XPLODE is robust to noise in the modifications, i.e., mistakes made by the user, and outperforms a post-processing approach that discovers all explanations and then finds the highest scoring one. (Section 6)

2. PRELIMINARIES

We consider a relation schema R defined over a set \mathcal{A} of k attributes, where each attribute $A_i \in \mathcal{A}$ is associated with a domain $\text{dom}(A_i)$. A tuple t over R is simply an element of $\text{dom}(A_1) \times \dots \times \text{dom}(A_k)$. A (database) instance D of R is a finite set of tuples over R . For a set X of attributes in \mathcal{A} and tuple t over R , $t[X]$ denotes the projection of that tuple on the attributes in X . We assume that each tuple $t \in D$ has a unique identifier tid , e.g., a natural number. We denote by $D[\text{tid}]$ the tuple t in D with identifier tid .

A *conditional functional dependency* (CFD) [14] φ over R is a pair $(X \rightarrow A, t_p)$, where (i) X is a set of attributes in \mathcal{A} , and A is a single attribute in \mathcal{A} ; (ii) $X \rightarrow A$ is a standard

functional dependency (FD); and (iii) t_p is a *pattern tuple* with attributes in X and A , where for each B in $X \cup \{A\}$, $t_p[B]$ is either a constant ‘ b ’ in $\text{dom}(B)$, or an unnamed variable ‘ $_$ ’. A CFD $\varphi = (X \rightarrow A, t_p)$ in which $t_p[A] = _$ is called *variable*, otherwise it is called *constant*. For constant CFDs, $t_p[X]$ may be assumed to consist of constants only. An FD is a (variable) CFD with t_p consisting solely of variables ‘ $_$ ’.

The semantics of a CFD $\varphi = (X \rightarrow A, t_p)$ on an instance D is defined as follows. A tuple $t \in D$ is said to *match* a pattern tuple t_p in attributes X , denoted by $t[X] \asymp t_p[X]$, if for all $B \in X$, either $t_p[B] = _$, or $t[B] = t_p[B]$. The tuple t *violates* a variable CFD $\varphi = (X \rightarrow A, t_p)$ if $t[X] \asymp t_p[X]$ and there exists another tuple t' in D such that $t[X] = t'[X]$ and $t[A] \neq t'[A]$. A tuple t *violates* a constant CFD $\varphi = (X \rightarrow A, t_p)$ if $t[X] = t_p[X]$ and $t[A] \neq t_p[A]$. The set of all tids of tuples in D that together violate a CFD φ is denoted by $\text{VIO}(\varphi, D)$. If $\text{VIO}(\varphi, D) = \emptyset$, then D *satisfies* φ , denoted by $D \models \varphi$.

The *support* of a CFD $\varphi = (X \rightarrow A, t_p)$ in D , denoted by $\text{supp}(\varphi, D)$, is defined as the number of tuples $t \in D$ such that $t[X] \asymp t_p[X]$, i.e., the support is the number of tuples in D that match the pattern of φ on the attributes in X . In line with the commonly used notion of confidence for approximate FDs [24, 21], we define the *confidence* of a CFD $\varphi = (X \rightarrow A, t_p)$ in D as $\text{conf}(\varphi, D) = 1 - \frac{|D'|}{\text{supp}(\varphi, D)}$, where $D' \subseteq D$ is a minimal subset such that $D \setminus D' \models \varphi$. This definition is suited to variable CFDs, where the set $\text{VIO}(\varphi, D)$ contains all tuples that *together* violate the CFD, but are not necessarily violations by themselves. For instance, if a violation set for a variable CFD contains two tuples with different A -values, the CFD can be made to hold by altering just one of the tuples. In other words, $|D'|$ is the minimum number of tuples that need to be altered or removed for φ to be satisfied. For a constant CFD, $|\text{VIO}(\varphi, D)| = |D'|$, and hence this definition of confidence reduces to the standard confidence of an association rule. We observe that $\text{conf}(\varphi, D) = 1$ means that $D \models \varphi$ and $\text{conf}(\varphi, D) = 0$ means that all tuples matching $t_p[X]$ need to be altered.

3. EXPLAINING REPAIRS

We here formalize the problem of discovering a *single* CFD that best explains an observed, possibly partial, repair of the data. Such repairs are represented by *modifications*. We describe what modifications are in Section 3.1. *Explaining* modifications in terms of CFDs, and differentiating between explanations based on a *scoring function*, is discussed in Section 3.2. Our problem statement is given in Section 3.3.

3.1 Modifications

We consider a setting in which no CFDs are provided. Instead, we have at our disposal two database instances, D_{rep} and D_{dirty} , where D_{rep} is obtained from the “dirty” instance D_{dirty} by applying a number of *modifications* to tuples. We assume that these instances have the same set of tids, such that tuples $D_{\text{rep}}[\text{tid}]$ and $D_{\text{dirty}}[\text{tid}]$ are both well-defined for every tid occurring in either instance. Every changed tuple in D_{rep} contains one or more modifications:

DEFINITION 1 (MODIFICATION). A *modification* \mathbf{m} is a quadruple $\mathbf{m} = (\text{tid}, A, v_d, v_c)$, where tid is the identifier of the tuple that is being changed, A is the attribute that is changed, v_d is the dirty value which was replaced, and v_c is the new, clean value, different from v_d . Given D_{dirty} and D_{rep} , a modification $\mathbf{m} = (\text{tid}, A, v_d, v_c)$ is *consistent*

with D_{dirty} and D_{rep} when, for tuples $s = D_{\text{dirty}}[\text{tid}]$ and $t = D_{\text{rep}}[\text{tid}]$, $s[A] = v_d$ and $t[A] = v_c$. \square

Given D_{dirty} and D_{rep} , we denote by $\mathfrak{M}(D_{\text{dirty}}, D_{\text{rep}})$ the set of all modifications that are consistent with D_{dirty} and D_{rep} . Observe that \mathfrak{M} can contain at most one modification for each tid and attribute. It should be clear that D_{dirty} and D_{rep} uniquely determine $\mathfrak{M}(D_{\text{dirty}}, D_{\text{rep}})$, as it is merely the “diff” of these two instances. We simply write \mathfrak{M} whenever the dirty and modified instances are clear from the context.

Given a set of modifications $M \subseteq \mathfrak{M}(D_{\text{dirty}}, D_{\text{rep}})$, we denote by $D_{\text{dirty}} \oplus M$ the version of D_{dirty} on which the modifications in M are applied. Consequently, $D_{\text{dirty}} \oplus \{\emptyset\} = D_{\text{dirty}}$ and $D_{\text{dirty}} \oplus \mathfrak{M}(D_{\text{dirty}}, D_{\text{rep}}) = D_{\text{rep}}$. We let $\sigma_M(D)$ denote the set of tuples in D that are involved in a modification, i.e., whose tids occur in a modification in M . Furthermore, $\sigma_M^{\text{tid}}(D)$ denotes the set of tids in $\sigma_M(D)$.

EXAMPLE 1. In our example, \mathfrak{M} consists of the modifications $\mathbf{m}_1 = (1, \text{CT}, \text{LA}, \text{MH})$, $\mathbf{m}_2 = (2, \text{CT}, \text{GLA}, \text{MH})$, and $\mathbf{m}_3 = (8, \text{CC}, 44, 01)$. Both $\sigma_{\mathfrak{M}}(D_{\text{dirty}})$ and $\sigma_{\mathfrak{M}}(D_{\text{rep}})$ consist of tuples t_1 , t_2 and t_8 , with tids 1, 2 and 8. \diamond

3.2 Explaining a Repair

Suppose that we are given a repair, represented by a set of modifications. These modifications are made by a user without any assumed knowledge of CFDs that may be required to hold on the data, capturing the semantics of “clean data”. As explained in the introduction, we want to recommend a CFD based on the current repair, such that this CFD may be used to detect further errors and suggest modifications.

(ε, δ)-CFDs. Intuitively, we relate candidate CFDs for “*explaining*” the current repair to D_{rep} , as this database instance is regarded to be cleaner than D_{dirty} . To narrow down such candidate CFDs, consider the following example.

EXAMPLE 2. Consider the partial repair in our running example corresponding to the single modification $\mathbf{m}_1 = (1, \text{CT}, \text{LA}, \text{MH})$. The CFD $\varphi_1 = (\text{NM} \rightarrow \text{CT}, (\text{Mike}, \text{MH}))$ could serve as an explanation for this modification as it is satisfied on the partial repair and relates to tuple t_1 . It is unlikely, however, that this CFD is useful. Indeed, φ_1 is only supported by a single tuple (t_1) and does not relate at all to, for example, modification $\mathbf{m}_2 = (2, \text{CT}, \text{GLA}, \text{MH})$ that will be made by the user to tuple t_2 . Consider next CFD $\varphi_2 = (\emptyset \rightarrow \text{CT}, (\text{MH}))$, stating that all cities should be MH. It is a well-supported CFD (its support is the entire database) and relates to \mathbf{m}_1 and \mathbf{m}_2 . It may not be useful, however, for further cleaning of the data. Indeed, φ_2 has a very low confidence: more than half of the data violates this CFD and hence too many tuples may be flagged as dirty. \diamond

This example illustrates the need for ensuring that explanations have sufficient *support*, as this excludes CFDs that are only supported by remaining errors and noise in D_{rep} . At the same time, since D_{rep} is only a partial repair, we should focus on explanations that only hold *approximately* in D_{rep} . All combined, this motivates the following definition.

DEFINITION 2 ((ε, δ)-CFD). A CFD φ on D_{rep} is an (ε, δ)-CFD if $\text{conf}(\varphi, D_{\text{rep}}) \geq 1 - \varepsilon$ and $\text{supp}(\varphi, D_{\text{rep}}) \geq \delta$. We denote by $\Sigma_{(\varepsilon, \delta)}(D_{\text{rep}})$ the set of all such (ε, δ)-CFDs. \square

From here on, the set $\Sigma_{(\varepsilon, \delta)}(D_{\text{rep}})$ is our candidate set of CFDs for explaining repairs. But what does it mean to “explain a repair”? We next address this question.

M-Repair Explanations. Our definition for explanations is based on the following intuitive condition:

“A CFD explains a repair if the repair improves the cleanliness of the data w.r.t. the CFD.”

We formalize this intuition by imposing three natural conditions on CFDs φ in $\Sigma_{(\varepsilon, \delta)}(D_{rep})$. Let $\mathfrak{M} = \mathfrak{M}(D_{dirty}, D_{rep})$.

1. The confidence of φ in D_{rep} should have increased compared to its confidence in D_{dirty} as the result of the modifications made to D_{dirty} . Since a confidence of 1 means that φ is no longer violated, an increase in confidence brings φ closer to being satisfied in the repair. Note that if the confidence of φ increases, then φ was necessarily violated in D_{dirty} .

2. We require that at least one of the violations of φ in D_{dirty} has a tid occurring in $\sigma_{\mathfrak{M}}^{tid}(D_{dirty})$, i.e., a tid of a modified tuple. This ensures that the increase in confidence, as required by the first condition, is the deliberate effect of resolving a violation of φ . Such a condition is necessary: in the running example, the CFD $(CC \rightarrow PN, (01, 2222222))$ is not violated on t_8 in the dirty data. However, modification m_3 in t_8 does increase the confidence of this CFD.

3. We require that φ is not violated in $\sigma_{\mathfrak{M}}(D_{rep})$, the part of the data that was specifically cleaned by the user. This ensures that if one uses φ later for further repairing, using any state-of-the-art CFD-based repairing algorithm, the tuples in $\sigma_{\mathfrak{M}}(D_{rep})$ will not be altered, i.e., they remain clean.

We next state these conditions more generally in terms of a set $M \subseteq \mathfrak{M}(D_{dirty}, D_{rep})$ of modifications:

DEFINITION 3 (M-REPAIR EXPLANATION). Consider instances D_{dirty} , D_{rep} , and modifications $M \subseteq \mathfrak{M}(D_{dirty}, D_{rep})$. A CFD $\varphi = (X \rightarrow A, t_p)$ is an *M-repair explanation* if

1. $\text{conf}(\varphi, D_{dirty} \oplus M) > \text{conf}(\varphi, D_{dirty})$;
2. $\text{VIO}(\varphi, D_{dirty}) \cap \sigma_M^{tid}(D_{dirty})$ is not empty; and
3. $\text{VIO}(\varphi, \sigma_M(D_{dirty} \oplus M))$ is empty.

We denote by $\text{Explain}_{(\varepsilon, \delta)}(D_{dirty} \oplus M)$ the set of *M-repair explanations* in $\Sigma_{(\varepsilon, \delta)}(D_{dirty} \oplus M)$. \square

Of particular interest is the case when $M = \mathfrak{M}$ and hence $D_{dirty} \oplus M = D_{rep}$. In this case, we also call CFDs in $\text{Explain}_{(\varepsilon, \delta)}(D_{rep})$ *global explanations*. The set of global explanations, however, can be quite sizeable, as we will show in the experimental section (Section 6). Worse still, explanations in this set do not necessarily relate to all modifications in \mathfrak{M} . In fact, a CFD may be an explanation because of just one of many modifications, as is illustrated next.

EXAMPLE 3. *When discovering global explanations in our example dataset, with $\varepsilon = 0.25$ and $\delta = 2$, there are 18 candidate CFDs. Among these, FD $([AC, CC] \rightarrow ZIP, (-, -, -))$ is a global explanation, for $\mathfrak{M} = \{m_1, m_2, m_3\}$, yet it is only related to one modification, namely m_3 . \diamond*

In order to distinguish between such “good” and “bad” explanations, we need to strengthen the connection between explanations and modifications. Instead of enforcing extra conditions on global explanations, we assign to them a quality metric. Observe that in the previous example, the CFD $([AC, CC] \rightarrow ZIP, (-, -, -))$ is a global explanation, i.e., it is an *M-repair explanation* for the entire set of modifications $M = \mathfrak{M}$, thanks to m_3 , but it is not an *M-repair explanation* for the subset $M = \{m_1, m_2\}$. We therefore introduce the concept of *locally explaining* modifications, on which we will base our quality metric.

Local Repair Explanations. Example 3 indicates that good global explanations should also explain the modifications involved *locally*. Intuitively, this means that the three conditions stated in Definition 3 should not only hold for the entire M , but also for subsets $M' \subseteq M$. Requiring these conditions to hold for all subsets of \mathfrak{M} is too strong, however, as there may not exist any global explanations with this property. We start by defining what it means for a CFD to locally explain a set of modifications.

DEFINITION 4 (LOCAL REPAIR EXPLANATION). Given a CFD φ in $\text{Explain}_{(\varepsilon, \delta)}(D_{rep})$, we say that φ *locally explains* a set $M \subseteq \mathfrak{M}$ if for every non-empty subset $M' \subseteq M$, φ is an M' -repair explanation. \square

In other words, when a global explanation φ locally explains M , it explains all repairs $D_{dirty} \oplus M'$ that can be obtained from applying modifications $M' \subseteq M$ to D_{dirty} . That is, in any order in which the modifications in M are applied, the cleanliness of the data w.r.t. φ improves at every step.

This notion gives rise to the following quality metric. Let φ be a global explanation in $\text{Explain}_{(\varepsilon, \delta)}(D_{rep})$. Then,

$$\text{score}(\varphi, \mathfrak{M}) := \max\{|M| \mid M \subseteq \mathfrak{M} \text{ and } \varphi \text{ locally explains } M\}.$$

If φ has a score close to $|\mathfrak{M}|$, then almost all modifications in \mathfrak{M} are both globally and locally explained. We are thus interested in global explanations with a high score.

EXAMPLE 4. *In the running example, the CFD $([CC, AC] \rightarrow ZIP, (-, -, -))$ has a score of 1, since it only explains modification m_3 . Indeed, it is easily verified that m_1 and m_2 do not improve the confidence of this CFD. Moreover, if a user would only supply m_1 and m_3 , this CFD could not be used to automatically clean the remainder of the data (i.e., apply m_2). On the other hand, the CFD $([CC, AC] \rightarrow CT, (-, -, -))$ can explain all 3 modifications, leading to a perfect score of 3. Even if only m_1 and m_3 were supplied by the user, this CFD would have the highest score of 2, and could automatically apply m_2 . This example strengthens our argument that user-supplied modifications can guide the CFD discovery process towards the CFD that is most useful for cleaning the remainder of the data. \diamond*

3.3 Problem Statement

We now have all ingredients for our problem statement:

PROBLEM: Repair Explanation Discovery

INPUT: Instances D_{dirty} and D_{rep} , modifications $\mathfrak{M}(D_{dirty}, D_{rep})$, thresholds ε and δ .

OUTPUT: A global explanation $\varphi \in \text{Explain}_{(\varepsilon, \delta)}(D_{rep})$ such that $\text{score}(\varphi, \mathfrak{M})$ is maximal.

In other words, we want to find the global explanation that also locally explains the largest subset of modifications of \mathfrak{M} . Any solution to this problem has somehow embedded in it the problem of discovering CFDs. Lattice traversal algorithms for CFD discovery [15, 7], exhibit an inherent exponential dependency in $|\mathcal{A}|d$, where $|\mathcal{A}|$ is the number of attributes and d is the maximum number of values occurring in any attribute. We will show that testing whether or not candidate CFDs satisfy the support and confidence thresholds and are global explanations imposes minimal overhead on the overall CFD discovery process. The score computation,

however, has a severe impact on the performance. Indeed, for each candidate global explanation it requires (worst case) to consider all subsets of \mathfrak{M} . Not all is lost, however. We will show in Section 5, that an efficient and good approximation of the scoring function can be computed.

4. DISCOVERING REPAIR EXPLANATIONS

We next describe an algorithm for discovering the best repair explanation, as defined in our problem statement. A solution to this problem is to first discover all CFDs which globally explain the observed repair, and then return the CFD with the highest score in a post-processing step. Clearly, this method does a lot of unnecessary work. Instead, we present an on-demand algorithm XPLODE, which returns the best explanation as soon as it is known. We begin in Section 4.1 with a detailed overview of XPLODE. Correctness of XPLODE is shown in Section 4.2, contingent on the availability of a *loose anti-monotonic* upper bound function on the scores of certain sets of CFDs. Examples of such upper bounds are provided as well. In Section 4.3, we briefly discuss how to modify XPLODE for returning multiple CFDs. Implementation details are presented in Section 4.4.

4.1 XPlode: Explanations On-Demand

We first provide a detailed overview of algorithm XPLODE (for eXplanations on-demand) and refer to Algorithm 1 for its pseudo-code. At the core of XPLODE is a traversal of a lattice, commonly used in FD and CFD discovery algorithms, in combination with the use of equivalence partitions to check whether CFDs are global explanations. These partitions, also used in FD and CFD discovery algorithms such as TANE [21] and CTANE [15], respectively, are well suited for the (ε, δ) -thresholds that we impose on confidence and support. We elaborate further on this partitioning technique in Section 4.4. For now, we assume that we can efficiently check whether or not a CFD φ is in $\text{Explain}_{(\varepsilon, \delta)}(D_{\text{dirty}} \oplus \mathfrak{M})$ for a given instance D_{dirty} and set \mathfrak{M} of modifications, and that $\text{score}(\varphi, \mathfrak{M})$ can be computed easily.

The description below focuses instead on how the best explanation can be found *during* the exploration of the lattice, rather than by post-processing $\text{Explain}_{(\varepsilon, \delta)}(D_{\text{dirty}} \oplus \mathfrak{M})$. The challenge is to quickly pinpoint a CFD in $\text{Explain}_{(\varepsilon, \delta)}(D_{\text{dirty}} \oplus \mathfrak{M})$ with guaranteed highest score.

To explain the workings of XPLODE, we need to introduce some concepts. First of all, the lattice that we traverse is the *power set lattice* of all attribute/value combinations (including the wildcard ‘.’ as a special value). As in CTANE [15], we denote the lattice elements by (X, t_p) , where X is a set of attributes and t_p is a pattern tuple over X . Hence, $(X, t_p) \subseteq (Y, s_p)$ iff $X \subseteq Y$ and $t_p = s_p[X]$. In this case, we also say that (Y, s_p) is a *child* of (X, t_p) and the children of (X, t_p) are thus obtained by expanding (X, t_p) with all possible attribute/value pairs (A, a) where $A \notin X$ and $a \in \text{dom}(A) \cup \{.\}$. Furthermore, since we are only interested in CFDs of high support, it suffices to only consider pairs (A, a) that have sufficient support in D_{rep} . That is, the number of tuples $t \in D_{\text{rep}}$ such that $t[A] \succ a$ should exceed δ .

Second, as in CTANE, an element (X, t_p) represents a set of *candidate* CFDs, denoted by $\text{CandCFD}(X, t_p)$, consisting of all CFDs of the form $(X \setminus \{A\} \rightarrow A, t_p)$, for $A \in X$.

Finally, let $\text{UB}(X, t_p)$ be an upper bound on $\text{score}(\varphi, \mathfrak{M})$ for any $\varphi \in \text{CandCFD}(X, t_p)$. As will be explained shortly, this upper bound function is used as a guide for the traversal

through the lattice, and serves to ensure that the highest-score global explanation can be identified *without the need for exploring the full lattice*.

Algorithm XPLODE relies on a traversal of the lattice such that its elements (X, t_p) are visited in *descending order* according to their upper bound $\text{UB}(X, t_p)$. To this aim, we keep generated elements in a priority queue Φ , initially containing (\emptyset, \emptyset) with upper bound $+\infty$ (line 3). During the run of the algorithm, we also maintain the global explanation φ_{max} with highest score seen so far, denoting its score by “**max**”. Initially, φ_{max} is set to nil and **max** = 0 (line 4). The algorithm ensures that, at any time, the queue Φ will only consist of elements (X, t_p) such that $\text{UB}(X, t_p) > \text{max}$. In other words, we only visit elements if a better explanation can possibly be found among its candidate CFDs.

Suppose XPLODE is currently exploring element (X, t_p) , i.e., the foremost element in the queue Φ with highest upper bound (line 6). For the current element (X, t_p) , a CFD φ in $\text{CandCFD}(X, t_p)$ is selected, that (i) is a global explanation; and (ii) has highest score among all other global explanations in $\text{CandCFD}(X, t_p)$, if such a CFD exists (line 7). The score of φ , and the scores of all CFDs in $\text{CandCFD}(X, t_p)$ have already been computed at an earlier stage of the algorithm, when (X, t_p) was generated. We expand on this later.

If φ exists and $\text{score}(\varphi, \mathfrak{M}) > \text{max}$, then φ is a better explanation than φ_{max} . In this case, φ_{max} is set to φ and **max** to $\text{score}(\varphi, \mathfrak{M})$ (line 9). Furthermore, the queue Φ is updated by removing each element with an upper bound smaller than or equal to the new **max**-value (line 10). This guarantees that all elements in Φ have an upper bound larger than the current **max**-value, as pointed out previously.

Finally, all children (Y, s_p) of (X, t_p) are generated and the scores of all candidate CFDs in $\text{CandCFD}(Y, s_p)$ are computed (lines 11-13). We can thus indeed assume, as we did earlier, that when XPLODE considers (Y, s_p) at a later stage, all scores of its candidate CFDs are available. Furthermore, if $\text{UB}(Y, s_p) > \text{max}$, then (Y, s_p) is inserted in Φ with upper bound value $\text{UB}(Y, s_p)$ (line 15). This guarantees that Φ contains only elements with a sufficiently high upper bound.

When the queue is empty, the algorithm terminates by returning φ_{max} (line 16). If $\varphi_{\text{max}} \neq \text{nil}$, the CFD φ_{max} is guaranteed to be a global explanation. In the next section, we identify sufficient conditions on the upper bound function such that φ_{max} is a global explanation with *maximal score*.

Remarks. (1) Multiple elements in Φ may hold the same maximal UB-value. We break ties by prioritizing on elements (X, t_p) with the highest-score CFD in $\text{CandCFD}(X, t_p)$. If ties persist, we pick (X, t_p) containing the highest number of wildcards in t_p . These choices can lead XPLODE quicker to elements that represent the best global explanation. (2) One may wonder why, when generating a child (Y, s_p) of (X, t_p) , we compute the scores of all its candidate CFDs (lines 11-13), and do not limit the score computation to candidate CFDs that are in fact global explanations. Indeed, CFDs that do not globally explain repairs do not impact the final result and consequently, their score could be simply set to 0. This would result in avoiding potentially expensive score computations. However, since scores are used for tie-breaking, experiments show that it is more efficient if we do compute all scores. Intuitively, more scores lead more quickly towards elements with higher scores, leading to a swifter discovery of an explanation with maximal score.

Algorithm 1 On-demand algorithm XPLoDE for obtaining the best explanation for a set \mathfrak{M} of modifications.

```

1: procedure XPLoDE( $D_{dirty}, D_{rep}, \mathfrak{M}, \varepsilon, \delta, \text{score}(\cdot), \text{UB}(\cdot)$ )
2:    $\Phi \leftarrow \text{PRIORITYQUEUE}(\{\})$ 
3:   Insert  $(\emptyset, \emptyset)$  into  $\Phi$  with upper bound  $+\infty$ 
4:    $\varphi_{\max} \leftarrow \text{nil}, \text{max} \leftarrow 0$ 
5:   while  $\Phi$  is not empty do
6:      $(X, t_p) \leftarrow \text{POP}(\Phi)$ 
7:     Let  $\varphi \in \text{CandCFD}(X, t_p)$  such that  $\varphi$  is a global
      explanation with highest score among all global ex-
      planations in  $\text{CandCFD}(X, t_p)$ 
8:     if  $\varphi$  exists and  $\text{score}(\varphi, \mathfrak{M}) > \text{max}$  then
9:        $\varphi_{\max} \leftarrow \varphi, \text{max} \leftarrow \text{score}(\varphi, \mathfrak{M})$ 
10:      Delete from  $\Phi$  all elements with  $\text{UB}$ -value  $\leq \text{max}$ 
11:      for all children  $(Y, s_p)$  of  $(X, t_p)$  do
12:        for all  $\psi \in \text{CandCFD}(Y, s_p)$  do
13:          Compute  $\text{score}(\psi, \mathfrak{M})$ 
14:        if  $\text{UB}(Y, s_p) > \text{max}$  then
15:          Insert  $(Y, s_p)$  into  $\Phi$  with value  $\text{UB}(Y, s_p)$ .
16:   return  $\varphi_{\max}$ .

```

4.2 Correctness and Upper Bound Functions

We next show how to guarantee that when XPLoDE outputs a global explanation, it is a global explanation of maximal score. The correctness of XPLoDE entirely relies on the upper bound function $\text{UB}(\cdot)$, as this function determines which elements are in the priority queue, and in what order. We first identify sufficient conditions on $\text{UB}(\cdot)$ to guarantee correctness. Examples of “good” upper bound functions are described at the end of this section.

Correctness. XPLoDE returns φ_{\max} when the priority queue Φ is empty. Clearly, every element that was ever generated has either been visited, or is not in the queue because its UB -value is below the score max of φ_{\max} . Observe that for any element (X, t_p) , its upper bound $\text{UB}(X, t_p)$ is larger than the score of any candidate CFD of (X, t_p) . This implies that none of the elements generated during execution, have a candidate CFD with a score higher than max .

Correctness of XPLoDE then requires that, when φ_{\max} is returned as the best global explanation, then any element (X, t_p) in the lattice whose upper bound value is larger than max must have been added to the queue at some prior stage. To this aim, we require that the upper bound function $\text{UB}(\cdot)$ is *loose anti-monotonic* [5]: For any (X, t_p) there exists a parent $(Y, s_p) = (X \setminus \{B\}, t_p[X \setminus \{B\}])$ such that $\text{UB}(X, t_p) \leq \text{UB}(Y, s_p)$. That is, every element in the lattice has at least one parent with a higher or equal UB -value. The proof of correctness is deferred to the appendix.

PROPOSITION 1. On input $D_{dirty}, D_{rep}, \mathfrak{M}, \varepsilon, \delta, \text{score}(\cdot)$, and $\text{UB}(\cdot)$, the algorithm XPLoDE returns the global explanation with maximal score, if it exists, provided that $\text{UB}(\cdot)$ is loose anti-monotonic and for any element (X, t_p) , $\text{UB}(X, t_p)$ is larger than the score of any of its candidate CFDs. \square

Loose Anti-monotonic Upper Bounds. The question is now whether there exist non-trivial¹ upper bound functions that satisfy the conditions in Proposition 1. We answer this affirmatively in this section. Defining $\text{UB}(X, t_p)$

¹One can choose a constant function $\text{UB}(\cdot)$. XPLoDE then performs an exhaustive breadth-first lattice traversal.

as the maximal *score* of the candidate CFDs of (X, t_p) does not suffice, however, as the following example illustrates.

EXAMPLE 5. We return to the running example, and consider the CFD $\varphi = (\text{CC} \rightarrow \text{CT}, (-, -))$, which locally explains all modifications $\{\mathfrak{m}_1, \mathfrak{m}_2, \mathfrak{m}_3\}$. Hence, its score is 3 (but it is not sufficiently confident, for $\varepsilon = 0.25$, to be a global explanation). This is a candidate CFD for the element $(\{\text{CC}, \text{CT}\}, (-, -))$ in the lattice, with parents $(\text{CC}, -)$ and $(\text{CT}, -)$. However, the CFDs $\varphi_1 = (\emptyset \rightarrow \text{CT}, (-))$ and $\varphi_2 = (\emptyset \rightarrow \text{CC}, (-))$ only have scores of 2 and 1, respectively. Indeed, φ_1 locally explains $\{\mathfrak{m}_1, \mathfrak{m}_2\}$ and φ_2 locally explains \mathfrak{m}_3 . No larger sets are locally explained by these CFDs. \diamond

Instead, we observe the following. For a CFD φ , we define $\text{ModVIO}(\varphi, D_{dirty}, \mathfrak{M})$ as the set of modifications in \mathfrak{M} that apply to tuples in $\text{VIO}(\varphi, D_{dirty})$. Here, a modification \mathfrak{m} applies to a tuple t when they share the same *tid*-value. It follows directly from Definition 4 that the set of modifications that a global explanation can also locally explain, consists at most of those modifications involved in violations of that explanation in the dirty database instance. Hence, $\text{score}(\varphi, \mathfrak{M}) \leq |\text{ModVIO}(\varphi, D_{dirty}, \mathfrak{M})|$. We can now define our upper bound function UB_0 ; we defer the proof of validity of the upper bound to the appendix.

DEFINITION 5 (UPPER BOUND). Let (X, t_p) be a lattice element, D_{dirty} a dirty dataset, and \mathfrak{M} a set of modifications.

$$\text{UB}_0(X, t_p) := \max_{\varphi \in \text{CandCFD}(X, t_p)} |\text{ModVIO}(\varphi, D_{dirty}, \mathfrak{M})|. \quad \square$$

PROPOSITION 2. The upper bound function $\text{UB}_0(\cdot)$ satisfies the conditions of Proposition 1. \square

We also introduce another upper bound, based on $\text{UB}_0(\cdot)$, with the difference that it also takes into account the attributes in ModVIO covered by explanations. More specifically, for a CFD φ , we define $\text{AttVIO}(\varphi, D_{dirty}, \mathfrak{M})$ as the set of attributes occurring in $\text{ModVIO}(\varphi, D_{dirty}, \mathfrak{M})$. Furthermore, we let λ be a parameter such that $0 \leq \lambda \cdot |\mathcal{A}| < 1$, where $|\mathcal{A}|$ is the number of attributes in the relation R . We define $\text{UB}_\lambda(X, t_p)$ as the maximum value of

$$|\text{ModVIO}(\varphi, D_{dirty}, \mathfrak{M})| + 1 - \lambda |\mathcal{X} \cup \text{AttVIO}(\varphi, D_{dirty}, \mathfrak{M})|,$$

where, as before, φ ranges over all CFDs in $\text{CandCFD}(X, t_p)$. The intuition behind the additional negative term, compared to $\text{UB}_0(\cdot)$, is to prioritize explanations to more general CFDs (containing a smaller number of attributes) during the execution of XPLoDE. The additional “+1” is to ensure that $\text{score}(\varphi, \mathfrak{M})$ remains smaller than $\text{UB}_\lambda(X, t_p)$ for every of its candidate CFDs φ , since $\lambda \cdot |\mathcal{A}| \leq 1$. Hence, $\text{UB}_\lambda(\cdot)$ only affects the priority among those CFDs explaining an identical number of modifications. The proof that $\text{UB}_\lambda(\cdot)$ is loose anti-monotonic is deferred to the appendix.

4.3 Discovering Multiple Explanations

As mentioned in the introduction, we have devised algorithm XPLoDE to be the core component of an interactive cleaning system. While finding the best explanation given a set of modifications is the crucial step in this system, in practice one might desire the algorithm to return multiple CFDs. We briefly discuss how XPLoDE can be altered to discover (a) the top- k explanations for the current set of modifications, and (b) a sequence of i modifications that incrementally explain the modifications. We believe a combination of these techniques covers most real-world scenario’s.

Discovering Top- k Explanations. Turning XPLODE into a top- k algorithm requires limited changes to the pseudocode shown in Algorithm 1. On line 4, we now initialize φ_{\max} as a list of length k . On line 7, the algorithm has to be changed such that all $\varphi \in \text{CandCFD}(\mathbf{X}, t_p)$ with score $> \max$ are processed, not just the highest-scoring φ . Finally, on line 9, the identified φ is *added* to (the list) φ_{\max} , and if more than k CFDs are present in φ_{\max} , the lowest-scoring one is removed. The value of max is subsequently set to the *lowest* score of those CFDs in φ_{\max} .

Incrementally Explaining Modifications. We next discuss how to discover CFDs that incrementally explain observed modifications. In other words, the best explanation is first discovered, and then the search is continued in order to find the best explanation for the modifications that *have not yet been explained*. To do this efficiently, we first make a change to the lattice elements: instead of associating a *score* with each CFD, we attach the *set of modifications* the CFD explains. This allows us to efficiently recompute the score of a CFD after removing already-explained modifications.

The main change to the algorithm is that we introduce a list *backup* to store generated lattice elements, *in addition* to the priority queue. After line 15 in the algorithm, each generated lattice element (\mathbf{Y}, s_p) is inserted into *backup* if $\text{UB}(\mathbf{Y}, s_p) > 0$, i.e., if a CFD in the element or its children can explain some modification. When the regular XPLODE algorithm finishes, because the best explanation φ_{\max} is found, we now remove all modifications explained by φ_{\max} from \mathfrak{M} . Subsequently, the list *backup* is examined, and the scores of its elements are updated. All elements (\mathbf{Y}, s_p) with $\text{UB}(\mathbf{Y}, s_p) = 0$ are removed, and the others are used to re-initialize the priority queue Φ . The algorithm then repeats, until all modifications are explained.

4.4 Implementation Details

We conclude this section by elaborating on how XPLODE checks the support and confidence thresholds and how global explanations are filtered out. These all crucially rely on so-called equivalence partitions, commonly used in (C)FD discovery algorithms.

Equivalence Partitions. These partitions, introduced in [11] and used in CTANE [15] for CFD discovery (and in its predecessor TANE [21] for FD discovery), are compact representations of sets of tuples that *agree* on sets of attributes.

More specifically, given (\mathbf{X}, t_p) , where \mathbf{X} is a set of attributes and t_p is a pattern tuple over \mathbf{X} , we say that two tuples s and t in D_{rep} are *equivalent relative to* (\mathbf{X}, t_p) if $s[\mathbf{X}] = t[\mathbf{X}] \succ t_p$. For a tuple $s \in D_{rep}$, $[s]_{(\mathbf{X}, t_p)}$ denotes the *equivalence class* consisting of the tids of all tuples $t \in D_{rep}$ that are equivalent with s relative to (\mathbf{X}, t_p) . The (*equivalence*) *partition of* (\mathbf{X}, t_p) , denoted by $\Pi(\mathbf{X}, t_p)$, is the collection of $[s]_{(\mathbf{X}, t_p)}$ for $s \in D_{rep}$. The *size* of $\Pi(\mathbf{X}, t_p)$, denoted by $|\Pi(\mathbf{X}, t_p)|$, is the number of equivalence classes in $\Pi(\mathbf{X}, t_p)$. We use $\|\Pi(\mathbf{X}, t_p)\|$ to denote the number of tids in $\Pi(\mathbf{X}, t_p)$. For instance, in the running example we have that $\Pi(\{\text{CC}, \text{CT}\}, (44, _)) = \{\{5, 6\}, \{7\}\}$ with size $|\Pi(\{\text{CC}, \text{CT}\}, (44, _))| = 2$ and $\|\Pi(\{\text{CC}, \text{CT}\}, (44, _))\| = 3$.

One of the key uses of equivalence partitions in discovery algorithms is to check the validity of CFDs [15]. It is easily verified that $D_{rep} \models (\mathbf{X} \rightarrow \mathbf{A}, t_p)$ if and only if $|\Pi(\mathbf{X}, t_p[\mathbf{X}])| = |\Pi(\mathbf{X} \cup \{\mathbf{A}\}, t_p)|$. That is, the number of equivalence classes relative to $(\mathbf{X}, t_p[\mathbf{X}])$ remains the same when adding $(\mathbf{A}, t_p[\mathbf{A}])$. Yet another way of phrasing this

is as follows. For an equivalence class $\text{eq} \in \Pi(\mathbf{X}, t_p[\mathbf{X}])$, we let $\text{Refine}(\text{eq}, (\mathbf{A}, t_p[\mathbf{A}]))$ be the set of equivalence classes in $\Pi(\mathbf{X} \cup \{\mathbf{A}\}, t_p)$ that subsume eq . Then, $D_{rep} \models \varphi$ if and only if $|\text{Refine}(\text{eq}, (\mathbf{A}, t_p[\mathbf{A}]))| = 1$ for every $\text{eq} \in \Pi(\mathbf{X}, t_p[\mathbf{X}])$. We use this formulation below.

Equivalence partitions can also be used to check support and confidence thresholds and for checking whether or not a CFD is a global explanation, as will be explained below. For this reason, we compute equivalence partitions during XPLODE’s lattice traversal. More precisely, just as in CTANE, we start by computing equivalence partitions for attribute/value pairs (\mathbf{A}, a) with $a \in \text{dom}(\mathbf{A}) \cup \{_ \}$ with support at least δ . Then, when children are generated, the equivalence partition for (\mathbf{X}, t_p) is obtained by intersecting the partitions of two parents of (\mathbf{X}, t_p) . Here, intersecting means intersecting every pair of equivalence classes from the two partitions. We implement this intersection as in TANE by means of a linear time algorithm based on lookup tables [21]. Only partitions of high support are retained. Since support is anti-monotonic, neither elements with insufficient support nor their children need to be considered.

Checking for (ε, δ) -CFDs. Consider a CFD $\varphi = (\mathbf{X} \rightarrow \mathbf{A}, t_p)$. Since the corresponding lattice element $(\mathbf{X} \cup \{\mathbf{A}\}, t_p)$ was added to the priority queue, it must have sufficient support. Due to the anti-monotonicity of support, the same is true for all its subsets, and hence $\text{supp}(\varphi, D_{rep}) \geq \delta$. Checking whether $\text{conf}(\varphi, D_{rep}) \geq 1 - \varepsilon$ can be done in terms of equivalence partitions along the same lines as the *error* of FDs is computed in TANE [21]. It is now easily verified (as in [21]) that $\text{conf}(\varphi, D_{rep})$ is equal to

$$\frac{1}{\|\Pi(\mathbf{X}, t_p[\mathbf{X}])\|} \left(\sum_{\text{eq} \in \Pi(\mathbf{X}, t_p[\mathbf{X}])} |\text{eq}| - \arg \max_{\text{eq}' \in \text{Refine}(\text{eq}, (\mathbf{A}, a))} |\text{eq}'| \right).$$

Indeed, intuitively, this expression tells that in order to resolve violating tuples with tids in eq by deleting a minimal number of tuples (as required by the definition of confidence), all tuples have to be removed that do belong to classes in $\text{Refine}(\text{eq}, (\mathbf{A}, a))$, except for one class of maximal size. We note that $\text{conf}(\varphi, D_{rep})$ is computed when element $(\mathbf{X} \cup \{\mathbf{A}\}, t_p)$ is considered. By contrast to CTANE, element $(\mathbf{X}, t_p[\mathbf{X}])$, needed to compute $\text{conf}(\varphi, D_{rep})$, may not have been visited yet due to the way the lattice is traversed. In this case, we compute $\Pi(\mathbf{X}, t_p[\mathbf{X}])$ on-the-fly from the equivalence classes $\Pi(\mathbf{B}, b)$ for $\mathbf{B} \in \mathbf{X}$ and $b = t_p[\mathbf{B}]$.

Checking for global explanations. In addition to satisfying support and confidence thresholds, for a CFD to be a global explanation, three more conditions (as stated in Definition 3) need to be verified. We next show how these checks can be done by using equivalence partitions. As before, let $\varphi = (\mathbf{X} \rightarrow \mathbf{A}, t_p)$. We start with condition (3), as it is the easiest one. Recall that this condition states that we wish to discover CFDs φ that are not violated on repaired tuples. In other words, $\text{VIO}(\varphi, \sigma_{\mathfrak{M}}(D_{rep}))$ should be empty. As already mentioned, checking for violations corresponds to finding equivalence classes $\text{eq} \in \Pi(\mathbf{X}, t_p[\mathbf{X}])$ such that $|\text{Refine}(\text{eq}, (\mathbf{A}, t_p[\mathbf{A}]))| > 1$. Since condition (3) only applies to tuples in $\sigma_{\mathfrak{M}}(D_{rep})$, it suffices to check whether there is an equivalence class $\text{eq} \in \Pi(\mathbf{X}, t_p[\mathbf{X}])$ for which there are tids in $\sigma_{\mathfrak{M}}(D_{rep})$ that belong to two different classes in $\text{Refine}(\text{eq}, (\mathbf{A}, t_p[\mathbf{A}]))$. Such a check can be easily integrated during the confidence computation of φ .

Conditions (1) and (2) in Definition 3 are a bit more challenging, as they require computations over D_{dirty} , whilst we only have equivalence partitions over D_{rep} . Suppose for the moment that we also have equivalence partitions over D_{dirty} at our disposal. We denote these by $\Pi_d(\mathbf{X}, t_p[\mathbf{X}])$ (with subscript “ d ” for dirty). Recall that condition (2) requires that $\text{VIO}(\varphi, D_{\text{dirty}}) \cap \sigma_{\mathfrak{M}}^{\text{tid}}(D_{\text{dirty}})$ is not empty. Given equivalence classes over D_{dirty} , this condition can be checked along the same lines as done for condition (3). Indeed, we compute $\text{conf}(\varphi, D_{\text{dirty}})$ and along the way we check for violations involving tids in $\sigma_{\mathfrak{M}}(D_{\text{dirty}})$, just as before. As a positive side-effect, condition (1) requires comparing $\text{conf}(\varphi, D_{\text{dirty}})$ and $\text{conf}(\varphi, D_{\text{rep}})$, both of which are now already computed.

Pulling back the equivalence partitions. It remains to explain how equivalence partitions $\Pi_d(\mathbf{Y}, s_p)$ are computed. Instead of recomputing them from scratch on D_{dirty} , we “pull them back” from $\Pi(\mathbf{Y}, s_p)$, the partition of the element in D_{rep} . A description and pseudocode of this technique is available in the appendix.

5. APPROXIMATING THE SCORE

As already observed in Section 3.3, the computation of $\text{score}(\varphi, \mathfrak{M})$, as defined in Definition 4, is quite expensive. Indeed, it requires the traversal of a power set lattice whose elements consist of *all subsets* of the modifications in \mathfrak{M} . In this section we propose an approximate scoring function, denoted by UC-score^2 , which is easy to compute. Moreover, experiments show that $\text{UC-score}(\varphi, \mathfrak{M})$ is a good approximation of $\text{score}(\varphi, \mathfrak{M})$. We show that $\text{UC-score}(\varphi, \mathfrak{M}) \leq \text{score}(\varphi, \mathfrak{M})$ and hence, $\text{UC-score}(\varphi, \mathfrak{M}) \leq \text{UB}_0(\mathbf{X}, t_p)$ (or $\text{UB}_\lambda(\mathbf{X}, t_p)$). This implies that XPLoDE can also be used for computing the global explanation with maximal UC-score.

5.1 Rationale Behind UC-score

We first observe that for constant CFDs φ , different violations are *independent* of each other (after all, constant CFDs concern violations consisting of single tuples only). Furthermore, let us call a set $M \subseteq \mathfrak{M}$ *valid* if every modification in M refers to a unique tuple. In other words, no two modifications in a valid set relate to the same tuple. Then, for constant CFDs and valid sets M , if each *single* modification $m \in M$ is locally explained by φ , then also the *entire set* M is locally explained by φ . No exhaustive enumeration of subsets of M is thus needed to check for local explainability. To obtain the best possible approximation of $\text{score}(\varphi, \mathfrak{M})$ in this way, it thus suffices to count the number of tids that occur in a modification in \mathfrak{M} that is locally explained by φ .

DEFINITION 6 (UC-score(\cdot), CONSTANT CFD). *Let φ be a constant CFD and \mathfrak{M} a set of modifications. We define*

$$\text{UC-score}(\varphi, \mathfrak{M}) := \max\{|M| \mid M \subseteq \mathfrak{M} \text{ is valid and each } m \in M \text{ is locally explained by } \varphi\}. \quad \square$$

As observed earlier, the $\text{UC-score}(\cdot)$ is equal to the size of the largest valid set M that is locally explained.

The situation for variable CFDs is quite different, however, due to *dependencies* between violations.

EXAMPLE 6. *Consider modification $m_3 = (8, \text{CC}, 44, 01)$ from the running example. A variable CFD that locally explains this modification is $\varphi = (\text{CC} \rightarrow \text{PN}, (-, -))$. Now,*

²The “UC” in UC-score refers to Unions of Constant CFDs.

assume a different modification, $m_4 = (3, \text{PN}, 222222, 111111)$. By itself, this modification is also explained by φ . When applying both modifications, however, tuples t_3, t_8 have $\text{CC} = 01$, but a different PN, violating the CFD φ . \diamond

To define an efficient, yet useful scoring function for variable CFDs, we will treat a variable CFD $\varphi = (\mathbf{X} \rightarrow \mathbf{A}, (t_p, -))$ as a *union* of a finite number of *constant* CFDs, say $\Sigma = \{\varphi_1, \dots, \varphi_m\}$ ³. Moreover, when we allow unions of constant CFDs to serve as the constraint language for global explanations, they inherit the nice properties of single constant CFDs, with some restrictions.

DEFINITION 7 (UC-score(\cdot), UNION OF CONSTANT CFDs). *Let φ be a CFD, $\Sigma_\varphi = \{\varphi_1, \dots, \varphi_m\}$ a union of constant CFDs, and \mathfrak{M} a set of modifications. We define*

$$\text{UC-score}(\Sigma_\varphi, \mathfrak{M}) := \max\{|M| \mid M \subseteq \mathfrak{M} \text{ is } \Sigma\text{-valid and each } m \in M \text{ is locally explained by } \varphi\}. \quad \square$$

It again holds that $\text{UC-score}(\Sigma, \mathfrak{M})$ is the size of the largest Σ -valid set of modifications that is locally explained by Σ , and can be computed as the number of tids that occur in a modification in \mathfrak{M} that is locally explained by Σ . This property is crucial for the efficient computation of $\text{UC-score}(\Sigma, \mathfrak{M})$. We explain the notion of Σ -valid set in the next section.

Consider now a variable CFD $\varphi = (\mathbf{X} \rightarrow \mathbf{A}, (t_p, -))$. We convert φ into a union Σ_φ of constant CFDs as follows: for each equivalence class $\text{eq} \in \Pi_d(\mathbf{X}, t_p)$, we denote by c_{eq} the projection on the \mathbf{X} -attributes of a tuple in eq (more precisely, a tuple with a tid in eq). Furthermore, we let a_{eq} be the most frequent \mathbf{A} -value in all tuples in eq . We then define $\varphi_{\text{eq}} = (\mathbf{X} \rightarrow \mathbf{A}, (c_{\text{eq}}, a_{\text{eq}}))$ and represent φ as the union $\Sigma_\varphi = \{\varphi_{\text{eq}} \mid \text{eq} \in \Pi_d(\mathbf{X}, t_p)\}$ of constant CFDs. Intuitively, the most frequent \mathbf{A} -value in each equivalence class is expected to reflect the correct value in that equivalence class. More importantly, recall that the confidence of φ (see Section 4.4) is computed by “removing tuples that do belong to classes in $\text{Refine}(\text{eq}, (\mathbf{A}, a))$, except for those tuples in the class of maximal size”. Thus, the most frequent \mathbf{A} -value directly relates to the confidence of the CFD. In conclusion, given a variable CFD φ , we define

DEFINITION 8 (UC-score(\cdot) OF A CFD). *Let φ be a CFD, $\Sigma = \{\varphi_1, \dots, \varphi_m\}$ a union of constant CFDs, and \mathfrak{M} a set of modifications. We define*

$$\text{UC-score}(\varphi, \mathfrak{M}) := \text{UC-score}(\Sigma_\varphi, \mathfrak{M}). \quad \square$$

In the remainder of this section we describe the crucial property underlying the definition of UC-score, show that $\text{UC-score}(\varphi, \mathfrak{M})$ is indeed smaller than or equal to $\text{score}(\varphi, \mathfrak{M})$, and verify that $\text{UC-score}(\varphi, \mathfrak{M})$ is easy to compute. The proofs of these properties can be found in the appendix.

5.2 Properties of UC-score

Consider a variable CFD $\varphi = (\mathbf{X} \rightarrow \mathbf{A}, (t_p, -))$ and let $\Sigma_\varphi = \{\varphi_{\text{eq}} \mid \text{eq} \in \Pi_d(\mathbf{X}, t_p)\}$ be the set of constant CFDs obtained from φ ⁴. We call a set $M \subseteq \mathfrak{M}$, Σ_φ -*valid* if it is valid and, in addition, for all tuples $t \in \sigma_M(D_{\text{dirty}} \oplus M)$ there either exists a constant CFD $\varphi_{\text{eq}} = (\mathbf{X} \rightarrow \mathbf{A}, (c_{\text{eq}}, a_{\text{eq}})) \in \Sigma_\varphi$ such that $t[\mathbf{X}] = c_{\text{eq}}$, or $t[\mathbf{X}] \neq t_p$. Intuitively, a set M of

³We represent a union of CFDs as a set of CFDs.

⁴To uniformly treat variable and constant CFDs, for a constant CFD φ we let Σ_φ be the singleton CFD $\{\varphi\}$.

modifications is Σ_φ -valid when in $\sigma_M(D_{\text{dirty}} \oplus M)$, either the violations of φ on D_{dirty} are repaired in accordance with the constant CFDs in Σ_φ , or these violations are repaired by invalidating the constants in the pattern tuple t_p of φ . By focusing on a set Σ_φ of constant CFDs, and Σ_φ -valid sets of modifications, we can indeed efficiently approximate **score**:

PROPOSITION 3. For any Σ_φ -valid set of modifications $M \subseteq \mathfrak{M}$, M is locally explained by Σ_φ if and only if each $m \in M$ is locally explained by Σ_φ . Furthermore, $\text{UC-score}(\Sigma_\varphi, M)$ is equal to the number of tids that occur in a modification in \mathfrak{M} that is locally explained by Σ_φ . \square

In other words, computing the score relative to M does not require an exhaustive exploration of all subsets of M , in contrast to the computation of $\text{score}(\varphi, \mathfrak{M})$ defined in Section 3.3. An important property is that:

PROPOSITION 4. For every global explanation φ and set \mathfrak{M} of modifications, $\text{UC-score}(\varphi, \mathfrak{M}) \leq \text{score}(\varphi, \mathfrak{M})$.

PROOF. (sketch) We show that if Σ_φ locally explains a Σ_φ -valid set M , then φ also locally explains M . Thus, $\text{UC-score}(\Sigma_\varphi, M) \leq \text{score}(\varphi, \mathfrak{M})$ for every Σ_φ -valid M . \square

Consequently, $\text{UC-score}(\varphi, \mathfrak{M}) \leq \text{score}(\varphi, \mathfrak{M}) \leq \text{UB}_0(X, t_p)$ (and $\text{UB}_\lambda(X, t_p)$) when φ is of the form $(X \setminus A \rightarrow A, t_p)$; hence XPLoDE finds the global explanation of highest UC-score.

5.3 Computation of UC-score

We conclude by explaining how $\text{UC-score}(\varphi, \mathfrak{M})$ can be efficiently computed. Proposition 3 tells that it suffices to count the number of tids that occur in a modification in \mathfrak{M} that is locally explained by Σ_φ . This is checked as follows:

PROPOSITION 5. Let Σ_φ be the set of constant CFDs corresponding to CFD $\varphi = (X \rightarrow A, t_p)$. Let $m = (\text{tid}, B, v_d, v_c) \in M$, with $M \subseteq \mathfrak{M}$ a Σ -valid set of modifications, $s = D_{\text{dirty}}[\text{tid}]$ and $t = (D_{\text{dirty}} \oplus m)[\text{tid}]$. Then Σ_φ locally explains m if and only if there exists a constant CFD $\varphi_{\text{eq}} = (X \rightarrow A, (c_{\text{eq}}, a_{\text{eq}})) \in \Sigma_\varphi$ such that $s[X] = c_{\text{eq}}$ and $s[A] \neq a_{\text{eq}}$ (s violates φ_{eq}), and:

1. either $t[A] = a_{\text{eq}}$ (t satisfies φ_{eq}); or
2. there exists another $\varphi_{\text{eq}'}$ $\in \Sigma_\varphi$ such that $t[X] = c_{\text{eq}'}$ and $t[A] = a_{\text{eq}'}$ (t satisfies some other CFD in Σ_φ); or
3. $t[X] \neq t_p$ (φ no longer applies to t). \square

Of course, for constant CFDs φ , $\Sigma_\varphi = \{\varphi\}$ and hence only cases (1) and (3) in the Proposition apply. Although Definition 3 also requires checking whether $\text{conf}(\Sigma_\varphi, D_{\text{dirty}}) < \text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus m)$, as part of the proof of Proposition 5 we show that this is implied by the conditions in its statement.

The pseudo-code for computing **UC-score**, shown in Algorithm 2, is based on Propositions 3 and 5. We first convert the CFD φ into its set Σ_φ , using function **CONVERTCFD**, as previously explained. Then, $\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})$ is computed. By Proposition 5 it suffices to only consider $m \in \mathfrak{M}$ that relate to tids in $\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})$. We partition modifications in \mathfrak{M} according to their tid. Let $\mathfrak{M}[\text{tid}]$ be the set of modifications in \mathfrak{M} that relate to tid. Since every modification occurs in exactly one attribute, we can further partition $\mathfrak{M}[\text{tid}]$ into $\mathfrak{M}[\text{tid}, X]$ and $\mathfrak{M}[\text{tid}, A]$, consisting of modifications on attributes in X and A , respectively. For modifications $m = (\text{tid}, B, v_d, v_c)$ in $\mathfrak{M}[\text{tid}, X]$, we increment the score on two occasions: (i) on line 7, if $t_p[B]$ is a constant for the attribute B in which a change happens, then clearly the

Algorithm 2 Computing the UC-score of a CFD φ .

```

1: procedure UC-SCORE( $D_{\text{dirty}}, D_{\text{rep}}, \mathfrak{M}, \varphi : (X \rightarrow A, t_p)$ )
2:    $\Sigma_\varphi \leftarrow \text{CONVERTCFD}(\varphi)$ 
3:    $\text{ucscore} \leftarrow 0$ 
4:   for all  $\text{tid} \in \text{VIO}(\Sigma_\varphi, D_{\text{dirty}})$  do
5:      $\mathfrak{M}[\text{tid}] \leftarrow \{m \in \mathfrak{M} \mid m \text{ relates to } \text{tid}\}$ 
6:     for all  $m = (\text{tid}, B, v_d, v_c) \in \mathfrak{M}[\text{tid}, X]$  do
7:       if  $t_p[B] \neq \cdot$  then
8:         Increment ucscore and go to next tid
9:       else if  $\exists \varphi_{\text{eq}'}, t = D_{\text{rep}}[\text{tid}], t[X] = c_{\text{eq}'}$  then
10:        if  $t[A] = a_{\text{eq}'}$  then
11:          Increment ucscore and go to next tid
12:        for all  $m \in \mathfrak{M}[\text{tid}, A]$  do
13:          if  $\exists \varphi_{\text{eq}}, s = D_{\text{dirty}}[\text{tid}], s[X] = c_{\text{eq}}$  then
14:            if  $t[A] = a_{\text{eq}}$  for  $t = D_{\text{rep}}[\text{tid}]$  then
15:              Increment ucscore and go to next tid
16:   return ucscore

```

change makes φ inapplicable to the tuple in question (condition 3 in Proposition 5); and (ii) on lines 9 and 10, if the change results in the tuple satisfying some constant CFD $\varphi_{\text{eq}'}$ $\in \Sigma_\varphi$ (condition 2 in Proposition 5). Finally, on lines 13–14, we perform a similar computation for modifications in $\mathfrak{M}[\text{tid}, A]$: We determine the constant CFD φ_{eq} that was violated in the tuple s in D_{dirty} , and increment the score if attribute A was modified such that tuple t satisfies φ_{eq} (condition 1 in Proposition 5). As soon as a modification is explained for a given tid t , it is counted, and the algorithm proceeds to the next tid.

Table 2: Statistics of the used datasets.

Dataset	#Tuples	#Attributes	%MinSupp
Abalone	8354	9	10%
Adult	97684	11	1%
Soccer	200000	10	10%
SP500	245148	7	1%

6. EXPERIMENTS

We experimentally validate our repair explanation method. Experiments illustrating the benefit of the optimizations discussed in Section 4, are postponed to the appendix.

6.1 Experimental Setup

Hardware. All our experiments were performed on an Intel Core i7 Processor (2.3GHZ) with 16GB of memory running OS X. All algorithms are implemented in C++ and run entirely in main memory. The code, datasets and CFDs used, are available for research purposes⁵.

Datasets. We use four datasets: **Abalone** and **Adult** are small datasets from the UCI Repository [26]; **Soccer** is a synthetic dataset about soccer players and their teams⁶, and finally, **SP500** is a real-world dataset about stock trading from [9]. To ensure that CFD violations can occur, we duplicate every tuple in these datasets. On the Adult dataset we only use constants CFDs, since mining general CFDs on this dataset is too time-consuming, because Adult has a higher number of attributes and many frequent constant patterns. Statistics of the data are shown in Table 2.

⁵<https://bit.ly/2MYzjcH>

⁶<http://www.db.unibas.it/projects/bart/>

Error Generation. We make use of the BART tool [1] for introducing violations in the datasets. BART takes a dataset and a set of data quality rules as input, and inserts a predefined percentage of violations into the data. The used percentages are reported in the **%Error** column in Table 4. To get the required quality rules, we used our implementation of CTANE [15] to discover 100% confident CFDs on the datasets, using the minimum support percentages shown in Table 2. As a note aside, support threshold δ is then simply $\frac{1}{100}(\#\text{Tuples} \times \% \text{MinSupp})$. These thresholds were set empirically, low enough to ensure that a reasonable number of 100% confident CFDs were found on the datasets, i.e., at least 50; yet high enough for the runtime to remain practical for experimenting. Combining the clean and dirty datasets, we generate *partial* repairs by starting from the dirty dataset, and replacing a subset of the dirty tuples with their clean variants. For each dataset, we obtain 3 different dirty datasets by using 3 different CFDs discovered on the clean data, denoted by CFD 1, CFD 2, and CFD 3. When considering a dirty dataset corresponding to a CFD i , then CFD i is the *target* CFD. In other words, it is the CFD we want to discover by repairing the corresponding dirty dataset. Obviously, CFD i is different for each dataset.

Falcon. We compare with FALCON [20], a system which discovers a SQL Update Statement (equivalent to a constant CFD) that explains a *single* modification. Since such a modification does not contain sufficient information to reliably discover the underlying CFD, FALCON employs *user interaction* to narrow down the search space. Feedback is received in the form of a user asserting the (in)validity of a given CFD. Using the fact that all generalizations of an invalid CFD are also invalid, and all specializations of a valid CFD are also valid, an efficient binary search algorithm limits the amount of user feedback required.

6.2 Experimental Results

Usefulness of Explaining from Repairs. Before evaluating the performance of our method for discovering CFDs that explain repairs, we first show that user-supplied modifications are indeed a good instrument to guide CFD discovery towards a CFD that is useful for repairing. We perform CFD discovery on the dirty data, and rank the discovered approximate CFDs by confidence and rule length, and compare this to a ranking using $\text{UC-score}(\cdot)$ on partially cleaned data, with 2 and 5 modifications, respectively.

In Table 3, we show the position of the target CFD in a list of CFDs ranked according to the different criteria (position 1 denotes the top of the ranking). The results clearly show that, when using confidence or rule size, the target is typically quite deep in the ranking. It is infeasible for a user to manually validate all of these CFDs one by one. This illustrates the need for integrating user feedback into the ranking function. Indeed, using information from modifications consistently quickly brings the target to the front, ratifying our approach. Repairing a handful of tuples automatically invalidates many CFDs, saving much user effort.

Explaining Full Repairs. To illustrate the efficacy of our method, we first perform explanation discovery on *full* repairs, i.e., using the fully clean data as D_{rep} . We are thus only interested in CFDs that are not violated on D_{rep} , and set $\varepsilon = 0$. The minimum support threshold is set according to the percentages from Table 2. These thresholds en-

sure that the target CFD is among the (ε, δ) -CFDs that are the candidates for global explanations. In this experiment, we first want to get an idea about the total number of global explanations. Therefore, we use the post-processing method: Run CTANE to find (ε, δ) -CFDs; then filter out all the global explanations. We find that the number of $(0, \delta)$ -CFDs that are global explanations is typically too large for manual inspection, ranging from around 40 on the Adult dataset up to 400 on the Soccer dataset. On partial repairs (when $\varepsilon > 0$), this number will only increase since the number of (ε, δ) -CFDs will increase considerably, up to 1500 on Abalone for $\varepsilon = 0.1$. Secondly, we want to validate that the target CFD is indeed the highest scoring (for $\text{UC-score}(\cdot)$) global explanation when considering full repairs. Indeed, on all datasets the target CFD was discovered. It shows that the $\text{UC-score}(\cdot)$ is a good measure for repair explanations.

Scoring Function. We have also evaluated whether the scoring function $\text{UC-score}(\cdot)$ is indeed a suitable surrogate for $\text{score}(\cdot)$, the exact scoring function. Since $\text{score}(\cdot)$ requires an expensive computation, which we implemented using a brute-force approach, we performed the experiment only on the smallest dataset, Abalone. We computed both the ranking difference and the absolute difference between both scoring functions. Due to space limitations, we postpone the results to the appendix. We learn that the top positions in the ranking are unaffected by the choice of scoring function. Moreover, the absolute error remains very small throughout the entire ranking. Hence, $\text{UC-score}(\cdot)$ offers a good approximation for $\text{score}(\cdot)$, and is unlikely to change which global explanation is returned by XPLODE, compared to $\text{score}(\cdot)$. We use $\text{UC-score}(\cdot)$ in all remaining experiments.

Explaining Partial Repairs. Our first experiment showed that the target CFD, the one used to dirty the data, can be recovered from a full repair. We now want to answer the question, “how many modifications are needed for a partial repair to recover the target CFD?”. We created partial repairs using an increasing number of modifications, and report the number and percentage of modifications needed until XPLODE returns the target CFD. On each dataset, we use the error percentage of the data as ε in the confidence threshold $1 - \varepsilon$; the support threshold δ is derived from the difference of (i) the minimum support percentage used to mine the CFDs and (ii) the error percentage of the data. For instance, on Abalone with 1% errors, we pick $\delta = (10\% - 1\%) \times 8354$ and $\varepsilon = 0.01$. As before, these guarantee that the target CFD is among the (ε, δ) -CFDs that are explored. Furthermore, we use $\text{UB}_\lambda(\cdot)$ as upper bound function in XPLODE. Experiments (not reported) show that this loose anti-monotonic function guides XPLODE more quickly towards the desired CFD than $\text{UB}_0(\cdot)$, since the λ penalty assigns a lower priority to lattice elements containing irrelevant attributes, i.e., attributes that do not occur in AttVIO .

The results for this experiment on all four datasets and CFD 1, CFD 2, CFD 3, are shown in Table 4. Here, the columns **%Error** and **#Error** contain, respectively, the percentage and absolute number of violations inserted by BART. The columns **%M(i)** and **#M(i)** display the percentage and number of modifications required before the target CFD is returned, for CFD i . We see that the percentage of modifications required to find the target CFD is typically low, and does not change much when the dirtiness of the data increases. This implies that our method has a greater benefit when cleaning dirtier data: if the target CFD can be

Table 3: Position of target CFD among all approximate CFDs according to various ranking criteria.

Dataset	CFD	Length (Asc)	Conf (Asc)	Conf (Desc)	UC-score ₂ (·)	UC-score ₅ (·)
Abalone	1	907	249	4617	18	2
	2	1825	1500	3703	12	3
	3	2492	305	4565	244	3
Adult	1	153006	143303	62202	256	10
	2	31948	141376	64517	15	1
	3	10064	191665	27548	2	1
Soccer	1	3896	806	11953	49	3
	2	1505	1424	11046	3	3
	3	1232	1385	12329	14	9
SP500	1	150	126	258	4	1
	2	171	127	248	42	2
	3	166	127	255	46	2

Table 4: Number and Percentage of Modifications required to retrieve the target CFD, for 3 different CFDs.

Dataset	%Error	#Error	%M(1)	#M(1)	%M(2)	#M(2)	%M(3)	#M(3)
Abalone	0.1	8	100%	8	50%	4	100%	8
	1	83	10%	8	5%	4	14%	12
	10	835	1%	8	1%	8	2%	18
Adult	0.1	97	20%	19	6%	6	30%	29
	0.5	488	5%	24	2%	12	0.6%	3
	1	976	2%	24	0.4%	4	5%	48
Soccer	0.1	200	9%	17	11%	22	1%	2
	1	2000	0.3%	7	1%	22	0.1%	2
	10	20000	0.2%	30	0.1%	20	0.1%	25
SP500	0.1	245	3%	7	3%	7	0.8%	2
	0.5	1225	0.5%	7	0.5%	7	0.1%	2
	1	2451	0.1%	3	0.1%	3	0.1%	3

found by manually cleaning 1% of the violations, then 99% of the violations may be cleaned automatically. Moreover, the number of attributes has a higher impact on the number of required modifications than the number of tuples.

Comparison with Falcon. We next compare XPLODE and FALCON [20], for the case where the target CFD is a constant CFD⁷. The experiment was done on the Soccer dataset, using 3 constant CFDs. Only modifications that relate to the consequent of the CFDs were considered, as FALCON supports only these. FALCON was able to recover each of the target CFDs using a single modification and 2 user questions, taking between 1 and 4 seconds. In comparison, XPLODE recovers the target CFD using 3 to 5 modifications, with a runtime around 4 seconds each time. Since the runtimes were obtained on different hardware, it is hard to compare efficiency. Nevertheless, the experiment suggests that XPLODE, although not specialized only towards constant CFDs, is comparable to FALCON.

We remark that this changes in the case of variable CFDs. Indeed, in this case FALCON finds a constant CFD for every constant pattern relating to the variable CFD. For each such CFD, a single modification and some user questions would be required. We consider again the Soccer dataset, this time using 3 variable CFDs. XPLODE can recover the single target CFD using, on average, 12 modifications. However, in order to capture all the errors in these datasets using constant CFDs, on average 55 constant CFDs are needed. This implies that FALCON would require this amount of user modifications, and outputs a large number of CFDs as well. A similar scenario occurs when considering modifications on attributes in the antecedent of CFDs, since FALCON assumes that modifications pertain only to the *consequent* of a rule.

⁷Due to IP restrictions, we were unable to obtain the code and perform an in-depth comparison. One of the authors of FALCON kindly performed an experiment instead.

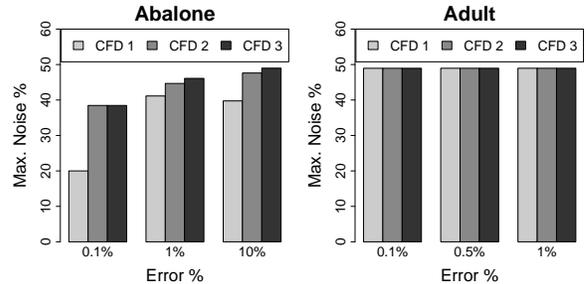


Figure 1: Noise-robustness of XPLODE.

Robustness to Noise. So far, we have run experiments using only modifications that belong to a single CFD, and as such, the returned CFD explains all the given modifications. However, when a user is manually cleaning data, without knowing the target CFD, corrections may be made in positions unrelated to the target CFD. We now verify whether XPLODE is robust to such “noise”.

We consider again a full repair, and add random modifications throughout the data, not connected to the target CFD. In Figure 1, we report how many random modifications can be added without distorting the output of XPLODE, up to 50% of the total number of modifications, i.e., as many random modifications as correct modifications. The results for SP500 and Soccer, available in the appendix, are similar to Abalone.

The results show our method is very robust to random noise, especially as more modifications are considered. On the Adult dataset, where we only consider constant CFDs, noise seems to have no impact. This makes sense: as variable CFDs can capture a larger variety of errors, they are also more likely to accidentally explain a random modification, whereas constant CFDs are unable to connect the random modifications to each other.

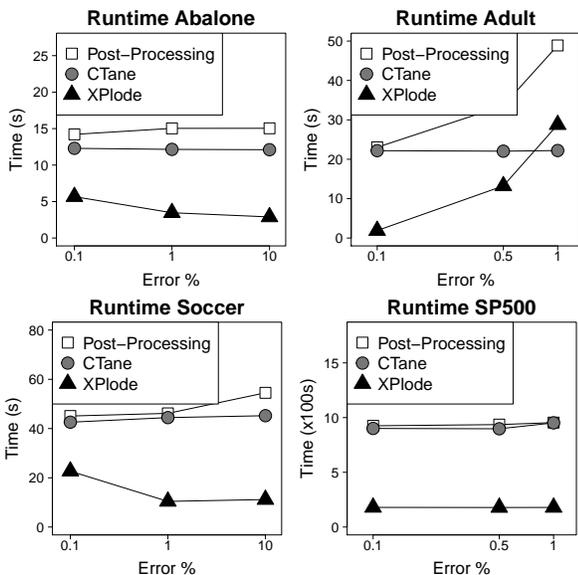


Figure 2: Runtime performance of XPLode, compared to post-processing and CTane.

Runtime Performance. Finally, we evaluate the runtime performance of XPLode on full repairs. We compare XPLode with two benchmarks: firstly, discovering all global explanations and then computing all the scores, and secondly, the discovery of *all* (ϵ, δ) -CFDs using the CTANE algorithm. Figure 2 shows the runtimes (in seconds) averaged over 3 independent runs with different CFDs. Algorithm XPLode clearly outperforms post-processing in every case, and is typically faster than a full CTANE execution. Overall, results indicate that the error percentage has little impact on runtimes. The only exception is the Adult dataset: here, the runtime deteriorates with a higher percentage of errors. We suspect that, for constant CFDs, checking for global explanations accounts for more of the total runtime. With more modifications, computing equivalence partitions in D_{dirty} becomes harder, as the difference between D_{dirty} and D_{rep} increases. Moreover, Adult shows a greater increase in the number of global explanations as dirtiness increases.

7. RELATED WORK

Our work is situated in the area of constraint-based data quality (see [13, 22] for surveys) and is based on CFDs [14].

Underlying XPLode is a constraint discovery process. By contrast to traditional (C)FD discovery algorithms [7, 15, 21, 30, 28] (see also [30, 27, 13, 22] for overviews) that aim to find *all* constraints that approximately hold on the data, we aim to find those CFDs that *explain a partial repair*. Traditional methods do not consider our notions of explanation and scoring function. To find explanations, these algorithms have to be combined with a post-processing step, which is expensive both in time and user interaction. In XPLode, post-processing is avoided by carefully integrating the notion of explanations in the discovery process. The resulting on-demand algorithm is similar in spirit to the method presented in [19], where constants patterns for a *fixed* FD are found that best describe the data. Our search strategy allows for general *loose anti-monotonic* functions [5], such as our scoring function, instead of the *strictly anti-monotonic* functions of support and confidence used in [19]. In addition, we explore the *entire* space of CFDs instead of requiring the

embedded FD to be fixed. We leverage equivalence partitions [11] as commonly used in constraint discovery [21, 15]. None of the traditional methods consider user interaction, which is *essential* for our method. Other discovery methods that leverage user interaction are [20, 34]: The FALCON system [20], described in Section 6.1, is most closely related. It finds *constant* CFDs based on a *single* modification and afterwards relies on a “user oracle” to (in)validate the proposed rules. The emphasis is on limiting calls to this oracle. By contrast, we *only* use information stemming from *multiple* modifications and have no need for a “black box” oracle. Our method is capable of handling *variable* CFDs, providing more flexibility and more compact explanations of modifications, since a variable CFD can represent a large number of constant CFDs. In addition, our method is on-demand. The UGUIDE system [34] aims to find a set of FDs such that their *violation set* overlaps maximally with the tuples holding *true* errors, whilst minimizing the number of violations that are not true errors. We aim to explain *repairs* rather than errors. UGUIDE asks users to either (in)validate FDs, or specify whether given cells or tuples are true errors or not. It is a best effort method given a budget on the number of user interactions. We treat both types of user interaction *uniformly* through modifications: by incorporating information on *how* errors should be repaired, we can *zoom in* directly to CFDs that are *useful for repairing* and not only for error detection. By contrast, UGUIDE initially uses *all* approximate FDs, which is costly and leads to users invalidating an excessive amount of spurious dependencies. Finally, we consider general CFDs rather than just FDs. Other forms of user interaction in data cleaning are considered in works such as [36, 35, 38]. In these works, the set of (valid) constraints (FDs, CFDs, or other) is assumed to be available already. User involvement is used to determine the right repairs relative to the constraints. Our work leverages user interaction in an earlier phase, to *find valid constraints for repairing*. Finally, we rely on users to make some initial repairs, after which existing data repairing algorithms [4, 12, 16, 17, 23, 25] can be used. Our approach is complementary to any of these repairing algorithms. In [6], a *description of errors* by means of a small set of conjunctive queries is discovered, without considering user interaction or repairing. Our repair explanations could also be of use to update constraints over evolving data [8, 3, 29].

8. CONCLUSION

We considered the problem of finding a CFD that best explains a partial repair. Our on-demand algorithm, XPLode, shows great promise for finding such explanations. Underlying the efficacy of the method is a scoring function and its efficient approximation, that counts the number of modifications explained. A search space traversal based on loose anti-monotonic bounds on the scores, leads to an improvement in efficiency, compared to a post-processing approach to the problem. Experiments show high precision in discovering the correct explanation, using only a small amount of modifications. Future work includes the investigation of alternative scoring functions, and extending XPLode towards discovering multiple explanations. Furthermore, the general concept of inferring constraints from a repair can be investigated in the context of other constraint formalisms [33, 37, 18, 10]. Suggesting modifications based on explanations also seems a natural extension of our framework.

9. REFERENCES

- [1] P. C. Arocena, B. Glavic, G. Mecca, R. J. Miller, P. Papotti, and D. Santoro. Messing up with BART: Error generation for evaluating data-cleaning algorithms. *PVLDB*, 9(2):36–47, 2015.
- [2] R. Bertens, J. Vreeken, and A. Siebes. Efficiently discovering unexpected pattern-co-occurrences. In *SDM*, pages 126–134, 2017.
- [3] G. Beskales, I. F. Ilyas, L. Golab, and A. Galiullin. On the relative trust between inconsistent data and inaccurate constraints. In *ICDE*, pages 541–552, 2013.
- [4] P. Bohannon, W. Fan, M. Flaster, and R. Rastogi. A cost-based model and effective heuristic for repairing constraints by value modification. In *SIGMOD*, pages 143–154, 2005.
- [5] F. Bonchi and C. Lucchese. Pushing tougher constraints in frequent pattern mining. In *PAKDD*, volume 5, pages 114–124, 2005.
- [6] A. Chalamalla, I. F. Ilyas, M. Ouzzani, and P. Papotti. Descriptive and prescriptive data cleaning. In *SIGMOD*, pages 445–456, 2014.
- [7] F. Chiang and R. J. Miller. Discovering data quality rules. *PVLDB*, 1(1):1166–1177, 2008.
- [8] F. Chiang and R. J. Miller. A unified model for data and constraint repair. In *ICDE*, pages 446–457, 2011.
- [9] X. Chu, I. F. Ilyas, and P. Papotti. Discovering denial constraints. *PVLDB*, 6(13):1498–1509, 2013.
- [10] X. Chu, I. F. Ilyas, P. Papotti, and Y. Ye. Ruleminer: Data quality rules discovery. In *ICDE*, pages 1222–1225, 2014.
- [11] S. S. Cosmadakis, P. C. Kanellakis, and N. Spyros. Partition semantics for relations. *Journal of Computer and System Sciences*, 33(2):203 – 233, 1986.
- [12] M. Dallachiesa, A. Ebaid, A. Eldawy, A. Elmagarmid, I. F. Ilyas, M. Ouzzani, and N. Tang. NADEEF: A commodity data cleaning system. In *SIGMOD*, pages 541–552, 2013.
- [13] W. Fan and F. Geerts. *Synthesis Lectures on Data Management: Foundations of data quality management*. Morgan & Claypool, 2012.
- [14] W. Fan, F. Geerts, X. Jia, and A. Kementsietsidis. Conditional functional dependencies for capturing data inconsistencies. *ACM TODS*, 33(2):6, 2008.
- [15] W. Fan, F. Geerts, J. Li, and M. Xiong. Discovering conditional functional dependencies. *IEEE TKDE*, 23(5):683–698, 2011.
- [16] H. Galhardas, D. Florescu, D. Shasha, E. Simon, and C. Saita. *Declarative data cleaning: Language, model, and algorithms*. PhD thesis, INRIA, 2001.
- [17] F. Geerts, G. Mecca, P. Papotti, and D. Santoro. That’s all folks!: Llunatic goes open source. *PVLDB*, 7(13):1565–1568, 2014.
- [18] L. Golab, H. Karloff, F. Korn, B. Saha, and D. Srivastava. Discovering conservation rules. *IEEE TKDE*, 26(6):1332–1348, 2014.
- [19] L. Golab, H. Karloff, F. Korn, D. Srivastava, and B. Yu. On generating near-optimal tableaux for conditional functional dependencies. *PVLDB*, 1(1):376–390, 2008.
- [20] J. He, E. Veltri, D. Santoro, G. Li, G. Mecca, P. Papotti, and N. Tang. Interactive and deterministic data cleaning: A tossed stone raises a thousand ripples. In *SIGMOD*, pages 893–907, 2016.
- [21] Y. Huhtala, J. Kärkkäinen, P. Porkka, and H. Toivonen. TANE: An efficient algorithm for discovering functional and approximate dependencies. *The computer journal*, 42(2):100–111, 1999.
- [22] I. F. Ilyas and X. Chu. *Trends in Cleaning Relational Data: Consistency and Deduplication*. Now Publishers Inc., 2015.
- [23] Z. Khayyat, I. F. Ilyas, A. Jindal, S. Madden, M. Ouzzani, P. Papotti, J.-A. Quiané-Ruiz, N. Tang, and S. Yin. Bigdancing: A system for big data cleansing. In *SIGMOD*, pages 1215–1230, 2015.
- [24] J. Kivinen and H. Mannila. Approximate inference of functional dependencies from relations. *Theor. Comput. Sci.*, 149(1):129–149, 1995.
- [25] S. Kolahi and L. V. Lakshmanan. On approximating optimum repairs for functional dependency violations. In *ICDT*, pages 53–62, 2009.
- [26] M. Lichman. UCI machine learning repository, 2013.
- [27] J. Liu, J. Li, C. Liu, and Y. Chen. Discover dependencies from data - A review. *IEEE Trans. Knowl. Data Eng.*, 24(2):251–264, 2012.
- [28] P. Mandros, M. Boley, and J. Vreeken. Discovering reliable approximate functional dependencies. In *KDD*, pages 355–363, 2017.
- [29] M. Mazuran, E. Quintarelli, L. Tanca, and S. Ugolini. Semi-automatic support for evolving functional dependencies. In *EDBT*, pages 293–304, 2016.
- [30] T. Papenbrock, J. Ehrlich, J. Marten, T. Neubert, J.-P. Rudolph, M. Schönberg, J. Zwiener, and F. Naumann. Functional dependency discovery: An experimental evaluation of seven algorithms. *PVLDB*, 8(10):1082–1093, 2015.
- [31] C. Pit-Claudel, Z. Mariet, R. Harding, and S. Madden. *Outlier Detection in Heterogeneous Datasets using Automatic Tuple Expansion*. Technical Report MIT-CSAIL-TR-2016-002, CSAIL, MIT, 32 Vassar Street, Cambridge MA 02139, February 2016.
- [32] J. Rammelaere, F. Geerts, and B. Goethals. Cleaning data with forbidden itemsets. In *ICDE*, pages 897–908, 2017.
- [33] S. Song and L. Chen. Discovering matching dependencies. In *CIKM*, pages 1421–1424, 2009.
- [34] S. Thirumuruganathan, L. Berti-Equille, M. Ouzzani, J.-A. Quiane-Ruiz, and N. Tang. UGuide: User-guided discovery of FD-detectable errors. In *SIGMOD*, pages 1385–1397, 2017.
- [35] M. Volkovs, F. Chiang, J. Szlichta, and R. J. Miller. Continuous data cleaning. In *ICDE*, pages 244–255, 2014.
- [36] J. Wang, T. Kraska, M. J. Franklin, and J. Feng. Crowder: Crowdsourcing entity resolution. *PVLDB*, 5(11):1483–1494, 2012.
- [37] J. Wang and N. Tang. Towards dependable data repairing with fixing rules. In *SIGMOD*, pages 457–468, 2014.
- [38] M. Yakout, A. K. Elmagarmid, J. Neville, M. Ouzzani, and I. F. Ilyas. Guided data repair. *PVLDB*, 4(5):279–289, 2011.

APPENDIX

A.1 Proof that X_{PLOBE} returns the best explanation

PROPOSITION 1. On input $D_{\text{dirty}}, D_{\text{rep}}, \mathfrak{M}, \varepsilon, \delta, \text{score}(\cdot)$, and $\text{UB}(\cdot)$, the algorithm X_{PLOBE} returns the global explanation with maximal score, if it exists, provided that $\text{UB}(\cdot)$ is loose anti-monotonic and for any element (X, t_p) , $\text{UB}(X, t_p)$ is larger than the score of any of its candidate CFDs.

PROOF. Let $\varphi_{\text{max}} \neq \text{nil}$ be the CFD returned by X_{PLOBE} , and let max be the score of φ_{max} . Assume, for the sake of contradiction, that there exists a global explanation $\varphi = (X \rightarrow A, t_p)$ with $\text{max} < \text{score}(\varphi, \mathfrak{M})$. Since $\text{UB}(X, t_p) \geq \text{score}(\varphi, \mathfrak{M}) > \text{max}$, this implies that $(X \cup \{A\}, t_p)$ was never added to the priority queue Φ .

Note, however, that since $\text{UB}(\cdot)$ is loose anti-monotonic, there exists a path in the lattice from $(X \cup \{A\}, t_p)$ to an attribute/value pair (B, b) , at the first level of the lattice, along which the UB -values never decrease. Clearly, (B, b) is added to the priority queue when (\emptyset, \emptyset) was expanded. This in turn implies that all elements on the path from (B, b) to (X, t_p) , all having a UB -value greater than max , must have been added to the queue before the algorithm can terminate. This contradicts our earlier observation. We may thus indeed conclude correctness. \square

A.2 Proof that $\text{UB}_0(\cdot)$ is a valid upper bound

PROPOSITION 2. The upper bound function $\text{UB}_0(\cdot)$ satisfies the conditions of Proposition 1.

PROOF. Consider an element (X, t_p) in the lattice. Let $\varphi = (X \setminus \{A\}, t_p)$ be the CFD in $\text{CandCFD}(X, t_p)$ such that $\text{UB}_0(X, t_p) = |\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})|$. Let $Y = X \setminus \{B\}$ for some $B \in X \setminus \{A\}$. We observe that for CFD $\psi = (Y \setminus \{A\} \rightarrow A, t_p[Y])$, we have that $\text{VIO}(\varphi, D_{\text{dirty}}) \subseteq \text{VIO}(\psi, D_{\text{dirty}})$ and hence also $|\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| \leq |\text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})|$. Furthermore, CFD ψ is in $\text{CandCFD}(Y, t_p[Y])$. As a consequence, $\text{UB}_0(Y, t_p[Y]) \geq |\text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})|$ which is larger or equal than $\text{UB}_0(X, t_p) = |\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})|$. It suffices to observe that $(Y, t_p[Y])$ is a parent of (X, t_p) . \square

A.3 Proof that $\text{UB}_\lambda(\cdot)$ is loose anti-monotonic

We use a similar argument as in the proof of Proposition 2, showing that $\text{UB}_0(\cdot)$ is loose anti-monotonic. That is, consider element (X, t_p) and a candidate CFD $\varphi = (Y \rightarrow A, t_p)$ for $Y = X \setminus \{A\}$ such that $\text{UB}_\lambda(X, t_p)$ is equal to

$$|\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| + 1 - \lambda |X \cup \text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})|.$$

Let $B \in Y$ and CFD $\psi = (Y \setminus B \rightarrow B, t_p[Y])$. We have $\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M}) \subseteq \text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})$. We distinguish between the following cases:

(a) $\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M}) = \text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})$, in which case also $\text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M}) = \text{AttVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})$. Since $Y = X \setminus A$, it is trivially verified that

$$|X \cup \text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| \geq |(X \setminus A) \cup \text{AttVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})|,$$

and hence we have that

$$\begin{aligned} & |\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| + 1 - \lambda |X \cup \text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| \\ &= |\text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})| + 1 - \lambda |X \cup \text{AttVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})| \\ &\leq |\text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})| + 1 - \lambda |Y \cup \text{AttVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})|. \end{aligned}$$

(b) $\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M}) \subset \text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})$ and thus $|\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| + 1 \leq |\text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})|$. Furthermore, $\text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})$ can be at most \mathcal{A} . It remains to verify that

$$-\lambda |X \cup \text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| \leq 1 - \lambda |\mathcal{A}|,$$

where $|X \cup \text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| \geq 0$. We obtain that $1 - \lambda |\mathcal{A}| \geq 0$, which is equivalent to our assumption that $\lambda |\mathcal{A}| \leq 1$. Thus, we may infer that also in this case

$$\begin{aligned} & |\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| + 1 - \lambda |X \cup \text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| \\ &\leq |\text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})| + 1 - \lambda |Y \cup \text{AttVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})|. \end{aligned}$$

Since ψ is a candidate CFD of $(Y, t_p[Y])$, which is a child of (X, t_p) , we have that $\text{UB}_\lambda(Y, t_p[Y]) \geq |\text{ModVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})| + 1 - \lambda |Y \cup \text{AttVIO}(\psi, D_{\text{dirty}}, \mathfrak{M})|$, which in turn is at least $|\text{ModVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})| + 1 - \lambda |X \cup \text{AttVIO}(\varphi, D_{\text{dirty}}, \mathfrak{M})|$, which was chosen to coincide with $\text{UB}_\lambda(X, t_p)$. Hence, $\text{UB}_\lambda(\cdot)$ is indeed loose anti-monotonic. \square

A.4 Proof of Proposition 3

We start by showing that Σ_φ locally explains a Σ_φ -valid set of modifications M if and only if Σ_φ locally explains each $\mathfrak{m} \in M$. Clearly, if Σ_φ locally explains M , then, by definition, Σ_φ is locally explain \mathfrak{m} for every $\mathfrak{m} \in M$. Suppose next that Σ_φ locally explains each $\mathfrak{m} \in M$. We need to show that Σ_φ is an M' -repair explanation for every subset M' of M . We start by showing some useful properties. First, by the definition of $\sigma_{M'}^{\text{tid}}(\cdot)$, we have

$$\sigma_{M'}^{\text{tid}}(D_{\text{dirty}}) = \bigcup_{\mathfrak{m} \in M'} \sigma_{\mathfrak{m}}^{\text{tid}}(D_{\text{dirty}}), \quad (1)$$

Similarly, by the definition of $\sigma_{M'}(\cdot)$ and the validity of M' , we have that

$$\sigma_{M'}(D_{\text{dirty}} \oplus M') = \bigcup_{\mathfrak{m} \in M'} \sigma_{\mathfrak{m}}(D_{\text{dirty}} \oplus \mathfrak{m}). \quad (2)$$

Furthermore, for a constant CFD $\varphi_{\text{eq}} \in \Sigma_\varphi$, $\text{VIO}(\varphi_{\text{eq}}, D_{\text{dirty}} \oplus M')$ consists of all tids corresponding to tuples t in $D_{\text{dirty}} \oplus M'$ such that $t \not\models \varphi_{\text{eq}}$. Since M' contains at most one modification per tuple, this implies that either t is unaffected by the modifications in M' , or t is a tuple in $D_{\text{dirty}} \oplus \mathfrak{m}$ for a (unique) modification $\mathfrak{m} \in M'$. Hence, in both cases, $t \in \text{VIO}(\varphi_{\text{eq}}, D_{\text{dirty}} \oplus \mathfrak{m})$ for some $\mathfrak{m} \in M'$. In other words,

$$\text{VIO}(\varphi_{\text{eq}}, D_{\text{dirty}} \oplus M') \subseteq \bigcup_{\mathfrak{m} \in M'} \text{VIO}(\varphi_{\text{eq}}, D_{\text{dirty}} \oplus \mathfrak{m}).$$

From this, we can infer that

$$\begin{aligned} \text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus M') &= \bigcup_{\varphi_{\text{eq}} \in \Sigma_\varphi} \text{VIO}(\varphi_{\text{eq}}, D_{\text{dirty}} \oplus M') \\ &\subseteq \bigcup_{\varphi_{\text{eq}} \in \Sigma_\varphi, \mathfrak{m} \in M'} \text{VIO}(\varphi_{\text{eq}}, D_{\text{dirty}} \oplus \mathfrak{m}) \\ &= \bigcup_{\mathfrak{m} \in M'} \text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathfrak{m}), \quad (3) \end{aligned}$$

where the first equality follows from the definition of $\text{VIO}(\cdot, \cdot)$, the second from our previous observation, and the third equality follows again from the definition of $\text{VIO}(\cdot, \cdot)$. We

also note that

$$|\text{supp}(\Sigma_\varphi, D_{\text{dirty}} \oplus M')| = |\text{supp}(\Sigma_\varphi, D_{\text{dirty}})| + \sum_{\mathbf{m} \in M'} (\text{supp}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) - \text{supp}(\Sigma_\varphi, D_{\text{dirty}})), \quad (4)$$

where each $\text{supp}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) - \text{supp}(\Sigma_\varphi, D_{\text{dirty}})$ is simply -1 if the modification reduces the support of Σ_φ , and 0 otherwise. In a similar fashion, we have that:

$$|\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus M')| = |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})| + \sum_{\mathbf{m} \in M'} (|\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m})| - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})|). \quad (5)$$

Both Equations (4) and (5) follow again from the fact that Σ_φ consists of constant CFDs and M' contains at most one modification per tuple.

Given Equations (1), (2), (3), (4) and (5), we next verify the three conditions of Def. 3 for M' .

▷ **Condition (1).** We need to verify that $\text{conf}(\Sigma_\varphi, D_{\text{dirty}}) < \text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus M')$. Let σ and ν denote $\text{supp}(\Sigma_\varphi, D_{\text{dirty}})$ and $|\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})|$, respectively; and let $\sigma_{\mathbf{m}}$ and $\nu_{\mathbf{m}}$ denote $\text{supp}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m})$ and $|\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m})|$, respectively. From Equations 4 and 5, we can infer that the confidence $\text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus M')$ is equal to

$$\frac{\sigma - \nu + \sum_{\mathbf{m} \in M'} ((\sigma_{\mathbf{m}} - \nu_{\mathbf{m}}) - (\sigma + \nu))}{\sigma + \sum_{\mathbf{m} \in M'} (\sigma_{\mathbf{m}} - \sigma)}.$$

By assumption, $\text{conf}(\Sigma_\varphi, D_{\text{dirty}}) < \text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m})$ for each $\mathbf{m} \in M'$. Assume that $\mathbf{m} = (\text{tid}, \mathbf{B}, v_d, v_c)$ and let $s = D_{\text{dirty}}[\text{tid}]$ and $t = (D_{\text{dirty}} \oplus \mathbf{m})[\text{tid}]$.

Note that it is not possible that $s \models \Sigma_\varphi$ because by assumption, $\text{VIO}(\Sigma_\varphi, D_{\text{dirty}}) \cap \sigma_{\mathbf{m}}(D_{\text{dirty}})$ is non-empty and the only tuple corresponding to $\sigma_{\mathbf{m}}(D_{\text{dirty}})$ is s . We thus have that $s \not\models \Sigma_\varphi$, which implies that there exists a $\varphi_{\text{eq}} \in \Sigma_\varphi$ such that $s[X] = c_{\text{eq}}$. Due to the fact that the CFDs in Σ_φ have mutually exclusive patterns c_{eq} , φ_{eq} is unique. We next distinguish two cases, depending on \mathbf{m} :

(a) $t[X] = c_{\text{eq}}$. In this case, for the confidence of Σ_φ to increase after modification \mathbf{m} is applied, the modification must be $\mathbf{m} = (\text{tid}, \mathbf{A}, s[\mathbf{A}], a_{\text{eq}})$. This modification does not affect the support of φ , brings down the number of violations of Σ_φ by one, and also does not affect supports of any other CFD in Σ_φ (due to mutual exclusive c_{eq} patterns). In other words, in this case

$$\begin{aligned} \sigma_{\mathbf{m}} &= \sigma \\ \nu_{\mathbf{m}} &= \nu - 1. \end{aligned}$$

(b) $t[X] \neq c_{\text{eq}}$. In this case, $t[X]$ and $s[X]$ can only differ in a single attribute, say $\mathbf{B} \in \mathbf{X}$. Hence, \mathbf{m} necessarily must be of the form $(\text{tid}, \mathbf{B}, c_{\text{eq}}[\mathbf{B}], t[\mathbf{B}])$. Furthermore, by assuming Σ_φ -validity the set M' of modifications, we again have two distinct options:

(b1) $t[X] = c_{\text{eq}'}$ for some $\text{eq}' \neq \text{eq}$. Since Σ_φ is not violated, i.e., $\text{VIO}(\Sigma_\varphi, \sigma_{\mathbf{m}}(D_{\text{dirty}} \oplus \mathbf{m}))$ is empty, it must hold that $t[\mathbf{A}] = a_{\text{eq}'}$. In other words, tuple t no longer violated Σ_φ , and hence the number of violations is reduced by one. Moreover, the support of c_{eq} is reduced by one while the support of $c_{\text{eq}'}$ increases by one, meaning that the support of Σ_φ

remains identical. We thus have that:

$$\begin{aligned} \sigma_{\mathbf{m}} &= \sigma \\ \nu_{\mathbf{m}} &= \nu - 1. \end{aligned}$$

(b2) There exist no $c_{\text{eq}'}$ in Σ_φ such that $t[X] = c_{\text{eq}'}$. By definition of Σ_φ -validity, this is only allowed if $t[X] \neq t_p$, where t_p is the pattern of the CFD φ . In other words, the CFD no longer applies to the tuple. Clearly, this reduces both the support of Σ_φ and the number of violations by one:

$$\begin{aligned} \sigma_{\mathbf{m}} &= \sigma - 1 \\ \nu_{\mathbf{m}} &= \nu - 1. \end{aligned}$$

In each of the cases (a), (b1) and (b2), the number of violations decreases with 1. Hence,

$$\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus M') = \text{VIO}(\Sigma_\varphi, D_{\text{dirty}}) - |M'|.$$

Let $\Delta(\sigma_{M'})$ denote the reduction in support of Σ_φ , i.e., the number of modifications belonging to case (b2), with $\Delta(\sigma_{M'}) \leq |M'|$. We can then write $\text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus M')$ as:

$$\frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})| + |M'| - \Delta(\sigma_{M'})}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - \Delta(\sigma_{M'})} \quad (6)$$

For $\Delta(\sigma_{M'}) = 0$, expression 6 is equal to

$$\frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})| + |M'|}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}})},$$

and for any $\Delta(\sigma_{M'}) > 0$, expression 6 is greater than or equal to

$$\frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})|}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - \Delta(\sigma_{M'})},$$

both of which are obviously strictly greater than

$$\text{conf}(\Sigma_\varphi, D_{\text{dirty}}) = \frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})|}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}})}$$

for any M' consisting of at least one element. It follows that condition (1) is satisfied.

▷ **Condition (2).** We need to verify that

$$\text{VIO}(\Sigma_\varphi, D_{\text{dirty}}) \cap \sigma_{M'}^{\text{tid}}(D_{\text{dirty}})$$

is non-empty. From Equation 1, this is equivalent to requiring

$$\bigcup_{\mathbf{m} \in M'} \text{VIO}(\Sigma_\varphi, D_{\text{dirty}}) \cap \sigma_{\mathbf{m}}^{\text{tid}}(D_{\text{dirty}})$$

to be non-empty. This follows from the fact that Σ_φ satisfied condition (2) for every $\mathbf{m} \in M'$.

▷ **Condition (3).** We need to verify that

$$\text{VIO}(\Sigma_\varphi, \sigma_{M'}(D_{\text{dirty}} \oplus M'))$$

is empty. Note that for constant CFDs this is equivalent to checking whether

$$\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus M') \cap \sigma_{M'}(D_{\text{dirty}} \oplus M')$$

is empty. Using Equations 1 and 3, we know that this set is included in

$$\bigcup_{\mathbf{m} \in M'} (\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) \cap \sigma_{\mathbf{m}}(D_{\text{dirty}} \oplus \mathbf{m})),$$

which is known to be empty, due the fact that Σ_φ satisfies condition (3) for each $\mathbf{m} \in M'$. It follows that

$$\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus M') \cap \sigma_{M'}^{\text{tid}}(D_{\text{dirty}} \oplus M')$$

is empty, and hence condition (3) is satisfied.

A.5 Proof of Proposition 4

We prove this by showing for any set of modifications $M \subseteq \mathfrak{M}$, if Σ_φ is an M -repair explanation then φ is also an M -repair explanation.

▷ **Condition 2:** $\text{VIO}(\varphi, D_{\text{dirty}}) \cap \sigma_M^{\text{tid}}(D_{\text{dirty}})$ is not empty: This follows from the fact that any violation of Σ_φ , is also a violation of φ . Indeed, let $\varphi_{\text{eq}} = (X \rightarrow A, (c_{\text{eq}}, a_{\text{eq}}))$ be a constant CFD violated in a tuple $t \in D_{\text{dirty}}$. Then $t[X] = c_{\text{eq}}$ and $t[A] \neq a_{\text{eq}}$. By construction of Σ_φ , it must hold that there exists a tuple $t' \in D_{\text{dirty}}$ such that $t'[X] = c_{\text{eq}}$ and $t'[A] = a_{\text{eq}}$. Clearly, t and t' together violate the CFD φ . It follows that $\text{VIO}(\Sigma_\varphi, D_{\text{dirty}}) \cap \sigma_M^{\text{tid}}(D_{\text{dirty}})$ is included in $\text{VIO}(\varphi, D_{\text{dirty}}) \cap \sigma_M^{\text{tid}}(D_{\text{dirty}})$. Since the former is non-empty, so is the latter.

▷ **Condition 3:** $\text{VIO}(\varphi, \sigma_M(D_{\text{dirty}} \oplus M))$ is empty We show this by contradiction. Assume there exist two tuples $t, t' \in \sigma_M(D_{\text{dirty}} \oplus M)$ that violate $\varphi : (X \rightarrow A, (t_p, -))$, i.e., $t[X] = t'[X]$ and $t[A] \neq t'[A]$. By construction Σ_φ , there must exist some $\varphi_{\text{eq}} : (X \rightarrow A, (c_{\text{eq}}, a_{\text{eq}}))$ such that $c_{\text{eq}} = t[X] = t'[X]$. Since φ_{eq} is not violated in $\sigma_M(D_{\text{dirty}} \oplus M)$, it must therefore also hold that $t[A] = t'[A] = a_{\text{eq}}$, contradicting the assumption that $t[A] \neq t'[A]$, and hence no such t and t' violating φ can exist in $\sigma_M(D_{\text{dirty}} \oplus M)$. It follows that $\text{VIO}(\Sigma_\varphi, \sigma_M(D_{\text{dirty}} \oplus M))$ is empty implies that $\text{VIO}(\varphi, \sigma_M(D_{\text{dirty}} \oplus M))$ is empty.

▷ **Condition 1:** $\text{conf}(\varphi, D_{\text{dirty}} \oplus M) > \text{conf}(\varphi, D_{\text{dirty}})$ In the proof of Proposition 3 (expression 6), it was shown that the confidence Σ_φ in $D_{\text{dirty}} \oplus M$ is equal to:

$$\frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})| + |M| - \Delta(\sigma_M)}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - \Delta(\sigma_M)}$$

Where $\Delta(\sigma_M)$ denotes the reduction in the support of Σ_φ as a result of the modifications M , with $\Delta(\sigma_{M'}) \leq |M|$.

Recall that the confidence of a CFD is defined based on a set D'_{dirty} , containing the minimum amount of tuples that need to be altered or removed from D_{dirty} for φ to be satisfied. Since the union of constants is constructed based on the most frequent value in each equivalence class, $|D'_{\text{dirty}}|$ is equal to $|\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})|$. Indeed, such a minimal set D'_{dirty} contains exactly those tuples that violate the constant CFDs in Σ_φ . We thus have that:

$$\text{conf}(\varphi, D_{\text{dirty}}) = \frac{\text{supp}(\varphi, D_{\text{dirty}}) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})|}{\text{supp}(\varphi, D_{\text{dirty}})}$$

Moreover, since every modification \mathbf{m} that is locally explained by Σ_φ must occur in a tuple violated by Σ_φ , it follows that every such modification pertains to the set D'_{dirty} . Clearly, when the tuple is modified to satisfy some $c_{\text{eq}} \in \Sigma_\varphi$, it now belongs to the most frequent value in its equivalence class, and is no longer a part of D'_{dirty} . It follows that $|(D_{\text{dirty}} \oplus M)'| = |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})| - |M|$.

The support of the CFD Σ_φ is reduced in only one case, namely when the modification is such that t_p no longer applies, i.e., $t[X] \neq t_p$, where t is the modified tuple. It is clear that such a modification also reduces the support of φ .

Hence, $\text{supp}(\varphi, D_{\text{dirty}} \oplus M) = \text{supp}(\varphi, D_{\text{dirty}}) - \Delta(\sigma_M)$.

We thus have that:

$$\begin{aligned} \text{conf}(\varphi, D_{\text{dirty}} \oplus M) &= \\ \frac{\text{supp}(\varphi, D_{\text{dirty}}) - \Delta(\sigma_M) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})| + |M|}{\text{supp}(\varphi, D_{\text{dirty}}) - \Delta(\sigma_M)}. \end{aligned}$$

Analogous to the proof of Proposition 3, it follows that

$$\text{conf}(\varphi, D_{\text{dirty}} \oplus M) > \text{conf}(\varphi, D_{\text{dirty}}). \quad \square$$

A.6 Proof of Proposition 5

Let $\mathbf{m} = (\text{tid}, B, v_d, v_c) \in \mathfrak{M}$, let $s = D_{\text{dirty}}[\text{tid}]$ and $t = (D_{\text{dirty}} \oplus \mathbf{m})[\text{tid}]$. Let Σ_φ be the union of constant CFDs constructed from $\varphi = (X \rightarrow A, (t_p, -))$. We only deal with the variable CFD case here, the constant CFD case is similar (and in fact easier).

⇒ We first show that the conditions in Proposition 5 are sufficient for Σ_φ to locally explain \mathbf{m} . To this aim we verify that conditions (1), (2) and (3) in Definition 3 are satisfied.

If there exists a constant CFD $\varphi_{\text{eq}} = (X \rightarrow A, (c_{\text{eq}}, a_{\text{eq}})) \in \Sigma_\varphi$ such that $s[X] = c_{\text{eq}}$ and $s[A] \neq a_{\text{eq}}$, then clearly Σ_φ is violated on s , and $\text{VIO}(\Sigma_\varphi, D_{\text{dirty}}) \cap \sigma_{\mathbf{m}}^{\text{tid}}(D_{\text{dirty}})$ is not empty, satisfying condition (2).

We now verify that conditions (1) and (3) are satisfied for each of the 3 cases in which Σ_φ is defined to explain \mathbf{m} :

Case 1. Assume $t[A] = a_{\text{eq}}$, i.e., \mathbf{m} changes s into the tuple t such that $t \models \varphi_{\text{eq}}$. Since φ_{eq} is the only CFD in Σ_φ that applies to tuple t , by construction of Σ_φ , it follows that Σ_φ is satisfied in the modified tuple, and hence $\text{VIO}(\varphi, \sigma_{\mathbf{m}}(D_{\text{dirty}} \oplus \mathbf{m}))$ is empty, satisfying condition (3) in Definition 3. Moreover, a violation has been resolved, while the support of Σ_φ remains unchanged. Indeed, s supported c_{eq} in D_{dirty} and t remains to do so in $D_{\text{dirty}} \oplus \mathbf{m}$. We thus have that

$$\begin{aligned} \text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) &= \\ &= \frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m})|}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m})} \\ &= \frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - (|\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})| - 1)}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}})}. \end{aligned}$$

Hence, also condition (1) in Definition 3, namely

$$\text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) > \text{conf}(\Sigma_\varphi, D_{\text{dirty}})$$

is satisfied.

Case 2. Assume that there exists another $\varphi_{\text{eq}'} \in \Sigma_\varphi$ such that $t[X] = c_{\text{eq}'}$ and $t[A] = a_{\text{eq}'}$, i.e., t satisfies some other CFD in Σ_φ . Again, $\varphi_{\text{eq}'}$ is by definition the only CFD in Σ_φ that applies to tuple t , implying that Σ_φ is satisfied in the modified tuple. Hence, $\text{VIO}(\varphi, \sigma_{\mathbf{m}}(D_{\text{dirty}} \oplus \mathbf{m}))$ is empty, satisfying condition (3). We have that

$$\text{supp}(\varphi_{\text{eq}'}, D_{\text{dirty}} \oplus \mathbf{m}) = \text{supp}(\varphi_{\text{eq}'}, D_{\text{dirty}}) - 1,$$

and

$$\text{supp}(\varphi_{\text{eq}'}, D_{\text{dirty}} \oplus \mathbf{m}) = \text{supp}(\varphi_{\text{eq}'}, D_{\text{dirty}}) + 1,$$

meaning that the support of Σ_φ again remains unchanged. Since one violation has been resolved, the analysis of confidence is identical to the previous case, and hence condition (1) in Definition 3 is also satisfied.

Case 3. Finally, assume that $t[X] \neq t_p$, i.e., φ no longer applies to t . It trivially follows that $\text{VIO}(\varphi, \sigma_{\mathbf{m}}(D_{\text{dirty}} \oplus \mathbf{m}))$ is empty, satisfying condition (3). The modification

has resolved a violation, while the support of Σ_φ has been reduced by one. We thus have that:

$$\begin{aligned} & \text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) \\ &= \frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) - |\text{VIO}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m})|}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m})} \\ &= \frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) + 1 - (|\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})| - 1)}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) + 1} \\ &= \frac{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) - (|\text{VIO}(\Sigma_\varphi, D_{\text{dirty}})|) + 2}{\text{supp}(\Sigma_\varphi, D_{\text{dirty}}) + 1}. \end{aligned}$$

And hence, it holds that

$$\text{conf}(\Sigma_\varphi, D_{\text{dirty}} \oplus \mathbf{m}) > \text{conf}(\Sigma_\varphi, D_{\text{dirty}}).$$

From the above, we conclude that if the conditions in Proposition 5 are satisfied, then the three conditions in Definition 3 are satisfied, and thus Σ_φ is a \mathbf{m} -repair explanation, or equivalently for single modifications, Σ_φ locally explains the modification \mathbf{m} .

□ We next verify that the conditions in Proposition 5 are also necessary for Σ_φ to explain a modification \mathbf{m} .

If there does not exist a constant CFD $\varphi_{\text{eq}} = (\mathbf{X} \rightarrow \mathbf{A}, (c_{\text{eq}}, a_{\text{eq}})) \in \Sigma_\varphi$ such that $s[\mathbf{X}] = c_{\text{eq}}$ and $s[\mathbf{A}] \neq a_{\text{eq}}$, then by construction of Σ_φ , the tuple s does not violate any CFD in Σ_φ . It follows that $\text{VIO}(\Sigma_\varphi, D_{\text{dirty}}) \cap \sigma_{\mathbf{m}}^{\text{tid}}(D_{\text{dirty}})$ is empty, and hence condition (2) is not satisfied. Therefore, Σ_φ cannot locally explain \mathbf{m} .

Alternatively, assume that there does exist a constant CFD $\varphi_{\text{eq}} = (\mathbf{X} \rightarrow \mathbf{A}, (c_{\text{eq}}, a_{\text{eq}})) \in \Sigma_\varphi$ such that $s[\mathbf{X}] = c_{\text{eq}}$ and $s[\mathbf{A}] \neq a_{\text{eq}}$. We show that, if none of the three cases in Proposition 5 hold in which Σ_φ is defined to explain \mathbf{m} , then condition (3) in Definition 3 cannot be satisfied. We distinguish between two cases, depending on the attribute \mathbf{B} which was modified:

- Assume $\mathbf{B} = \mathbf{A}$, i.e., the modification relates to the consequent of the CFD φ . In this case, only case (1) is relevant. If this does not hold, then $t[\mathbf{A}] \neq a_{\text{eq}}$. However, no changes were made to the other values in t , meaning that φ_{eq} still applies to t , and is thus still violated. It follows that Σ_φ is violated after the modification, and $\text{VIO}(\varphi, \sigma_{\mathbf{m}}(D_{\text{dirty}} \oplus \mathbf{m}))$ is not empty. Hence, Σ_φ does not locally explain \mathbf{m} .
- Assume $\mathbf{B} \in \mathbf{X}$, i.e., the modification relates to the antecedent of the CFD φ . Both scenario's (2) and (3) are relevant. If neither scenario applies, then $t[\mathbf{X}] \asymp t_p$, and thus the CFD φ still applies. It remains to show that the modified tuple t necessarily violates Σ_φ . Note that, since \mathbf{m} is assumed to be Σ -valid, there must exist exactly one CFD $\varphi_{\text{eq}'} = (\mathbf{X} \rightarrow \mathbf{A}, (c_{\text{eq}'}, a_{\text{eq}'})) \in \Sigma_\varphi$ such that $t[\mathbf{X}] = c_{\text{eq}'}$. However, since scenario (2) does not hold, $t[\mathbf{A}] \neq a_{\text{eq}'}$. Clearly, t violates $\varphi_{\text{eq}'}$, and hence $\text{VIO}(\varphi, \sigma_{\mathbf{m}}(D_{\text{dirty}} \oplus \mathbf{m}))$ is again not empty, violating condition (3).

From the above, it follows that if the conditions in Proposition 5 do not apply, then Σ_φ does not locally explain \mathbf{m} . This proves the proposition. □

A.7 Pullback Procedure

We show the pseudocode of the PULLBACK algorithm, discussed in Section 4.4, in Algorithm 3. The PULLBACK procedure goes as follows: For an equivalence class $\text{eq} \in \Pi(\mathbf{Y}, s_p)$, let Δeq be the set of tids in eq that underwent changes in attributes in \mathbf{Y} , consistent with the pattern tuple s_p . That

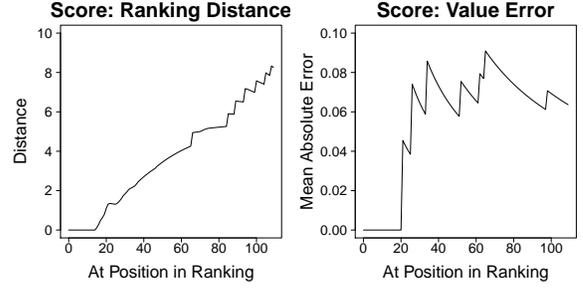


Figure 3: Comparison of $\text{UC-score}(\cdot)$ and $\text{score}(\cdot)$ on the Abalone dataset, using 20 modifications.

is, Δeq contains all tids in eq such that there is a modification $\mathbf{m} = (\text{tid}, \mathbf{B}, v_d, v_c)$ in \mathfrak{M} such that $\mathbf{B} \in \mathbf{Y}$, $v_c \asymp s_p[\mathbf{B}]$ and $\text{tid} \in \text{eq}$. We initialize $\Pi_d(\mathbf{Y}, s_p)$ by adding an equivalence class $\text{eq}_d = \text{eq} \setminus \Delta\text{eq}$, for each $\text{eq} \in \Pi(\mathbf{Y}, s_p)$. Since each tid in eq_d corresponds to a tuple in D_{dirty} that has not been changed in positions relevant to (\mathbf{Y}, s_p) , we have that for any two tids , say tid and tid' , in eq_d , $s[\mathbf{Y}] = s'[\mathbf{Y}] \asymp s_p$ for $s = D_{\text{dirty}}[\text{tid}]$ and $s' = D_{\text{dirty}}[\text{tid}']$. In other words, eq_d is a proper equivalence class for D_{dirty} .

Those tuples $t \in D_{\text{dirty}}$ corresponding to a $\text{tid} \in \Delta\text{eq}$, for some $\text{eq} \in \Pi(\mathbf{Y}, s_p)$, still need to be assigned to their correct equivalence class in $\Pi_d(\mathbf{Y}, s_p)$, or even to a new equivalence class, in case no correct class can be found. More precisely, for each $\text{eq} \in \Pi(\mathbf{Y}, s_p)$ and each $\text{tid} \in \Delta\text{eq}$, we look up $t[\mathbf{Y}]$ for $t = D_{\text{dirty}}[\text{tid}]$, and add tid to the equivalence class eq_d corresponding to tuples $s \in D_{\text{dirty}}$ with $s[\mathbf{Y}] = t[\mathbf{Y}]$, if such eq_d exists. Otherwise, we create a new equivalence class for tuples in D_{dirty} agreeing with t on its \mathbf{Y} -attributes.

Algorithm 3 Pulling back an equivalence class $\Pi(\mathbf{Y}, s_p)$ in D_{rep} to its counterpart $\Pi_d(\mathbf{Y}, s_p)$ in D_{dirty}

```

1: procedure PULLBACK( $\Pi(\mathbf{Y}, s_p), D_{\text{dirty}}, \mathfrak{M}$ )
2:    $\Pi_d(\mathbf{Y}, s_p) \leftarrow \emptyset$ 
3:   for all  $\text{eq} \in \Pi(\mathbf{Y}, s_p)$  do
4:     Compute  $\Delta\text{eq}$ , the set of modified  $\text{tids}$  in  $\text{eq}$ 
5:      $\text{eq}_d \leftarrow \text{eq} \setminus \Delta\text{eq}$ 
6:      $\Pi_d(\mathbf{Y}, s_p) \leftarrow \Pi_d(\mathbf{Y}, s_p) \cup \{\text{eq}_d\}$ 
7:    $\Delta \leftarrow \{\Delta\text{eq} \mid \text{eq} \in \Pi(\mathbf{Y}, s_p)\}$ 
8:   for all  $\text{tid} \in \Delta$  do
9:      $\text{eq}_d \leftarrow \text{FINDEQ}(\text{tid}, \Pi_d(\mathbf{Y}, s_p), D_{\text{dirty}})$ 
10:    if  $\text{eq}_d \neq \emptyset$  then
11:       $\text{eq}_d \leftarrow \text{eq}_d \cup \{\text{tid}\}$ 
12:    else
13:      Create a new  $\text{eq}_d$ , initially empty
14:       $\text{eq}_d \leftarrow \{\text{tid}\}$ 
15:       $\Pi(\mathbf{Y}, s_p) \leftarrow \Pi(\mathbf{Y}, s_p) \cup \{\text{eq}_d\}$ 
16:   return  $\Pi_d(\mathbf{Y}, s_p)$ 

```

A.8 Experiments on Scoring Function

We now further evaluate whether the scoring function $\text{UC-score}(\cdot)$ is indeed a suitable surrogate for $\text{score}(\cdot)$, the exact scoring function. Since $\text{score}(\cdot)$ requires an expensive computation, which we implemented using a brute-force approach, we perform the experiment only on the smallest dataset, Abalone. We mined all global explanations on the data with a 1% error percentage, using a partial repair containing 20 modifications. Both $\text{UC-score}(\cdot)$ and $\text{score}(\cdot)$ report values ranging from 1 to 20 for the global explanations. Again, since

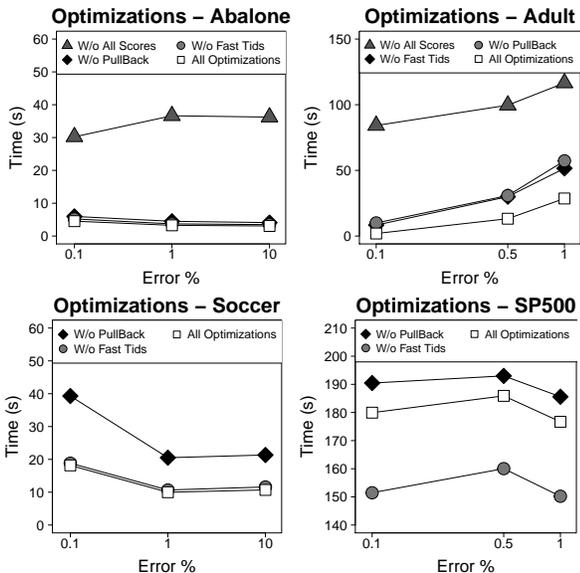


Figure 4: Impact of optimizations on runtime.

we need all global explanations, the naive post-processing method is used. We rank all global explanations, using both scoring functions. At each point in the ranked list, we compute the Mean Spearman’s Footrule Distance and the Mean Absolute Error of the scores, for all explanations up to this point in the ranking. The results are shown in Figure 3. We learn that the top positions in the ranking are unaffected by the choice of scoring function. Moreover, the absolute error remains very small throughout the entire ranking. The sea-saw pattern is a result of averaging: when a difference is encountered, the error rises, and then smoothens out as additional positions, without error, are considered. Hence, $UC\text{-score}(\cdot)$ offers a good approximation for $score(\cdot)$, and is unlikely to change which global explanation is returned by XPLoDE, compared to using $score(\cdot)$. We therefore use $UC\text{-score}(\cdot)$ in all remaining experiments.

A.9 Experiments on Impact of Optimizations

We evaluate the impact of three optimizations discussed throughout Section 4. We compare the runtimes with each of these optimizations turned off in turn, to the runtime of XPLoDE with all optimizations turned on. The runtimes were again obtained using a full repair, as in the runtime experiment. The results are shown in Figure 4. The worst runtime is obtained when disabling “All scores”, which is the computation of scores for CFDs that are not global explanations (discussed under Remarks at the end of Section 4.1). Computing these scores leads to better informed tiebreaking among elements with equal upper bounds, at the cost of having to compute the score. Clearly, computing all scores leads to a major improvement in runtime. For this reason, we do not test without this optimization on the larger datasets, Soccer and SP500. The other two optimizations concern the method of obtaining equivalence partitions by XPLoDE. Recall that, when computing the confidence of a CFD, the equivalence partition of a subset may not be available. We evaluate as “Fast Tids” the computation of these equivalence partitions on-the-fly from the partitions of singletons, as opposed to computing them from the data. This optimization has a positive impact on runtime, on all datasets except SP500. Finally, we evaluate the use of the

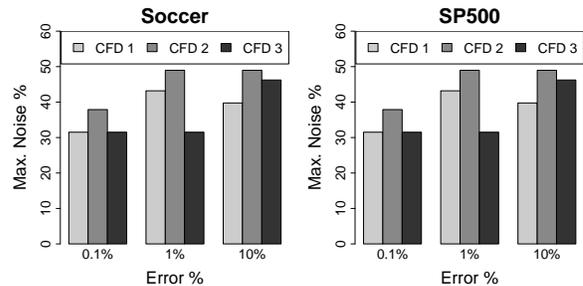


Figure 5: Noise-robustness of XPLoDE.

PULLBACK method, to convert an equivalence partition on D_{rep} to its counterpart on D_{dirty} . PULLBACK offers the least benefit of our optimizations, but still leads to a clear and consistent improvement. The percentage of errors does not seem to change the relative influence of the optimizations.

A.10 Noise-robustness of XPLoDE

The plots for the noise-robustness of XPLoDE on Soccer and SP500 are shown in Figure 5.

A.11 Additional Details On Experiments

We show the CFDs used in the experiments in Table 5.

Dataset	CFDs
Abalone	(viscera, whole, rings=10) → height
	(whucked, whole, height, diameter) → rings
	(shell, whole, length, rings=9) → height
Adult	(maritalstatus=Never-married, race=White, sex=Male, relationship=Own-child, age=18-21) → income=LessThan50K
	(maritalstatus=Never-married, age=18-21, education=HS-grad) → income=LessThan50K
	(education=Some-college, relationship=Husband) → sex=Male
Soccer	(season, position, surname, stadium=King Power Stadium) → team
	(city, birthplace, season=2013) → position
	(position, surname, city=Solna) → season
SP500	(g, e, d) → f
	(g, f, b) → d
	(g, f, c) → d

Table 5: CFDs used in the experiments.