Blocking Anonymity Threats Raised by Frequent Itemset Mining

Maurizio Atzori[†‡] Francesco Bonchi[†] Fosca Giannotti[†] Dino Pedreschi[‡]

[†]Pisa KDD Laboratory
ISTI - CNR, Area della Ricerca di Pisa
Via Giuseppe Moruzzi, 1 - 56124 Pisa, Italy

[‡]Pisa KDD Laboratory
Computer Science Dep., University of Pisa
Largo Pontecorvo, 3 - 56127 Pisa, Italy

Abstract

In this paper we study when the disclosure of data mining results represents, per se, a threat to the anonymity of the individuals recorded in the analyzed database. The novelty of our approach is that we focus on an objective definition of privacy compliance of patterns without any reference to a preconceived knowledge of what is sensitive and what is not, on the basis of the rather intuitive and realistic constraint that the anonymity of individuals should be guaranteed. In particular, the problem addressed here arises from the possibility of inferring from the output of frequent itemset mining (i.e., a set of itemsets with support larger than a threshold σ), the existence of patterns with very low support (smaller than an anonymity threshold k)[3]. In the following we develop a simple methodology to block such inference opportunities by introducing distortion on the dangerous patterns.

1. Introduction

Consider a medical institution where the physicians collect, for research purposes, data about their patients. Having unrestricted access to the data, they can perform real mining on all available information using traditional mining tools – not necessarily the privacy preserving ones. This way they maximize the outcome of the knowledge discovery process, without any concern about privacy of the patients which are recorded in the data. But the anonymity of patients becomes a hot issue when the physicians want to share their discoveries (e.g., association rules holding in the data) with their scientific community. It is generally believed that data mining results do not violate the *anonymity* of the individuals recorded in the source database. In fact, data mining models and patterns, in order to ensure a required statistical significance, represent a large number of individuals and thus conceal individual identities: this is the case of the minimum support threshold in association rule mining. We have recently shown that the above belief is ill-founded [3].

Example 1 *Consider the following association rule:*

$$a_1 \wedge a_2 \wedge a_3 \Rightarrow a_4 \quad [sup = 80,\ conf = 98.7\%]$$

where sup and conf are the usual interestingness measures of support and confidence as defined in [2]. Since the given rule holds for a number of individuals (80), which seems large enough to protect individual privacy, one could conclude that the given rule can be safely disclosed. But, is this all the information contained in such a rule? Indeed, one can easily derive the support of the premise of the rule:

$$sup(\{a_1, a_2, a_3\}) = \frac{sup(\{a_1, a_2, a_3, a_4\})}{conf} \approx \frac{80}{0.987} = 81.05$$

Given that the pattern $a_1 \wedge a_2 \wedge a_3 \wedge a_4$ holds for 80 individuals, and that the pattern $a_1 \wedge a_2 \wedge a_3$ holds for 81 individuals, we can infer that in our database there is just one individual for which the pattern $a_1 \wedge a_2 \wedge a_3 \wedge \neg a_4$ holds.

In [3] we say that the two itemsets $\{a_1, a_2, a_3\}$ and $\{a_1, a_2, a_3, a_4\}$ represent an *inference channel* (a simple one), for the pattern $a_1 \wedge a_2 \wedge a_3 \wedge \neg a_4$.

By shifting the concept of *k-anonymity* [6] from data to patterns, we have formally characterized the notion of a threat to anonymity in the context of pattern discovery, and provided a methodology to efficiently and effectively identify such threats that might arise from the disclosure of a set of extracted patterns. In this paper we study a methodology to eliminate the threats to anonymity by introducing distortion on the dangerous patterns in a controlled way, by measuring the effects of the distortion.

2. Problem Definition

Definition 1 *A binary database $\mathcal{D} = (\mathcal{I}, \mathcal{T})$ consists of a finite set of binary variables $\mathcal{I} = \{i_1, \ldots, i_p\}$, also known as* items*, and a finite multiset $\mathcal{T} = \{t_1, \ldots, t_n\}$ of p-dimensional binary vectors (*transactions*) recording the values of the items. A* pattern *for the variables in $\mathcal{I}$ is a propositional sentence built by AND ($\wedge$), OR ($\vee$) and NOT ($\neg$) logical connectives, on variables in $\mathcal{I}$. The domain of all possible patterns is denoted $\mathcal{P}at(\mathcal{I})$.*

$\mathcal{D}$

	a	b	c	d	e	f	g	h
t_1	1	1	1	1	1	1	1	1
t_2	1	1	1	1	1	0	1	0
t_3	1	1	1	1	1	0	0	0
t_4	1	1	1	1	1	1	1	0
t_5	1	1	1	1	1	0	0	0
t_6	1	1	1	1	1	0	0	0
t_7	1	1	0	1	1	0	0	0
t_8	1	0	0	0	1	1	1	0
t_9	0	0	1	1	1	1	1	0
t_{10}	0	0	1	1	1	0	0	0
t_{11}	0	0	1	1	1	1	1	1
t_{12}	1	1	0	0	0	1	1	0

(a)

Notation: patterns

$sup_{\mathcal{D}}(a \vee f) = 11$
$sup_{\mathcal{D}}(e \wedge (\neg b \vee \neg d)) = 4$
$sup_{\mathcal{D}}(h \wedge \neg b) = 1$

Notation: itemsets

$sup_{\mathcal{D}}(abc) = 6$
$sup_{\mathcal{D}}(abde) = 7$
$sup_{\mathcal{D}}(cd) = 9$

(b)

$\mathcal{F}(\mathcal{D}, 8) = \{\langle\emptyset, 12\rangle, \langle a, 9\rangle, \langle b, 8\rangle, \langle c, 9\rangle, \langle d, 10\rangle, \langle e, 11\rangle, \langle ab, 8\rangle, \langle ae, 8\rangle, \langle cd, 9\rangle, \langle ce, 9\rangle, \langle de, 10\rangle, \langle cde, 9\rangle\}$

(c)

$\mathcal{C}l(\mathcal{D}, 8) = \{\langle\emptyset, 12\rangle, \langle a, 9\rangle, \langle e, 11\rangle, \langle ab, 8\rangle, \langle ae, 8\rangle, \langle de, 10\rangle, \langle cde, 9\rangle\}$

(d)

$\mathcal{MC}h(3, \mathcal{C}l(\mathcal{D}, 8)) = \{\langle \mathcal{C}_{\emptyset}^{cde}, 1\rangle, \langle \mathcal{C}_{a}^{ab}, 1\rangle, \langle \mathcal{C}_{a}^{ae}, 1\rangle, \langle \mathcal{C}_{e}^{cde}, 1\rangle, \langle \mathcal{C}_{de}^{cde}, 1\rangle\}$

(e)

Figure 1. Example: (a) the binary database $\mathcal{D}$; (b) different notation used for patterns and itemsets; the set of frequent ($\sigma = 8$)(c), and of closed (d) itemsets ; (e) the set of maximal inference channels for $k = 3$.

Definition 2 *Given a database $\mathcal{D}$, a transaction $t \in \mathcal{D}$ and a pattern p, we write $p(t)$ if t makes p true. The support of p in $\mathcal{D}$ is given by the number of transactions which make p true: $sup_{\mathcal{D}}(p) = |\{t \in \mathcal{D} \mid p(t)\}|$.*

Definition 3 *The set of all itemsets $2^{\mathcal{I}}$, is a pattern class consisting of all possible conjunctions of the form $i_1 \wedge i_2 \wedge \ldots \wedge i_m$. Given a database $\mathcal{D}$ and a support threshold σ, the set of σ-frequent itemsets in $\mathcal{D}$ is denoted $\mathcal{F}(\mathcal{D}, \sigma) = \{\langle X, sup_{\mathcal{D}}(X)\rangle \mid X \in 2^{\mathcal{I}} \wedge sup_{\mathcal{D}}(X) \geq \sigma\}$.*

The problem addressed in this paper is given by the possibility of inferring from *the output of frequent itemset mining*, i.e, $\mathcal{F}(\mathcal{D}, \sigma)$, *the existence of patterns with very low support*, that represent a threat for the anonymity of the individuals about which they are true (Figure 1(b) shows the different notation used for *general patterns* and for *itemsets*).

Definition 4 *A set S of pairs $\langle X, n\rangle$, where $X \in 2^{\mathcal{I}}$ and $n \in \mathbb{N}$, and a database $\mathcal{D}$ are said to be σ-compatible if $S \subseteq \mathcal{F}(\mathcal{D}, \sigma)$. Given a pattern p we say that $S \models sup(p) > x$ (respectively $S \models sup(p) < x$) if, for all databases $\mathcal{D}$ σ-compatible with S, we have that $sup_{\mathcal{D}}(p) > x$ (respectively $sup_{\mathcal{D}}(p) < x$).*

In general [4], the support of a pattern $p = i_1 \wedge \cdots \wedge i_m \wedge \neg a_1 \wedge \cdots \wedge \neg a_n$ can be inferred if we know the support of itemsets $I = \{i_1, \ldots, i_m\}$, $J = I \cup \{a_1, \ldots, a_n\}$, and every itemset X such that $I \subset X \subset J$.

$$sup_{\mathcal{D}}(p) = \sum_{I \subseteq X \subseteq J} (-1)^{|X \setminus I|} sup_{\mathcal{D}}(X) \qquad (1)$$

The right-hand side of Equation (1) is denoted $f_I^J(\mathcal{D})$.

Definition 5 *Given a database $\mathcal{D}$ and two itemsets $I \subseteq J \in 2^{\mathcal{I}}$, if $0 < f_I^J(\mathcal{D}) < k$, then the set of itemsets $\{X | I \subseteq X \subseteq J\}$ is an inference channel. We denote such inference channel $\mathcal{C}_I^J$ and we write $sup_{\mathcal{D}}(\mathcal{C}_I^J) = f_I^J(\mathcal{D})$.*

Example 2 *Consider $\mathcal{D}$ in Figure 1, and suppose $k = 3$. We have that $sup_{\mathcal{D}}(b \wedge \neg d \wedge \neg e) = f_b^{bde}(\mathcal{D}) = sup_{\mathcal{D}}(b) - sup_{\mathcal{D}}(bd) - sup_{\mathcal{D}}(be) + sup_{\mathcal{D}}(bde) = 8 - 7 - 7 + 7 = 1$. Therefore, $\mathcal{C}_b^{bde}$ is an inference channel of support 1.*

In [3] we have shown that, since a generic pattern $p \in \mathcal{P}at(\mathcal{I})$ can be considered without loss of generality in *normal disjunctive form*, we can conclude that all possible threats to anonymity are due to inference channels of the form $\mathcal{C}_I^J$. However two cases can be distinguished: (i) I and J are both frequent itemsets; (ii) J is not frequent. For lack of space we focus only on the first and most essential form of channels.

Definition 6 *The set of all $\mathcal{C}_I^J$ holding in $\mathcal{F}(\mathcal{D}, \sigma)$, together with their supports, is denoted $\mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma)) = \{\langle \mathcal{C}_I^J, f_I^J(\mathcal{D})\rangle \mid \langle J, sup_{\mathcal{D}}(J)\rangle \in \mathcal{F}(\mathcal{D}, \sigma)\}$.*

In our previous paper [3], we have studied the problem of how to detect the inference channels holding in a set of frequent itemsets that we want to disclose (i.e., computing $\mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$). In this paper we study how to transform (sanitize) the collection to remove the inference channels.

Problem 1 *Given a collection of frequent itemsets $\mathcal{F}(\mathcal{D}, \sigma)$, and the set of all its inference channels $\mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$, transform $\mathcal{F}(\mathcal{D}, \sigma)$ in a collection of frequent itemsets $\mathcal{O}^k$, which can be safely disclosed. $\mathcal{O}^k$ is the output of our problem, and it must satisfy the following conditions: (i) $\mathcal{C}h(k, \mathcal{O}^k) = \emptyset$; (ii) $\exists \mathcal{D}' : \mathcal{O}^k = \mathcal{F}(\mathcal{D}', \sigma)$.*

The second condition constraints the output collection of itemsets to be *"realistic"*; i.e., to be compatible with at least a database. This requirement is due to the fact that disclosing an output which is not compatible with any database could represent a threat. In fact, a malicious adversary could recognize that the set

of itemsets disclosed is not "real", and he could exploit this leak by reconstructing the missing patterns, starting from those ones present in the output. We call this kind of threat *inverse mining attacks.*

3. Avoiding Redundant Distortion

In this Section we show how, using a condensed representation of $\mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$, we can avoid redundant distortion when blocking inference channels. Consider the two inference channels $\langle \mathcal{C}_{ad}^{acd}, 1\rangle$ and $\langle \mathcal{C}_{abd}^{abcd}, 1\rangle$ holding in the database in Fig. 1(a): one is more specific than the other, but they both uniquely identify transaction t_7. It is easy to see that many other families of equivalent, and thus redundant, inference channels can be found. The theory of *closed itemsets* [5] can help us with this problem.

Definition 7 *The set of frequent closed itemsets is* $\mathcal{C}l(\mathcal{D}, \sigma) = \{\langle X, sup_{\mathcal{D}}(X)\rangle \in \mathcal{F}(\mathcal{D}, \sigma) | \nexists Y \supset X s.t. \langle Y, sup_{\mathcal{D}}(X)\rangle \in \mathcal{F}(\mathcal{D}, \sigma)\}$. *An itemset* $I \in \mathcal{C}l(\mathcal{D}, \sigma)$ *is said to be maximal iff* $\nexists J \supset I$ *s.t.* $\langle J, s\rangle \in \mathcal{C}l(\mathcal{D}, \sigma)$.

Analogously to what happens for the pattern class of itemsets, if we consider the pattern class of conjunctive patterns we can rely on the *anti-monotonicity property of frequency.*

Proposition 1 *Given* $\mathcal{C}_I^J$ *and* $\mathcal{C}_H^L$ *we say that* $\mathcal{C}_I^J \preceq \mathcal{C}_H^L$ *when* $I \subseteq H$ *and* $(J \setminus I) \subseteq (L \setminus H)$. *It holds that:* $\mathcal{C}_I^J \preceq \mathcal{C}_H^L \Rightarrow \forall \mathcal{D}\ .\ f_I^J(\mathcal{D}) \geq f_H^L(\mathcal{D})$.

Definition 8 *An inference channel* $\mathcal{C}_I^J$ *is said to be* maximal *w.r.t.* $\mathcal{D}$ *and* σ, *if* $\forall H, L$ *such that* $I \subseteq H$ *and* $(J \setminus I) \subseteq (L \setminus H)$, $f_H^L = 0$. *The set of maximal inference channels is denoted* $\mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$.

Proposition 2
$\mathcal{C}_I^J \in \mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma)) \Rightarrow I \in \mathcal{C}l(\mathcal{D}, \sigma) \wedge J$ *is maximal.*

In [3] we have proved that from $\mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$ we can reconstruct every channel in $\mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$ and its support without accessing the database.

Proposition 3 *Given* $\langle \mathcal{C}_I^J, f_I^J(\mathcal{D})\rangle \in \mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$, *let* M *be any maximal itemset such that* $M \supseteq J$. *The following equation holds:*

$$f_I^J(\mathcal{D}) = \sum_{c(X)} f_{c(X)}^M(\mathcal{D}) \tag{2}$$

where $c(I) \subseteq c(X) \subseteq M$ *and* $c(X) \cap (J \setminus I) = \emptyset$.

As the set of all closed frequent itemsets $\mathcal{C}l(\mathcal{D}, \sigma)$ contains all the information of $\mathcal{F}(\mathcal{D}, \sigma)$ in a more compact representation, analogously the set $\mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$ represents, without redundancy, all the information in $\mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$. Removing the redundancy existing in $\mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$, we also implicitly avoid redundant sanitization, and thus we dramatically reduce the distortion needed to block all the inference channels. In fact, to block an inference channel $\mathcal{C}_I^J \in \mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$ we have two main options:

- making the inference channel anonymous enough, i.e., forcing $f_I^J(\mathcal{D}) \geq k$;
- making the inference channel disappear, i.e., forcing $f_I^J(\mathcal{D}) = 0$.

The following two propositions show that, whichever option we choose, we can just block the channels in $\mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$, obtaining to block all the inference channels in $\mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$.

Proposition 4 *Given a database* $\mathcal{D}$, *consider a database* $\mathcal{D}'$ *s.t.* $\forall \langle \mathcal{C}_H^L, f_H^L(\mathcal{D})\rangle \in \mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$ *it holds that* $f_H^L(\mathcal{D}') \geq k$. *Then from Proposition 1 it follows that* $\forall \langle \mathcal{C}_I^J, f_I^J(\mathcal{D})\rangle \in \mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$, $f_I^J(\mathcal{D}') \geq k$.

Proposition 5 *Given a database* $\mathcal{D}$, *consider a database* $\mathcal{D}'$ *s.t.* $\forall \langle \mathcal{C}_H^L, f_H^L(\mathcal{D})\rangle \in \mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$ *it holds that* $f_H^L(\mathcal{D}') = 0$. *Then from Proposition 3 it follows that* $\forall \langle \mathcal{C}_I^J, f_I^J(\mathcal{D})\rangle \in \mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$, $f_I^J(\mathcal{D}') = 0$.

In the next Section we exploit the properties above to reduce the distortion needed to sanitize our output.

4. Suppressive Sanitization

The basic idea of Suppressive Sanitization is to hide inference channels, pushing their support to 0: this can be done by removing transactions t s.t. $I \subseteq t \wedge (J \setminus I) \cap t = \emptyset$. Unfortunately, we can not simulate such suppression of transactions simply by decreasing the support of the itemset I by f_I^J for each $\mathcal{C}_I^J \in \mathcal{C}h(k, \mathcal{F}(\mathcal{D}, \sigma))$, since we would lose database-compatibility due to the other items appearing in the dangerous transactions. Consider for instance a transaction $I \cup \{x, y, z\}$: removing it we reduce the support of I, but as uncontrolled side effect, we also reduce the support of the itemset $I \cup \{x\}$. Therefore, in order to maintain database-compatibility, we must take into account these other items carefully. One naïve way of achieving this is to really access the database, suppress the dangerous transactions, and reduce the support of all itemsets contained in the suppressed transactions accordingly. But this is not enough: while introducing distortion to block the "real" inference channels holding in $\mathcal{F}(\mathcal{D}, \sigma)$, transforming it in $\mathcal{O}^k$, we could possibly create some new "fake" inference channels (not existing in the original database and thus not violating the anonymity of real individuals). We do not allow this possibility: although fake, such inference channels could be the starting point for a backward reasoning of a malicious adversary, in other terms, could open the door to inverse mining attacks. The unique solution here is to perform again the detection algorithm [3], and if necessary, to block the novel inference

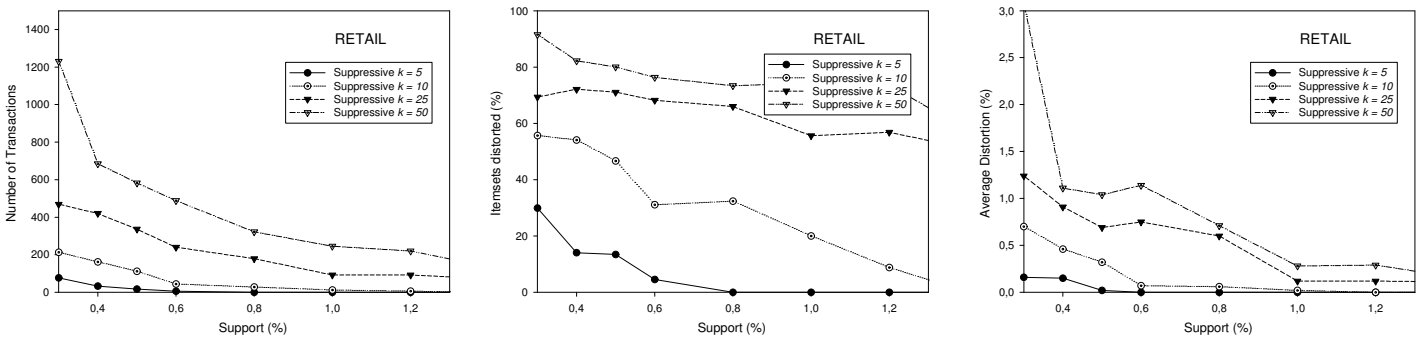

Figure 2. Distortion empirical evaluation.

channels found. Obviously, this process can make some frequent itemsets become infrequent.

Algorithm 1 implements the suppressive sanitization which access the database $\mathcal{D}$ on the basis of the information in $\mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$, and adjust $\mathcal{C}l(\mathcal{D}, \sigma)$ with the information found in $\mathcal{D}$: the following pseudo-code outputs a sanitized version of $\mathcal{C}l(\mathcal{D}, \sigma)$ but nothing prevents us from disclosing a sanitized version of $\mathcal{F}(\mathcal{D}, \sigma)$.

Example 3 *Consider* $\mathcal{C}l(\mathcal{D}, 8)$ *and* $\mathcal{MC}h(3, \mathcal{C}l(\mathcal{D}, 8))$ *in Fig.1(d) and (e): we got 5 maximal channels* $\langle \mathcal{C}_{\emptyset}^{cde}, 1\rangle, \langle \mathcal{C}_{a}^{ab}, 1\rangle, \langle \mathcal{C}_{a}^{ae}, 1\rangle, \langle \mathcal{C}_{e}^{cde}, 1\rangle, \langle \mathcal{C}_{de}^{cde}, 1\rangle$ *due to transactions* t_{12}, t_8 *and* t_7*. The suppression of these 3 transactions reduces the support of some closed itemsets: at the end of the suppression phase (line 8 of Algorithm 1) we got that* $\mathcal{C}l(\mathcal{D}, 8) = \{\langle \emptyset, 9\rangle, \langle a, 6\rangle, \langle e, 9\rangle, \langle ab, 6\rangle, \langle ae, 6\rangle, \langle de, 9\rangle, \langle cde, 9\rangle\}$*. Compacting* $\mathcal{C}l(\mathcal{D}, 8)$ *means to remove from it itemsets which, due to the transactions suppression, are no longer frequent or no longer closed (lines 9 – 12), i.e.,* $\mathcal{C}l(\mathcal{D}, 8) = \{\langle cde, 9\rangle\}$*. At this point Algorithm 1 invokes the detection algorithm (see [3]) to find out the maximal channels in the new* $\mathcal{C}l(\mathcal{D}, 8)$*, and if necessary, starts a new suppression phase (in our example this is not the case).*

The plots in Fig.2 report the following three measures of distortion (recorded on the well-known RETAIL dataset [1] for different values of σ and k):

1. Absolute number of transactions virtually suppressed.
2. The fraction of itemsets in $\mathcal{F}(\mathcal{D}, \sigma)$ which have their support changed in $\mathcal{O}^k$:
$$\frac{|\{\langle I, sup_{\mathcal{D}}(I)\rangle \in \mathcal{F}(\mathcal{D}, \sigma) \mid sup_{\mathcal{O}^k}(I) \neq sup_{\mathcal{D}}(I)\}|}{|\mathcal{F}(\mathcal{D}, \sigma)|}$$
where $sup_{\mathcal{O}^k}(I) = s$ if $\langle I, s\rangle \in \mathcal{O}^k$; 0 otherwise.
3. The average distortion w.r.t. the original support of itemsets:
$$\frac{1}{|\mathcal{F}(\mathcal{D}, \sigma)|} \sum_{I \in \mathcal{F}(\mathcal{D}, \sigma)} \frac{|sup_{\mathcal{O}^k}(I) - sup_{\mathcal{D}}(I)|}{sup_{\mathcal{D}}(I)}$$

Algorithm 1 Suppressive Sanitization

Input: $\mathcal{C}l(\mathcal{D}, \sigma), \mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma)), \mathcal{D}$
Output: $\mathcal{O}^k$

1: **while** $\mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma)) \neq \emptyset$ **do**
2: *//Scan the database*
3: **for all** $t \in \mathcal{D}$ **do**
4: **if** $\exists \langle \mathcal{C}_I^J, f_I^J\rangle \in \mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$ s.t. $I \subseteq t$ and $(J \setminus I) \cap t = \emptyset$ **then**
5: *//Transaction suppression*
6: **for all** $\langle X, sup_{\mathcal{D}}(X)\rangle \in \mathcal{C}l(\mathcal{D}, \sigma)$ s.t. $X \subseteq t$ **do**
7: $sup_{\mathcal{D}}(X) \leftarrow sup_{\mathcal{D}}(X) - 1$;
8: *//Compact* $\mathcal{C}l(\mathcal{D}, \sigma)$
9: **for all** $\langle X, s\rangle \in \mathcal{C}l(\mathcal{D}, \sigma)$ **do**
10: **if** $\exists \langle Y, s\rangle \in \mathcal{C}l(\mathcal{D}, \sigma)$ s.t. $Y \supset X$ **or** $s < \sigma$ **then**
11: $\mathcal{C}l(\mathcal{D}, \sigma) \leftarrow \mathcal{C}l(\mathcal{D}, \sigma) \setminus \langle X, s\rangle$;
12: **detect** $\mathcal{MC}h(k, \mathcal{C}l(\mathcal{D}, \sigma))$ **in** $\mathcal{C}l(\mathcal{D}, \sigma)$;
13: $\mathcal{O}^k \leftarrow \mathcal{C}l(\mathcal{D}, \sigma)$;

Note that the last measure is really tough with itemsets which become infrequent during the sanitization process: for these itemsets we count a maximum distortion of 1 ($sup_{\mathcal{O}^k}(I) = 0$). However, while the number of itemsets distorted is usually very large, the average distortion on itemsets is very low: this means that quite all itemsets are touched by the sanitization, but their supports are changed just a little.

References

[1] `http://fimi.cs.helsinki.fi/data/`.
[2] R. Agrawal, T. Imielinski, and A. N. Swami. Mining association rules between sets of items in large databases. In *Proceedings of the 1993 ACM SIGMOD.*
[3] M. Atzori, F. Bonchi, F. Giannotti, and D. Pedreschi. k-anonymous patterns. In *Proceedings of the PKDD'05.*
[4] T. Calders and B. Goethals. Mining all non-derivable frequent itemsets. In *Proceedings of the 6th PKDD*, 2002.
[5] N. Pasquier, Y. Bastide, R. Taouil, and L. Lakhal. Discovering frequent closed itemsets for association rules. In *Proc. ICDT '99*, 1999.
[6] L. Sweeney. k-anonymity: a model for protecting privacy. *International Journal on Uncertainty Fuzziness and Knowledge-based Systems*, 10(5), 2002.