

ρ -uncertainty Anonymization by Partial Suppression

Xiao Jia¹, Chao Pan¹, Xinhui Xu¹, Kenny Q. Zhu¹, and Eric Lo²

¹ Shanghai Jiao Tong University, China
kzhu@cs.sjtu.edu.cn **

² Hong Kong Polytechnic University, China

Abstract. We present a novel framework for set-valued data anonymization by partial suppression regardless of the amount of background knowledge the attacker possesses, and can be adapted to both space-time and quality-time trade-offs in a “pay-as-you-go” approach. While minimizing the number of item deletions, the framework attempts to either preserve the original data distribution or retain mineable useful association rules, which targets statistical analysis and association mining, two major data mining applications on set-valued data.

1 Introduction

Set-valued data sources are valuable in many data mining and data analysis tasks. For example, retail companies may want to know what items are top sellers (e.g., milk), or whether there is an association between the purchase of two or more items (e.g., people who buy flour also buy milk). According to our observation there are two main categories of set-valued data analysis: one is *statistical analysis* such as computing max, min and average values; the other is *mining of association rules* between items. In many cases, analysis tasks are *outsourced* to other external companies or individuals, or simply *published* to the general masses for scientific and public research purposes.

Publishing set-valued, and especially transactional data, can pose significant privacy risks. Set-valued transactions consist of one or more data items, which can be divided into two categories: *non-sensitive* and *sensitive*. Privacy is in general associated with the sensitive items. Table 1(a) shows an example of retail transactions in which each record (row) represents a set of items purchased in a single transaction by an individual. All the items are non-sensitive, except the *condom* which is sensitive. An individual’s privacy is breached if he or she can be *re-identified*, or associated with a record in the data which contains one or more sensitive items. Past research has shown that such breach is possible through *linking attacks* [7], e.g. linking milk with condom in Table 1(a).

The privacy model we want to achieve is called ρ -uncertainty, where no sensitive association rules can be inferred with a confidence higher than ρ [3]. The

** Kenny Q. Zhu is the contact author and is supported by NSFC grants 61100050, 61033002, 61373031 and Google Faculty Research Award.

Table 1. A Retail Dataset and Anonymization Results

(a) Original Dataset		(b) Global Suppression	
TID	Transaction	TID	Transaction
1	bread, milk, <i>condom</i>	1	bread, milk, <i>condom</i>
2	bread, milk	2	bread, milk
3	milk, <i>condom</i>	3	milk, <i>condom</i>
4	flour, fruits	4	flour, fruits
5	flour, <i>condom</i>	5	flour, <i>condom</i>
6	bread, fruits	6	bread, fruits
7	fruits, <i>condom</i>	7	fruits, <i>condom</i>

(c) Our Approach 1		(d) Our Approach 2	
TID	Transaction	TID	Transaction
1	bread, milk, <i>condom</i>	1	bread, milk, <i>condom</i>
2	bread, milk	2	bread, milk
3	milk, <i>condom</i>	3	milk, <i>condom</i>
4	flour, fruits	4	flour, fruits
5	flour , <i>condom</i>	5	flour, <i>condom</i>
6	bread, fruits	6	bread, fruits
7	fruits, <i>condom</i>	7	fruits, <i>condom</i>

most popular approach to achieve ρ -uncertainty is called “global suppression” [3] in which once an occurrence of an item t is determined to be removed from one record, all occurrences of t are removed from the whole dataset. We instead opt to *partially* suppress the data set so only *some* occurrences of item t are deleted. Table 1 shows the example dataset and three anonymized datasets produced by global suppression and our approaches. The original dataset is not safe because sensitive rules such as $\{bread, milk\} \rightarrow condom$ can be inferred with confidence great than $1/3$ ($1/2$ in this case), which is our threshold. Table 1(b) is the anonymized dataset where all the occurrences of the sensitive item *condom* are deleted due to global suppression. Table 1(c) shows the anonymized dataset produced by our first approach, which is optimized to preserve data distribution, so different items (*condom* and flour) are deleted to make the dataset safe. Table 1(d) is the result of our second approach, which is optimized to preserve important data association in the original dataset, so only two occurrences of the item *condom* are deleted for safety. Even in such a small dataset, both our approaches outperform global suppression in the number of deletions (2 vs. 4) while retaining useful information at the same time.

To the best of our knowledge, the partial suppression technique has not been studied in the context of set-valued data anonymization before. We choose to solve the set-valued data anonymization problem by partial suppression because global suppression tends to delete more items than necessary, and the removal of all occurrences of the same item not only changes the data distribution significantly but also makes mining association rules about the deleted items impos-

sible. The problem of anonymization by suppression (global or partial) is very challenging [1, 18], exactly because, (i) the number of possible inferences from a given dataset is exponential, and (ii) the size of the search space, i.e. the number of ways to suppress the data is also exponential to the number of data items. We therefore propose two heuristics in this paper to anonymize input data in two different ways, giving rise to the two kinds of output in Table 1.

The main contributions of this paper are as follows.

1. To the best of our knowledge, we are the first to propose an effective *partial* suppression framework for anonymizing set-valued data (Section 2 and 3).
2. We adopt a “pay-as-you-go” approach based on divide-and-conquer, which can be adapted to achieve both space-time and quality-time trade-offs (Section 3 and 4). Our two heuristics can be adapted to either preserve data distribution or retain useful association in the data (Section 3).
3. Experiments show that our algorithm outperforms the peers in preserving the original data distribution (more than 100 times better than peers) or retaining mineable useful association rules while reducing the item deletions by large margins (Section 4).

2 Problem Definition

This section introduces the privacy model and data utility before formally describing the problem of partial suppression.

2.1 Privacy Model

$X \rightarrow Y$ is a *sensitive association rule* iff Y contains at least one sensitive item. The principle privacy model in this paper maintains that a table is *safe* iff no sensitive rules can be inferred with a confidence higher than threshold ρ [3]. It is easy to show that, if all sensitive association rules with one-item consequent from T are safe then all sensitive association rules are safe and hence T is safe.

Formally, we define quasi-identifier (also *qid*) to be any itemset (including sensitive items) drawn from any record in table T . A *qid* q is safe w.r.t. ρ iff $\text{conf}(q \rightarrow e) \leq \rho$, for any sensitive item e in T . We say T is safe w.r.t. ρ iff q is safe w.r.t. ρ for any *qid* q in T . A *suppressor* is a function $S : T \mapsto T'$ where T' is a suppressed table which is safe w.r.t. ρ . There are many different ways to suppress a table. The goal is to find a suppressor that maximizes the *utility* of the suppressed table.

2.2 Data Utility

In this paper, we identify two major uses of an anonymized table: *statistical analysis* and *association rule mining*. In the first case, we want the anonymized table to have a distribution as close to the original table as possible; in the second case, we would like the anonymized data to retain all non-sensitive association rules

while introducing few or no spurious rules. In both scenarios, the common goal is to minimize the *information loss*, i.e., the total number of items suppressed.

With these two scenarios in mind, we define two variants of an objective function $f(T, T')$ as:

$$f(T, T') = \begin{cases} NS(T, T') \cdot KL(T' \parallel T) & \text{(data distribution)} \\ \frac{NS(T, T')}{J(nr(T), nr(T'))} & \text{(rule mining)} \end{cases} \quad (1)$$

where

$$NS(T, T') = \frac{\sum_{e \in D(T)} (sup_T(e) - sup_{T'}(e))}{\sum_{e \in D(T)} sup_T(e)} \quad (2)$$

$$KL(P \parallel Q) = \sum_i Q(i) \log \frac{Q(i)}{P(i)} \quad (3)$$

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (4)$$

Here $D(T)$ denotes the domain of items in T , $sup_T(e)$ denotes the support of item e in T , $nr(T)$ denotes the set of all non-sensitive associations rules mineable from T with sufficient support and confidence, the functions NS , KL and J represent total number of suppressions (normalized to 1), K-L divergence[11] and Jaccard similarity[10], respectively. K-L divergence measures the distance between two probability distributions, while Jaccard similarity measures the similarity between two sets.

2.3 Optimal Partial Suppression Problem

The optimal partial suppression problem is to find a *Partial Suppressor* S which anonymizes an input set-valued table T to minimize the objective function:

$$\min_S f(T, S(T))$$

such that $S(T)$ is *safe* w.r.t. to our privacy model.

3 Partial Suppression Algorithm

The Optimal Suppression Problem defined in Section 2 is an NP-hard problem. We therefore present the partial suppression algorithm as a heuristic solution to the Optimal Suppression Problem. To simplify the discussion of the algorithm, we make the following definitions.

Definition 1 (Number of Suppressions). *To disable an unsafe rule $q \rightarrow e$, the number of items of type $t \in q \cup \{e\}$ that need to be suppressed is*

$$N_s(t, q \rightarrow e) = \begin{cases} sup(q \cup \{e\}) - sup(q)\rho & t = e \\ \frac{sup(q \cup \{e\}) - sup(q)\rho}{1 - \rho} & t \in q \end{cases}$$

In other words, for each sensitive rule r , we need to delete $\min_t N_s(t, r)$ items of type t to make it safe. In this work, we select these items randomly for deletion.

Definition 2 (Leftover Items). *The leftover of item type t is defined as*

$$\text{leftover}(t) = \text{sup}_{T'}(\{t\}) / \text{sup}_T(\{t\})$$

T is the original data and T' is the intermediate suppressing result. The ratio shows the percentage of remaining items of type t in the intermediate result T' .

The key intuition of our algorithm is that although the total number of “bad” sensitive association rules maybe, in the worst case, exponential in the original data, incremental “invalidation” of some of the rules through partial suppression of a small number of affected items can massively reduce the number of these bad rules, which leads to quick convergence to a solution, that is, a safe data set. Next we present the basic algorithm of this framework.

3.1 The Basic Algorithm

PARTIALSUPPRESSOR (Algorithm 1) presents the top-level algorithm. The partial suppressor iterates over the table T , and for each record $T[i]$, the algorithm first generates *qids* from $T[i]$ and sanitizes the unsafe ones. The suppressor terminates when the whole table is scanned and there is no unsafe *qid*.

Algorithm 1: PARTIALSUPPRESSOR(T, b_{\max})

```

1:  $T_0 \leftarrow T$  (original table)
2: loop
3:   Initialize the sup of all qids to 0
4:   while  $|B| < b_{\max}$  and  $i \leq |T|$  do
5:     Fill  $B$  with qids generated by  $T[i]$ 
6:     Update sup of all qids
7:      $i \leftarrow i + 1$ 
8:   end while
9:   if  $B$  contains an unsafe qids then
10:    SANITIZEBUFFER( $T_0, T, B$ )
11:    safe  $\leftarrow$  false
12:   end if
13:   if  $i \geq |T|$  and safe then
14:     break
15:   else if  $i \geq |T|$  then
16:      $i \leftarrow 1$ 
17:     safe  $\leftarrow$  true
18:   end if
19: end loop

```

A *qid* is a combination of different items, and the number of distinct *qids* to be enumerated is exponential. We therefore introduce a *qid* buffer of capacity b_{\max} to balance the space consumption with the generation time. The value of b_{\max} is significant. Small b_{\max} values cause repetitive generation of *qids*, while large b_{\max} values cause useless generation of *qids* which do not exist by the time to process them in the queue.

3.2 Buffer Sanitization

Each time *qid* buffer B is ready, SANITIZEBUFFER (Algorithm 2) is invoked to start processing *qids* in B and make all of them safe. $D_S(T)$ denotes the domain of all sensitive items in T . We first partition *qids* in B into two groups, *safe* and *unsafe*. Then in each iteration (Lines 2-18), SANITIZEBUFFER picks the “best” (according to heuristic functions H) unsafe sensitive association rule to sanitize (Lines 6 and 8). SUPPRESSIONPOLICY in SANITIZEBUFFER uses one of the the following two heuristic function.

Algorithm 2: SANITIZEBUFFER(T_0, T, B)

```

1:  $\mathcal{P} \leftarrow$  SUPPRESSIONPOLICY()
2: repeat
3:   pick an unsafe qid  $q$  from  $B$ 
4:    $E \leftarrow \{e \mid \text{conf}(q \rightarrow e) > \rho \wedge e \in D_S(T)\}$ 
5:   if  $\mathcal{P} = \text{Distribution}$  then
6:      $(d, q, e) \leftarrow \arg \max_{d \in q \cup E, q, e \in E} H_{\text{dist}}(d, q, e, T_0, T)$ 
7:   else if  $\mathcal{P} = \text{Mine}$  then
8:      $(d, q, e) \leftarrow \arg \min_{d \in q \cup E, q, e \in E} H_{\text{mine}}(d, q, e)$ 
9:   end if
10:   $X \leftarrow q \cup \{e\}$ 
11:   $k \leftarrow N_s(d, q \rightarrow e)$ 
12:  while  $k > 0$  do
13:    pick a record  $R$  from  $T$  where  $R \subseteq \mathcal{C}(X)$ 
14:     $R \leftarrow R - \{d\}$ 
15:    Update sup of qids contained in  $R$ 
16:     $k \leftarrow k - 1$ 
17:  end while
18: until there is no unsafe qid in  $B$ 

```

Preservation of Data Distribution Consider an unsafe sensitive association rule $q \rightarrow e$ where $\text{conf}(q \rightarrow e) > \rho$, and $q \in B$. To reduce $\text{conf}(q, e)$ below ρ , we suppress a number of items of type $t \in q \cup \{e\}$ from $\mathcal{C}(q \cup \{e\})$.³ We hope to minimize $KL(T \parallel T_0)$ (see Equation (3)). From Equation (3), we observe that by suppressing some items of type t where $T(t) > T_0(t)$,⁴ the KL divergence tends to decrease, thus we define the following heuristic function

$$H_{\text{dist}}(t, q, e, T_0, T) = \frac{T(t) \log \frac{T(t)}{T_0(t)}}{N_s(t, q \rightarrow e)}. \quad (5)$$

The maximizing this function aims at suppressing item type t which maximally recovers the original data distribution and minimizes the number of deletions.

³ We define $\mathcal{C}(X) = \{T[i] \mid X \subseteq T[i], 1 \leq i \leq |T|\}$.

⁴ We denote the probability of item type t in T as $T(t)$, which is computed by $\frac{\text{sup}_T(t)}{|T|}$.

Preservation of Useful Rules A spurious rule ($q \rightarrow e$) is introduced when the denominator of $conf(q \rightarrow e)$, $sup(q)$, is sufficiently small so that the confidence appears large enough. However, if $sup(q)$ is too small, the rule would not have enough support and can be ignored. Therefore, our objective is to suppress those items which have been suppressed before to minimize the support of the potential spurious rules. Therefore, we seek to minimize

$$H_{mine}(t, q, e) = leftover(t) \cdot N_s(t, q \rightarrow e)$$

3.3 Optimization with Divide-and-Conquer

When data is very large we can speed up by a divide-and-conquer (DnC) framework that partitions the input data dynamically, runs PARTIALSUPPRESSOR on them individually and combines the results in the end. This approach is correct in the sense that if each suppressed partition is safe, so is the combined data. This approach also gives rise to the parallel execution on multi-core or distributed environments which provides further speed-up (this will be shown in Section 4).

Algorithm 3: DNCSPLITDATA(T, t_{\max})

```

1: if  $Cost(T) > t_{\max}$  then
2:   Split  $T$  equally into  $T_1, T_2$ 
3:   DNCSPLITDATA( $T_1, t_{\max}$ )
4:   DNCSPLITDATA( $T_2, t_{\max}$ )
5: else
6:   PARTIALSUPPRESSOR( $T, t_{\max}$ )
7: end if

```

Algorithm 3 splits the input table whenever the estimated cost of suppressing that table is greater than t_{\max} . Cost is estimated as:

$$Cost(T) = \frac{|T| \cdot 2^{\frac{N}{|T|}}}{|D(T)|} \quad (6)$$

where N is the total number of items in T .

4 Experimental Results

We conducted a series of experiments on 4 main datasets in Table 2. BMS-POS and BMS-WebView are introduced in [20] and are commonly used for data mining. Retail is the retail market basket data [2]. Syn is the synthetic data in which each item is generated with equal probability and the max record length is 50. We randomly designate 40% of the item types in each dataset as

Table 2. Five Original Datasets

Dataset	Description	Recs	Dom. Size	Sensitive items	Non-Sens. items
BMS-POS (POS)	Point-of-sale data from a large electronics retailer	515597	1657	1183355	2183665
BMS-WebView (WV)	Click-stream data from e-commerce web site	77512	3340	137605	220673
Retail	Retail market basket data	88162	16470	340462	568114
Syn	Synthetic data with max record length = 50	493193	5000	828435	1242917

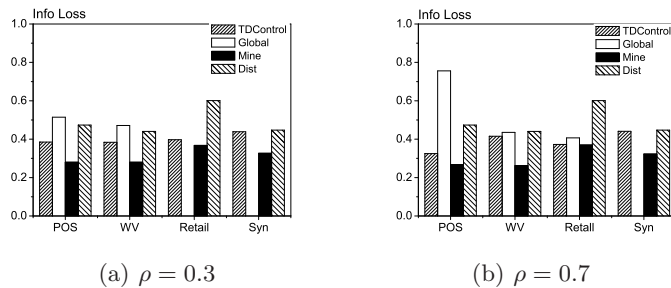


Fig. 1. Comparisons in Information Loss

sensitive items and the rest as non-sensitive. To evaluate the performance of rule mining, we produce four additional datasets by truncating all records in the original datasets to 5 items only, and denote such datasets as “cutoff = 5”.

We compare our algorithm with the global suppression algorithm (named Global) and generalization algorithm (named TDCControl) of Cao *et al.* [3].⁵ Our algorithm has two variants, *Dist* and *Mine*, which optimize for data distribution and rule mining, respectively. Experiments that failed to complete in 2 hours is marked as “N/A” or an empty place in the bar charts. We run all the experiments on Linux 2.6.34) with an Intel 16-core 2.4GHz CPU and 8GB RAM.

In what follows, we first present results in data utility, then the performance of the algorithms, and then the effects of changing various parameters in our algorithm. Finally, we compare with a permutation method which utilizes a similar privacy model but with different optimization goals. Unless otherwise noted, we use the following default parameters: $b_{max} = 10^6$, $t_{max} = 500$.

4.1 Data Utility

We compare the algorithms in terms of information loss, data distribution, and association rule mining.

⁵ The source code of these algorithms was directly obtained from Cao.

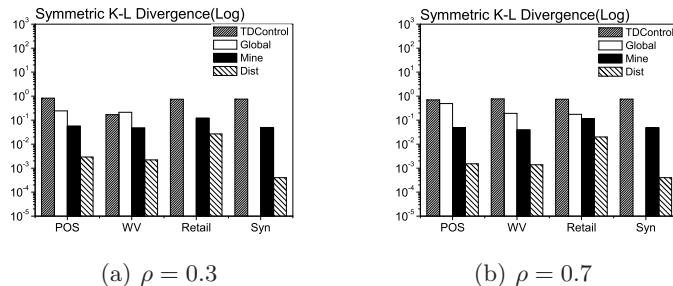


Fig. 2. Comparisons in Symmetric K-L Divergence

Figure 1 shows that *Mine* is uniformly better among the other four techniques. It suppresses only about 26% items in POS and WV and about 35% items in Retail, while the other four techniques incur on average 10% more losses than *Mine* and up to 75% losses in the worst case. We notice that *Dist* performs worse than *Global* even though it tries to minimize the information loss at each iteration. The reason is that it also tries to retain the data distribution. Further, we argue that for applications that require data statistics, the distribution, that is, summary information, is more useful than the details, hence losing some detailed information is acceptable. Note that *Global* and *TDCControl* failed to complete in some datasets, because these methods don't scale very well.

To determine the similarity between the item frequency distribution of original data and that of the anonymized data, we use the Kullback-Leibler divergence (also called relative entropy) as our standard. To prevent zero denominators, we modified Equation (4) to a symmetric form [6] defined as

$$\mathcal{S}(H_1||H_2) = \frac{1}{2}KL(H_1||H_1 \oplus H_2) + \frac{1}{2}KL(H_2||H_1 \oplus H_2)$$

where $H_1 \oplus H_2$ represents the union of distributions H_1 and H_2 . Figure 2 shows that *Dist* outperforms the peers as its output has the highest resemblance to the original datasets. On the contrary, *TDCControl* is the worst performer since generalization algorithm creates a lot of new items while suppressing too many item types globally. Since the symmetric relative entropy of *Dist* is very small, y-axis is in logarithmic scale to improve visibility. Therefore, the actual difference in K-L divergence is two or three orders of magnitude.

The most common criticism of partial suppression is that it changes the support of good rules in the data and introduces spurious rules in rule mining. In this experiment, we test the algorithms on data sets with the max record length=5 (cutoff=5), and check the rules mined from the anonymized data with support equals to 0.05% ⁶ and confidence equals to 70% and 30%. Figure 3 gives the results. Both *TDCControl* and *Global* perform badly in this category, with negligible number of original rules remaining after anonymization. Conversely, all

⁶ We choose this support level just to reflect a practical scenario.

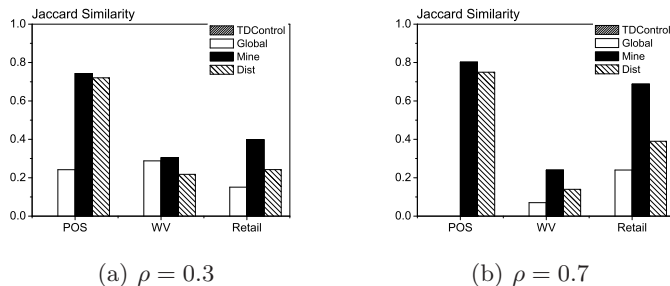


Fig. 3. Association Rules Mining with Support 0.05%

of the partial suppression algorithms manage to retain most of the rules and the Jaccard Similarity reaches 80% in some datasets which shows our heuristic works very well. Specifically, *Mine* performs the best among partial algorithms. The rules generated from TDControl are all in general form which is totally different from the original one. To enable comparison, we specialize the more general rules from the result of TDControl into rules of original level of abstraction in the generalization hierarchy. For example, we can specialize a rule {dairy product \rightarrow grains} into: milk \rightarrow wheat, milk \rightarrow rice, yogurt \rightarrow wheat, etc. Take WV as an example, there are 4 rules left in the result of TDControl when the ρ is 0.7 and the number becomes 28673 after specialization, which makes the results almost invisible.

4.2 Performance

Next we evaluate the time performance and scalability of our algorithms.

Table 3. Comparison in Time Performance ($\rho = 0.7, t_{max} = 300$)

Algorithm	POS	WV	Retail	Syn
TDControl	183	30	156	476
Global	1027	81	646	N/A
<i>Dist</i>	395	151	171	130
<i>Mine</i>	1554	478	256	132

From Table 3, TDControl is the clear winner for two of the four datasets. *Mine* does not perform well in BMS-POS. The reason is that *Mine* incurs the least information loss among all the competing methods. This means most of the original data remains unsuppressed. Given the large scale of BMS-POS, checking whether the dataset is safe in each iteration is therefore more time consuming than other methods or in other datasets. Results for Global are not available for Syn because it runs out of memory.

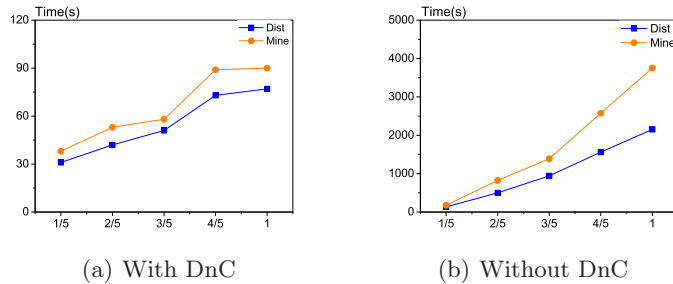


Fig. 4. Scale-up with Input Data ($\rho = 0.7$)

Next experiment illustrates the scalability of our algorithm w.r.t. data size. We choose *Retail* as our target dataset here because *Retail* has the maximum average record length of 10.6. We run partial algorithms on 1/5, 2/5 through 5/5 of *Retail* respectively. Figure 4 shows the time cost of our algorithm increases reasonably with the input data size with or without DnC. Furthermore, increased level of partitioning causes the algorithm to witness superlinear speedup in Figure 4(a). In particular, the dataset is automatically divided into 4, 8, 16 and 32 parts at 1/5, 2/5, 3/5 and the whole of the data, respectively.

4.3 Effects of Parameters on Performance

In this section, we study the effects of t_{max} , b_{max} on the quality of solution (in terms of information loss) and time performance.

We choose *Retail* as the target dataset again since *Retail* is the most time-consuming dataset that can terminate within acceptable time without DnC strategy. The value of t_{max} determines the size of a partition in DnC. Here, we evaluate how partitioning helps with time performance and its possible effects on suppression quality. Figure 5(a) shows the relationship between partitions and information loss. The lines of *Dist* is flat, indicating that increasing t_{max} doesn't cost us the quality of the solution. *Mine* shows a slight descending tendency at first and then tends to be flat. We argue that a reasonable t_{max} will not cause our result quality to deteriorate. On the other hand, Figure 5(b) shows that time cost increases dramatically with the increase of t_{max} . The reason is that partitioning decreases the cost of enumerating *qids* which is the most time-consuming part in our algorithm. Moreover, parallel processing is also a major reason for the acceleration.

Next experiment (See Figure 6) illustrates the impact of varying b_{max} on performance. We choose *WV* as our target dataset since the number of distinct *qids* are relatively smaller than other datasets and our algorithm can terminate even when we set a small b_{max} .

Note first that varying b_{max} has no effect on the information loss which indicates that this parameter is purely for performance tuning. At lower values,

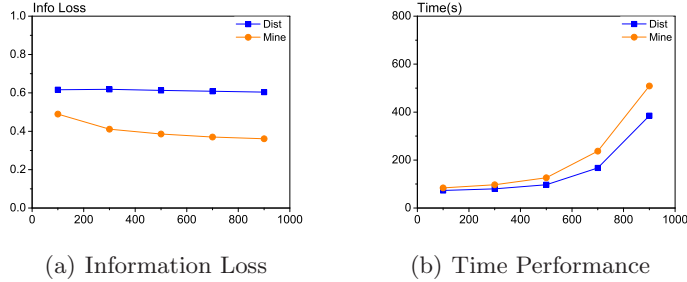


Fig. 5. Variation of t_{max} ($\rho = 0.7$)

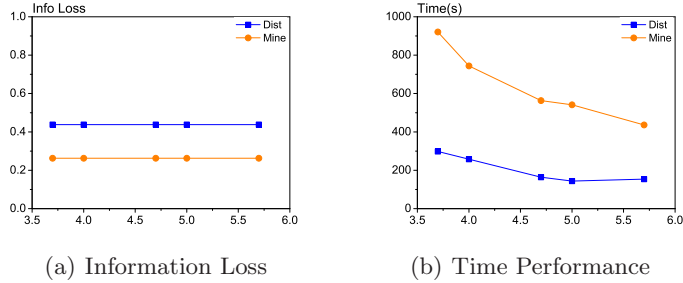


Fig. 6. Variation of Buffer Size b_{max} ($\rho = 0.7$)

increasing b_{max} gives almost exponential savings in running time. But as b_{max} reaches a certain point, the speedup saturates, which suggests that given the fixed size of the data, when B is large enough to accommodate all $qids$ at once after some iterations, further increase in b_{max} is not useful. The line for *Mine* hasn't saturated because *Mine* suppresses fewer items and retains more $qids$, hence requires a much larger buffer.

4.4 A Comparison to Permutation Method

In this section, we compare our algorithms with a permutation method [8] which we call M . The privacy model of M states that the probability of associating any transaction $R \in T$ with any sensitive item $e \in D_S(T)$ is below $1/p$, where p is known as a privacy degree. This model is similar to ours when $\rho = 1/p$, which allows us to compare three variants of our algorithm against M where $p = 4, 6, 8, 10$ on dataset WV which was reported in [8]. Figure 7(a) shows the result on K-L divergence. All variants of our algorithm outperform M in preserving the data distribution. Figure 7(b) shows timing results. Even though M is faster, our algorithms terminate within acceptable time.

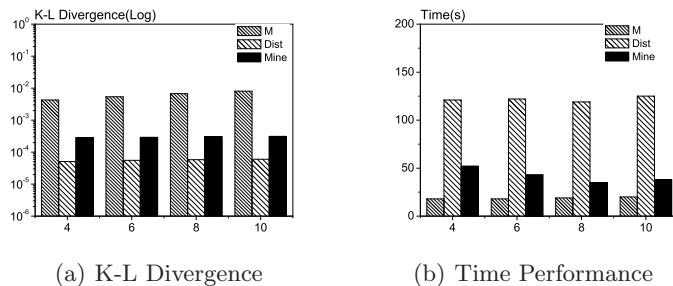


Fig. 7. Comparison with Permutation

5 Related Work

Privacy-preserving data publishing of relational tables has been well studied in the past decade since the original proposal of k -anonymity by Sweeney *et al.* [12]. Recently, privacy protection of set-valued data has received increasing interest. The original set-valued data privacy problem was defined in the context of association rule hiding [1, 15, 16], in which the data publisher wishes to “sanitize” the set-valued data (or *micro-data*) so that all sensitive or “bad” associate rules cannot be discovered while all (or most) “good” rules remain in the published data. Subsequently, a number of privacy models including (h, k, p) -coherence [18], k^m -anonymity [14], k -anonymity [9] and ρ -uncertainty [3] have been proposed. k^m -anonymity and k -anonymity are carried over directly from relational data privacy, while (h, k, p) -coherence and ρ -uncertainty protect the privacy by bounding the confidence and the support of any sensitive association rule inferrable from the data. This is also the privacy model this paper adopts.

A number of anonymization techniques were developed for these models. These generally fall in four categories: *global/local generalization* [14, 9, 3], *global suppression* [18, 3], *permutation* [8] and *perturbation* [19, 4]. Next we briefly discuss the pros and cons of these anonymization techniques.

Generalization replaces a specific value by a generalized value, e.g., “milk” by “dairy product”, according to a generalization hierarchy [7]. While generalization preserves the correctness of the data, it compromises accuracy and preciseness. Worse still, association rule mining is impossible unless the data users have access to the same generalization taxonomy and they agree to the target level of generalization. For instance, if the users don’t intend to mine rules involving “dairy products”, then all generalizations to “dairy products” are useless.

Global suppression is a technique that deletes all items of some types so that the resulting dataset is safe. The advantage is that it preserves the support of existing rules that don’t involve deleted items and hence retains these rules [18], and also it doesn’t introduce additional/spurious association rules. The obvious disadvantage is that it can cause unnecessary information loss. In the past, partial suppression has not been attempted mainly due to its perceived side effects of changing the support of inference rules in the original data [18, 3, 15, 16]. But

our work shows that partial suppression introduces limited amount of new rules while preserving many more original ones than global suppression. Furthermore, it preserves the data distribution much better than other competing methods.

Permutation was introduced by Xiao *et al.* [17] for relational data and was extended by Ghinita *et al.* [8] for transactional data. Ghinita *et al.* propose two novel anonymization techniques for sparse high-dimensional data by introducing two representations for transactional data. However the limitation is that the quasi-identifier is restricted to contain only *non-sensitive items*, which means they only consider associations between quasi-identifier and sensitive items, and not *among* sensitive items. Manolis *et al.* [13] introduced “disassociation” which also severs the links between values attributed to the same entity but does not set a clear distinction between sensitive and non-sensitive attributes. In this paper, we consider all kinds of associations and try best to retain them.

Perturbation is developed for statistical disclosure control [7]. Common perturbation methods include *additive noise*, *data swapping*, and *synthetic data generation*. Their common criticism is that they damage the data integrity by adding noises and spurious values, which makes the results of downstream analysis unreliable. Perturbation, however, is useful in non-deterministic privacy model such as differential privacy [5], as attempted by Chen *et al.* [4] in a probabilistic top-down partitioning algorithm based on a context-free taxonomy.

The most relevant work to this paper is by Xu *et al.* [18] and Cao *et al.* [3]. The (h, k, p) -coherence model by Xu *et al.* requires that the attacker’s prior knowledge to be no more than p public (non-sensitive) items, and any inferable rule must be supported by at least k records while the confidence of such rules is at most $h\%$. They believe private items are essential for research and therefore only remove public items to satisfy the privacy model. They developed an efficient greedy algorithm using global suppression. In this paper, we do not restrict the size or the type of the background knowledge, and we use a partial suppression technique to achieve less information loss and also better retain the original data distribution.

Cao *et al.* [3] proposed a similar ρ -uncertainty model which is used in this paper. They developed a global suppression method and a top-down generalization-driven global suppression method (known as TDCControl) to eliminate all sensitive inferences with confidence above a threshold ρ . Their methods suffer from same woes discussed earlier for generalization and global suppression. Furthermore, TDCControl assumes that data exhibits some monotonic property under a generalization hierarchy. This assumption is questionable. Experiments show that our algorithm significantly outperforms the two methods in preserving data distribution and useful inference rules, and in minimizing information losses.

6 Conclusion

We proposed a partial suppression framework including two heuristics which produce anonymized data that is highly useful to data analytics applications. Compared to previous approaches, this framework generally deletes fewer items

to satisfy the the same privacy model. We showed that the first heuristic can effectively limit spurious rules while maximally preserving the useful rules mineable from the original data. The second heuristic which minimizes the K-L divergence between the anonymized data and the original data helps preserve the data distribution, which is a feature largely ignored by the privacy community in the past. Finally the divide-and-conquer strategy effectively controls the execution time with limited compromise in the solution quality.

References

1. M. Atallah, E. Bertino, A. Elmagarmid, M. Ibrahim, and V. Verykios. Disclosure limitation of sensitive rules. In *KDEX*, 1999.
2. T. Brijs, G. Swinnen, K. Vanhoof, and G. Wets. Using association rules for product assortment decisions: A case study. In *Knowledge Discovery and Data Mining*, pages 254–260, 1999.
3. J. Cao, P. Karras, C. Raïssi, and K.-L. Tan. ρ -uncertainty: inference-proof transaction anonymization. *VLDB*, pages 1033–1044, 2010.
4. R. Chen, N. Mohammed, B. C. M. Fung, B. C. Desai, and L. Xiong. Publishing set-valued data via differential privacy. *VLDB*, pages 1087–1098, 2011.
5. C. Dwork. Differential privacy: A survey of results. In *TAMC*, pages 1–19, 2008.
6. K. Fisher, D. Walker, K. Q. Zhu, and P. White. From dirt to shovels: fully automatic tool generation from ad hoc data. In *POPL*, pages 421–434, 2008.
7. B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu. Privacy-preserving data publishing: A survey of recent developments. *ACM Comput. Surv.*, 2010.
8. G. Ghinita, P. Kalnis, and Y. Tao. Anonymous publication of sensitive transactional data. *TKDE*, pages 161–174, 2011.
9. Y. He and J. F. Naughton. Anonymization of set-valued data via top-down, local generalization. *VLDB*, pages 934–945, 2009.
10. P. Jaccard. The distribution of the flora in the alpine zone. *New Phytologist*, 11:37–50, 1912.
11. S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 21(1):79–86, 1951.
12. L. Sweeney. k-anonymity: a model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, pages 557–570, 2002.
13. M. Terrovitis, J. Liagouris, N. Mamoulis, and S. Skiadopoulos. Privacy preservation by disassociation. *PVLDB*, 2012.
14. M. Terrovitis, N. Mamoulis, and P. Kalnis. Privacy-preserving anonymization of set-valued data. *VLDB*, pages 115–125, 2008.
15. V. S. Verykios, A. K. Elmagarmid, E. Bertino, Y. Saygin, and E. Dasseni. Association rule hiding. *TKDE*, pages 434–447, 2004.
16. Y.-H. Wu, C.-M. Chiang, and A. L. P. Chen. Hiding sensitive association rules with limited side effects. *TKDE*, pages 29–42, 2007.
17. X. Xiao and Y. Tao. Anatomy: simple and effective privacy preservation. In *PVLDB*, pages 139–150, 2006.
18. Y. Xu, K. Wang, A. W.-C. Fu, and P. S. Yu. Anonymizing transaction databases for publication. In *KDD*, pages 767–775, 2008.
19. Q. Zhang, N. Koudas, D. Srivastava, and T. Yu. Aggregate query answering on anonymized tables. In *ICDE*, pages 116–125, 2007.
20. Z. Zheng, R. Kohavi, and L. Mason. Real world performance of association rule algorithms. In *KDD*, pages 401–406, 2001.