

Privacy-Aware Data Cleaning-as-a-Service

Yu Huang^a, Mostafa Milani^b, Fei Chiang^a

^aDept. of Computing and Software, McMaster University, Hamilton, Canada

^bDept. of Computer Science, Western University, London, Canada

Abstract

Data cleaning is a pervasive problem for organizations as they try to reap value from their data. Recent advances in networking and cloud computing technology have fueled a new computing paradigm called *Database-as-a-Service*, where data management tasks are outsourced to large service providers. In this paper, we consider a *Data Cleaning-as-a-Service* model that allows a client to interact with a data cleaning provider who hosts curated, and sensitive data. We present *PACAS*: a Privacy-Aware data Cleaning-As-a-Service model that facilitates interaction between the parties with client query requests for data, and a service provider using a data pricing scheme that computes prices according to data sensitivity. We propose new extensions to the model to define *generalized data repairs* that obfuscate sensitive data to allow data sharing between the client and service provider. We present a new semantic distance measure to quantify the utility of such repairs, and we re-define the notion of consistency in the presence of generalized values. The PACAS model uses (X,Y,L)-anonymity that extends existing data publishing techniques to consider the semantics in the data while protecting sensitive values. Our evaluation over real data show that PACAS safeguards semantically related sensitive values, and provides lower repair errors compared to existing privacy-aware cleaning techniques.

1. Introduction

Data cleaning is a pervasive problem motivated by the fact that real data is rarely error free. Organizations continue to be hindered by poor data quality as they wrangle with their data to extract value. Recent studies estimate that up to 80% of the data analysis pipeline is consumed by data preparation tasks such as data cleaning [1]. A wealth of data cleaning solutions have been proposed to reduce this effort: constraint based cleaning that use dependencies as a benchmark to repair data values such that the data and dependencies are consistent [2], statistical based cleaning, which propose updates to the data according to expected statistical distributions [3], and leveraging master data as a source of ground truth [4].

Recent advances in networking and cloud infrastructure have motivated a new computing paradigm called *Database-as-a-Service* that lowers the cost and increases access to a suite of data management services. A *service provider* provides the necessary hardware and software platforms to support a variety of data management tasks to a *client*. Companies such as Amazon, Microsoft, and Google each provide storage platforms, accelerated computational capacity, and advanced data

Email addresses: huang223@mcmaster.ca (Yu Huang), mostafa.milani@uwo.ca (Mostafa Milani), fchiang@mcmaster.ca (Fei Chiang)

analytics services. However, the adoption of data cleaning services has been limited due to privacy restrictions that limit data sharing. Recent data cleaning efforts that use a curated, master data source share a similar service model to provide high quality data cleaning services to a client with inconsistent data [4, 5]. However, these techniques largely assume the master data is widely available, without differentiating information sensitivity among the attribute values.




Figure 1: (a) DGH^{med} and (b) VGH^{med} .




Figure 2: (a) DGH^{age} and (b) VGH^{age} .

Example 1.1. A data broker gathers longitudinal data from hospital and doctors' records, prescription and insurance claims. Aggregating and curating these disparate datasets lead to a valuable commodity that clients are willing to pay for gleaned insights, and to enrich and clean their individual databases. Table 1 shows the curated data containing patient demographic, diagnosis and medication information. The schema consists of patient gender (**GEN**), age (**AGE**), zip code (**ZIP**), diagnosed illness (**DIAG**), and prescribed medication (**MED**).

A client such as a pharmaceutical firm, owns a stale and inconsistent subset of this data as shown in Table 2. For example, given a functional dependency (FD) $\varphi : [\text{GEN}, \text{DIAG}] \rightarrow [\text{MED}]$ on Table 2, it states that a person's gender and diagnosed condition determine a prescribed medication. That is, for any two tuples in Table 2, if they share equal values in the $[\text{GEN}, \text{DIAG}]$ attributes, then they should also have equal values in the **MED** attribute. We see that tuples $t_1 - t_5$ falsify φ . Error cells that violate φ are highlighted in light and dark gray, and values in bold are inaccurate according to Table 1.

If the client wishes to clean her data with the help of a data cleaning service provider (i.e., the data broker and its data), she must first match the inconsistent records in Table 2 against the curated records in Table 1. This generates possible fixes (also known as repairs) to the data in Table 2. The preferred repair is to update $t_2[\text{MED}]$, $t_3[\text{MED}]$ to *ibuprofen* (from m_6), and $t_4[\text{DIAG}] = \textit{migraine}$ (from m_3). However, this repair discloses sensitive patient information about diagnosed illness and medication from the service provider. It may be preferable to disclose a more general value that

ID	GEN	AGE	ZIP	DIAG	MED
m_1	male	51	P0T2T0	osteoarthritis	ibuprofen
m_2	female	45	P2Y9L8	tendinitis	addaprin
m_3	female	32	P8R2S8	migraine	naproxen
m_4	female	67	V8D1S3	ulcer	tylenol
m_5	male	61	V1A4G1	migraine	dolex
m_6	male	79	V5H1K9	osteoarthritis	ibuprofen

Table 1: Curated medical records (R_{SP})

ID	GEN	AGE	DIAG	MED
t_1	male	51	osteoarthritis	ibuprofen
t_2	male	79	osteoarthritis	intropes
t_3	male	45	osteoarthritis	addaprin
t_4	female	32	ulcer	naproxen
t_5	female	67	ulcer	tylenol
t_6	male	61	migrane	dolex
t_7	female	32	pylorospasm	appaprtin
t_8	male	37	hypertension	dolex

Table 2: Dirty client records w.r.t. φ .

ID	GEN	AGE	ZIP	DIAG	MED
g_1	*	[31,60]	P*	osteoarthritis	ibuprofen
g_2	*	[31,60]	P*	tendinitis	addaprin
g_3	*	[31,60]	P*	migraine	naproxen
g_4	*	[61,90]	V*	ulcer	tylenol
g_5	*	[61,90]	V*	migraine	dolex
g_6	*	[61,90]	V*	osteoarthritis	ibuprofen

Table 3: Public table.

is semantically similar to the true value to protect individual privacy. For example, instead of disclosing the medication *ibuprofen*, the service provider discloses *Non-steroid anti-inflammatory drug (NSAID)*, which is the family of drugs containing *ibuprofen*. In this paper, we explore how to compute such generalized repairs in an interactive model between a service provider and a client to protect sensitive data and to improve accuracy and consistency in client data.

State-of-the-Art. Existing work in data privacy and data cleaning have been limited to imputation of missing values using decision trees [6], or information theoretic techniques [7], and studying trade-offs between privacy bounds and query accuracy over differentially private relations [8]. In the *data cleaning-as-a-service* setting, which we consider in this paper, using differential privacy poses the following limitations: (i) differential privacy provides guarantees assuming a limited number of interactions between the client and service provider; (ii) queries are limited to aggregation queries; and (iii) data randomization decreases data utility.

Given the above limitations, we explore the use of *Privacy Preserving Data Publishing* (PPDP) techniques, that even though do not provide the same provable privacy guarantees as differential

privacy, do restrict the disclosure of sensitive values without limiting the types of queries nor the number of interactions between the client and service provider. PPDP models prevent re-identification and break attribute linkages in a published dataset by hiding an individual record among a group of other individuals. This is done by removing identifiers and generalizing *quasi-identifier* (QI) attributes (e.g., GEN, AGE, ZIP) that together can re-identify an individual. Well-known PPDP methods such as *k-anonymity* require the group size to be at least k such that an individual cannot be identified from $k - 1$ other individuals [9, 10]. Extensions include (X, Y) -*anonymity* that break the linkage between the set X of QI attributes, and the set Y of sensitive attributes by requiring at least k distinct sensitive values for each unique X [11]. For example, Table 3 is k -anonymous for $k = 3$ by generalizing values in the QI attributes. It is also (X, Y) -anonymous for sensitive attribute MED for $k = 3$ since there are three distinct medications for each value in the QI attributes X , e.g., values $(*, [31, 60], P*)$ of X co-occur with *ibuprofen*, *addaprin*, and *naproxen* of Y .

To apply PPDP models in data cleaning, we must also address the highly contextualized nature of data cleaning, where domain expertise is often needed to interpret the data to achieve correct results. It is vital to incorporate these domain semantics during the cleaning process, and into a privacy model during privacy preservation. Unfortunately existing PPDP models only consider syntactic forms of privacy via generalization and suppression of values, largely ignoring the data semantics. For example, upon closer inspection of Table 3, the values in the QI attributes in records $g_1 - g_3$ are associated with the same family of medication, since *ibuprofen*, *addaprin*, and *naproxen* all belong to the NSAID class, which are *analgesic* drugs. In past work, we introduced a privacy-preserving framework that uses an extension of (X, Y) -anonymity to incorporate a *generalization hierarchy*, which capture the semantics of an attribute domain¹ [12]. In Table 3, the medications in $g_1 - g_3$ are modeled as synonyms in such a hierarchy. In this paper, we extend our earlier work [12] to define a semantic distance metric between values in a generalization hierarchy, and to incorporate generalized values as part of the data repair process.

Example 1.2. To clean Table 2, a client requests the correct value(s) in tuples $t_1 - t_3$ from the data cleaning service provider. The service provider returns *ibuprofen* to the client to update $t_2[\text{MED}]$ and $t_3[\text{MED}]$ to $m_6[\text{MED}] = \text{ibuprofen}$. However, disclosing the specific *ibuprofen* medication violates the requirements of the underlying privacy model. Hence, the service provider must determine when such cases arise, and offer a revised solution to disclose a less informative, generalized value, such as *NSAID* or *analgesic*. The model must provide a mechanism for the client and service provider to negotiate such a trade-off between data utility and data privacy.

Technical Challenges.

1. Data cleaning with functional dependencies (FDs), and record matching (between the service provider and client) is an NP-complete problem. In addition, it is approximation-hard, i.e., the problem cannot be approximately solved with a polynomial-time algorithm and a constant approximation ratio [4]. We extend this data cleaning problem with privacy restrictions on the service provider, making it as hard as the initial problem. In our earlier work, we analyzed the complexity of this problem, and proposed a data cleaning framework realizable in practice [12]. In this paper, we extend this framework to consider a new type of repair with generalized values that further protects sensitive data values.

¹ k -anonymity and extensions, e.g., l -diversity and t -closeness also do not consider the underlying data semantics.

2. Given a generalization hierarchy as shown in Figures 1 and 2, proposed repairs from a service provider may contain general values that subsume a set of specific values at the leaf level. In such cases, we study and propose a new (repair) semantics for data consistency with respect to (w.r.t.) an FD involving general values.
3. By proposing generalized values as repair values, we lose specificity in the client data instance. We study the trade-off between generalized versus specific repair values. Resolving inconsistencies w.r.t. a set of FDs requires traversing the space of possible fixes (which may now include general values). Quantifying the individual utility of a generalized value as a repair value, while respecting data cleaning budgets from the client is our third challenge.

Our Approach and Contributions. We build upon our earlier work that introduced *PACAS*, a Privacy-Aware data Cleaning-As-a-Service framework that facilitates data cleaning between a client and a service provider [12]. The interaction is done via a data pricing scheme where the service provider charges the client for each disclosed value, according to its adherence to the privacy model. *PACAS* includes a new privacy model that extends (X,Y)-anonymity to consider the data semantics, while providing stronger privacy protection than existing PPDP methods. In this work, we present extensions to *PACAS* to resolve errors (FD violations) by permitting updates to the data that are generalizations of the true value to avoid disclosing details about sensitive values. Our goal is to provide repair recommendations that are semantically equivalent to the true value in the service provider data. We make the following contributions:

1. We extend *PACAS* the existing privacy-preserving, data cleaning framework that identifies FD errors in client data, and allows the client to purchase clean, curated data from a service provider [12]. We extend the repair semantics to introduce *generalized repairs* that update values in the client instance to generalized values.
2. We re-define the notion of consistency between a relational instance (with generalized values) and a set of FDs. We propose an entropy-based measure that quantifies the semantic distance between two values to evaluate the utility of repair candidates.
3. We extend our *SafeClean* algorithm that resolves FD errors by using external data purchased from a service provider. *SafeClean* proposes repairs to the data by balancing its data privacy requirements against satisfying query purchase requests for its data at a computed price [12]. Given a cleaning budget, we present a new budget allocation algorithm that improves upon previous, fixed allocations, to consider allocations according to the number of errors in which a database value participates.
4. We evaluate the effectiveness of *SafePrice* and *SafeClean* over real data showing that by controlling data disclosure via data pricing, we can effectively repair FD violations while guaranteeing (X,Y,L)-anonymity. We also show that our *SafeClean* achieves lower repair error than comparative baseline techniques.

We provide notation and preliminaries in Section 2. In Section 3, we define generalized relations, and their consistency w.r.t. a set of FDs, and present a new semantic distance function. We present our problem definition, and review the *PACAS* system in Section 4. We discuss how data privacy is enforced via data pricing in Section 5, and describe new extensions to our data cleaning algorithm that consider generalized values in Section 6. We present our experimental results in Section 7, related work in Section 8, and conclude in Section 9.

Table 4: Summary of notation and symbols.

Symbol	Description
R, \mathcal{R}	relation and relational schema
A, B	relational attributes
X, Y, Z	sets of relational attributes
D, \mathfrak{R}	database instance and database schema
φ, ϕ	functional dependency and matching dependency
$\text{Dom}^A, \text{dom}^A$	domain of attribute A
$\text{dom}^A(l)$	sub-domain of attribute A in level l
$\text{DGH}^A, \text{VGH}^A$	domain and value generalization hierarchies
\leq	generalization relation for levels
\preceq	generalization relation for values, tuples, and relations
Q, G	Simple query and generalized query (GQ)
l, L	level and sequence of levels
$\mathcal{S}, \mathcal{C}^Q$	support set, and conflict set
\mathcal{B}, B_i	total and the budget for the i -th iteration
l_{\max}	generalization level
c, e	database cell and error cell
δ	distance function for values, tuples and relations

2. Preliminaries

2.1. Relations and Dependencies

A *relation* (table) R with a *schema* $\mathcal{R} = \{A_1, \dots, A_n\}$ is a finite set of n -ary tuples $\{t_1, \dots, t_N\}$. A *database* (*instance*) D is a finite set of relations R_1, \dots, R_m with schema $\mathfrak{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_m\}$. We denote by small letters x, y, z as variables. Let A, B, C refer to single attributes and X, Y, Z as sets of attributes. A *cell* $c = t[A_i]$ is the i -th position in tuple t with its value denoted by $c.\text{value}$. We use c to refer to the value $c.\text{value}$ if it is clear from the context. A *functional dependency* (FD) φ over a relation R with set of attributes \mathcal{R} is denoted by $\varphi : X \rightarrow Y$, in which X and Y are subsets of \mathcal{R} . We say φ holds over R , $R \models \varphi$, if for every pair of tuples $t_1, t_2 \in R$, $t_1[X] = t_2[X]$ implies $t_1[Y] = t_2[Y]$. Table 4 summarizes our symbols and notations.

A *matching dependency* (MD) ϕ over two relations R and R' with schemata $\mathcal{R} = \{A_1, A_2, \dots\}$ and $\mathcal{R}' = \{A'_1, A'_2, \dots\}$ is an expression of the following form:

$$\bigwedge_{i \in [1, n]} R[X_i] \approx R'[X'_i] \rightarrow R[Y] \Leftarrow R'[Y'], \quad (1)$$

where (X_i, X'_i) and (Y, Y') are comparable pairs of attributes in R and R' . The MD ϕ states that for a pair of tuples (t, t') with $t \in R$ and $t' \in R'$, if $t[X'_i]$ values are similar to values $t'[X'_i]$ according to the similarity function \approx , the values of $t[Y]$ and $t'[Y']$ are identical [5].

2.2. Generalization

Generalization replaces values in a private table with less specific, but semantically consistent values according to a generalization hierarchy. To generalize attribute A , we assume a set of levels $\mathcal{L}^A = \{l_0^A, \dots, l_h^A\}$ and a partial order \leq^A , called a *generalization relation* on \mathcal{L}^A . Levels l_i^A are assigned with disjoint domain-sets $\text{dom}(l_i^A)$. In \leq^A , each level has at most one parent. The

domain-set $\text{dom}(l_n^A)$ is the maximal domain set and it is a singleton, and $\text{dom}(l_0^A)$ is the *ground domain set*. The definition of \leq^A implies the existence of a totally ordered hierarchy, called the *domain generalization hierarchy*, DGH^A . The domain set $\text{dom}(l_i^A)$ generalizes $\text{dom}(l_j^A)$ iff $l_j^A \leq l_i^A$. We use h^A to refer to the number of levels in DGH^A . Figures 1(a) and 2(a) show the DGH s for the medication and age attributes, resp. A *value generalization relationship* for A , is a partial order \preceq^A on $\text{Dom}^A = \bigcup \text{dom}(l_i^A)$. It specifies a *value generalization hierarchy*, VGH^A , that is a tree whose leaves are values of the ground domain-set $\text{dom}(l_0^A)$ and whose root is the single value in the maximal domain-set $\text{dom}(l_n^A)$ in DGH^A . For two values v and v' in Dom^A , $v' \preceq^A v$ means v' is more specific than v according to the VGH. We use \preceq rather than \preceq^A when the attribute is clear from the context. The VGH for the MED and AGE attributes are shown in Figures 1(b) and 2(b), respectively. According to VGH of MED, $\text{ibuprofen} \preceq \text{NSAID}$.

A value is *ground* if there is no value more specific than it, and it is *general* if it is not ground. In Figure 1, *ibuprofen* is ground and *analgesic* is general. For a value v , its base denoted by $\text{base}(v)$ is the set of ground values u such that $u \preceq v$. We use $0 \leq \text{level}(v) \leq h^A$ to refer to the level of v according to VGH^A .

A *general relation* (table) is a relation with some general values and a ground relation has only ground values. A *general database* is a database with some general relations and a ground database has only ground relations. The generalization relation \preceq^A trivially extends to tuples. We give an extension of the generalization relation to general relations and databases in Section 3.

The generalization hierarchies (DGH and VGH) are either created by the data owners with help from domain experts or generated automatically based on data characteristics [13]. The automatic generation of hierarchies for categorical attributes [14] and numerical attributes [15, 16, 17] apply techniques such as histogram construction, binning, numeric clustering, and entropy-based discretization [13].

2.3. Generalized Queries

We review *generalized queries* (GQs) to access values at different levels of the DGH [12].

Definition 1 (Generalized Queries). A generalized query (GQ) with schema \mathcal{R} is a pair $G = \langle Q, L \rangle$, where Q is an n -ary select-projection-join query over \mathcal{R} , and $L = \{l_1, \dots, l_n\}$ is a set of levels for each of the n values in Q according to the DGHs in \mathcal{R} . The set of answers to G over R , denoted as $G(R)$, contains n -ary tuples t with values at levels in $l_i \in L$, such that $\exists t' \in Q(R)$ and t generalizes t' , i.e. $t' \preceq t$.

Intuitively, answering a GQ involves finding the answers of $Q(R)$, and then generalizing these values to levels that match L . For a fixed size DGH, the complexity of answering G is the same as answering Q .

Example 2.1. Consider a GQ $G = \langle Q, L \rangle$ with level $L = \{l_0^{\text{GEN}}, l_1^{\text{MED}}\}$ and query $Q(R_{SP}) = \Pi_{\text{GEN}, \text{MED}}(\sigma_{\text{DIAG}=\text{migraine}}(R_{SP}))$ posed on relation R_{SP} is Table 1. The query Q requests the gender and medication of patients with a migraine. The answers to $Q(R_{SP})$ are $\{(female, \text{naproxen}), (male, \text{dolex})\}$. The GQ G asks for the same answers but at the levels specified by L . Therefore the answers to G are $\{(female, \text{NSAID}), (male, \text{acetaminophen})\}$, which are generalized according to L and Figure 1.

2.4. Privacy-Preserving Data Publishing

The most well-known PPDP privacy model is k -anonymity that prevents re-identification of a single individual in an anonymized data set [9, 10].

Definition 2 (X-group and k -anonymity). A relation R is k -anonymous if every QI-group has at least k tuples. An X -group is a set of tuples with the same values in X .

As an example, Table 3 has two QI-groups, $\{g_1, g_2, g_3\}$ and $\{g_4, g_5, g_6\}$, and it is k -anonymous with $k = 3$. k -anonymity is known to be prone to attribute linkage attacks where an adversary can infer sensitive values given QI values. Techniques such as (X,Y)-anonymity aim to address this weakness by ensuring that for each QI value in X , there are at least k different values of sensitive values in attribute(s) Y in the published (public) data [18].

Definition 3 ((X,Y)-anonymity). A table R with schema \mathcal{R} and attributes $X, Y \subseteq \mathcal{R}$ is (X,Y)-anonymous with value k if for every $t \in R$, there are at least k values in $Q^t(R)$, $|Q^t(R)| \geq k$, where $Q^t(R) = \Pi_Y(\sigma_{X=t[X]}(R))$.

The query $Q^t(R)$ in Definition 3 returns distinct values of attributes Y that appear in R with the values $t[X]$. Therefore, if the size of $Q^t(R)$ is greater than k for every tuple t , that means every values of X are linked with at least k values of Y . Note that k -anonymity is a special case of (X,Y)-anonymity when X is the set of QI attributes and Y is the set of sensitive attributes that are also a key in R [18].

Example 2.2. For $X = \{\text{GEN}, \text{AGE}, \text{ZIP}\}$, $Y = \{\text{MED}\}$ with $k = 3$, Table 3 is (X,Y)-anonymous since each $g_i \in R$, $Q^{g_i}(R)$ is either $\{\text{ibuprofen}, \text{addaprin}, \text{naproxen}\}$ or $\{\text{tylenol}, \text{dolex}, \text{ibuprofen}\}$, which means the values of X in each tuple g_i are linked to $k = 3$ distinct values of Y ($Q^{g_i}(R)$ represents the set of distinct values of Y that are linked to the values of $g_i[X]$ in R). Table 3 is not (X,Y)-anonymous with $X = \{\text{DIAG}\}$ and $Y = \{\text{MED}\}$, since for g_1 , $Q^{g_1}(R) = \{\text{ibuprofen}\}$ with size $1 \leq k$.

The (X,Y)-anonymity model extends k -anonymity; if Y is a key for R and X, Y are QI and sensitive attributes, respectively, then (X,Y)-anonymity reduces to k -anonymity.

The l -diversity privacy model extends k -anonymity with a stronger restriction over the X -groups [19]. A relation is considered l -diverse if each X -group contains at least l “well-represented” values in the sensitive attributes Y . Well-representation is normally defined according to the application semantics, e.g., *entropy l -diversity* requires the entropy of sensitive values in each X -group to satisfy a given threshold [19]. When well-representation requires l sensitive values in each Y attribute, (X,Y)-anonymity reduces to l -diversity.

2.4.1. (X,Y,L)-anonymity.

In earlier work, we introduced (X,Y,L)-anonymity, which extends (X,Y)-anonymity to consider the semantic closeness of values [12]. Consider the following example showing the limitations of (X,Y)-anonymity.

Example 2.3. (X,Y)-anonymity prevents the linkage between sets of attributes X and Y in a relation R by making sure that every values of Y appears with at least k distinct values of X in R . For example, Table 3 is (X,Y)-anonymous with $X = \{\text{GEN}, \text{AGE}, \text{ZIP}\}$, $Y = \{\text{MED}\}$, and $k = 3$ because the values of X appear with three different values; $(*, [31, 60], P*)$ co-occurs with *ibuprofen*, *addaprin*, *naproxen* and $(*, [61, 90], V*)$ appears with *tylenol*, *dolex*, *ibuprofen*.

(X,Y)-anonymity ignores how close the values of Y are according to their VGHs. For example, although *ibuprofen*, *addaprin*, *naproxen* are distinct values, they are similar medications with the

same parent NSAID according to the VGH in Figure 1. This means the value $(*, [31, 60], P*)$ of attributes X is linked to NSAID which defies the purpose of (X, Y) -anonymity.

(X, Y, L) -anonymity resolves this issue by adding a new parameter L for levels of Y . In (X, Y, L) -anonymity, every values of X has to appear with at least k values of Y in the levels of L . For example, Table 3 is not (X, Y, L) -anonymous if $L = \{l_1^{\text{MED}}\}$ because $(*, [31, 60], P*)$ appears with Y values that all roll up to NSAID at level l_1^{MED} . Thus, to achieve (X, Y, L) -anonymity we need to further suppress values in Table 3 to prevent the linkage between $(*, [31, 60], P*)$ and NSAID.

Definition 4 ((X,Y,L)-anonymity). Consider a table R with schema \mathcal{R} and attributes $X, Y \subseteq \mathcal{R}$, and a set of levels L corresponding to attribute DGHs from Y . R is (X, Y, L) -anonymous with value k if for every $t \in R$, there are at least k values in $G^t(R)$, where $G^t = \langle Q^t, L \rangle$ is a GQ with $Q^t(R) = \Pi_Y(\sigma_{X=t[X]}(R))$.

In Definition 4, the GQ $G^t(R)$ returns the set of Y values in levels L that appear with the X values of the tuples t . If the size of $G^t(R)$ is greater than k , then every value of X appears with at least k values of Y in the level L , which is the objective of (X, Y, L) -anonymity, as shown in Example 2.3.

2.5. Data Pricing

High quality data is a valuable commodity that has lead to increased purchasing and selling of data online. Data market services provide data across different domains such as geographic and locational (e.g., AggData [20]), advertising (e.g., Oracle [21], Acxiom [22]), social networking (e.g., Facebook [23], Twitter [24]), and people search (e.g. Spokeo [25]). These vendors have become popular in recent years, and query pricing has been proposed as a fine-grained and user-centric technique to support the exchange and marketing of data [26].

Given a database instance D and query Q , a *pricing function* returns a non-negative real number representing the price to answer Q [26]. A pricing function should have the desirable property of being *arbitrage-free* [26]. Arbitrage is defined using the concept of *query determinacy* [27]. Intuitively, Q determines Q' if for every database D , the answers to Q' over D can be computed from the answers to Q over D . Arbitrage occurs when the price of Q is less than that of Q' , which means someone interested in purchasing Q' can purchase the cheaper Q instead, and compute the answer to Q' from Q . For example, $Q(R) = \Pi_Y(R)$ determines $Q'(R) = \Pi_Y(\sigma_{Y < 10}(R))$ because the user can apply $Y < 10$ over $Q(R)$ to obtain $Q'(R)$. Arbitrage occurs if Q is cheaper than Q' which means a user looking for $Q'(R)$ can buy Q and compute answers to Q' . An arbitrage-free pricing function denies any form of arbitrage and returns consistent pricing without inadvertent data leakage [26]. A pricing scheme should also be *history-aware* [26]. A user should not be charged multiple times for the same data purchased through different queries. In such a pricing scheme, the data seller must track the history of queries purchased by each buyer and price future queries according to historical data.

3. Generalized Relations

Using general values as repair values in data cleaning requires us to re-consider the definition of consistency between a relational instance and a set of FDs. In this section, we introduce an entropy-based measure that allows us to quantify the utility of a generalized value as a repair value, and present a revised definition of consistency that has not been considered in past work [12].

3.1. Measuring Semantic Distance

By replacing a value v' in a relation R with a *generalized* value v , there is necessarily some information loss. We present an entropy-based penalty function that quantifies this loss [28]. We then use this measure as a basis to define a distance function between two general values.

Definition 5 (Entropy-based Penalty [28]). Consider an attribute A in a ground relation R . Let X_A be a random variable to randomly select a value from attribute A in R . Let VGH^A be the value generalization hierarchy of the attribute A (the values of A in R are ground values from VGH^A). The entropy-based penalty of a value v in VGH^A denoted by $E(v)$ is defined as follows:

$$E(v) = P(X_A \in \text{base}(v)) \times H(X_A | X_A \in \text{base}(v)),$$

where $P(X_A \in \text{base}(v))$ is the probability that the value of X_A is a ground value and a descendant of v , $X_A \in \text{base}(v)$. The value $H(X_A | X_A \in \text{base}(v))$ is the entropy of X_A conditional to $X_A \in \text{base}(v)$.

Intuitively, the entropy $H(X_A | X_A \in \text{base}(v))$ measures the uncertainty of using the general value v . $E(v)$ measures the information loss of replacing values in $\text{base}(v)$ with v . Note that $E(v) = 0$ if v is ground because $H(X_A | X_A \in \text{base}(v)) = 0$, and $E(v)$ is maximum if v is the root value $*$ in VGH^A . $E(v)$ is monotonic whereby if $v \preceq v'$, then $E(v) \leq E(v')$. Note that the conditional entropy $H(X_A | X_A \in \text{base}(v))$ is not a monotonic measure [28].

Example 3.1. Consider the general value $v = g_1[\text{AGE}] = [31, 60]$ in Table 1. Three ground values 51, 45 and 32 in $\text{base}(v)$ appear in Table 1, which has total 6 records, so $P(X_{\text{AGE}} \in \text{base}(v)) = \frac{3}{6} = \frac{1}{2}$. According to Table 1, the conditional entropy is $H(X_A | X_{\text{AGE}} \in \text{base}(v)) = 3 \times (-\frac{1}{3} \times \log \frac{1}{3}) = 1.58$ because 51, 45 and 32 each appear exactly once in the table. Therefore, the entropy-based penalty of v is $E(v) = \frac{1}{2} \times 1.58 = 0.79$.

Definition 6 (Semantic Distance Function, δ). The semantic distance between v and v' in VGH^A is defined as follows:

$$\delta(v, v') = \begin{cases} |E(v') - E(v)| & \text{if } v \preceq v' \text{ or } v' \preceq v \\ \delta(v, a) + \delta(a, v') & \text{otherwise} \end{cases}$$

in which $a = \text{lca}(v, v')$ is the least common ancestor of v and v' .

Intuitively, if v is a descendant of v' or vice versa, i.e. $v \preceq v'$ or $v' \preceq v$, their distance is the the difference between their entropy-based penalties, i.e., $|E(v') - E(v)|$. This is the information loss incurred by replacing a more informative child value v with its ancestor v' when $v \preceq v'$. If v and v' do not align along the same branch in the VGH , i.e. $v \not\preceq v'$ and $v' \not\preceq v$, $\delta(v, v')$ is the total distance between v and v' as we travel through their least common ancestor a , i.e. $\delta(v, a) + \delta(a, v')$.

Example 3.2. (Ex. 3.1 continued) According to Definition 6, $\delta([31, 60], 51) = |E([31, 60]) - E(51)| = |0.79 - 0| = 0.79$ because $51 \preceq [31, 60]$. Similarly, $\delta([31, 60], 45) = 0.79$. Also, $\delta(45, 51) = \delta(45, [31, 60]) + \delta([31, 60], 51) = 1.58$ because 45 and 51 do not belong to the same branch, and $[31, 60]$ is their least common ancestor.

The $\delta(v'v)$ distance captures the semantic closeness between values in the VGH . We extend the definition of δ to tuples by summing the distances between corresponding values in the two tuples. The δ function naturally extends to sets of tuples and relations. We use the $\delta(v'v)$ distance measure to define repair error in our evaluation in Section 7.

3.2. Consistency in Generalized Relations

A (generalized) relation R may contain generalized values that are syntactically equal but are semantically different and represent different ground values. For example, $g_1[\text{AGE}]$ and $g_2[\text{AGE}]$ in Table 3 are syntactically equal (containing [31,60]), but their true values in Table 1 are different, $m_1[\text{AGE}] = 51$ and $m_2[\text{AGE}] = 45$, respectively. In contrast, two syntactically different general values may represent the same ground value. For example, although *analgesic* and *NSAID* are different general values, they may both represent *ibuprofen*. The consistency between traditional FDs and a relation R is based on syntactic equality between values. We present a revised definition of consistency with the presence of generalized values in R .

Given a generalized relation R' with schema \mathcal{R} , and FD $\varphi : A \rightarrow B$ with attributes A, B in \mathcal{R} , R' satisfies φ ($R' \models \varphi$), if for every pair of tuples t_1, t_2 in R' , if $t_1[A]$ and $t_2[A]$ are equal ground values then $t_1[B]$ is an ancestor of $t_2[B]$ or vice-versa, i.e. $t_1[B] \preceq t_2[B]$ or $t_2[B] \preceq t_1[B]$. Since a general value v encapsulates a set of distinct ground values in $\text{base}(v)$, relying on syntactic equality between two (general) values, $t_1[B]$ and $t_2[B]$ is not required to determine consistency, since they may be syntactically different but represent the same ground value. Our definition requires that $t_1[B]$ be an ancestor of $t_2[B]$ (or vice-versa), which means they represent the same ground value. This definition extends to FDs $X \rightarrow Y$ with a set of attributes X, Y in \mathcal{R} . The definition reduces to the classical FD consistency when R' is ground. Assuming VGHs have fixed size, consistency checking in the presence of generalized values is in quadratic time according to $|R'|$.

Example 3.3. Consider an FD $\varphi : [\text{GEN}, \text{DIAG}] \rightarrow [\text{MED}]$ that states two patients of the same gender that are diagnosed with the same disease should be prescribed the same medication. According to the classic definition of FD consistency, t_1 and t_2 in Table 2 are inconsistent as t_1 and t_2 are males with osteoarthritis, i.e. $t_1[\text{GEN}, \text{DIAG}] = t_2[\text{GEN}, \text{DIAG}] = \{\text{male}, \text{osteoarthritis}\}$, but are prescribed different medications, i.e. $t_1[\text{MED}] = \text{ibuprofen} \neq \text{intropes} = t_2[\text{MED}]$. Under our revised definition of consistency for a generalized relation, if we update $t_2[\text{MED}]$ to *NSAID*, which is the ancestor of *ibuprofen*, then t_1 and t_2 are now consistent. In our revised consistency definition, $t_1[\text{MED}] = \text{ibuprofen} \preceq \text{NSAID} = t_2[\text{MED}]$, indicating that the general value $t_2[\text{MED}] = \text{NSAID}$ may represent the ground value $t_1[\text{MED}] = \text{ibuprofen}$. However, if we were to update $t_2[\text{MED}]$ to *vasodilators*, which is not the ancestor of $t_1[\text{MED}] = \text{ibuprofen}$, then t_1 and t_2 remain inconsistent under the generalized consistency definition, since the general value $t_2[\text{MED}] = \text{vasodilators}$ cannot represent $t_1[\text{MED}] = \text{ibuprofen}$.

4. PACAS Overview

We formally define our problem, and then review the PACAS framework [12], including extensions for generalized repairs.

4.1. Problem Statement

Consider a client, CL , and a service provider, SP , with databases D_{CL}, D_{SP} containing single relations R_{CL}, R_{SP} , respectively. Our discussion easily extends to databases with multiples relations. We assume a set of FDs Σ defined over R_{CL} that is falsified. We use FDs as the benchmark for error detection, but our framework is amenable to other error detection methods. The shared generalization hierarchies are generated by the service provider (applying the techniques mentioned in Section 2.2). The problem of privacy-preserving data cleaning is twofold, defined separately for

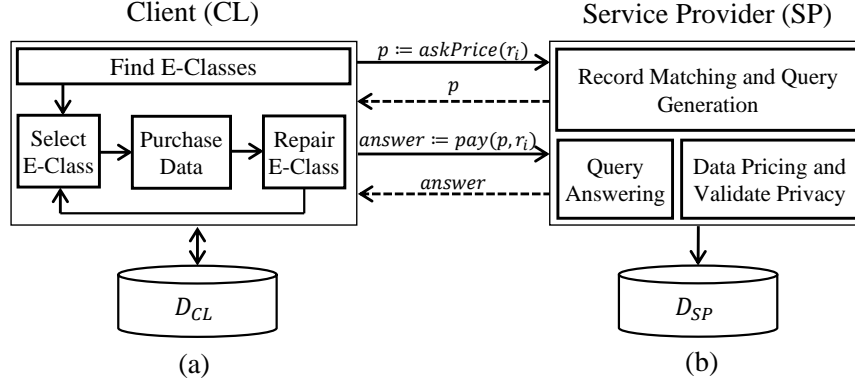


Figure 3: Framework overview.

CL and SP . We assume a generalization level l_{\max} , which indicates the maximum level that values in our repaired database can take from the generalization hierarchy.

Client-side: For every cell $c \in R_{CL}$ with value $c.value$, let $c.value^*$ be the corresponding accurate value in R_{SP} . A cell c is considered dirty if $c.value \neq c.value^*$. We assume the client can initiate a set of requests r_1, \dots, r_n to R_{SP} in which each request r_i is of the form $r_i = (t, A, l)$, that seeks the clean value of database cell $t[A]$ at level l in R_{SP} . We assume $\sum_i (price(r_i)) \leq \mathcal{B}$ for a fixed cleaning budget \mathcal{B} . Let R_{CL}^* be the clean version of R_{CL} where for each cell, $c.value = c.value^*$. The problem is to generate a set of requests r_1, \dots, r_n , where the answers (possibly containing generalized values) are used to compute a relation R'_{CL} such that: (i) $R'_{CL} \models \Sigma$, (ii) $\delta(R'_{CL}, R_{CL}^*)$ is minimal, and (iii) for each $c.value$, its level $l \leq l_{\max}$.

In our implementation, we check consistency $R'_{CL} \models \Sigma$ using the consistency definition in Section 3.2, and measure the distance $\delta(R'_{CL}, R_{CL}^*)$ using the semantic distance function δ (Defn. 6).

Service-side: The problem is to compute a pricing function $price(r_i)$ that assigns a price to each request r_i such that R_{SP} preserves (X,Y,L)-anonymity.

4.2. Solution Overview

Figure 3 shows the PACAS system architecture consisting of two main units that execute functions for CL and SP . Figure 3(a) shows the CL unit containing four modules. The first module finds equivalence classes in R_{CL} . An *equivalence class* (eq) is a set of cells in R_{CL} with equal values in order for R_{CL} to satisfy Σ [2]. The next three modules apply an iterative repair for the eqs. These modules select an eq, purchase accurate value(s) of dirty cell(s) in the class, and then repair the cells using the purchased value. If the repairs contain generalized values, the CL unit must verify consistency of the generalized relation against the defined FDs. The cleaning iterations continue until \mathcal{B} is exhausted or all the eqs are repaired. We preferentially allocated the budget \mathcal{B} to cells in an eq based on the proportion of errors in which the cells (in an eq) participate. Figure 3(b) shows the SP unit. The *Record Matching* module receives a request $r_i = (t, A, l)$ from CL , identifies a matching tuple t' in R_{SP} , and returns $t'[A]$ at the level l according to the VGH. The *Data Pricing and Validate Privacy* module computes prices for these requests, and checks whether answering these requests will violate (X,Y,L)-anonymity in R_{SP} . If so, the request is not safe to be answered. The *Query Answering* module accepts payment for requests, and returns the corresponding answers to CL .

5. Limiting Disclosure of Sensitive Data

The *SP* must carefully control disclosure of its curated and sensitive data to the *CL*. In this section, we describe how the *SP* services an incoming client request for data, validates whether answering this request is safe (in terms of violating (X,Y,L)-anonymity), and the data pricing mechanism that facilitates this interaction.

5.1. Record Matching and Query Generation

Given an incoming *CL* request $r_i = (t, A, l)$, the *SP* must translate r_i to a query that identifies a matching (clean) tuple t' in R_{SP} , and returns $t'[A]$ at level l . To answer each request, the *SP* charges the *CL* a price that is determined by the level l of data disclosure and the adherence of $t'[A]$ to the privacy model. To translate a request r to a GQ G_r , we assume a schema mapping exists between R_{SP} and R_{CL} , with similarity operators \approx to compare the attribute values. This can be modeled via *matching dependencies* (MDs) in *SP* [5].

Example 5.1. Consider a request $r=(t_2, \text{MED}, l_1^{\text{MED}})$ that requests the medication of the patient in tuple t_2 at level l_1^{MED} . To translate r into a generalized query $G_r = \langle Q_r, L_r \rangle$, we use the values of the QI attributes GEN and AGE in tuple t_2 . We define a query $Q_r(R_{SP}) = \Pi_{\text{MED}}(\sigma_{\text{GEN}=t_2[\text{GEN}] \wedge \text{AGE}=t_2[\text{AGE}]}(R_{SP}))$ requesting the medications of patients in R_{SP} with the same gender and age as the patient in tuple t_2 . The GQ G_r level is equal to the level of the request, i.e. $L_r = \{l_1^{\text{MED}}\}$. The query Q_r is generated from an assumed matching dependency (MD) $R_{CL}[\text{GEN}] = R_{SP}[\text{GEN}] \wedge R_{CL}[\text{AGE}] = R_{SP}[\text{AGE}] \rightarrow R_{CL}[\text{MED}] = R_{SP}[\text{MED}]$ that states if the gender and age of two records in *SP* and *CL* are equal, then the two records refer to patients with the same prescribed medication.

5.2. Enforcing Privacy

Recall from Section 2.4 that (X,Y,L)-anonymity extends (X,Y)-anonymity to consider the attribute domain semantics. Given a GQ, the *SP* must determine whether it is safe to answer this query, i.e., decide whether disclosing the value requested in r_i violates (X,Y,L)-anonymity (defined in Defn. 4). In this section, we formalize the privacy guarantees of (X,Y,L)-anonymity, and in the next section, review the *SafePrice* data pricing algorithm that assigns prices to GQs to guarantee (X,Y,L)-anonymity.

(X,Y,L)-anonymity applies tighter privacy restrictions compared to (X,Y)-anonymity by tuning the parameter L (determined by the data owner). At higher levels of L , it becomes increasingly more difficult to satisfy the (X,Y,L)-anonymity condition, while (X,Y,L)-anonymity reduces to (X,Y)-anonymity at $l_i = 0$ for every $l_i \in L$. We formally state this in Theorem 5.1. (X,Y,L)-anonymity is a semantic extension of (X,Y)-anonymity. Similar extensions can be defined for PPDP models such as (X,Y)-privacy, l -diversity, t -closeness. Unfortunately, all these models do not consider the data semantics in the attribute DGHs and VGHs [11].

Theorem 5.1. Consider a relation R and let X and Y be two subsets of attributes in R . Let L_1 and L_2 be levels of the attributes in Y with $L_1 \leq L_2$, i.e. every level of L_1 is lower than or equal to its corresponding level in L_2 . The following holds for relation R with a fixed $k > 1$:

1. If R is (X, Y, L_2) -anonymous it is also (X, Y, L_1) -anonymous, but not vice versa (i.e. R can be (X, Y, L_1) -anonymous but not (X, Y, L_2) -anonymous).
2. For any levels L of attributes in Y , if R is (X, Y, L) -anonymous, it is also (X, Y) -anonymous.




Figure 4: Possible relations \mathcal{I} , conflict set \mathcal{C}_Q and admissible relations \mathcal{I}_Q for query Q .

Proof of Theorem 5.1. In the first item, if R is (X, Y, L_2) -anonymous, then for every t , $|G_1^t(R)| \geq k$ in which $G_1^t = \langle Q^t, L_1 \rangle$. Considering $G_2^t = \langle Q^t, L_2 \rangle$, we can claim $|G_2^t(R)| \geq k$, which proves R is also (X, Y, L_1) -anonymous. The claim holds because if Q^t has n answers in levels of L_1 , it will have fewer or equal number of answers in the higher level L_2 since different values in levels of L_1 may be replaced with the same ancestors in the levels of L_2 . The second item holds because if R is (X, Y, L) -anonymous it is also (X, Y, L^\perp) -anonymous where L^\perp contains the bottom levels of attributes in Y ; due to $L^\perp \leq L$ and item 1. If R is (X, Y, L^\perp) -anonymous, then every value of X appears with k ground values and we can conclude it is (X, Y) -anonymous. \square

5.3. Pricing Generalized Queries.

PPDP models have traditionally been used in non-interactive settings where a privacy-preserving table is published once. We take a user-centric approach and let users dictate the data they would like published. We apply PPDP in an interactive setting where values from relation R_{SP} are published incrementally according to (user) CL requests, while verifying that the disclosed data satisfies (X, Y, L) -anonymity [12]. We adopt a data pricing scheme that assigns prices to GQs by extending the baseline data pricing algorithm defined in [29] to guarantee (X, Y, L) -anonymity. We review this baseline pricing algorithm next.

Baseline Data Pricing. The baseline pricing model computes a price for a query Q over relation R according to the amount of information revealed about R when answering Q [29].

Given a query Q over a relation R (a database with single relation R), the baseline pricing model determines the price of Q based on the amount of information revealed about R by answering Q . Let \mathcal{I} be a set of *possible relations* that the buyer believes to be R , representing his initial knowledge of R . As the buyer receives answers to $Q(R)$, he gains new knowledge, allowing him to eliminate relations R' from \mathcal{I} , which provide a different answer $Q(R') \neq Q(R)$. This set of eliminated instances R' is called the *conflict set of Q* denoted as \mathcal{C}_Q , and intuitively represents the amount of information that is revealed by answering Q (Figure 4). As more queries are answered, the size of \mathcal{I} is reduced. We can apply a set function that uses \mathcal{C}_Q to compute a price for Q . We can use the *weighted cover set* function with predefined weights assigned to the relations in \mathcal{I} . Query prices are computed by summing the weights for relations in \mathcal{C}_Q , which has been shown to give arbitrage-free prices [29]. In practice, the set \mathcal{I} is usually infinite making it infeasible to implement. To circumvent this problem, a smaller, finite subset \mathcal{S} called the *support set* is used to generate arbitrage-free prices [29]. The support set is defined as the neighbors of R , generated from R via tuple updates, insertions, and deletions. The values that are used to generate the support set are from the same domain of the original relation R .

Several optimizations are applied to the baseline pricing, and its effectiveness is experimentally justified [29]. Most importantly, the relations in the support set can be modeled by update op-

Algorithm 1: $price(Q, D, \mathcal{S}, w)$ [29]

Input : A query Q , a database D , support set \mathcal{S} , weight function w
Output: Price to answer Q

```
1  $p \leftarrow 0$ ;  
2 for  $D' \in \mathcal{S}$  do  
3   | if  $Q(D') \neq Q(D)$  then  $p \leftarrow p + w(D')$  ;  
4 end  
5 return  $p$ ;
```

erations. That is, we generate each relation in \mathcal{S} from R by applying its corresponding update operation and use the resulting relation to compute the value of the weighted function as the final price. We roll-back the update to restore R , and continue this process to compute the weighted function using other relations in \mathcal{S} . We avoid storing all databases in \mathcal{S} to enable more efficient price computations.

Algorithm 1 provides pseudocode of the baseline algorithm. The algorithm takes query Q , database D , a support set \mathcal{S} , and a weight function w , and computes the price to answer Q over R . The baseline algorithm is history-aware as input \mathcal{S} excludes databases that were already considered by past queries.

SafePrice Algorithm. We describe the *SafePrice* algorithm (originally introduced in [12]) that enforces (X,Y,L)-anonymity over R (equivalently R_{SP} in our framework). We first present the definition of a *safe query*, i.e., criteria for a GQ to preserve (X,Y,L)-anonymity.

Definition 7. (Safe Query) Consider a GQ G over a relation R with schema \mathcal{R} , $X, Y \subseteq \mathcal{R}$, and levels L corresponding to attributes in Y . Let $\mathcal{I}_G \subseteq \mathcal{I}$ be the set of relations R'' such that $G(R) = G(R'')$. G is safe (or preserves (X,Y,L)-anonymity of R) with value k , if for every tuple $t \in R$, there are at least k tuples in the set of answers $\{t'' \mid \exists R'' \in \mathcal{I}_G, t'' \in G^t(R'')\}$ where $G^t = \langle Q^t, L \rangle$ is a GQ with $Q^t(R'') = \Pi_Y(\sigma_{X=t[X]}(R''))$.

In Definition 7, \mathcal{I}_G represents the set of relations that the buyer believes R is drawn from after observing the answer $G(R)$. If there are at least k tuples in the answer set of G^t over \mathcal{I}_G , this indicates that the buyer does not have enough information to associate the values in X to less than k values of Y at level L , thus preserving (X,Y,L)-anonymity. If the *SP* determines that a GQ is safe, he will assign a finite price relative to the amount of disclosed information about R .

Algorithm 2 presents the *SafePrice* details. The given support set \mathcal{S} represents the user's knowledge about relation R after receiving the answer to past purchased queries. We maintain a set \mathcal{S}_G that represents all admissible relations after answering G and captures the user's posterior knowledge about R after answering G . We use \mathcal{S}_G to check whether answering G is safe. This is done by checking over instances in \mathcal{S}_G whether values in X are associated with at least k values of Y using the query G^t . The price for a query is computed by summing the weights of the inadmissible relations in the conflict set $\mathcal{C}_G = \mathcal{S} \setminus \mathcal{S}_G$ (Line 4). Similar to the baseline pricing, we use the support set \mathcal{S} and admissible relations \mathcal{S}_G rather than \mathcal{I} and \mathcal{I}_G , respectively. We iterate over tuples $t \in R$ (Line 5), and check whether values in X are associated with less than k values in Y over relations in \mathcal{S}_G . If so, we return an infinite price reflecting that query G is not safe (Line 8).

Proposition 1. If the input GQ is not safe, Algorithm 2 preserves (X,Y,L)-anonymity by returning an infinite price.

Proof sketch: According to Definition 7, if G violates (X,Y,L)-anonymity then for $t \in R$, there are less than k answers to G^t over relations in \mathcal{I}_G . Since $\mathcal{S}_G \subseteq \mathcal{I}_G$, there will be less than k answers to G^t over relations in \mathcal{S}_G , meaning Algorithm 2 assigns infinite price to G in Line 8. Note that if Algorithm 2 returns an infinite price, it does not imply G is unsafe (this only occurs when $\mathcal{S} = \mathcal{I}$). ■

Algorithm 2: $\text{SafePrice}(G, R, \mathcal{S}, w)$

Input : G, R, \mathcal{S}, w
Output: Price of G

```

1  $p \leftarrow 0; \mathcal{S}_G \leftarrow \emptyset;$ 
2 for  $T \in \mathcal{S}$  do
3   if  $G(T) = G(R)$  then  $\mathcal{S}_G \leftarrow \mathcal{S}_G \cup \{T\};$ 
4   else  $p \leftarrow p + w(T);$ 
5 for  $t \in R$  do
6    $A \leftarrow \emptyset;$ 
7   for  $R'' \in \mathcal{S}_G$  do  $A \leftarrow A \cup G^t(R'');$ 
8   if  $|A| < k$  then return  $\infty;$ 
9 return  $p;$ 
```

Given the interactive setting between the CL and the SP , we must ensure that all (consecutively) disclosed values guarantee (X,Y,L)-anonymity over R . We ensure that SafePrice is history-aware by updating the support set \mathcal{S} after answering each GQ. The SP uses the pricing function in Algorithm 2 to update \mathcal{S} to reflect the current relations R' that have been eliminated by answering the latest GQ. The SP implements $\text{AskPrice}(r_i, R_{SP})$ (cf. Figure 3) by translating r_i to a GQ, and then invoking SafePrice .

We assume that requesting a query price is free. We acknowledge that returning prices might leak information about the data being priced and purchased. This problem is discussed in the data pricing literature, particularly to incentivize data owners to return trustful prices when prices reveal information about the data (cf. [30] for a survey on this issue). We consider this problem as a direction of future work.

5.4. Query Answering

A CL data request is executed via the $\text{Pay}(p, r_i, R_{SP})$ method, where she purchases the value in r_i at price p . The *Query Answering* module executes $\text{Pay}(p, r_i, R_{SP})$ via SP accepts payment, translates r_i to a GQ, and returns the answer over R_{SP} to CL . Lastly, SP updates the support set \mathcal{S} to ensure SafePrice has an accurate history of disclosed data values.

We note that all communication between CL and SP is done via the AskPrice and Pay methods (provided by SP). We assume there is a secure communication protocol between the CL and SP , and that all data transfer is protected and encrypted. The model is limited to a single SP that sells data at non-negotiable prices. We intend to explore general cases involving multiple SP providers and price negotiation as future work.

6. Data Cleaning with Generalized Values

Existing constraint-based data repair algorithms that propose updates to the data to satisfy a set of data dependencies assume an open-access data model with no data privacy restrictions

[2, 31, 32, 33, 34]. In these repair models, inadvertent data disclosure can occur as the space of repair candidates is not filtered nor transformed to obscure sensitive values. A privacy-aware data repair algorithm must address the challenge of providing an accurate repair to an error while respecting data generalizations and perturbations to conceal sensitive values.

In earlier work, we introduced *SafeClean*, a data repair algorithm that resolves errors in a relation R_{CL} using data purchased from a service provider R_{SP} [12]. The key distinctions of *SafeClean* from past work include: (i) *SafeClean* interacts with the service provider, SP , to purchase trusted values under a constrained budget \mathcal{B} . This eliminates the overhead of traversing a large search space of repair candidates; and (ii) *SafeClean* tries to obtain values with highest utility from SP for repairing CL . However, the existing *SafeClean* algorithm grounds all purchased values thereby not allowing generalized values in R_{CL} . In this paper, we extend *SafeClean* to consider generalized values in R_{CL} by re-defining the notion of consistency between a relation and a set of FDs (Section 3.2), and revising the budget allocation algorithm to requests according to the number of errors in which cells in an equivalence participate (Section 6.4). In contrast, the existing *SafeClean* provides fixed budget allocations. We give an overview of our cleaning algorithm and subsequently describe each component in detail.

6.1. Overview

For a fixed number of FDs Σ , the problem of finding minimal-cost data repairs to R_{CL} such that R_{CL} satisfies Σ is NP-complete [2]. Due to these intractability results, we necessarily take a greedy approach that cleans cell values in R_{CL} that maximally reduce the overall number of errors w.r.t. Σ . This is in similar spirit to existing techniques that have used various weighted cost functions [2, 35] or conflict hypergraphs [33] to model error interactions among data dependencies in Σ .

Given R_{CL} and Σ , we identify a set of error cells that belong to tuples that falsify some $\sigma \in \Sigma$. For an error cell $e \in \mathcal{E}$, we define $\text{eq}(e)$ as the equivalence class to which e belongs. An equivalence class is a set of database cells with the same value such that Σ is satisfied [2]. We use eqs for two reasons: (i) by clustering cells into eqs, we determine a repair value for a *group of cells* rather than an individual cell, thereby improving performance; and (ii) we utilize every cell value within an eq to find the best repair.

The *SafeClean* algorithm repairs FD errors by first finding all the eqs in R_{CL} . The algorithm then iteratively selects an eq, and purchases the true value of a cell, and updates all dirty cells in the same class to the purchased value. The SP may decide that returning a generalized value is preferred to protect user sensitive data in R_{SP} . If so, then the CL must validate consistency of its data including these generalized values. At each iteration, we repair the eq class with cells that participate in the largest number of FD errors. At each iteration, *SafeClean* assigns a portion of the budget \mathcal{B} that is proportional to the number of errors relative to the total number of errors in R_{CL} (this new extension to *SafeClean* is discussed in Section 6.4). *SafeClean* continues until all the eqs are repaired, or the budget is exhausted.

SafeClean Algorithm. Algorithm 3 gives details of *SafeClean*’s overall execution. The algorithm first generates the set of eqs via *GenerateEQs* in Line 2. Equivalence classes containing only one value are removed since there is no need for repair. In Line 7, *SafeClean* selects an equivalence class eq_i with cells participating in the largest number of violations for repair (further details in Section 6.3). To repair the error cells in eq_i , *SafeClean* generates a request r_i using *GenerateRequest* that requests a repair value for a cell in eq_i (Line 8). This request is made at the lowest possible level (less than l_{max}) at a price allowable within the given budget. The algorithm assigns a fraction of the remaining budget, i.e. $\alpha_i \times B$ ($B \leq \mathcal{B}$ is the remaining budget) to purchase data at each

Algorithm 3: *SafeClean*($R_{CL}, R_{SP}, \Sigma, \mathcal{B}$)

Input : $R_{CL}, R_{SP}, \Sigma, \mathcal{B}$
Output: Clean R'_{CL}

```
1  $R'_{CL} \leftarrow R_{CL}$ ;  
2  $EQ \leftarrow \text{GenerateEQs}(R'_{CL}, \Sigma)$ ;  
3  $B \leftarrow \mathcal{B}$ ;  
4 for  $eq \in EQ$  do  
5 | if Resolved( $eq$ ) then  $EQ \leftarrow EQ \setminus \{eq\}$ ;  
6 while  $B > 0$  and  $EQ \neq \emptyset$  do  
7 |  $eq \leftarrow \text{Select}(EQ)$ ;  
8 |  $r_i \leftarrow \text{GenerateRequest}(eq, \alpha_i \times B, R_{SP})$ ;  
9 | if  $r_i \neq \text{null}$  then  
10 | |  $p_i \leftarrow \text{AskPrice}(r_i, R_{SP})$ ;  
11 | |  $u_i \leftarrow \text{Pay}(p_i, r_i, R_{SP})$ ;  
12 | |  $B \leftarrow B - p_i$ ;  
13 | | ApplyRepair( $eq, u_i$ )  
14 |  $EQ \leftarrow EQ \setminus eq$ ;  
15 return  $R'_{CL}$ 
```

iteration. This fraction depends on the number of violations in eq_i (cf. Section 6.4 for details). If such a request can be satisfied, the value(s) are purchased and applied (Lines 10-13). If there is insufficient budget remaining to purchase a repair value, then the eq cannot be repaired. In either case, *SafeClean* removes eq_i from EQ and continues with the next eq (Line 14). *SafeClean* terminates when \mathcal{B} is exhausted, or there are no eqs are remaining. We present details of eq generation, selection, request generation, and data purchase/repair in the following sections.

6.2. Generating Equivalence Classes

The eqs are generated by *GenerateEQs* in Algorithm 4 that takes as input R_{CL} and Σ , and returns the set of eqs EQ . For every cell $c_i \in R_{CL}$, the procedure initializes the eq of c_i as $eq(c_i) = \{c_i\}$, and adds it to the set of eqs EQ (Line 2). We then iteratively merge the eqs of any pair of cells $c_1 = t_1[B], c_2 = t_2[B]$ if there is a FD $\varphi : A \rightarrow B \in \Sigma$, $t_1[A] = t_2[A]$, and both $t_1[A]$ and $t_2[A]$ are ground. The procedure stops and returns EQ when no further pair of eqs can be merged.

Algorithm 4: *GenerateEQs*(R_{CL}, Σ)

Input : R_{CL}, Σ ,
Output: The set of equivalence classes EQ

```
1  $EQ \leftarrow \emptyset$ ;  
2 for  $c_i \in R_{CL}$  do  $EQ \leftarrow EQ \cup \{\{c_i\}\}$  ;  
3 for every  $t_1, t_2 \in R_{CL}$  and  $\varphi : A \rightarrow B \in \Sigma$  do  
4 | if  $t_1[A] = t_2[A]$  then merge( $EQ(t_1[B]), EQ(t_2[B])$ ) ;  
5 end  
6 return  $EQ$ ;
```

Example 6.1. Given FD $\varphi : [\text{GEN}, \text{DIAG}] \rightarrow [\text{MED}]$ in Table 2, since t_1 , t_2 and t_3 have the same attribute values on GEN and DIAG, we merge $t_1[\text{MED}]$, $t_2[\text{MED}]$ and $t_3[\text{MED}]$ into the same EQ_1 . Similarly, we cluster t_4 and t_5 into the same EQ_2 .

6.3. Selecting Equivalence Classes for Repair

In each iteration of Algorithm 3, we repair the cells in an eq eq_i that will resolve the most errors in R_{CL} w.r.t. Σ . To achieve this goal, we choose eq_i as the eq with cells participating in the largest number of errors. For a cell $c_j \in R_{CL}$ and an FD $\varphi : A \rightarrow B \in \Sigma$, let $\mathcal{E}(R_{CL}, \varphi, c_j)$ be the set of errors $\{t_1, t_2\}$ w.r.t. φ and $c_j \in \{t_1[A], t_1[B], t_2[A], t_2[B]\}$. For an eq eq_i , let $\mathcal{E}(R_{CL}, \varphi, eq_i) = \bigcup_{c_j \in eq_i} \mathcal{E}(R_{CL}, \varphi, c_j)$ and $\mathcal{E}(R_{CL}, \Sigma, eq_i) = \bigcup_{\varphi \in \Sigma} \mathcal{E}(R_{CL}, \varphi, eq_i)$. The eq eq_i returned by *Select* in Line 2 of Algorithm 3 is the eq with the largest number of errors in $\mathcal{E}(R_{CL}, \Sigma, eq_i)$. In other words, this is the number of tuple pairs that every cell in eq_i participates, summed over all FDs in Σ .

Example 6.2. We continue Example 6.1 with EQ_1 and EQ_2 . According to our error definition, given FD $\varphi : [\text{GEN}, \text{DIAG}] \rightarrow [\text{MED}]$, EQ_1 is involved in three errors with tuples $(\{t_1, t_2\}, \{t_1, t_3\})$ and $\{t_2, t_3\}$, while EQ_2 has one error $(\{t_3, t_4\})$. Hence, our algorithm will select EQ_1 to repair first.

6.4. Data Request Generation

To repair cells in eq_i , we generate a request $r_i = \langle c, l \rangle$ by *GenerateRequest* (Algorithm 5) that requests accurate and trusted value of a cell $c \in eq_i$ at level l from R_{SP} . There are two restrictions on this request: (i) the price must be within the budget $\alpha_i \times B$, and (ii) the level l must be $\leq l_{max}$. The value α_i is defined as follows:

$$\alpha_i = \frac{\mathcal{E}(R_{CL}, \Sigma, eq_i)}{\sum_{eq_j \in EQ} \mathcal{E}(R_{CL}, \Sigma, eq_j)}.$$

In this definition, the allocated budget $\alpha_i \times B$ to each iteration is proportional to the number of FD violations in eq_i , and also depends on the total number of errors in R_{CL} . This allocation model improves upon previous work that decreases the budget allocated to the i^{th} -request by a factor of $\frac{1}{i}$, and does not adjust the allocation to the number of errors in which a cell participates [12]. Note that if the price paid for r_i (i.e. p_i) is less than this allocated budget, the remaining budget carries to the next iteration through B .

If there is no such request, *GenerateRequest* returns null, indicating that eq_i cannot be repaired with the allocated budget (Line 10). If there are several requests that satisfy (i) and (ii), we follow a greedy approach and select the request at the lowest level with the price to break ties (Line 5). For a cell $c_i \in eq_i$, *LowestAffordableLevel* in Line 3 finds the lowest level in which the value of c_i can be purchased from R_{SP} considering the restrictions in (i) and (ii). Our greedy algorithm spends most of the allocated budget $\alpha_i \times B$ in the current iteration. An alternative approach is to select requests with the highest acceptable level (l_{max}) to preserve the budget for future iterations.

6.5. Purchase Data and Repair

To repair the cells in eq_i , Algorithm 3 invokes *Pay*(r_i, R_{SP}) to purchase the trusted value u_i and replaces the value of every cell in eq_i with u_i in *ApplyRepair*. The algorithm then removes the eq eq_i from EQ and continues to repair the next eq. The algorithm stops when there is no eq to repair or the budget is exhausted.

Algorithm 5: *GenerateRequest*(C, b, R_{SP})

Input : eq C , budget b and R_{SP} .
Output: Request r_i

```
1  $l \leftarrow l_h; p \leftarrow \infty; c \leftarrow null;$   
2 for  $c_i \in C$  do  
3    $l_i \leftarrow \text{LowestAffordableLevel}(c_i, b, l_{max}, R_{SP});$   
4    $p_i \leftarrow \text{AskPrice}(\langle c_i, l_i \rangle, R_{SP});$   
5   if  $l_i < l$  or ( $l_i == l$  and  $p_i \leq p$ ) then  
6      $c \leftarrow c_i; l \leftarrow l_i; p \leftarrow p_i;$   
7   end  
8 end  
9 if  $c \neq null$  then return  $\langle c, l \rangle;$   
10 return  $null;$ 
```

Example 6.3. Continuing from Example 6.2, since EQ_1 has the largest number of errors, we purchase the trusted value from R_{SP} to repair EQ_1 first. If the purchased value is general value $NSAID$, we update all cells in EQ_1 , i.e., $t_1[MED]$, $t_2[MED]$ and $t_3[MED]$ to $NSAID$ to resolve the inconsistency.

6.6. Complexity Analysis

We analyze the complexity of *SafeClean*'s modules:

- Identifying errors involves computing equivalence classes, and selecting an equivalence class for repair. In the worst case, this is quadratic in the number of tuples in R_{CL} .
- For each data request to resolve an error, executing *SafePrice* and *pay* rely on R_{SP} , the support set \mathcal{S} , and the complexity of GQ answering. We assume the size of \mathcal{S} is linear w.r.t the size of R_{SP} . The complexity of running GQs is the same as running SQL queries. Thus, all procedures run in polynomial time to the size of R_{SP} .
- In *applyRepair*, for each returned value from SP , we update the error cells in R_{CL} for each equivalence class and each FD, taking time on the order of $\mathcal{O}(|\mathcal{E}||A||\Sigma|)$ for attribute domain size $|A|$. We must update the affected cells in the equivalence class, and their dirty scores; both taking time bounded by the number of cells in R_{CL} . Hence, Algorithm 3 runs in polynomial time to the size of R_{SP} and R_{CL} .

7. Experiments

Our evaluation focuses on the following objectives:

1. We evaluate the impact of generalized values on the repair error, and the runtime performance. In addition, we study the proportion of generalized repair values that are recommended for varying budget levels.
2. The efficiency and effectiveness of *SafePrice* to generate reasonable prices that allow *SafeClean* to effectively repair the data.
3. We evaluate the repair error and scalability of *SafeClean* as we vary the parameters k, l, e and \mathcal{B} to study the repair error to runtime tradeoff.

Table 5: Data characteristics.

	Clinical	Census	Food
$ R_{CL} $	345,000	300,000	30,000
n	29	40	11
$ \Pi_A(R) $	61/73	17/250	248/70
$ \text{VGH}^A $	5/5	5/6	5/5

Table 6: Parameter values (defaults in bold).

Sym.	Description	Values
\mathcal{B}	budget	0.2, 0.4, 0.6, 0.8
$ \mathcal{S} $	support set size	6, 8, 10 , 12, 14
l	generalization level	0, 1, 2 , 3, 4
k	#tuples in X-group	1, 2, 3 , 4, 5
e	error rate	0.05, 0.1 , 0.15, 0.2, 0.25

4. We compare *SafeClean* against PrivateClean, a framework that explores the link between differential privacy and data cleaning [8]. We study the repair error to runtime tradeoff between these two techniques, and show that data randomization in PrivateClean significantly hinders error detection and data repair, leading to increased repair errors.

7.1. Experimental Setup

We implement *PACAS* using Python 3.6 on a server with 32 Core Intel Xeon 2.2 GHz processor with 64GB RAM. We describe the datasets, baseline comparative algorithm, metrics and parameters. The datasets, their schema, and source code implementation can be found at [36].

Datasets. We use three real datasets. Table 5 gives the data characteristics, showing a range of data sizes in the number of tuples ($|R_{CL}|$), number of attributes (n), number of unique values in the sensitive attribute A ($|\Pi_A(R)|$), and the height of the attribute VGH^A ($|\text{VGH}^A|$). We denote sensitive attributes with an asterisk (*).

Clinical Trials (Clinical). The Linked Clinical Trials database describes patient demographics, diagnosis, prescribed drugs, symptoms, and treatment [37]. We select the **country**, **gender**, **source** and **age** as QI attributes. We define two FDs: (i) $\varphi_1 : [\text{age}, \text{overall_status}, \text{gender}] \rightarrow [\text{drug_name}^*]$; and (ii) $\varphi_2 : [\text{overall_status}, \text{timeframe}, \text{measure}] \rightarrow [\text{condition}^*]$. We construct attribute value generalization hierarchies (VGH) with five levels on attributes **drug name** and **condition**, respectively using external ontologies (Bioprotal Medical Ontology [38], the University of Maryland Disease Ontology [39], and the Libraries of Ontologies from the University of Michigan [40]). The average number of children per node in the VGH is eight.

Census. The U.S. Census Bureau provides population characteristics such as education level, years of schooling, occupation, income, and age [41]. We select **sex**, **age**, **race**, and **native country** as QI attributes. We define two FDs: (i) $\varphi_3 : [\text{age}, \text{education-num}] \rightarrow [\text{education}^*]$; and (ii) $\varphi_4 : [\text{age}, \text{industry code}, \text{occupation}] \rightarrow [\text{wage-per-hour}^*]$. We construct VGH on attributes **wage-per-hour** and **education** by stratifying wage and education levels according to hierarchies from US statistics [42], and the US Department of Education [43]. The average number of children per node is five.

Food Inspection (Food). This dataset contains violation citations of inspected restaurants in New York City describing the **address**, **borough**, **zipcode**, **violation code**, **violation description**, **inspection type**, **score**, **grade**. We define **inspection type**, **borough**, **grade** as QI attributes. We define two FDs: (i) $\sigma_5 : [\text{borough}, \text{zipcode}] \rightarrow [\text{address}^*]$; and (ii) $\sigma_6 : [\text{violation code}, \text{inspection type}] \rightarrow [\text{violation description}^*]$. We construct attribute VGH on **address** and **violation description** by classifying streets into neighborhoods, districts, etc, and extracting topic keywords from the description and classifying the violation according to the Food Service Establishment Inspection Code [44]. The average number of children per node in the VGH is four.

For each dataset, we manually curate a clean instance R_{SP} according to the defined FDs, verify with external sources and ontologies, and is used as the ground truth. To create a dirty instance

R_{CL} , we duplicate R_{SP} to obtain R_{CL} , and use BART, an error generation benchmarking tool for data cleaning applications to inject controlled errors in the FD attributes [45]. We use BART to generate constraint-induced errors and random errors, and define error percentages ranging from 5% to 25%, with respect to the number of tuples in a table. We use BART’s default settings for all other parameters.

Comparative Baseline. The closest comparative baseline is PrivateClean, a framework that explores the link between differential privacy and data cleaning [8]. PrivateClean provides a mechanism to generate ϵ -differentially private datasets on numerical and discrete values from which a data analyst can apply data cleaning operations. PrivateClean proposes randomization techniques for discrete and numeric values, called Generalized Randomized Response (GRR), and applies this across all attributes. To guarantee a privatized dataset with discrete attributes is ϵ -differentially private at confidence $(1 - \alpha)$, a sufficient number of distinct records is needed. In our comparison evaluation, we set $\alpha = 0.05$, and the degree of privacy $p = 0.5$ (applied uniformly across the attributes and equivalent to ϵ for discrete attributes [8]). Since PrivateClean does not directly propose a data cleaning algorithm, but rather data cleaning operators, we apply the well-known Greedy-Repair [2] FD repair algorithm to generate a series of updates to correct the FD violations. We use the `transform()` operation to represent each of these updates in PrivateClean, and use the source code provided by the authors in our implementation. We choose the Greedy-Repair algorithm for its similarity to our repair approach; namely, to repair cells on an equivalence class basis, and to minimize a cost function that considers the number of data updates and the distance between the source and target values.

Metrics. We compute the average runtime over four executions. To measure the quality of the recommended repairs, we define the *repair error* of a cell as the distance between a cell’s true value and its repair value. We use the semantic distance measure, $\delta(v, v')$, in Section 3.1, that quantifies the distance between two values v (suppose a true value), and v' (a repair value) in the attribute value generalization hierarchy, and considers the distribution of v and v' in the relation. We assume a cell’s value, and its repair are both in the generalization hierarchy. The repair error of a relational instance is computed as the sum of the repair errors across all cells in the relation. We use the absolute value of the repair error in our experiments (denoted as δ). Similar to other error metrics (such as mean squared error), lower error values are preferred.

Parameters. Unless otherwise stated, Table 6 shows the range of parameter values we use, with default values in bold. We vary the following parameters to evaluate algorithm performance: (i) the budget \mathcal{B} ; (ii) the size of the support set $|\mathcal{S}|$ in the pricing function; (iii) l (the lower bound of generalization); (iv) k , the number of tuples in an X -group; and (v) the error rate e in R_{CL} .

7.2. Generalized Values

We measure the proportion of generalized values that are returned for varying levels of the budget \mathcal{B} . Since a repair value may be returned at any level of the VGH, we compute the semantic distance between the generalized value v and the corresponding ground value v' in the CL . We normalize the distance between (v, v') into four ranges: $[0-0.25]$, $[0.25-0.5]$, $[0.5-0.75]$, $[0.75-1]$, where a distance of zero indicates v and v' are both ground values.

Figure 5 shows the relative proportions for varying \mathcal{B} values over the clinical dataset. As expected, the results show that the proportion of generalized values at the highest levels of the VGH occur for low \mathcal{B} values since we can only afford (cheaper) generalized values under a constrained budget. In contrast, for \mathcal{B} values close to 0.9, close to 85% of the repair values are specific, ground values with a distance range of $[0, 0.25]$, while the remaining 15% are generalized values at the next




Figure 5: Ratio of gen. values.




Figure 6: Error vs. gen. values.




Figure 7: Runtime vs. gen. values.

level. We observe that for $\mathcal{B} > 0.6$ approximately 70% of the total repair values are very close to the ground value (where distance is at $[0-0.25]$), indicating higher quality repairs.

We evaluate the impact on the runtime to repair error tradeoff when an increasing number of generalized values occur in the repaired relation. We control the number of generalized values indirectly via the number of error cells under a constrained budget $\mathcal{B} = 0.1$ where it is expected that close to all repair recommendations will be general values. Figure 6 and Figure 7 show the repair error and runtime curves over the clinical data, respectively. As expected, we observe for an increasing number of generalized values, the repair error increases as the constrained budget leads to more values returned at the highest * generalized value, thereby increasing the distance between ground-truth relation and repaired relation. In contrast, the increased number of generalized values leads to lower runtimes due to the increased number of unsatisfied query requests.

7.3. SafePrice Efficiency and Effectiveness

The *SafePrice* algorithm relies on a support set \mathcal{S} to determine query prices by summing the weights of discarded instances from the conflict set \mathcal{C} . These discarded instances represent the knowledge gained by *CL*. We vary the size of the initial \mathcal{S} to determine its influence on the repair error (δ), and the overall runtime. Figure 8 shows a steady decrease in the repair error for increasing $|\mathcal{S}|$. As $|\mathcal{S}|$ grows, the *SP* is less restrictive to answer GQs and fewer requests are declined at lower levels. As more requests are answered at these lower levels, the repair error decreases. Figure 9 shows that the *SafePrice* runtime scales linearly with increasing $|\mathcal{S}|$, making it feasible to implement in practice. From Figures 8 and 9, we determine that *SafePrice* achieves an average 6% reduction in the repair error at a cost of 16m runtime. This is expected due to the additional time needed to evaluate the larger space of instances to answer GQs. Comparing across the three datasets, the data sizes affect runtimes as a larger number of records must be evaluated during query pricing. This is reflected in longer runtimes for the larger clinical and census datasets.




Figure 8: Repair error vs. $|\mathcal{S}|$.

7.4. SafePrice Parameter Sensitivity

We vary the parameters k , GQ level l , error rate e , budget \mathcal{B} , and measure their influence on *SafeClean* repair error and runtime over all three datasets. We expect that enforcing more stringent privacy requirements through larger k and l values will result in larger repair errors. Figures 10 to 13 do indeed reflect this intuition.




Figure 9: Runtime vs. $|S|$.




Figure 10: Repair error vs. k .




Figure 11: Runtime vs. k .




Figure 12: Repair error vs. l .




Figure 13: Runtime vs. l .




Figure 14: Repair error vs. e .




Figure 15: Runtime vs. e .




Figure 16: Repair error vs. B .




Figure 17: Runtime vs. B .

In Figure 10, *SafeClean* experiences larger repair errors for increasing k as generalizations to conceal sensitive values become increasingly dependent on the attribute domain and its VGH i.e., (X,Y,L) -anonymity indicates there must be at least k values at a given level l in the VGH. Otherwise, the data request is denied. Figure 11 shows that for increasing k , runtimes decrease linearly as query requests to satisfy more stringent k become more difficult. On average, we observe that an approximate 10% improvement in runtime leads to a 7% increase in the repair error for each increment of k . Figures 12 and 13 show the repair error and runtime, respectively, as we vary the query level parameter l . The repair error increases, particularly after $l = 3$ as more stringent privacy requirements are enforced, i.e., l distinct values are required at each generalization level. This makes satisfying query requests more difficult, leading to unrepaired values and lower runtimes, as shown in Figure 13.

Figure 14 shows the repair error δ increases as we scale the error rate e . For a fixed budget, increasing the number of FD errors leads to a decreasing budget for each FD error. This makes some repairs unaffordable for the *CL*, leading to unrepaired values and an increased number of generalized repair values. This situation can be mitigated if we increase the budget B . As expected, Figure 15 shows that the *SafeClean* runtime increases for an increasing number of FD violations, due to the larger overhead to compute more equivalence classes, compute prices to answer queries, and to check consistency in the *CL*.

Figures 16 and 17 show the repair error and runtime, respectively, as we vary the budget B .

For increasing budget allocations, we expect the *SP* to recommend more ground (repair) values, and lower repair error values, as shown in Figure 16. Given the larger budget allocations, the *SP* is able to answer a larger number of query requests, and must compute their query prices, thereby increasing algorithm runtime. We observe that an average 14% reduction in the repair error leads to an approximate 7% increase in runtime.

7.5. Comparative Evaluation

We compare *SafeClean* to PrivateClean, a framework for data cleaning on locally differentially private relations [8]. Section 7.1 describes the baseline algorithm parameter settings and configuration. Despite the differing privacy models, our evaluation aims to explore the influence of increasing error rates on the repair error δ , and algorithm runtimes using the Clinical dataset. We hope these results are useful for practitioners to understand qualitative and performance trade-offs between the two privacy models. For *SafeClean*, we measure total time of the end-to-end process from error detection to applying as many repairs as the budget allows. In PrivateClean, we measure the time to privatize the R_{CL} , error detection, running the Greedy-Repair FD repair algorithm [2], and applying the updates via the Transform operation. For PrivateClean, we measure the repair error $\delta(v, v')$ for source value v and target (clean) value v' , as recommended by Greedy-Repair, where both v, v' are ground values.

Figure 18 shows the comparative repair error between *SafeClean* and PrivateClean as we vary the error rate e . *SafeClean* achieves an average -41% lower repair error than PrivateClean. This poor performance by PrivateClean is explained by the underlying data randomization used in differential privacy, which provides strong privacy guarantees, but poor data utility, especially in data cleaning applications. As acknowledged by the authors, identifying and reasoning about errors over randomized response data is hard [8]. This randomization may update rare values to be more frequent, and similarly, common values to be more rare. This leads to more uniform distributions (where stronger privacy guarantees can be provided), but negatively impact error detection and cardinality-minimal data repair techniques that rely on attribute value frequencies to determine repair values [2]. We envision that *SafeClean* is a compromise towards providing an (X,Y,L)-anonymous instance while preserving data utility.

SafeClean's lower repair error comes at a runtime cost, as shown in Figure 19. As we scale the error rate, *SafeClean*'s runtime scales linearly due to the increased overhead of data pricing. Recall the pricing mechanism in Section 5.3, the price of query Q is determined by the query answer over the relation D and its neighbors D' in the support set \mathcal{S} . In contrast, PrivateClean does not incur such an overhead due to its inherent data randomization. There are opportunities to explore optimizations to lower *SafeClean*'s overhead to answer query requests and compute prices to answer each query. In practice, optimizations can be applied to improve overall performance, including decreasing the parameter k according to application requirements, and considering distributed (parallel) executions of query processing over partitions of the data. *SafeClean* aims to provide a balanced approach towards achieving data utility for data cleaning applications while protecting sensitive values via data pricing and (X,Y,L)-anonymity.




Figure 18: Comparative repair error.




Figure 19: Comparative runtime.

8. Related Work

Data Privacy and Data Cleaning. We extend the PACAS framework introduced in [12], which focuses on repairs involving ground values. While the prior PACAS framework prices and returns generalized values, the *CL* does not allow generalized values in its relation R_{CL} . Instead, the *CL* instantiates a grounding process to ground any generalized values returned by the *SP*. In this work, we remove this limitation by re-defining the notion of consistency between R_{CL} and the set of FDs Σ to include generalized values. We have also extended the budget allocation scheme to consider allocations according to the proportion of errors in which cells (in an equivalence class) participate. This improves the PACAS framework to be more adaptive to the changing number of unresolved errors instead of only assigning fixed allocations. Our revised experiments, using repair error δ as a measure of accuracy, show the influence of generalized repairs along varying parameters towards improved efficiency and effectiveness.

There has been limited work that considers data privacy requirements in data cleaning. Jaganathan and Wright propose a privacy-preserving protocol between two parties that imputes missing data using a lazy decision-tree imputation algorithm [6]. Information-theoretic metrics have been used to quantify the information loss incurred by disclosing sensitive data values [57, 7]. As mentioned previously, PrivateClean provides a framework for data cleaning on local differentially private relations using a set of generic user-defined cleaning operations showing the trade-off between privacy bounds and query accuracy [8]. While PrivateClean provides stronger privacy bounds than PACAS, these bounds are only applicable under fixed query budgets that limit the number of allowed queries (only aggregation queries), which is difficult to enforce in interactive settings. Our evaluation has shown the significant repair error that PrivateClean incurs due to the inherent randomization needed in local differential privacy.

Ge et. al., introduce APEX, an accuracy-aware, data privacy engine for sensitive data exploration that allow users to pose adaptive queries with a required accuracy bound. APEX returns query answers satisfying the bound and guarantees that the entire data exploration process is differentially private with minimal privacy loss. Similar to PACAS, APEX allows users to interact with private data through declarative queries, albeit specific aggregate queries (PACAS has no such restriction). APEX enable users to perform specific tasks over the privatized data, such as entity resolution. While similar in spirit, APEX uses the given query accuracy bound to tune differential privacy mechanisms to find a privacy level ϵ that satisfies the accuracy bound. In contrast, PACAS applies generalization to target predicates in the given query requests, without randomizing the entire dataset, albeit with looser privacy guarantees, but with higher data utility.

Dependency Based Cleaning. Declarative approaches to data cleaning (cf. [46] for a survey) have focused on achieving consistency against a set of dependencies such as FDs and inclusion dependencies and their extensions. Data quality semantics are declaratively specified with the dependencies, and data values that violate the dependencies are identified as errors, and repaired [5, 47, 35, 48, 49, 58]. There are a wide range of cleaning algorithms, based on user/expert feedback [50, 51, 52], master database [4, 5], knowledge bases or crowdsourcing [53], probabilistic inference [54, 55]. Our repair algorithm builds upon these existing techniques to include data disclosure requirements of sensitive data and suggesting generalized values as repair candidates.

Data Privacy. PPDP uses generalization and suppression to limit data disclosure of sensitive values [9, 10]. The generalization problem has been shown to be intractable, where optimization, and approximation algorithms have been proposed (cf. [11] for a survey). Extensions have proposed tighter restrictions to the baseline k -anonymity model to protect against similarity and skewness attacks by considering the distribution of the sensitive attributes in the overall population in the

table [11]. Our extensions to (X,Y)-anonymity to include semantics via the VGH can be applied to existing PPDP techniques to semantically enrich the generalization process such that semantically equivalent values are not inadvertently disclosed.

In differential privacy, the removal, addition or replacement of a single record in a database should not significantly impact the outcome of any statistical analysis over the database [56]. An underlying assumption requires that the service provider know in advance the set of queries over the released data. This assumption does not hold for the interactive, service-based data cleaning setting considered in this paper. PACAS aims to address this limitation by marrying PPDP and data pricing in interactive settings.

Data Pricing. We use the framework by Deep and Koutris that provides a scalable, arbitrage-free pricing for SQL queries over relational databases [29]. Recent work has considered pricing functions to include additional properties such as being *reasonable* (differentiated query pricing based on differentiated answers), and *predictable, non-disclosive* (the inability to infer unpaid query answers from query prices) and *regret-free* (asking a sequence of queries during multiple interactions is no more costly than asking at once) [26]. We are investigating extensions of *SafePrice* to include some of these properties.

9. Conclusions

We present PACAS, a data cleaning-as-a-service framework that preserves (X,Y,L)-anonymity in a service provider with sensitive data, while resolving errors in a client data instance. PACAS anonymizes sensitive data values by implementing a data pricing scheme that assigns prices to requested data values while satisfying a given budget. We propose generalized repair values as a mechanism to obfuscate sensitive data values, and present a new definition of consistency with respect to functional dependencies over a relation. To adapt to the changing number of errors in the database during data cleaning, we propose a new budget allocation scheme that adjusts to the current number of unresolved errors. We believe that PACAS provides a new approach to privacy-aware cleaning that protects sensitive data while offering high data cleaning utility, as data markets become increasingly popular. As next steps, we are investigating: (i) optimizations to improve the performance of the data pricing modules; and (ii) extensions to include price negotiation among multiple service providers and clients.

References

- [1] H. Bowne-Anderson, What data scientists really do, Harvard Business Review.
- [2] P. Bohannon, W. Fan, M. Flaster, R. Rastogi, A cost-based model and effective heuristic for repairing constraints by value modification, in: Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD), 2005, p. 143?154.
- [3] T. Dasu, J. M. Loh, Statistical distortion: Consequences of data cleaning, Proceedings of the VLDB Endowment 5 (11) (2012) 1674?1683.
- [4] W. Fan, J. Li, S. Ma, N. Tang, W. Yu, Towards certain fixes with editing rules and master data, The VLDB Journal 21 (2) (2012) 213?238.
- [5] W. Fan, X. Jia, J. Li, S. Ma, Reasoning about record matching rules, Proceedings of the VLDB Endowment 2 (1) (2009) 407?418.

- [6] G. Jagannathan, R. N. Wright, Privacy-preserving imputation of missing data, *Data and Knowledge Engineering* 65 (1) (2008) 40?56.
- [7] F. Chiang, D. Gairola, Infoclean: Protecting sensitive information in data cleaning, *Journal of Data and Information Quality (JDIQ)* 9 (4) (2018) 1–26.
- [8] S. Krishnan, J. Wang, M. J. Franklin, K. Goldberg, T. Kraska, Privateclean: Data cleaning and differential privacy, in: *Proceedings of the 2016 International Conference on Management of Data (SIGMOD)*, 2016, p. 937?951.
- [9] L. Sweeney, K-anonymity: A model for protecting privacy, *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10 (5) (2002) 557–570.
- [10] P. Samarati, Protecting respondents’ identities in microdata release, *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 13 (6) (2001) 1010–1027.
- [11] B. C. M. Fung, K. Wang, R. Chen, P. S. Yu, Privacy-preserving data publishing: A survey of recent developments, *ACM Computing Surveys* 42 (4) (2010) 14:1–14:53.
- [12] Y. Huang, M. Milani, F. Chiang, PACAS: Privacy-aware, data cleaning-as-a-service, in: *2018 IEEE International Conference on Big Data (Big Data)*, 2018, pp. 1023–1030.
- [13] J. Han, M. Kamber, J. Pei, *Data Mining: Concepts and Techniques*, 3rd Edition, Morgan Kaufmann Publishers Inc., 2011.
- [14] S. Lee, S.-Y. Huh, R. D. McNeil, Automatic generation of concept hierarchies using wordnet, *Expert Systems with Applications* 35 (3) (2008) 1132–1144.
- [15] U. M. Fayyad, K. B. Irani, Multi-interval discretization of continuous-valued attributes for classification learning, in: *International Joint Conference on Artificial Intelligence (IJCAI)*, 1993, p. 1022?1027.
- [16] R. Kerber, Chimerge: Discretization of numeric attributes, in: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 1992, p. 123?128.
- [17] H. Liu, R. Setiono, Feature selection via discretization, *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 9 (4) (1997) 642–645.
- [18] K. Wang, B. C. M. Fung, Anonymizing sequential releases, in: *Conference on Knowledge Discovery and Data Mining (KDD)*, 2006, p. 414?423.
- [19] A. Machanavajjhala, D. Kifer, J. Gehrke, M. Venkitasubramaniam, *l*-diversity: Privacy beyond *k*-anonymity, *ACM Transactions on Knowledge Discovery from Data (TKDD)* 1 (1) (2007) 1–52.
- [20] AggData locational data, <https://www.aggdata.com/> (2009).
- [21] Oracle Data Marketplace, <https://www.oracle.com/ca-en/data-cloud/> (2018).
- [22] Acxiom: Data, Identity Solution, <https://www.acxiom.com/> (2018).
- [23] Facebook Consumer Behaviour and Marketing, <https://www.facebook.com/business/insights> (2018).

- [24] Twitter Social Data, <https://developer.twitter.com/en/enterprise> (2018).
- [25] Spokeo, <https://www.spokeo.com/purchase> (2006).
- [26] M. Balazinska, B. Howe, D. Suciu, Data markets in the cloud: An opportunity for the database community, *Proceedings of the VLDB Endowment (PVLDB)* 4 (12) (2011) 1482–1485.
- [27] A. Nash, L. Segoufin, V. Vianu, Views and queries: Determinacy and rewriting, *ACM Transactions on Database Systems (TODS)* 35 (3) (2010) 1–41.
- [28] A. Gionis, T. Tassa, k-anonymization with minimal loss of information, *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 21 (2) (2009) 206–219.
- [29] S. Deep, P. Koutris, QIRANA: A framework for scalable query pricing, in: *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2017, pp. 699–713.
- [30] A. Roth, Buying private data at auction: The sensitive surveyor’s problem, *ACM SIGecom Exchanges* 11 (1) (2012) 1–8.
- [31] F. Chiang, R. J. Miller, Active repair of data quality rules, in: *Proceedings of the International Conference on Information Quality (ICIQ)*, 2011, pp. 174–188.
- [32] G. Beskales, I. F. Ilyas, L. Golab, Sampling the repairs of functional dependency violations under hard constraints, in: *Proceedings of the VLDB Endowment (PVLDB)*, 2010, pp. 197–207.
- [33] L. Golab, I. F. Ilyas, G. Beskales, A. Galiullin, On the relative trust between inconsistent data and inaccurate constraints, in: *IEEE International Conference on Data Engineering (ICDE)*, 2013, pp. 541–552.
- [34] M. Dallachiesa, A. Ebaid, A. Eldawy, A. Elmagarmid, I. F. Ilyas, M. Ouzzani, N. Tang, Nadeef: a commodity data cleaning system, in: *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2013, pp. 541–552.
- [35] M. Volkovs, F. Chiang, J. Szlichta, R. J. Miller, Continuous data cleaning, in: *IEEE International Conference on Data Engineering (ICDE)*, 2014, pp. 244–255.
- [36] PACAS: Experimental datasets and source code, <https://pacas123.github.io/PACAS/>.
- [37] Linked Clinical Trials, <https://old.datahub.io/dataset/linkedct> (2016).
- [38] Bioportal Medical Ontology, <https://bioportal.bioontology.org/ontologies> (2018).
- [39] Disease Ontology, <http://disease-ontology.org/> (2018).
- [40] University of Michigan Library, Libraries of ontologies, <https://guides.lib.umich.edu/ontology/ontologies> (2017).
- [41] Census Income Database, <https://archive.ics.uci.edu/ml/machine-learning-databases/census-income-mld/census-income.html> (2017).
- [42] US Data and Statistics, <https://www.usa.gov/statistics> (2018).

- [43] US Department of Education, <https://www.ed.gov/> (2017).
- [44] New York State Food Services Inspection, <https://www.health.ny.gov> (2017).
- [45] P. Arocena, B. Glavic, G. Mecca, R. J. Miller, P. Papotti, D. Santoro, Messing up with bart: error generation for evaluating data-cleaning algorithms, *Proceedings of the VLDB Endowment (PVLDB)* 9 (2) (2015) 36–47.
- [46] L. Bertossi, L. Bravo, Generic and declarative approaches to data quality management, in: S. Sadiq (Ed.), *Handbook of Data Quality: Research and Practice*, Springer Berlin Heidelberg, 2013, pp. 181–211.
- [47] F. Chiang, R. J. Miller, A unified model for data and constraint repair, in: *IEEE International Conference on Data Engineering (ICDE)*, 2011, pp. 446–457.
- [48] S. Kolahi, L. Lakshmanan, On approximating optimum repairs for functional dependency violations, in: *International Conference on Database Theory (ICDT)*, 2009, pp. 53–62.
- [49] I. F. Ilyas, X. Chu, Trends in cleaning relational data: Consistency and deduplication, *Foundations and Trends in Databases* 5 (4) (2015) 281–393.
- [50] C. Gokhale, S. Das, A. Doan, J. F. Naughton, N. Rampalli, J. Shavlik, X. Zhu, Corleone: Hands-off crowdsourcing for entity matching, in: *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2014, pp. 601–612.
- [51] J. Wang, T. Kraska, M. J. Franklin, J. Feng, Crowder: Crowdsourcing entity resolution, *Proceedings of the VLDB Endowment (PVLDB)* 5 (11) (2012) 1483–1494.
- [52] M. Yakout, A. K. Elmagarmid, J. Neville, M. Ouzzani, I. F. Ilyas, Guided data repair, *Proceedings of the VLDB Endowment (PVLDB)* 4 (5) (2011) 279–289.
- [53] X. Chu, J. Morcos, I. F. Ilyas, M. Ouzzani, P. Papotti, N. Tang, Y. Ye, KATARA: A data cleaning system powered by knowledge bases and crowdsourcing, in: *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2015, pp. 1247–1261.
- [54] T. Rekatsinas, X. Chu, I. F. Ilyas, C. Ré, Holoclean: Holistic data repairs with probabilistic inference, *Proceedings of the VLDB Endowment* 10 (11) (2017) 1190–1201.
- [55] Z. Yu, X. Chu, Piclean: A probabilistic and interactive data cleaning system, in: *Proceedings of the International Conference on Management of Data (SIGMOD)*, 2019, pp. 2021–2024.
- [56] C. Dwork, Differential privacy, in: *International Colloquium on Automata, Languages and Programming (ICALP)*, 2006, pp. 1–12.
- [57] D. Huang, D. Gairola, Y. Huang, Z. Zheng, F. Chiang, PARC: Privacy-Aware Data Cleaning, *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*, 2016, pp. 2433–2436.
- [58] S. Baskaran, A. Keller, F. Chiang, L. Golab, J. Szlichta, Efficient Discovery of Ontology Functional Dependencies, in: *Proceedings of the ACM International Conference on Information and Knowledge Management, CIKM*, 2017, pp. 1847–1856.