

Cleaning data with Swipe

Toon Boeckling, Antoon Bronselaer

March 2024

Abstract

The repair problem for functional dependencies is the problem where an input database needs to be modified such that all functional dependencies are satisfied and the difference with the original database is minimal. The output database is then called an *optimal repair*. If the allowed modifications are value updates, finding an optimal repair is NP-hard. A well-known approach to find *approximations* of optimal repairs builds a Chase tree in which each internal node resolves violations of one functional dependency and leaf nodes represent repairs. A key property of this approach is that controlling the branching factor of the Chase tree allows to control the trade-off between repair quality and computational efficiency. In this paper, we explore an extreme variant of this idea in which the Chase tree has only one path. To construct this path, we first create a partition of attributes such that classes can be repaired sequentially. We repair each class only once and do so by fixing the order in which dependencies are repaired. This principle is called *priority repairing* and we provide a simple heuristic to determine priority. The techniques for attribute partitioning and priority repair are combined in the Swipe algorithm. An empirical study on four real-life data sets shows that Swipe is one to three orders of magnitude faster than multi-sequence Chase-based approaches, whereas the quality of repairs is comparable or better. Moreover, a scalability analysis of the Swipe algorithm shows that Swipe scales well in terms of an increasing number of tuples.

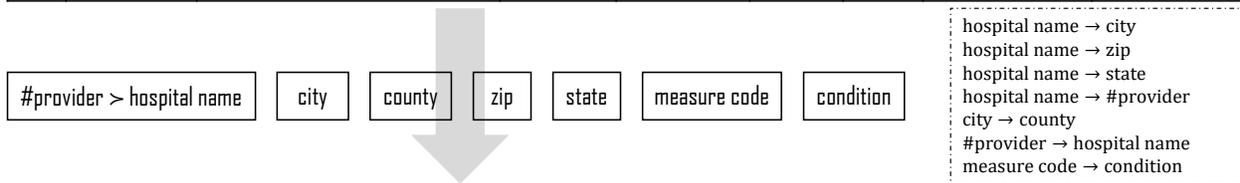
1 Introduction

We study the problem of repairing an inconsistent database through value modifications in the case where constraints are functional dependencies (FDs). More precisely, starting from a database that violates some FDs (i.e., a dirty database) we seek to generate a database that satisfies all FDs (i.e., a repair) and that originates from the dirty database by sequentially changing the value of a cell in a table. Beyond the scope of toy problems, there usually exist many possible repairs for a given dirty database and most of them are perceived as “not good” because they simply make too much changes to the dirty database. To deal with this, one of the most prevalent approaches, is to define a *cost model* that assigns a positive cost to each change. Introducing this cost model gives rise to the *repair problem for FDs*, which is the problem of finding a repair that minimizes the accumulative cost of all changes one must make to produce that repair. Such a repair is called a *minimal-cost repair*. The following example illustrates the repair problem for FDs and will be used as a running example throughout the paper.

Example 1 *Figure 1 (top) shows a snippet of the Hospital data set [13], for which the seven FDs shown must be satisfied. If we assume that any change to an attribute value has cost 1, then a minimal-cost repair can be obtained by changing the values marked in grey (which happen to be here the actual errors in the data set). This repair has cost 6 and it can be verified by the reader that no repair of cost 5 exists. It can be seen from this example that minimal-cost repairs are not unique, since several other repairs of cost 6 exist. For example, the violation of measure code \rightarrow condition by tuples 1 and 5 can also be resolved by changing the value of measure code for tuple 5.*

The Llunatic Chase algorithm Although we can easily find a minimal-cost repair in Example 1, it is known that the repair problem for FDs as defined here, cannot be solved efficiently. Deciding whether there exists a repair with a cost lower than a constant C , is NP-complete and finding constant-factor approximations of minimal-cost repairs is NP-hard [30]. To deal with this, algorithms have been proposed to search for *approximate* solutions [8, 30, 5, 32]. In the current paper, we focus on one particular approach that relies on the Chase algorithm and has been implemented in the open-source framework Llunatic [24, 25]. We note and emphasize here that Llunatic can in fact deal with constraints that are more expressive than FDs. Yet, we limit the explanation of Llunatic here to FDs only, as these constraints are the focus of the current paper. Originally, the Chase algorithm was designed to work as a proof engine for dependency implication [3, 1]. Informally, the algorithm takes as input (i) a set of tuples and (ii) a set of FDs and then builds a sequence of *Chase steps*. In each such Chase step, one FD that is currently violated on the given set of tuples, becomes satisfied by equating the right-hand side attributes of the FD. The

tid	#provider	hospital name	city	county	zip	state	measure code	condition
1	10006	eliza coffee memorial hospital	florence	laudxrdalx	35631	al	ami-4	heart attack
2	10006	eliza coffee memorial hospital	florence	laudexdale	35631	al	hf-2	heart failure
3	10006	eliza coffee memorial hospital	florence	lauderdale	3563x	al	ami-2	heart attack
4	1000x	eliza coffee memorial hospital	florence	lauderdale	35631	al	pn-3b	pneumonia
5	1003x	huntsville hospital	huntsville	madison	35801	al	ami-4	heart attackx
6	10039	huntsville hospital	huntsville	madison	35801	al	ami-5	heart attack



tid	#provider	hospital name	city	county	zip	state	measure code	condition
1	10006	eliza coffee memorial hospital	florence	lauderdale	35631	al	ami-4	heart attackx
2	10006	eliza coffee memorial hospital	florence	lauderdale	35631	al	hf-2	heart failure
3	10006	eliza coffee memorial hospital	florence	lauderdale	35631	al	ami-2	heart attack
4	10006	eliza coffee memorial hospital	florence	lauderdale	35631	al	pn-3b	pneumonia
5	1003x	huntsville hospital	huntsville	madison	35801	al	ami-4	heart attackx
6	1003x	huntsville hospital	huntsville	madison	35801	al	ami-5	heart attack

Figure 1: An example cleaning scenario with hospital data [13] (top) and seven FDs (middle right). Actual errors in the data are marked in grey. A partition of attributes is shown (middle left) over which FDs are forward repairable. The first partition class shows a priority model over its attributes. For this partition, a repair obtained by using majority voting with random tie breaking as repair function is shown (bottom). Correct changes are shown in green bold font and incorrect changes are shown in red bold font.

output of a Chase step is then used as input for the next Chase step and this process continues until all FDs are satisfied. In the original algorithm, tuples contained *labelled variables* and conflict resolution among variables with different labels is done by a consistent choice (e.g., if labels are integer-valued indices, we can always choose the variable with the lower index). In this setting, it is well-known that the Chase algorithm terminates and satisfies the *Church Rosser* property [1]. The latter property states that the output of the Chase algorithm does not depend on the order in which FDs are selected.

In order to use the Chase algorithm as a repair tool in the presence of constants, a modification of the Chase algorithm called the *Llunatic Chase* has been proposed [25]. In this algorithm, resolution of conflicts is based on a *partial order* that models preferences of different constants and variables (e.g., null values). If a conflict cannot be resolved by means of this partial order, special labelled variables (i.e., lluns) are introduced, which can later be used to query for human input. While the Llunatic Chase has been shown to terminate, it does not satisfy the Church Rosser property and the outcome of a sequence of Chase steps is thus dependent on the order in which FDs are selected for repair. To deal with this, Llunatic considers a *Chase tree* in which each path is a sequence of Chase steps and leaf nodes are repairs. From all repairs in a Chase tree, a repair can then be selected. For example, one could consider a cost model and choose the repair with minimal cost.

Because it is infeasible to generate the complete Chase tree, Llunatic uses a *cost manager* that allows the definition of a wide range of pruning strategies. Examples of these strategies include (i) limiting the number of outgoing branches of a node by means of a branching threshold, (ii) limiting the amount of leaf nodes by a potential solutions threshold and (iii) limiting modifications to forward repairs only [25]. The cost manager thus offers a trade-off between repair quality (i.e., generating more repairs implies a higher chance of a finding a high quality repair) and computational cost (i.e., generating more repairs implies larger Chase trees).

Single-path Chase trees with Swipe The key contribution of the current paper is to investigate an extreme case of cost management. More specifically, we investigate a degenerate variant of the Llunatic Chase where a Chase tree represents a single sequence of Chase steps and generates only one repair. In addition, we restrict to forward repairs only, which means that violations of an FD are always restored by equating values of the right-hand side attributes in that FD. The main problem we face in this scenario, is how to best select the next FD to repair. To solve this problem, we will first partition the attributes involved in a given set of FDs such that the partition is *forward repairable*. This means we can iterate over the partition classes sequentially and for each class C , apply forward repairing for those FDs of which the right-hand side attribute is in C . Once a partition of attributes is available, the next step is to visit each partition class and for each class, consider a *priority model* for the FDs we must repair for that class. The combination of (i) the partition of attributes and (ii) the priority models for each class gives us the ability to produce a (degenerate) Chase tree with a single-sequence of Chase steps, where each Chase step uses forward repairing. We can now summarize the key contributions of this paper as follows:

- We present the *Swipe* algorithm to repair FD violations. This algorithm is grounded on two key ideas: (i) the notion of a partition of attributes that is forward repairable and (ii) the notion of a priority model for FDs. We show that for any set of FDs there always exists a partition that is forward repairable and we present a simple algorithm to construct one that is maximally refined. In addition, we provide a simple but effective heuristic to create a priority model for FDs when a class from a forward repairable partition is given.
- Similar to other FD repair algorithms, we use equivalence relations on the set of tuples to keep track of repair steps previously done. In order to do this in an efficient way, the *Swipe* algorithm uses *disjoint set forests*. These data structures model equivalence classes in a tree-based manner and have the property that the asymptotic complexity of merging two classes, is constant.
- We study the theoretical properties of *Swipe*. We prove that it always terminates and when it does, it produces a repair. Moreover, we show that for unary FDs (i.e., FDs with a singleton left-hand side), each FD must be repaired at most once whenever resolution of conflicts is based on choice.
- We empirically study the trade-off between repair quality and computational cost. On one hand, we demonstrate that the generation of repairs with *Swipe* is one to three orders of magnitude faster than with *Llunatic*. On the other hand, the repair quality in terms of F -score of correctly repaired attribute values, is shown to be comparable or better. This provides first evidence for the fact that the construction of a single sequence of Chase steps can lead to good repairs.

The remainder of this paper is organized as follows. In Section 2, we revise the vast body of literature on FD repairing and highlight the most important connections to existing methods. In Section 3, we introduce the basic concepts and notations related to the relational model and functional dependencies used throughout the paper. In Section 4, we first formalize the notion of an attribute partition that is (forward) repairable (Section 4.1) and then develop an algorithm to produce such a partition for a given set of FDs (Section 4.2). Next, we show how repair of a partition class is done (Section 4.3) and finally combine our results to formulate the *Swipe* algorithm (Section 4.4). We provide an experimental analysis in Section 5 and summarize the main contributions in Section 6.

2 Related work

In this section, we provide a concise overview of existing solutions for repairing violations of FDs and contrast them to the proposed algorithm in this paper. For a more detailed discussion, the interested reader is referred to overview papers [28] and monographs [21, 29]. We make a distinction between (i) traditional approaches where searching repairs is cast into an (*approximate*) *optimization problem* and (ii) more recent approaches where *learning* techniques are used to generate repairs.

The traditional approach towards repairing violations of FDs is to consider a cost model that encodes a positive cost to changes made to the original data. The goal is then to find repairs that can be produced with a minimal sum of costs. In the setting of Consistent Query Answering (CQA) [2] one usually considers the deletion of tuples as an elementary change [41, 32]. In the current paper, we focus on the problem where each elementary change is the replacement of the value of one attribute in one row, with some other value. Finding an optimal repair (i.e., a repair with minimal cost) in this setting, has been shown to be max-SNP hard ([30], Theorem 5). Several strategies have therefore been proposed to search for approximate optimal repairs efficiently, including greedy search algorithms [8, 7, 15, 30], sampling approaches [5], usage of MAX-SAT solvers [16] and alignment with knowledge graphs [14]. Many of these approaches make use of the notion of equivalence classes over the set of tuples in order to keep track of which rows must receive the same value for some attribute. This technique has been found necessary to ensure termination of repair algorithms [8]. The way in which equivalence classes are used, can differ between approaches. Equivalence classes could for example be gradually built as a repair is constructed row by row [8]. Another approach that has been used, initializes equivalence classes in a greedy way prior to repair, thereby

determining equivalences that already exist in the dirty relation [5]. The algorithm we present also uses equivalence relations on rows, although without this notion of greedy initialization. Moreover, we explicitly use the notion of disjoint set forests to manage equivalence relations on rows efficiently.

When we evaluate search-based approaches, several experimental studies show that they either deal with scalability issues (even for moderate-sized data sets) [16, 8, 7] or produce repairs of low quality [37]. This provides empirical evidence that it is challenging to find a good trade-off between computational cost and repair quality. One particular method that facilitates this trade-off, is the open-source framework Llunatic [24, 25], which uses a variant of the Chase algorithm. As mentioned in the introduction, this approach produces a Chase tree where each path is a sequence of repair steps and leaf nodes are actual repairs, associated with a certain cost. To control the space and time complexity, the Chase tree can be pruned in several ways. The algorithm we present here can be thought of as a degenerate variant of such a Chase algorithm that considers a single sequence of Chase steps. By doing so, we minimize the computational cost and try to maximize quality of repairs as much as possible. Informally, the key idea is hereby to repair small groups of attributes of a dirty relation one after the other. This idea is similar to the Fellegi-Holt approach to repair tuple-level constraints [22, 9, 6, 10]. Interestingly, in the case of tuple-level constraints, it can be shown that if the set of constraints is satisfiable, there always exists a partition composed of *singleton* classes, such that we can repair sequentially. This result follows from Theorem 1 in [22] for nominal edit rules and has been extended for other types of tuple-level constraints [17, 6]. As a consequence, it is possible for these constraints to repair one attribute at a time for each separate row. A similar property does *not* hold for FDs: we cannot ensure the existence of a partition that is sequentially repairable in the sense defined here and that consists solely of singleton classes [26].

Over the past decade, there has been a growing interest in using learning approaches to deal with the repair problem. The most well-known approach is HoloClean [36], which combines integrity constraints with statistical information about a given data set to compose a probabilistic model that generates repairs. This learning-based technique has gained a lot of attention [33] and inspired others to increase the role of learning. More specifically, learning approaches have also been used for error detection [34, 27, 35] where a sample of clean data is used to learn how errors should be characterized. In these approaches, there is no need anymore to use explicit integrity constraint formalisms. Rather, there is a model that is being trained to recognize errors. Although these approaches are promising for the future, their experimental evaluations have shown they face scalability issues as well, displaying run times of over 10 seconds for data sets with 1000 tuples. We believe there is a requirement for tools that allow to produce repairs much faster than that. Evidence for that requirement can be found in a recent survey paper investigating the landscape of commercial tools for data quality monitoring [20]. Interestingly, none of those tools have adopted methods for automated error detection or repair, whereas we are convinced that these methods are key in measurement and improvement of data quality.

3 Preliminaries

Let \mathbb{A} be a countable set of attributes. For each $a \in \mathbb{A}$, let A denote the *domain* of a . We define a (*relational*) *schema* as a non-empty and finite set of attributes $\mathcal{R} = \{a_1, \dots, a_k\}$ and a *relation* R with schema \mathcal{R} as a finite set $R \subseteq A_1 \times \dots \times A_k$. Elements of R are called *tuples*. We assume that \mathcal{R} always has a special attribute tid that is unique for each $r \in R$. For a relation R with schema \mathcal{R} , a set of attributes $X \subseteq \mathcal{R}$ and a predicate P , we denote the *projection* of R over X by $R[X]$ and the *selection* over P by R_P . For convenience of notation, if X is a singleton set $\{a\}$, we denote $R[\{a\}]$ by $R[a]$. In a similar way, if R contains only one tuple r , we use the notation $r[X]$. Note that the projection $R[X]$ is a set and therefore contains no duplicate tuples. In some cases, it is useful to restrict the schema to attributes X but keep duplicates. We denote $R[\llbracket X \rrbracket]$ as the *multiset* of values obtained after restricting the schema to attributes X without removal of duplicates.

A *functional dependency* (FD) ϕ defined over a schema \mathcal{R} is an expression of the form $X \rightarrow a$ such that $X \subseteq \mathcal{R}$ and $a \in \mathcal{R}$. In this expression, X is called the left-hand side and a is called the right-hand side of the FD¹. For an FD ϕ , we will use the notation $\text{LHS}(\phi)$ for the left-hand side of ϕ and $\text{RHS}(\phi)$ for the right-hand side of ϕ . A relation R with schema \mathcal{R} *satisfies* ϕ (denoted by $R \models \phi$) if:

$$\forall r_1 \in R : \forall r_2 \in R : r_1[X] = r_2[X] \Rightarrow r_1[a] = r_2[a].$$

Similarly, R *satisfies* a set of FDs Φ (denoted by $R \models \Phi$) if it satisfies all $\phi \in \Phi$. For a set of FDs Φ and a set $Z \subseteq \mathcal{R}$, the *projection* of Φ over Z is denoted by $\Phi[Z]$ and defined as:

$$\Phi[Z] = \{\phi \mid \phi \in \Phi \wedge \text{LHS}(\phi) \subseteq Z \wedge \text{RHS}(\phi) \in Z\}.$$

An FD ϕ defined over \mathcal{R} is said to be *implied* by a set of FDs Φ , if for any relation R with schema \mathcal{R} , we have $R \models \Phi \Rightarrow R \models \phi$ (i.e., satisfaction of Φ implies satisfaction of ϕ). We use the notation $\Phi \models \phi$ to say that ϕ is logically implied by the set Φ . Two sets of FDs Φ and Φ' are said to be *equivalent*, denoted by $\Phi \equiv \Phi'$ if:

$$(\forall \phi \in \Phi : \Phi' \models \phi) \wedge (\forall \phi' \in \Phi' : \Phi \models \phi').$$

¹We can assume the right-hand side is a single attribute without loss of generality.

An FD $X \rightarrow a$ is called *minimal* if it is not implied by any $\{X' \rightarrow a\}$ where $X' \subset X$. Finally, for a given set of FDs Φ , a *minimal cover* of Φ is any set $\underline{\Phi}$ such that (i) $\Phi \equiv \underline{\Phi}$, (ii) all $\phi \in \underline{\Phi}$ are minimal and (iii) no proper subset of $\underline{\Phi}$ is equivalent with Φ .

Table 1: Overview of symbols and key concepts.

Symbol	Meaning
\mathbb{A}	Countable set of attributes
$a \in \mathbb{A}$	An attribute with domain A
\mathcal{R}	Relational schema $\{a_1, \dots, a_k\}$
tid	Tuple identifier attribute
R	Relation with schema \mathcal{R}
$r \in R$	A tuple in relation R
$R[X]$	Projection of R over $X \subseteq \mathcal{R}$
$R[[X]]$	Projection of R over $X \subseteq \mathcal{R}$ with preservation of duplicates
ϕ	Functional dependency (FD) of the form $X \rightarrow a$
Φ	A set of FDs
$\Phi \models \phi$	FD ϕ is implied by Φ
$\underline{\Phi}$	A minimal cover of Φ
$\underline{\Phi}[Z]$	Projection of Φ over Z
\mathcal{P}	A partition of \mathcal{R} of the form $[C_1 \dots C_m]$
\mathcal{R}_i	For $\mathcal{P} = [C_1 \dots C_m]$ and $1 \leq i \leq m$, $\mathcal{R}_i = \{C_1 \cup \dots \cup C_i\}$
R_i^*	Partial repair: a relation with schema \mathcal{R} that satisfies $\Phi[\mathcal{R}_i]$
$\mathcal{P} \rightsquigarrow \Phi$	Φ is sequentially repairable over \mathcal{P}
$\mathcal{P} \rightsquigarrow_F \Phi$	Φ is forward repairable over \mathcal{P}
\mathcal{P}^+	The preorder obtained from a set of FDs Φ
$\text{DSF}(a)$	A disjoint set forest (DSF) for a
ρ_a	A repair function for a

4 Sequential repairing

4.1 Basic definitions

Consider a schema \mathcal{R} and a set of FDs Φ defined over \mathcal{R} . Let $\mathcal{P} = [C_1 \dots C_m]$ denote a *partition* of \mathcal{R} where $m \leq k$ and for any $i \in \{1, \dots, m\}$, $C_i \subseteq \mathcal{R}$ are mutually disjoint *partition classes*. We assume in this paper that partition classes C_i are *ordered* by their index i and we refer to this order as the *natural* order of the partition. In the context of such a partition \mathcal{P} , for any $i \in \{1, \dots, m\}$, we denote by R_i^* a relation with schema \mathcal{R} such that $R_i^* \models \Phi[\mathcal{R}_i]$ where $\mathcal{R}_i = \{C_1 \cup \dots \cup C_i\}$. In the context of our repairing approach, we will call such a relation a *partial repair*. We now say that Φ is *sequentially repairable* over \mathcal{P} , denoted by $\mathcal{P} \rightsquigarrow \Phi$, if:

$$\forall i \in \{1, \dots, m\} : \forall R_{i-1}^* : \exists R_i^* : R_i^*[\mathcal{R}_{i-1}] = R_{i-1}^*[\mathcal{R}_{i-1}] \quad (1)$$

where we adopt the convention that $R_0^* = R$. In words, if Φ is sequentially repairable over \mathcal{P} then any partial repair R_{i-1}^* can be transformed into a partial repair R_i^* by modifying only the values for attributes in C_i . By repetition, we eventually obtain a relation R_m^* that satisfies Φ . The relevance of sequential repairability lies in the fact that if $\mathcal{P} \rightsquigarrow \Phi$, then classes from \mathcal{P} can be repaired *one at a time* in order of increasing i . We now make two refinements to the notion of sequential repairability in order to make it a useful repair instrument.

First, it easy to see that for any schema \mathcal{R} and any set Φ defined over \mathcal{R} , there always exists at least one (trivial) partition for which Φ is sequentially repairable. Indeed, since FDs are always satisfiable, we have that $[\mathcal{R}] \rightsquigarrow \Phi$. This partition is not very useful because it simply tells us we can repair violations of Φ . However, if we are able to *refine* this partition (i.e., splitting classes into disjoint sub-classes) and still satisfy the condition of sequential repairability, we can separate the treatment of different FDs and exploit this property in a repair algorithm. For that reason, we are interested here in those partitions \mathcal{P} that are *maximally refined* without losing the property of sequential repairability for Φ .

Second, if $\mathcal{P} \rightsquigarrow \Phi$, it is possible that FDs must be satisfied by either equating the values for the right-hand side attribute or differentiating the values for the left-hand side attributes. For example, if $\mathcal{R} = \{a, b\}$ and $\Phi = \{a \rightarrow b\}$ then we have $[\{a\}, \{b\}] \rightsquigarrow \Phi$ as well as $[\{b\}, \{a\}] \rightsquigarrow \Phi$. That is, for given values of a , we can choose values for b and make sure equal

values for a imply equal values for b . Similarly, for given values of b , we can choose values for a that do not cause a violation of $a \rightarrow b$. In this paper, we restrict repairing to the case where violations of FDs are resolved by changing the values of the right-hand side attribute. This operation is also known as *forward repairing* [25]. We provide the following definition.

Definition 1 For a schema \mathcal{R} and a set of FDs Φ defined over \mathcal{R} , let \mathcal{P} be a partition of \mathcal{R} . Then Φ is forward repairable over partition \mathcal{P} , denoted by $\mathcal{P} \rightsquigarrow_F \Phi$, if and only if:

$$\forall C_i \in \mathcal{P} : \forall \phi \in \Phi[\mathcal{R}_i] \setminus \Phi[\mathcal{R}_{i-1}] : \text{RHS}(\phi) \in C_i \quad (2)$$

In words, Φ is forward repairable if for every class C_i , every FD that was not considered in the scope of previous classes, has a right-hand side attribute that is an element of C_i . We now have the following result.

Proposition 1 For a schema \mathcal{R} and FDs Φ defined over \mathcal{R} , we have $(\mathcal{P} \rightsquigarrow_F \Phi) \Rightarrow (\mathcal{P} \rightsquigarrow \Phi)$.

In this paper, we seek to repair violations of FDs by first building a partition of attributes that allows forward repairing and then apply a repair algorithm to each class of this partition, following the natural order of that partition. This repair algorithm uses a *priority* model that dictates the order in which FDs are repaired. We now demonstrate the main ideas behind this algorithm in the context of the running example.

Example 2 In Figure 1, the attribute partition (middle left) is forward repairable for the FDs. This partition is also maximally refined w.r.t. forward repairability. That means we cannot split classes without breaking the condition of forward repairability. To construct a repair R^* , we visit the classes of \mathcal{P} in their natural order. The first class is $C_1 = \{\text{hospital name}, \#\text{provider}\}$ and we have:

$$\Phi[C_1] = \{\text{hospital name} \rightarrow \#\text{provider}, \#\text{provider} \rightarrow \text{hospital name}\}.$$

We then sort these FDs (this sorting is made more precise later in the paper) and visit them in order. Suppose for now we first visit $\text{hospital name} \rightarrow \#\text{provider}$, then we find two violations: one for tuples $\{1, 2, 3, 4\}$ and one for tuples $\{5, 6\}$. We then do two things. First, we register that these two groups of tuples must have the same value for $\#\text{provider}$ in a special data structure. This data structure is used to keep track of decisions made during repair of visited FDs. Second, we resolve the violation by applying a function that maps the bag of values observed for $\#\text{provider}$ in the violating tuples to a single value. In this example, we do this by majority voting with random tie breaking. Concretely, this means we change $\#\text{provider}$ for $\text{tid} = 4$ into 10006. For the other violation, we have a random choice between the values so assume we choose 1003x. We then visit $\#\text{provider} \rightarrow \text{hospital name}$ and find no violations. Since all FDs in $\Phi[C_1]$ are now satisfied, we have constructed the first partial repair R_1^* . We then proceed to $C_2 = \{\text{city}\}$. The set $\Phi[C_1 \cup C_2]$ now has one additional FD, which is $\text{hospital name} \rightarrow \text{city}$ and it has no violations. Moreover, the FDs from the previous step are ensured to remain satisfied so we now have a second partial repair R_2^* for which $R_2^* \models \Phi[C_1 \cup C_2]$. This procedure continues until we have visited $C_7 = \{\text{condition}\}$. After this last step, the original data has been iteratively transformed into a repair R^* that satisfies all FDs. An example repair under the assumption of majority voting with random tie breaking to fix violations, is shown in Figure 1 (bottom).

From the example given above, two important questions arise. The first question deals with the generation of a maximally refined partition that is forward repairable for a given set of FDs. We will treat this question in Section 4.2. The second question deals with how violations of FDs are resolved for a single class. We will treat this question in Section 4.3. After that, we combine all these ideas in Section 4.4 to present the Swipe algorithm.

4.2 Attribute partition building

This section presents a method to construct, for given \mathcal{R} and Φ over \mathcal{R} , a partition \mathcal{P} that (i) is maximally refined and (ii) satisfies $\mathcal{P} \rightsquigarrow_F \Phi$. To do so, note that Definition 1 (Eq. (2)) implies that for any FD $X \rightarrow a$ in Φ , attributes in X must be part of a class that does *not lie after* the class that contains a . We will encode this information in a preorder relation constructed from Φ .

As a first step, we verify that Φ does not contain any trivial or implied FDs. This is to avoid that trivial FDs introduce unnecessary constraints on the order of attributes. In the scope of the running example in Figure 1, introducing the trivial FD $\{\text{county}, \text{zip}\} \rightarrow \text{county}$ would create a constraint on the order of the classes that contain county and zip . This constraint is unnecessary since the FD is always satisfied. The same observation holds for FDs that are not minimal (e.g., $\{\text{hospital name}, \text{state}\} \rightarrow \text{zip}$). To ensure that we encode only necessary requirements, we transform Φ into a *minimal cover* $\underline{\Phi}$. From such a minimal cover, we compose a binary relation $P \subseteq \mathcal{R} \times \mathcal{R}$ that satisfies the following criteria:

1. $\forall a \in \mathcal{R} : (a, a) \in P$,

2. $\forall (X \rightarrow a) \in \Phi : \forall b \in X : (b, a) \in P$.

Here, $(b, a) \in P$ expresses the requirement that b should not occur after a in the sequence of partition classes. Next, P is transformed into its *transitive closure* denoted by P^+ [40]. The obtained relation P^+ is the finest preorder (reflexive and transitive) that contains P . We say that P^+ is the preorder obtained from Φ . This preorder contains an equivalence relation \equiv_{P^+} on \mathcal{R} . More specifically, we have $a \equiv_{P^+} a'$ whenever both $(a, a') \in P^+$ and $(a', a) \in P^+$. Consider now any partition $\mathcal{P} = [C_1 \dots C_m]$ such that:

$$\forall a \in \mathcal{R} : \forall a' \in \mathcal{R} : a \equiv_{P^+} a' \Leftrightarrow (\exists C_i \in \mathcal{P} : a \in C_i \wedge a' \in C_i) \quad (3)$$

and:

$$\forall C_i \in \mathcal{P} : \forall C_j \in \mathcal{P} : i < j \Rightarrow (\forall a \in C_i : \forall a' \in C_j : (a', a) \notin P^+) \quad (4)$$

A partition \mathcal{P} that satisfies these two constraints is said to be induced by P^+ . A partition induced by P^+ can be constructed readily by first assigning each equivalence class from \equiv_{P^+} to a class C_i and then sort the obtained classes to enforce Eq.(4). This construction shows that a partition induced by P^+ is a topological sort of the quotient set \mathcal{R}/\equiv_{P^+} . We can now state the following result.

Theorem 1 For a schema \mathcal{R} and a set of FDs Φ defined over \mathcal{R} , let \mathcal{P} be a partition induced by P^+ . Then (i) $\mathcal{P} \rightsquigarrow_F \Phi$ and (ii) $\neg(\mathcal{P}' \rightsquigarrow_F \Phi)$ for any refinement \mathcal{P}' of \mathcal{P} .

P^+	hospital name	#provider	city	county	zip	state	measure code	condition
hospital name	x	x	x	*	x	x		
#provider	x	x	*	*	*	*		
city			x	x				
county				x				
zip					x			
state						x		
measure code							x	x
condition								x

Figure 2: The construction of a preorder P^+ for the FDs from Figure 1 (middle). Elements derived directly from the FDs are marked with ‘x’ and elements added to compute the transitive closure are marked with ‘*’. The preorder is the union of two weak orders, marked by bold lines.

Example 3 For the FDs shown in Figure 1 (middle), Figure 2 shows the construction of the preorder P^+ . First, we derive P from a minimal cover of the given FDs as explained. In our example, the available FDs are already a minimal cover and thus we proceed with the FDs as given. In Figure 2, the entries in P are marked by ‘x’ symbols. We then compute the transitive closure P^+ and the additional entries to get P^+ are marked by ‘*’ symbols. Because $\text{hospital name} \equiv_{P^+} \text{\#provider}$, both attributes must belong to the same partition class. Attributes determined by hospital name appear in singleton classes after the class that contains hospital name and #provider. Finally, condition should be after measure code.

4.3 Priority repair

Assume a partition $\mathcal{P} = [C_1 \dots C_m]$ for \mathcal{R} such that $\mathcal{P} \rightsquigarrow_F \Phi$ for a given set of FDs Φ . Definition 1 then ensures we can visit classes C_i in their natural order and generate a partial repair R_i^* by changing only attributes from C_i . That is, we are given a partial repair R_{i-1}^* and we must modify attributes C_i to fix violations of FDs $\Phi[R_i]$ and obtain a new partial repair R_i^* . To do this, we propose a technique called *priority repairing*. The main principles of this technique are the following:

- FDs are inspected and repaired in a specific order, dictated by a *priority* model. The key idea is to repair FDs that contain less reliable attributes in their right-hand side first and thereby try to maximize the accuracy of the repairs. We propose a simple heuristic to estimate the reliability of an attribute by means of the estimated number of changes that an attribute requires.
- In order to keep track of tuples that must receive the same value for $a \in C_i$, we use a *disjoint set forest* (DSF) [23] to represent equivalence classes on tuples. A DSF has the advantage that the necessary operations on equivalence classes have a near-constant time complexity [38]. This mitigates the complexity of *FD revision*, where one FD must be repaired again because of possible new violations caused by the repair of another FD.
- The value assigned to an equivalence class for $a \in C_i$ is determined by a *repair function* ρ_a . This is a function that maps the bag of values of a involved in a violation onto a new value. We have a particular interest in those repair functions that are *preservative*. These functions simply choose one of the values involved in a violation.

In the remainder of the section, we detail and formalize these main principles.

Priority model In order to repair class C_i , we must ensure that the FDs over $\Phi[\mathcal{R}_i]$ are satisfied. Clearly, any FD from $\Phi[\mathcal{R}_{i-1}]$ contains no attributes from C_i and must therefore not be considered. We can therefore restrict ourselves to the set $\Phi_i := \Phi[\mathcal{R}_i] \setminus \Phi[\mathcal{R}_{i-1}]$. In this set there can be FDs ϕ for which:

$$\text{LHS}(\phi) \cap C_i = \emptyset. \quad (5)$$

Those FDs have a left-hand side that contains only attributes that are clean at the time of repairing class C_i because these attributes are all part of a class C_j with $j < i$. Such FDs are called *pilot* FDs for class C_i and we repair them first, with two clear advantages. First, because we rely on *clean* data to make the first changes to attributes $a \in C_i$, we expect these changes to be accurate. Second, pilot FDs never require revision as the values of their left-hand side attributes never change. Therefore, they need to be considered only once.

Example 4 In Example 2, there are no pilot FDs for class C_1 . For all other classes, we have only pilot FDs because the left-hand side of the FDs contains no attributes that are part of the class.

If an FD is not a pilot FD for class C_i , it contains at least one attribute from C_i in the left-hand side. Those FDs are further sorted by means of a simple *priority model* that encodes the order in which attributes must be repaired. This priority model is a simple total order $>$ on attributes C_i where $a > a'$ indicates that a must be repaired prior to a' . In other words, we require that $X \rightarrow a$ is repaired *before* $X' \rightarrow a'$ whenever $a > a'$. To construct a priority model, we propose to rank attributes within class C_i by increasing reliability. To make such a ranking, the most obvious method would be to ask for input of a human supervisor, which is not uncommon in repair approaches [24, 25]. However, such input might not be easy to collect, either because there is no supervisor, or a supervisor cannot easily rank attributes. In that case, we propose to use an *estimate* of the number of changes we will need to make to values of a as follows.

Suppose we have an FD $\phi := X \rightarrow a$ and a relation R , then for any $x \in R[X]$, the bag $R_{X=x}[[a]]$ gives us the values for attribute a of those tuples for which $X = x$. If there are *different* elements in this bag, we must change values for a in order to satisfy $X \rightarrow a$. If we denote the element with *highest* multiplicity² in this bag by $\text{mv}(x, a)$, then changing the values of a for tuples $R_{X=x \wedge a \neq \text{mv}(x, a)}$ allows to resolve the conflict with a *minimal* amount of changes. We can now collect these tuples for all values $x \in R[X]$:

$$\text{VioFD}(X \rightarrow a) := \bigcup_{x \in R[X]} R_{X=x \wedge a \neq \text{mv}(x, a)}, \quad (6)$$

and finally aggregate over all FDs $\phi \in \Phi$ where $\text{RHS}(\phi) = a$:

$$\text{Vio}(a) := \bigcup_{\phi \in \Phi \wedge \text{RHS}(\phi) = a} \text{VioFD}(X \rightarrow a) \quad (7)$$

The set $\text{Vio}(a)$ basically gives us the tuples in R that require a change to attribute a if (i) we are interested in keeping the number of changes as small as possible and (ii) we would consider each FD independently. The more changes an attribute requires, the less reliable that attribute is. We therefore use $|\text{Vio}(a)|$ to build our priority model such that:

$$a > a' \Leftrightarrow |\text{Vio}(a)| \geq |\text{Vio}(a')| \quad (8)$$

The rationale of this heuristic is that since we now must sort non-pilot FDs, there are attributes from C_i that appear in the left-hand side of these FDs. It is likely that some of these attributes will themselves contain errors (i.e., because they appear in

²If there are multiple such values, we choose one at random.

the right-hand side of another FD). We thus seek to first repair an FD that has attributes in the left-hand side, that are *maximally* reliable. We now make two observations about this heuristic.

First, the assumption about FD independence, is a naive one. When repairing two FDs with the same right-hand side attribute a we must merge equivalence classes induced by both FDs (this is explained later) and our estimate neglects this. We accept this error in our estimate at the benefit of simple and efficient computation. Moreover, we will show empirically that classes C_i are in practical scenarios often singleton set, in which case the estimate must not be computed because all FDs are then by definition pilot FDs. If we do have multiple attributes in C_i , then empirical results show that the heuristic leads to sequences of FDs with repair quality superior to other possible sequences.

A second observation is that $|\text{Vio}(a)|$ can be computed at different times with different results. One could for example compute $|\text{Vio}(a)|$ on the original relation R , before any modifications are done. The advantage of doing so would be that estimates are unbiased and not dependent on changes made by the Swipe algorithm. An alternative strategy is to postpone the computation of $\text{Vio}(a)$ up to the moment where the class that contains a is visited. In other words, we then compute $\text{Vio}(a)$ from the partial repair R_{i-1}^* . Variations of this strategy are possible. We can for example estimate reliability before repair of C_i starts or postpone it to the point where pilot FDs have been repaired. Empirically, we have however not observed any significant differences between these variations on the data sets we tested. In the remainder of this paper, we favor a strategy where estimation of reliability is done on *partial repairs*.

Example 5 In Example 2, we have $C_1 = \{\text{hospital name}, \#\text{provider}\}$. Both attributes appear once in the right-hand side of an FD. It can be verified that $\text{Vio}(\text{hospital name}) = \emptyset$ and $\text{Vio}(\#\text{provider})$ is either $\{4, 5\}$ or $\{4, 6\}$. It follows that $\#\text{provider} > \text{hospital name}$ and thus that we repair $\text{hospital name} \rightarrow \#\text{provider}$ first and then $\#\text{provider} \rightarrow \text{hospital name}$.

Example 5 demonstrates (i) the estimation of reliability of attributes and (ii) the implied order on FDs. Note that the two sorting principles we use (i.e., pilot FDs first and then least reliable right-hand side first) provides us with a *partial* order on the FDs we must repair. Any total order that *contains* this partial order can be used in Swipe.

Tuple equivalence Having determined the order in which FDs are treated, we turn our attention to the actual repair process. In order to repair attributes in C_i , we consider for each $a \in C_i$, an equivalence relation on the set of tuples from R_{i-1}^* . We use these equivalence classes to keep track of which tuples must receive the same value for a in order to satisfy FDs ϕ for which $\text{RHS}(\phi) = a$. More specifically, whenever two tuples are in the same equivalence class for attribute a , we require they must receive the same value for a . Initially, each tuple is put in a separate singleton class. As we fix violations of FDs of the form $X \rightarrow a$, we require that tuples with the same value for X also have the same value for a . In other words, if we observe tuples in R_{i-1}^* with equal values for X , we must *merge* the corresponding classes in which these tuples appear. In order to do this efficiently, we use a *disjoint set forest* [23] for each attribute a , denoted by $\text{DSF}(a)$. With this data structure, each equivalence class is represented by a *tree* on tuple identifiers (i.e., values for tid). Each $\text{DSF}(a)$ basically supports three operations:

- Operation `makeset(id)` adds a new singleton set (i.e., $\{\text{id}\}$) to the forest. Such a singleton set is represented by a one-node tree. This operation is used to initialize each $\text{DSF}(a)$.
- Operation `find(id)` finds the root tid of the tree in which id resides and is used to determine if two values are in the same class or not.
- Operation `union(id1, id2)` merges the classes in which id_1 and id_2 reside. This is done by finding the root nodes of the trees and in case the root nodes are distinct, attach one root as a child of the other. During the repair of an FD $X \rightarrow a$, this operation is used to merge equivalence classes for tuples that have equal values for X .

A key result is that the three elementary operations of a disjoint set forest can be implemented such that their *asymptotic* time complexity is very efficient. More specifically, one can provide an implementation such that the asymptotic time complexity of `makeset(id)` and `union(id1, id2)` is $O(1)$, while that of `find(id)` is $O(\alpha(n))$ where $\alpha(\cdot)$ is an inverse functional of the Ackermann function and n is the number of elements in the class where id is located [38, 39, 4]. The asymptotic efficiency of these operations is key because it allows us to provide an efficient strategy for the update of $\text{DSF}(a)$ during the repair of an FD $X \rightarrow a$. More specifically, if we want to repair $X \rightarrow a$, then we must update $\text{DSF}(a)$ such that after the update we have:

$$\forall r \in R_{i-1}^* : \forall r' \in R_{i-1}^* : r[X] = r'[X] \Rightarrow \text{find}(r[\text{tid}]) = \text{find}(r'[\text{tid}]) \quad (9)$$

In words, after the update is done, we want that any two tuples from R_{i-1}^* with equal values for X are equivalent in $\text{DSF}(a)$ and thus must receive the same value for a .

The pseudo-code shown in Algorithm 1 provides a simple way to do so by using a map H with constant-time complexity for `get` and `put` operations. This map keeps, for each value in $R_{i-1}^*[X]$, the root node of the tree in $\text{DSF}(a)$ in which tuples with this value reside. We then iterate over tuples $r \in R_{i-1}^*$ and check if H currently contains $r[X]$. If this is not the case, we

Algorithm 1 Update DSF (a) for $X \rightarrow a$ and R_{i-1}^*

```
1: procedure UPDATE( $R_{i-1}^*, X \rightarrow a, \text{DSF}(a)$ )
2:    $H \leftarrow []$  ▷ Initialize empty map
3:   for  $r \in R_{i-1}^*$  do
4:     if  $H[r[X]] = \text{NULL}$  then
5:        $H[r[X]] \leftarrow \text{DSF}(a).\text{find}(r[\text{tid}])$  ▷ Register root
6:     else
7:        $\text{DSF}(a).\text{union}(r[\text{tid}], H[r[X]])$  ▷ Merge classes
8:        $H[r[X]] \leftarrow \text{DSF}(a).\text{find}(r[\text{tid}])$ 
9:     end if
10:  end for
11: end procedure
```

register the root node of the tree where $r[\text{tid}]$ is currently located (line 5). Else, we merge the current root node with the tree where $r[\text{tid}]$ is currently located and update the map H (lines 7 - 8). The update procedure described here requires in each step one find operation and possibly one union operation on $\text{DSF}(a)$, which implies an asymptotic complexity of $O(|R_{i-1}^*| \cdot \alpha(n))$, with n the size of the largest equivalence class in $\text{DSF}(a)$.

Example 6 In Example 5 it is shown that during the visit of C_1 , we first repair $\text{hospital name} \rightarrow \#\text{provider}$. This is the first FD we visit, which means $\text{DSF}(\#\text{provider})$ contains a singleton class for each tuple. Then application of

update ($R, \text{hospital name} \rightarrow \#\text{provider}, \text{DSF}(\#\text{provider})$)

modifies $\text{DSF}(\#\text{provider})$ in such a way that tuples with equal values for hospital name are in the same class. In other words, after the update, $\text{DSF}(\#\text{provider})$ has classes $\{\{1, 2, 3, 4\}, \{5, 6\}\}$.

Fixing violations with repair functions After updating $\text{DSF}(a)$ for some FD $X \rightarrow a$ with Algorithm 1, $\text{DSF}(a)$ represents classes of tuple identifiers from R_{i-1}^* that must have the same value for a . In other words, $\text{DSF}(a)$ corresponds with set $\{\mathcal{E}_1, \dots, \mathcal{E}_\ell\}$ such that $\forall i \in \{1, \dots, \ell\} : \mathcal{E}_i \subseteq R_{i-1}^*[\text{tid}]$. For each such set \mathcal{E} , we can get the corresponding tuples from R_{i-1}^* and verify whether or not these tuples have equal values for a . If not, the tuples violate $X \rightarrow a$ and we fix this violation by application of a repair function defined by $\rho_a : \mathcal{B}(A) \rightarrow A$, where $\mathcal{B}(A)$ is the set of all *bags* (or *multisets*) with elements from A . A repair function maps a bag of values from A onto a single value³ from the domain A . A particular class of repair functions are those that are *preservative*, which means that $\rho_a(B) \in B$ for any bag $B \in \mathcal{B}(A)$ [12]. Our interest in preservative repair functions is not coincidental. First of all, preservative repair functions possess important properties: they are *idempotent* ($\rho\{v, \dots, v\} = v$), they ensure that for any repair R^* , we have $\forall a \in \mathcal{R} : R^*[a] \subseteq R[a]$ (i.e., they do not introduce new values in the repair) and they play an important role in notion of cardinality-minimal repairs [5]. Second, preservative repair functions have some well-known representatives like majority voting, weighted voting and median or min/max selection for domains equipped with an order relation. Finally, we will show in the following that they play a key role in the avoidance of revisions for unary FDs.

Algorithm 2 provides the pseudo-code for the resolution of violations of $X \rightarrow a$. It first updates the $\text{DSF}(a)$ by using Algorithm 1. Then, for each class of tuple identifiers in the updated $\text{DSF}(a)$, it gets the tuples from R^* with identifiers in this class and stores these tuples in a variable $R_{\mathcal{E}}^*$ (line 5). These tuples are inspected for their values for a (line 6). If a violation is observed, the repair function for a is applied to the bag of values $R_{\mathcal{E}}^*[a]$ and the result is stored in the variable v_{fix} (line 7). This value is then assigned to attribute a for all the tuples of R_{i-1}^* involved in the violation (line 8)

Example 7 In Example 6 we showed that $\text{DSF}(\#\text{provider})$ has two equivalence classes. Suppose we choose as repair function majority voting with random tie breaking then fixing $\text{hospital name} \rightarrow \#\text{provider}$ must change the value of $\#\text{provider}$ for tuple 4 into 10006. In addition, we must change also tuple 5 or 6.

Priority repair We now have all components in place to present an algorithm that performs priority repair on a class of attributes C_i . The pseudo-code for this algorithm is shown in Algorithm 3. The algorithm receives a class of attributes C_i and the corresponding set of FDs $\Phi_i := \Phi[\mathcal{R}_i] \setminus \Phi[\mathcal{R}_{i-1}]$, which contains only those FDs that are possibly still violated in R_{i-1}^* . The algorithm starts with sorting the FDs according to the principles explained before: we prioritize pilot FDs (see Eq. (5)) followed by the other FDs, which are sorted by their right-hand side attribute based on the cardinality of the set $\text{Vio}(\cdot)$. FDs are added to a stack that maintains this order *at all time* (line 2).

³In the literature on repairing FDs, many existing approaches allow labelled variables during repair, but we restrict ourselves to constant values only.

Algorithm 2 Fix violations of $X \rightarrow a$

```
1: procedure FIX( $R_{i-1}^*, X \rightarrow a, \text{DSF}(a)$ )
2:   update ( $R_{i-1}^*, X \rightarrow a, \text{DSF}(a)$ ) ▷ Algorithm 1
3:   fixes  $\leftarrow 0$  ▷ Count violations
4:   for  $\mathcal{E} \in \text{DSF}(a)$  do ▷ For each class in DSF( $a$ )
5:      $R_{\mathcal{E}}^* \leftarrow \{r \mid r \in R_{i-1}^* \wedge r[\text{tid}] \in \mathcal{E}\}$ 
6:     if  $|R_{\mathcal{E}}^*[a]| > 1$  then ▷ Violation check
7:        $v_{\text{fix}} \leftarrow \rho_a(R_{\mathcal{E}}^*[a])$  ▷ Apply repair function
8:       exec update  $R_{i-1}^*$  set  $a = v_{\text{fix}}$  where  $\text{tid} \in \mathcal{E}$  ▷ Fix violation
9:       fixes++
10:    end if
11:  end for
12:  return fixes
13: end procedure
```

Algorithm 3 Repair FDs Φ_i by changing attributes C_i to obtain partial repair R_i^*

Require: C_i is a class of \mathcal{P} such that $\mathcal{P} \rightsquigarrow_F \Phi$

```
1: procedure PRIORITYREPAIR( $R_{i-1}^*, \Phi_i, C_i$ )
2:   init ( $\mathbb{S}, \Phi_i$ ) ▷ Sorted stack of FDs
3:    $\forall a \in C_i$  : init ( $\text{DSF}(a)$ ) ▷ Initialize DSF with singleton classes
4:   while  $\mathbb{S} \neq \emptyset$  do
5:      $\phi \leftarrow \text{poll}(\mathbb{S})$ 
6:     fixes  $\leftarrow \text{fix}(R_{i-1}^*, \phi, \text{DSF}(\text{RHS}(\phi)))$  ▷ Fix step
7:     if fixes  $\neq 0$  then ▷ Revision step
8:       for  $\phi' \in (\Phi_i \setminus \mathbb{S})$  do
9:         if  $\text{RHS}(\phi) \in \text{LHS}(\phi')$  then
10:          add ( $\mathbb{S}, \phi'$ )
11:        end if
12:      end for
13:    end if
14:  end while
15:  return  $R_{i-1}^*$ 
16: end procedure
```

Next, we initialize a DSF structure for each $a \in C_i$. This initialization creates a singleton class for each $r \in R_{i-1}^*$ and thus requires $|C_i| \cdot |R_{i-1}^*|$ makeset operations (line 3). We then start a loop that continues until the stack with FDs is empty. In each iteration, we poll the first FD from the stack and perform two steps: the fix step and the revision step. In the *fix* step, we retrieve the current DSF structure for $\text{RHS}(\phi)$ and we apply Algorithm 2 to fix any violations (line 6). In the *revision* step, we verify if the fix step made changes to R_{i-1}^* . If this is the case, then the values for $\text{RHS}(\phi)$ in R_{i-1}^* have been modified. We then check if there are FDs ϕ' that have $\text{RHS}(\phi)$ in the left-hand side and add those FDs back to the stack to check if new violations have been introduced (line 9). This continues until no new violations are found.

Algorithm 3 can now be attributed with the following properties. First of all, for each attribute a , the DSF structure is initialized only once and during each update (Algorithm 1), we change $\text{DSF}(a)$ only by merging classes. In particular, this means that $\text{DSF}(a)$ before update is a partition refinement of $\text{DSF}(a)$ after update. As such, any FDs that were previously repaired and have a in the right-hand side, remain satisfied. For FDs that have a in the left-hand side, this might not be the case and these FDs are added to the stack for revision (only if $\text{fixes} > 0$). During such a revision, $\text{DSF}(a)$ can again only change by merging classes, which means that in consecutive revisions, the number of classes in $\text{DSF}(a)$ is monotonically decreasing. This leads us to the formulation of the following theorem.

Theorem 2 Let $\mathcal{P} = [C_1 \dots C_m]$ be a partition of \mathcal{R} such that $\mathcal{P} \rightsquigarrow_F \Phi$ for a set of FDs Φ . Then for any class C_i from \mathcal{P} and any partial repair R_{i-1}^* , we have (i) execution of **priorityRepair** (R_{i-1}^*, Φ_i, C_i) (Algorithm 3) terminates and (ii) after termination R_{i-1}^* is modified into a partial repair R_i^* .

Theorem 2 states that even in case of revisions, we eventually end up with a partial repair R_i^* . Nevertheless, the presence of revisions comes with a cost as it requires a new update of the DSF structure and potential modifications to the repair. It is therefore clear that avoiding revisions has a computational advantage. In that regard, we note that whenever $|C_i| = 1$, all FDs

are necessarily pilot FDs and thus never require revision. Moreover, in the case of unary FDs, revision is not needed if the repair function for the single left-hand side attribute is preservative.

Proposition 2 *Let $\mathcal{P} = [C_1 \dots C_m]$ be a partition of \mathcal{R} such that $\mathcal{P} \rightsquigarrow_F \Phi$ for a set of FDs Φ . Then for any class C_i from \mathcal{P} and any partial repair R_{i-1}^* , a unary FD $a' \rightarrow a$ can be ignored in the revision step of **priorityRepair** (R_{i-1}^*, Φ_i, C_i) if $\rho_{a'}$ is preservative.*

The proposition above shows that if we use preservative repair functions, then repair of unary FDs is extremely efficient as it is revision-free.

4.4 The Swipe algorithm

The ability to (i) construct a maximally refined partition that is forward repairable and (ii) repair classes of such a partition by means of a priority model can now be used to construct a repair algorithm for FDs that we call the Swipe algorithm. The pseudo-code to repair a dirty relation R in this way is shown in Algorithm 4. The algorithm starts by computing a minimal cover of FDs in order to remove any redundant information that is present (line 2). Based on such a minimal cover, we construct a partition $\mathcal{P} = [C_1 \dots C_m]$ induced by the preorder P^+ (line 3) as explained in Section 4.2. We initialize the repair with the original dirty relation R . This repair is then modified step by step. More precisely, we iterate over the classes C_i of \mathcal{P} in the natural order and in each step apply the priority repair algorithm, where we give the current partial repair R_{i-1}^* as input, together with those FDs that involve at least one attribute from C_i (line 6). During the execution of the priority repair algorithm, R_{i-1}^* is modified by assigning values to attributes in C_i in such a way we get a new partial repair R_i^* . Because the partition is forward repairable, we know that such a modification is possible in each iteration. Eventually, we obtain a repair $R^* := R_m^*$ that satisfies all FDs.

Algorithm 4 Swipe

```

1: procedure SWIPE( $R, \Phi$ )
2:    $\underline{\Phi} \leftarrow$  minimalCover ( $\Phi$ )                                      $\triangleright$  Minimal cover
3:    $\underline{\mathcal{P}} \leftarrow$  buildFromPreorder ( $\underline{\Phi}$ )                              $\triangleright$  Construct  $\mathcal{P}$  from preorder  $P^+$  on  $\Phi$ 
4:    $R_0^* \leftarrow R$                                                   $\triangleright$  Initialize repair
5:   for  $i \in \{1, \dots, |\mathcal{P}|\}$  do
6:      $R_i^* \leftarrow$  priorityRepair( $R_{i-1}^*, \underline{\Phi}[\mathcal{R}_i] \setminus \underline{\Phi}[\mathcal{R}_{i-1}], C_i$ )  $\triangleright$  Priority repair on  $C_i$  for new FDs
7:   end for
8:   return  $R_m^*$ 
9: end procedure

```

5 Experimental evaluation

In this section, we report the results of an empirical study that scrutinizes the efficiency and accuracy of the Swipe algorithm. Because the intent of the paper is to optimize the trade-off between repair quality and computational cost in a Chase-like procedure, we are primarily interested in comparing the Swipe algorithm (a single-sequence Chase procedure) with Llunatic (a multi-sequence Chase procedure). The goal is hereby to seek evidence for the hypothesis that the Swipe algorithm allows to find good repairs of FD violations in an efficient manner and that exploring many repairs does not necessarily improve the quality of the final repair. To that end, the study assesses both the quality of repairs in terms of precision, recall and F -score (Section 5.3) as well as the scalability of the Swipe algorithm (Section 5.4). All experiments reported below were executed on a machine running Ubuntu 20.04.1 with 64GB of RAM memory, 12 cores (i9-10920X/3.50GHz) and two 500 GB Solid State Disks in mirror (WDS500G3XHC-00SJG0).

5.1 Data sets

Real-life data sets The study involves four real-life data sets, which differ in size, number of errors and number of considered constraints. Two of them (Hospital [16, 36, 25] and Flight [31, 25, 10]) are commonly used in experiments on data quality, while the others have been composed and used in previous work by the authors. Each data set is accompanied by a gold standard that contains the correct tuples for a sample of the data set.

Table 2: Summary of data sets used in the experiments. Repairable attributes are attributes that occur in at least one FD. Cells in error are cells for which the value is different in R_{gold} compared to R .

	Hospital	Allergen	Eudract	Flight
# rows	1000	1160	86670	776067
# rows w. ground truth	1000	206	3133	70951
# attributes	19	22	9	5
# repairable attributes	15	22	9	5
# FDs	15	21	11	4
# cells in error	509 (2.68%)	358 (8.28%)	2962 (11.82%)	123799 (43.62%)
# cells with NULL	0	0	1722 (6.87%)	59837 (21.08%)

Table 2 provides a summary of the data set descriptions, including the size of the gold standard and the number of constraints considered in our experiments. More details on the origin of these data sets and their gold standard, as well as full downloads, are openly available⁴. A short description for each data set is included below.

- **Hospital.** A benchmark data set often used in literature on data cleaning [16, 36]. The original data [16] stems from the U.S. Department of Health and Human Services⁵. We use the version consisting of 1000 tuples that has been used before in data repairing experiments [36]. The running example from Figure 1 is a sample from this data set.
- **Allergen.** A data set containing data regarding allergens of food products as found on two different websites [10]. The attributes indicate the presence ('2'), traces ('1'), or absence ('0') of allergens in a product which can be identified by its barcode. A gold standard was compiled by manual consultation of pictures of a sample of products [10].
- **Eudract.** A data set with data about clinical trials conducted in Europe [11, 10] and obtained from the Eudract register⁶. A gold standard was compiled based on the information of the same studies in other registries.
- **Flight.** A data set that describes flights annotated with the departing and arrival airports as well as their expected and actual time of departure and arrival. The original data set was used in the context of data fusion [31, 18, 19] and is available online⁷. In the current study, we use the version and gold standard proposed in recent work [10].

Artificial data sets In order to assess scalability, artificial data sets (i.e., an artificial relation R paired with a set of constraints Φ) are used in Section 5.4. The process to generate these data sets takes two parameters: the number of attributes (i.e., $|\mathcal{R}|$) and the number of tuples (i.e., $|R|$). For a given number of tuples and attributes, we generate tuples by randomly assigning values to attributes under the condition that $|A| = 10$ for each attribute a . By choosing the domain size sufficiently small, we ensure that there are many violations of FDs that need repairing. Clearly, if we would increase domain sizes $|A|$ we would obtain less violations and observe faster run times. Once relations are generated, we generated a set Φ such that $|\Phi| = |R|$, meaning that the number of FDs is equal to the number of attributes. This reflects the fact that for larger schemata, one can expect more dependencies that must be satisfied. To generate an FD, we first choose the size of the left-hand side by sampling from a uniform distribution $\text{unif}\{1, \lceil |\mathcal{R}|/10 \rceil\}$. This reflects that for larger schemata, it becomes more likely to have FDs with multiple attributes in the left-hand side. We then sample for that FD the required number of attributes uniformly from \mathcal{R} to compose the left-hand side. The right-hand side attribute is chosen from the remaining attributes. When a dirty relation R and a set of FDs has been generated, we apply Swipe to obtain a repair R^* .

5.2 Repair techniques

We compare the Swipe algorithm with Llunatic [25]. The data sets were stored in a PostgreSQL database in separate schemas. Note that we used PostgreSQL version 9.5 because Llunatic relies on the ‘with oids’ option during table creation, which is abandoned as of version 12⁸.

⁴<https://gitlab.com/antoonbronselaer/swipe-reproducibility>

⁵<http://www.hospitalcompare.hhs.gov>

⁶<https://eudract.ema.europa.eu/>

⁷<http://lunadong.com/fusionDataSets.htm>

⁸See <https://www.postgresql.org/docs/12/release-12.html>, section E.13.2

Table 3: Configuration details of LLunatic (FT) for each data set.

Data set	LLunatic (FT)
Hospital	Similarity cost manager
	Higher branching threshold
	More potential solutions
	Backward repairs allowed
Allergen	Mimic max repair function
Eudract	Higher branching threshold
	More potential solutions
	Backward repairs allowed
Flight	Similarity cost manager
	Higher branching threshold
	More potential solutions
	Backward repairs allowed

LLunatic. LLunatic is a general purpose Chase engine that builds a Chase tree where nodes correspond to (sequences of) repair steps and leaf nodes represent repairs. The implementation allows configuration of (i) pruning strategies to keep the search efficient and (ii) a partial order that indicates preferences of values for repair. We use the most recent version (2.0) available at GitHub⁹. We made a local build of the LLunatic code in order to be able to run it from the command line on an lxc container. We encoded the available FDs for each data set in separate configuration files. As LLunatic allows for a wide range of configuration options, we decided to create two configurations for each data set:

- The first configuration is the standard one (S) and is the same for each data set. It searches for forward repairs only and we set both the branching threshold and the potential solutions threshold to 2. This means that the Chase tree is binary and we stop searching after two repairs have been found. The partial order to resolve conflicts is based on the frequency of values. The main idea of this standard configuration is that it resembles the search strategy of Swipe. More specifically, it resolves conflicts in the same way as the default method in Swipe (i.e., majority voting), it uses forward repairing only and it constructs few sequences. The main difference is thus that Swipe will construct a specific order on the FDs and uses no variables.
- The second configuration is a fine-tuned configuration (FT) that is tailored to each data set individually. First of all, we allowed backward repairing for Hospital, Eudract and Flight to see if this helps in finding better repairs. We did not allow this for Allergen as there we wanted to focus on mirroring the behaviour of Swipe when using the max repair function. Second, to account for typographical errors that are present in the attributes of Hospital and Flight, we used a cost manager that uses Smith-Waterman string similarity between values to guide resolution. For Hospital, Eudract and Flight, we also increased the branching threshold such that a larger part of the search space is explored and more repairs are searched for.

A summary of the main differences of the FT configurations is given for each data set in Table 3. The full configuration files are available online and can be used for reproduction purposes⁴.

Swipe. The Swipe algorithm was implemented in Java (Version 8 or higher) as part of the ledc framework¹⁰, which is an open-source framework that bundles various techniques for data quality under one umbrella project. The implementation of Swipe contains a set of basic repair functions. We consider three of them in our experiments:

- The first one is a *majority* voting function (mv) with random tie breaking. This repair function was used in the examples throughout this paper and is considered a default choice.
- The second one is a *weighted* voting function (wv) with random tie breaking. The weight of a value in the bag is based on the number of NULL values that occur in the row where that value is taken. More specifically, if a value from the bag is taken from a tuple with N NULL values, the weight for that value is equal to $(|\mathcal{R}| - N)^4$. We include this repair function as it has been shown effective in previous work [10].

⁹<https://github.com/donatelloosantoro/LLunatic>

¹⁰<https://gitlab.com/ledc/ledc-fundy>

- The third repair function simply chooses the *maximal* value (max) and can be used in a scenario where attribute domains are equipped with an order relation that is relevant to the repair. This is true for all datasets included in the experiment, except for the Hospital dataset, where using the natural order simply sorts the attribute values alphabetically.

For each data set, we run Swipe three times and in each run use a different repair function. That repair function is then used for *all* attributes in the data set. Note that each of these three functions is *preservative*, which means Proposition 2 applies in each scenario we test here.

5.3 Precision and recall

In order to examine the quality of repairs, we compare, for each data set, the given relation R with a repair R^* and a gold standard R_{gold} and compute the following measures:

- Precision (P): the number of correctly repaired cells divided by the number of repaired cells.
- Recall (R): the number of correctly repaired cells divided by the number of erroneous cells.
- F -score (F): harmonic mean of precision and recall.

Hereby, a *repaired* cell is a cell for which the value is different in R^* as compared to R . A *correctly repaired* cell is a repaired cell where in addition the values are the same in R^* and R_{gold} . An *erroneous* cell is a cell for which the value is different in R_{gold} compared to R . As Llunatic uses a multi-sequence Chase method, it provides multiple repairs. We report here the results for the repair that has the lowest cost. Moreover, in the case of Llunatic, R^* may contain named variables (llun-values) that indicate that the cell should be assigned some constant value, but the value self remains unspecified. Llunatic offers the option to call on human intervention to fill in the correct constant. To report repair quality in the presence of variables, we propose to consider two extreme cases: one in which all variables are assigned the *correct* constant and one in which all variables are assigned an *erroneous* constant. These cases provide a best and worst outcome of precision, recall and F -value whenever all variables would be assigned with constant values at random. As such, for Llunatic, we will report *intervals* of repair quality rather than single values. Recent studies on Llunatic have used measures that correspond to our proposed upper bound [25]. However, the inclusion of the lower bound introduces an important dimension to the analysis of results in the sense that a large amount of variables implies a big gap between lower and upper bound. This is important to recognize as many variables imply many human interventions after the repair has been produced.

Table 4: Overview of precision (P), recall (R) and F -score obtained for Llunatic (two configurations) and Swipe (three repair functions). Best results are marked in bold font. n/a indicates the repair function was not applicable on that dataset.

	Data set	Llunatic (S)	Llunatic (FT)	Swipe (mv)	Swipe (wv)	Swipe (max)
P	Hospital	[0.93,0.96]	[0.11,0.12]	0.96	0.96	n/a
	Allergen	[0.00,0.57]	[0.64,0.64]	0.21	0.21	0.64
	Eudract	[0.80,0.84]	[0.80,0.84]	0.83	0.84	0.23
	Flight	[0.67,0.72]	[0.80,0.80]	0.71	0.79	0.56
R	Hospital	[0.83,0.85]	[0.29,0.31]	0.89	0.89	n/a
	Allergen	[0.00,0.54]	[0.30,0.30]	0.10	0.10	0.30
	Eudract	[0.31,0.32]	[0.31,0.32]	0.36	0.37	0.32
	Flight	[0.66,0.71]	[0.54,0.54]	0.68	0.79	0.58
F	Hospital	[0.88,0.90]	[0.16,0.17]	0.92	0.92	n/a
	Allergen	[0.00,0.55]	[0.41,0.41]	0.13	0.13	0.41
	Eudract	[0.44,0.47]	[0.44,0.47]	0.51	0.52	0.27
	Flight	[0.67,0.72]	[0.65,0.65]	0.69	0.79	0.57

Table 4 provides an overview of the repair quality obtained by the two configurations of Llunatic and the three configurations of Swipe on all data sets. For each combination of data set and measure of quality, we indicate the best performing approach in bold font. For recall and F -score on the Allergen data set, we did not indicate a best approach as it cannot be indicated with certainty. When inspecting the results in Table 4, several interesting observations can be made. A first observation is that results are sensitive to the exact configurations. This is especially true for the Allergen dataset, which provides allergen information for each product coming from only two sources. This low number of sources makes voting procedures less suitable. The resolution of conflicts by means of the max repair function, is a superior choice here. This behaviour is also

Table 5: Mean execution times (in seconds) of the repair methods.

Data set	Llunatic (S)	Llunatic (FT)	Swipe (best)
Hospital	20.31	167.09	0.20
Allergen	1.52	264.82	0.28
Eudract	19.38	98.09	7.15
Flight	910.05	3987.90	17.80

encoded in the FT scenario of Llunatic and it was confirmed that in this case, Llunatic and Swipe produced exactly the same results.

Interestingly, for other datasets, the voting scenarios of Swipe provide robust and good results. When a data set has no NULL values (Hospital and Allergen), both voting scenarios are equivalent. When NULL values are present (Eudract and Flight), weighted voting produces better results and turns out to be a good baseline choice. The max function provide good results only in particular scenarios like the Allergen dataset.

For Llunatic, the sensitivity to configuration choices seems to be higher than for Swipe. For the Hospital data set, there is a large difference in repair quality between the two configurations, most likely due to the usage of backward repairing in the second configuration. In general, we found that giving Llunatic the possibility to include backward repairs (this is the case in the fine-tuned configurations) seems to decrease repair quality on the data sets considered here. Generalizing this observation should be done with caution, but the results reported here make a strong case for at least questioning the usefulness of backward repair on real-life data sets.

A second observation is that when we compare the repair quality of Swipe with that of the configurations of Llunatic, it can be seen that Swipe produces repairs that are sometimes comparable and usually better than the best possible outcomes of Llunatic. The improvement in F -score is hereby mostly attributed to an improvement in recall, whereas precision behaves comparable. In this regard, it is interesting to see that Llunatic allows the usage of variables (i.e., lluns) when generating repairs. Opposed to that, repair functions used in the Swipe algorithm always produce constants. The results in Table 4 show that the more aggressive strategy where we always choose some constant is usually paying off. Moreover, it can be seen that the number of variables might become very high. This happens for example with the Allergen data set when using the standard configuration of Llunatic. A high number of variables can be problematic as it requires much human intervention in the repair process. This might however be a necessity in order to obtain good repairs. Indeed, we see that for the Allergen data set, the upper bound of recall and F -score are much higher for the standard configuration of Llunatic than for the other approaches.

An interesting result is found for the Eudract data set, which is the only data set where we encounter a partition class with more than 2 attributes. For this class with 3 attributes, the Swipe algorithm estimates the reliability of the involved attributes and uses that estimate to rank order FDs. Table 4 shows that using this order of FDs eventually produces a repair with $P = 0.84$ and $R = 0.37$. To investigate how good this order is, we repeated the experiment but forced Swipe to use a different priority model for this class. More precisely, we considered every possible priority order for the involved attributes in order to compute a best and worst case for precision and recall. We then found that precision ranged between 0.69 and 0.84 and recall ranged between 0.31 and 0.37. These additional results shows that the order in which we repair FDs can indeed make a large difference in the repair quality. However, the sequence based on estimated reliability leads to the best possible repair quality. This is true for the Eudract data set and also for the Hospital data set, where we have one partition class with 2 attributes. Also in the latter case, it is confirmed that the chosen order of FDs is the best one, although the difference in quality is less pronounced.

5.4 Scalability

Run time comparison In this section, we investigate the run time efficiency and scalability of Swipe. First, we measure the mean run time over five runs of Swipe and Llunatic on the different data sets and report the mean execution times (in seconds). We do not report other scalability parameters like CPU usage or main memory consumption, although we have investigated these parameters and came to similar conclusions. In case of Llunatic, we differentiate between the standard configuration and the configuration that was fine-tuned for each data set (Table 3). For Swipe, we took for each dataset, the best performing repair function. The results are reported in Table 5.

The results from Table 5 show that Swipe outperforms Llunatic on each data set in terms of execution time. This comes as no surprise as Swipe is designed explicitly as a simplification of the multi-sequence Chase algorithm that underlies Llunatic. Yet, we observe that the gain in execution time is considerable. Swipe is two to three orders of magnitude faster than the fine-tuned configuration of Llunatic. When we compare to the standard Llunatic configuration, the difference reduces to one order of magnitude. This provides us with some interesting insights. First of all, the cost manager of Llunatic does what it

is supposed to do: restricting repair operations to forward repair only and limiting the branching factor significantly reduces the execution time of Llunatic. However, Table 4 shows that this gain in computational efficiency also leads to lesser quality of the repairs. In that sense, the Swipe algorithm makes more considerate choices in reducing computational cost with better repair quality as a result.

Size of classes in \mathcal{P} The execution time of Swipe is strongly influenced by the size of classes in \mathcal{P} . For smaller partition classes we usually require less revisions. To that extent, we computed the sizes of the partition classes for each of the four data sets and came to the conclusion that over all data sets, there are 47 classes of which 45 were singletons. One class had size 2 (Hospital) and one class had size 3 (Eudract). This observation confirms that in real-life data sets, it is reasonable to apply the partitioning approach and to separate the repairs of different FDs.

Scalability in terms of $|R|$ and $|\mathcal{R}|$ In order to gain further understanding of the scalability of Swipe, we experimented with randomly generated data and constraints. These data and constraints are generated as explained in Section 5.1. The repair functions are defaulted to majority voting with random tie breaking for all attributes. We generate repairs of the given dirty data sets with Swipe and then measure the time to generate a repair in milliseconds. For a fixed number of tuples and fixed number of attributes, we repeat the procedure 10 times and report the mean of the measured run times.

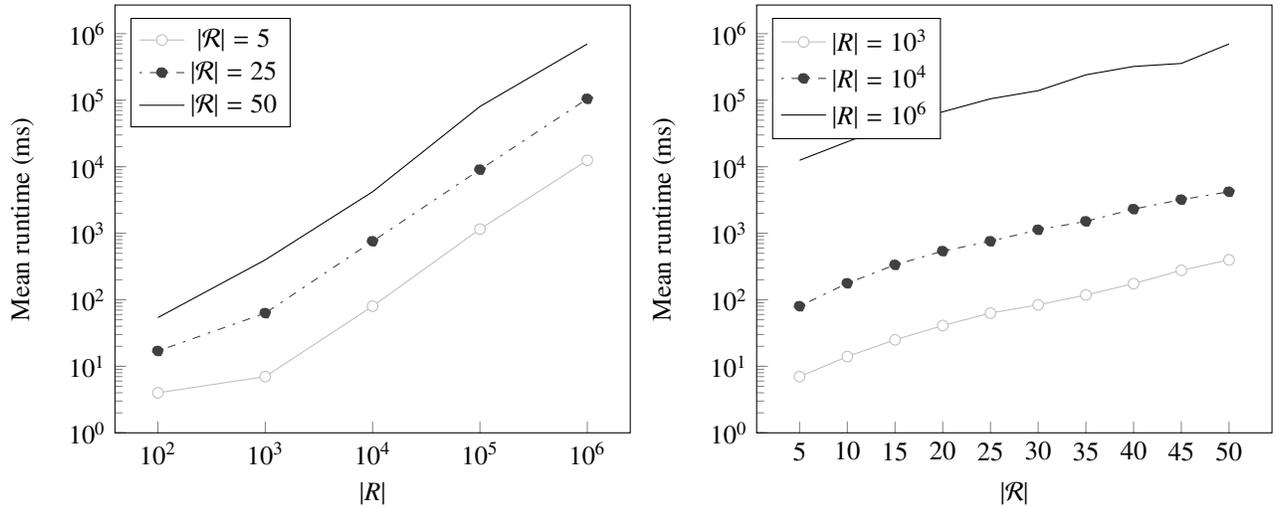


Figure 3: Mean run time (ms) of 10 executions of Swipe in function of changing $|R|$ (left) and $|\mathcal{R}|$ (right).

Figure 3 shows the obtained mean run times for a varying number of tuples (left) and a varying the number of attributes (right). More precisely, the left panel shows the evolution of the mean run time as the number of tuples increases with powers of 10, starting at $|R| = 10^2$ and ending at $|R| = 10^6$. This trend is shown three times: one time for a small schema ($|\mathcal{R}| = 5$), one time for a medium-sized schema ($|\mathcal{R}| = 25$) and one time for a large schema ($|\mathcal{R}| = 50$). These results show that the run time of Swipe scales linearly in terms of increasing $|R|$ for different schema sizes. Note in addition that the shift in run time from the small schema (5 attributes) to the large schema (50 attributes) is two orders of magnitude, suggesting quadratic behavior in terms of $|\mathcal{R}|$. This is confirmed in the right panel, that shows the evolution of the mean run time as the number of attributes increases with 5, starting at $|\mathcal{R}| = 5$ and ending at $|\mathcal{R}| = 50$. This trend is again shown three times: one time for a small relation ($|R| = 10^3$), one time for a medium-sized relation ($|R| = 10^4$) and one time for a large relation ($|R| = 10^6$). These results show that Swipe scales quadratic in terms of increasing $|\mathcal{R}|$ for different relation sizes. In addition, for an increase in $|R|$ with one order of magnitude, we observe shifts of one order of magnitude in run time, confirming the linear behaviour in terms of $|R|$. These results show that Swipe is very efficient in repairing large relations when the number of attributes is relatively small. For medium-sized schemata and number of FDs, Swipe can still repair relations relatively fast, especially if we compare to other state-of-the-art repair methods.

6 Conclusion

In this paper, we have introduced the Swipe algorithm to repair violations of functional dependencies (FDs). This algorithm is a degenerate variant of the Chase-based approach towards FD repairing. It hinges on two key principles. The first principle is

to use a partition of attributes that is forward repairable. We have provided an algorithm that, for a given set of FDs, constructs the most refined partition that meets this requirement of forward repairability. The second principle is that of priority repairing, which fixes the order in which FDs are treated. We have shown a simple heuristic to build such a priority model based on the estimated reliability of attributes. From a theoretical point of view, we have shown that Swipe is ensured to terminate and that there are easy to meet conditions under which unary FDs are revision-free. Empirical results show that Swipe provides an excellent trade-off between repair quality and computational efficiency. Future improvements of Swipe can focus on applying the principles developed here, to more expressive constraints like conditional functional dependencies.

References

- [1] ABITEBOUL, S., HULL, R., AND VIANU, V., Eds. *Foundations of Databases: The Logical Level*, 1st ed. Addison-Wesley Longman Publishing Co., Inc., 1995.
- [2] ARENAS, M., BERTOSSI, L., AND CHOMICKI, J. Consistent query answers in inconsistent databases. In *Proceedings of the Eighteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (New York, NY, USA, 1999), PODS '99, Association for Computing Machinery, p. 68–79.
- [3] BEERI, C., AND VARDI, M. A proof procedure for data dependencies. *Journal of the ACM* 31 (1984), 718–741.
- [4] BEN-AMRAM, A., AND YOFFE, S. A simple and efficient union-find-delete algorithm. *Theoretical Computer Science* 412 (02 2011), 487–492.
- [5] BESKALES, G., ILYAS, I., AND GOLAB, L. Sampling the repairs of functional dependency violations under hard constraints. *PVLDB* 3 (09 2010), 197–207.
- [6] BOECKLING, T., DE TRÉ, G., AND BRONSELAER, A. Cleaning data with selection rules. *IEEE Access* 10 (2022), 125212–125229.
- [7] BOHANNON, P., FAN, W., GEERTS, F., JIA, X., AND KEMENTSIETSIDIS, A. Conditional functional dependencies for data cleaning. In *Proceedings of the IEEE International Conference on Data Engineering* (Istanbul, Turkey, 2007), IEEE, pp. 746–755.
- [8] BOHANNON, P., FLASTER, M., FAN, W., AND RASTOGI, R. A cost-based model and effective heuristic for repairing constraints by value modification. In *SIGMOD Conference* (Baltimore, Maryland, USA, 2005), ACM, p. 143–154.
- [9] BOSKOVITZ, A. *Data Editing and Logic: The covering set method from the perspective of logic*. PhD thesis, The Australian National University, 2008.
- [10] BRONSELAER, A., AND ACOSTA, M. Parker: Data fusion through consistent repairs using edit rules under partial keys. *Information Fusion* 100 (2023), 101942.
- [11] BRONSELAER, ANTOON AND BOECKLING, TOON AND PATTYN, FILIP. Dynamic repair of categorical data with edit rules. *EXPERT SYSTEMS WITH APPLICATIONS* 201 (2022), 15.
- [12] BRONSELAER, ANTOON AND DE TRÉ, GUY. Aspects of object merging. In *North American Fuzzy Information Processing Society, 2010 Annual meeting* (Toronto, Canada, 2010), IEEE, pp. 27–32.
- [13] CHU, X., ILYAS, I., AND PAPOTTI, P. Holistic data cleaning: Putting violations into context. In *Proceedings - International Conference on Data Engineering* (Brisbane, QLD, Australia, 04 2013), IEEE, pp. 458–469.
- [14] CHU, X., MORCOS, J., ILYAS, I. F., OUZZANI, M., PAPOTTI, P., TANG, N., AND YE, Y. Katara: A data cleaning system powered by knowledge bases and crowdsourcing. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data* (Melbourne, Victoria, Australia, 2015), SIGMOD '15, Association for Computing Machinery, p. 1247–1261.
- [15] CONG, G., FAN, W., GEERTS, F., JIA, X., AND MA, S. Improving data quality: Consistency and accuracy. In *VLDB 2007* (Vienna, Austria, 2007), ACM, p. 315–326.
- [16] DALLACHIESA, M., EBAID, A., ELDAWY, A., ELMAGARMID, A., ILYAS, I. F., OUZZANI, M., AND TANG, N. Nadeef: A commodity data cleaning system. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data* (New York, New York, USA, 2013), SIGMOD '13, ACM, p. 541–552.

- [17] DE WAAL, T., PANNEKOEK, J., AND SCHOLTUS, S. *Handbook of statistical Data Editing and Imputation*. Wiley, 2001.
- [18] DONG, X. L., BERTI-EQUILLE, L., AND SRIVASTAVA, D. Integrating conflicting data: the role of source dependence. In *Proceedings of the VLDB conference* (Lyon, France, 2009), ACM, pp. 550–561.
- [19] DONG, X. L., SAHA, B., AND SRIVASTAVA, D. Less is more: Selecting sources wisely for integration. In *Proceedings of the VLDB conference* (Trento, Italy, 2012), VLDB Endowment, pp. 37–48.
- [20] EHRLINGER, L., AND WÖSS, W. A survey of data quality measurement and monitoring tools. *Frontiers in Big Data 5* (2022), 28.
- [21] FAN, W., AND GEERTS, F. *Foundations of Data Quality Management*. Morgan & Claypool Publishers, 2012.
- [22] FELLEGI, I., AND HOLT, D. A systematic approach to automatic edit and imputation. *Journal of the American Statistical Association 71*, 353 (1976), 17–35.
- [23] GALLER, B., AND FISCHER, M. An improved equivalence algorithm. *Communications of the ACM 7* (1964), 301–303.
- [24] GEERTS, F., MECCA, G., PAPOTTI, P., AND SANTORO, D. The LLUNATIC data-cleaning framework. *Proc. VLDB Endow.* 6, 9 (2013), 625–636.
- [25] GEERTS, F., MECCA, G., PAPOTTI, P., AND SANTORO, D. Cleaning data with Llunatic. *The VLDB Journal 29* (2019), 867–892.
- [26] GINSBURG, S., AND ZAIDAN, S. M. Properties of functional-dependency families. *Journal of the ACM 29*, 3 (1982), 678–698.
- [27] HEIDARI, A., MCGRATH, J., ILYAS, I. F., AND REKATSINAS, T. Holodetect: Few-shot learning for error detection. In *Proceedings of the 2019 International Conference on Management of Data* (New York, NY, USA, 2019), SIGMOD '19, Association for Computing Machinery, p. 829–846.
- [28] ILYAS, I. F., AND CHU, X. Trends in cleaning relational data: Consistency and deduplication. *Foundations and Trends in Databases 5*, 4 (2015), 281–393.
- [29] ILYAS, I. F., AND CHU, X. *Data Cleaning*. Association for Computing Machinery, New York, NY, USA, 2019.
- [30] KOLAH, S., AND LAKSHMANAN, L. V. S. On approximating optimum repairs for functional dependency violations. In *Proceedings of the 12th International Conference on Database Theory* (St. Petersburg, Russia, 2009), ICDT '09, Association for Computing Machinery, p. 53–62.
- [31] LI, X., DONG, X. L., LYONS, K., MENG, W., AND SRIVASTAVA, D. Truth finding on the deep web: Is the problem solved? In *Proceedings of the VLDB conference* (Trento, Italy, 2013), VLDB Endowment, pp. 97–108.
- [32] LIVSHITS, E., KIMELFELD, B., AND ROY, S. Computing optimal repairs for functional dependencies. In *Proceedings of the ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems* (Houston, Texas, USA, 2018), ACM, p. 225–237.
- [33] MAHDAVI, M., AND ABEDJAN, Z. Baran: Effective error correction via a unified context representation and transfer learning. In *Proceedings of the VLDB Endowment* (Tokio, Japan, 2020), vol. 13, VLDB Endowment, p. 1948–1961.
- [34] MAHDAVI, M., ABEDJAN, Z., CASTRO FERNANDEZ, R., MADDEN, S., OUZZANI, M., STONEBRAKER, M., AND TANG, N. Raha: A configuration-free error detection system. In *Proceedings of the 2019 International Conference on Management of Data* (Amsterdam, Netherlands, 2019), SIGMOD '19, Association for Computing Machinery, p. 865–882.
- [35] PHAM, M., KNOBLOCK, C. A., CHEN, M., VU, B., AND PUJARA, J. Spade: A semi-supervised probabilistic approach for detecting errors in tables. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21* (Montreal, Canada, 2021), International Joint Conferences on Artificial Intelligence Organization, pp. 3543–3551.
- [36] REKATSINAS, T., CHU, X., ILYAS, I., AND RÉ, C. Holoclean: Holistic data repairs with probabilistic inference. In *Proceedings of the VLDB Endowment* (Munich, Germany, 2017), VLDB Endowment, pp. 1190–1201.
- [37] REZIG, E. K., OUZZANI, M., AREF, W. G., ELMAGARMID, A. K., AND MAHMOOD, A. R. Pattern-driven data cleaning, 2017.
- [38] TARJAN, R. E. Efficiency of a good but not linear set union algorithm. *Journal of the ACM 22* (1975), 215–255.

- [39] TARJAN, R. E., AND VAN LEEUWEN, J. Worst-case analysis of set union algorithms. *Journal of the ACM* 31 (1984), 245–281.
- [40] WARREN, H. A modification of warshall’s algorithm for the transitive closure of binary relations. *Communications of the ACM* 18, 4 (1975), 218–220.
- [41] WIJSEN, J. Project-join-repair: An approach to consistent query answering under functional dependencies. In *Flexible Query Answering Systems* (Berlin, Heidelberg, 2006), H. L. Larsen, G. Pasi, D. Ortiz-Arroyo, T. Andreasen, and H. Christiansen, Eds., Springer Berlin Heidelberg, pp. 1–12.

A Proofs

Proof 1 (Proof of Proposition 1) If \mathcal{P} satisfies the condition from Definition 1, we can assign, for each $C_i \in \mathcal{P}$ with partial repair R_{i-1}^* and for each $\phi \in \Phi[\mathcal{R}_i] \setminus \Phi[\mathcal{R}_{i-1}]$, equal values to RHS(ϕ) in order to satisfy ϕ .

Proof 2 (Proof of Theorem 1) We show first that \mathcal{P} allows sequential forward repairability for Φ (i) and then that \mathcal{P} cannot be refined without losing this property (ii).

(i) If \mathcal{P} is induced by P^+ for a given set Φ , then consider an arbitrary class C_i from \mathcal{P} and any FD $\phi \in \Phi[\mathcal{R}_i] \setminus \Phi[\mathcal{R}_{i-1}]$. If $\text{RHS}(\phi) \notin C_i$, then \mathcal{P} cannot be induced by P^+ because it would violate either Eq. (3) or Eq. (4). It follows that we must have $\text{RHS}(\phi) \in C_i$ and thus $\mathcal{P} \rightsquigarrow_F \Phi$.

(ii) Suppose there is a refinement \mathcal{P}' of \mathcal{P} for which $\mathcal{P}' \rightsquigarrow_F \Phi$. This would imply that there is a class from \mathcal{P} that can be split into two disjoint subclasses without breaking forward repairability. In turn, this would imply that there is an equivalence class from \equiv_{P^+} that can be split into two subclasses. Let us denote the equivalence relation in which this split is done by $\equiv_{P'}$. Now, if $\equiv_{P'}$ still contains P then P^+ cannot be the transitive closure of P and we have a contradiction. Alternatively, if $\equiv_{P'}$ does not contain P , then there is at least one $(b, a) \in P$ that is not accounted for and \mathcal{P}' is not forward repairable. In both cases, we obtain a conclusion that is in contradiction with the premise and it follows no refinement of \mathcal{P} is forward repairable.

Proof 3 (Proof of Theorem 2) We first prove that Algorithm 3 terminates (i) and then that after termination we have obtained a partial repair R_i^* (ii).

(i) To see that **priorityRepair** (R_{i-1}^*, Φ_i, C_i) terminates, note that each time an FD ϕ is considered, either DSF(RHS(ϕ)) does not change or its number of classes decreases. If for all $a \in C_i$, DSF(a) does not change anymore, the algorithm terminates. Else, we must reach a point in which each DSF(a) has only one class. In that case, for each $a \in C_i$, each tuple in R_{i-1}^* gets the same value for a . It follows that in this case, there are no more violations, from which it follows that all DSF(a) remain unchanged and the algorithm stops.

(ii) Each time an FD $X \rightarrow a$ is polled from the stack, the fix step updates DSF(a) and potentially changes R_{i-1}^* in attribute a . In the revision step that follows, any FD $\phi' \in \Phi_i$ that is not on the stack, was polled and repaired before and there are three options. First, if $a \notin \text{LHS}(\phi')$ and $a \neq \text{RHS}(\phi')$, then clearly this FD is still satisfied. Second, if $a = \text{RHS}(\phi')$, then because DSF(a) changes only by merging classes, any two tuples with the same value for LHS(ϕ') are still in the same class in DSF(a) and thus will also have the same value after fixing $X \rightarrow a$. Hence, ϕ' remains satisfied. Third, if $a \in \text{LHS}(\phi')$, we put ϕ' back on the stack for revision. It follows that after application of **fix** $(R_{i-1}^*, X \rightarrow a, \text{DSF}(a))$, any FD from Φ_i not on the stack is currently satisfied by R_{i-1}^* . At the same time, at termination time, the stack is empty and thus all FDs are satisfied.

Proof 4 (Proof of Proposition 2) For an FD $a' \rightarrow a$ during a revision step, we can make a sequence of observations:

1. If $a' \notin C_i$ then we have $\forall \phi \in \Phi_i : \text{RHS}(\phi) \neq a'$ and consequently $a' \rightarrow a$ is never considered during any revision step. We therefore assume that $a' \in C_i$.
2. $a' \rightarrow a$ is considered in a revision step only if it is not an element of \mathbb{S} . This means it was fixed already and thus that any two tuples with equal values for a' are in the same class of DSF(a).
3. $a' \rightarrow a$ is considered in a revision step if the preceding fix step involves an FD $X \rightarrow a'$ and is therefore of the form **fix** $(R_{i-1}^*, X \rightarrow a', \text{DSF}(a'))$. In the remainder of this proof, we denote the partial repair before this fix step as R_{before}^* and after this fix step as R_{after}^* . Clearly, $R_{\text{after}}^* \models X \rightarrow a'$.
4. From (1) and (2) we have that $X \rightarrow a'$ mentioned in (3) is not a pilot FD. Thus, FDs are polled from the stack in order induced by $>$.
5. The fix step in (3) is preceded by a poll of $X \rightarrow a'$, so (2) implies $a > a'$.
6. From (2) and (5) we have that $R_{\text{before}}^* \models a' \rightarrow a$. In addition, transitivity of FDs implies $R_{\text{before}}^* \models X \rightarrow a$.

Suppose now R_{after}^* fails $a' \rightarrow a$, then there must exist two rows $r_1 \in R_{\text{before}}^*$ and $r_2 \in R_{\text{before}}^*$ that are transformed into $r_1^* \in R_{\text{after}}^*$ and $r_2^* \in R_{\text{after}}^*$, respectively, for which we have

$$r_1^*[a'] = r_2^*[a'] \wedge r_1^*[a] \neq r_2^*[a].$$

In addition, we know that:

$$r_1^*[a'] \neq r_1[a'] \vee r_2^*[a'] \neq r_2[a']$$

because otherwise $r_1[a'] = r_1^*[a'] = r_2^*[a'] = r_2[a']$ from which $R_{\text{before}}^* \models a' \rightarrow a$ that $r_1^*[a] = r_2^*[a]$ and this contradicts our construction of r_1 and r_2 . We can then distinguish between two cases.

[1] If $r_1^*[X] = r_2^*[X]$ then we have that $R_{\text{after}}^* \not\models X \rightarrow a$ and because $a' \notin X$, we have $R_{\text{before}}^*[X] = R_{\text{after}}^*[X]$. As such, it follows that $R_{\text{before}}^* \not\models X \rightarrow a$, which contradicts with (6).

[2] If $r_1^*[X] \neq r_2^*[X]$ then because $\rho_{a'}$ is preservative, there must exist rows $r_3 \in R_{\text{before}}^*$ and $r_4 \in R_{\text{before}}^*$ such that on one hand:

$$r_1[X] = r_3[X] \wedge r_1^*[a'] = r_3^*[a'] = r_3[a']$$

and on the other hand

$$r_2[X] = r_4[X] \wedge r_2^*[a'] = r_4^*[a'] = r_4[a'].$$

That is, r_1 and r_2 received their values for a' from respectively r_3 and r_4 and this is only possible if r_1 and r_3 have the same value for X and r_2 and r_4 have the same value for X . Since we assumed $r_1^*[a'] = r_2^*[a']$, it follows that $r_3[a'] = r_4[a']$ and because $R_{\text{before}}^* \models a' \rightarrow a$ we must also have $r_3[a] = r_4[a]$. Finally, since $R_{\text{before}}^* \models X \rightarrow a$, we also find that $r_1[a] = r_3[a]$ and $r_2[a] = r_4[a]$ by which we find that $r_1[a] = r_2[a]$. Again, this contradicts our construction of r_1 and r_2 .