



A Frequent Itemset Hiding Toolbox

Aris Gkoulalas-Divanis^{1(✉)}, Vasileios Kagklis², and Elias C. Stavropoulos³

¹ IBM Watson Health, Cambridge, MA, USA
gkoulala@us.ibm.com

² Computer Engineering and Informatics Department, University of Patras,
Patras, Greece
kagklis@ceid.upatras.gr

³ Educational Content, Methodology and Technology Laboratory,
Hellenic Open University, Patras, Greece
estavop@eap.gr

Abstract. Advances in data collection and storage technologies have given way to the establishment of transactional databases among companies and organizations, as they allow enormous amounts of data to be stored efficiently. Useful knowledge can be mined from these data, which can be used in several ways depending on the nature of the data. Quite often companies and organizations are willing to share data for the sake of mutual benefit. However, the sharing of such data comes with risks, as problems with privacy may arise. Sensitive data, along with sensitive knowledge inferred from this data, must be protected from unintentional exposure to unauthorized parties. One form of the inferred knowledge is frequent patterns mined in the form of frequent itemsets from transactional databases. The problem of protecting such patterns from being discovered, is known as the frequent itemset hiding problem. In this paper we present a toolbox, which provides several implementations of frequent itemset hiding algorithms. Firstly, we summarize the most important aspects of each algorithm. We then introduce the architecture of the toolbox and its novel features. Finally, we provide experimental results on real world datasets, demonstrating the capabilities of the toolbox and the convenience it offers in comparing different algorithms.

Keywords: Privacy preserving data mining · Knowledge hiding · Frequent itemset hiding · Sensitive knowledge

1 Introduction

Nowadays, transactional databases are being used more and more by organizations, as they support efficient storage of large volumes of data. By using data mining techniques on such data, modern companies can extract useful information that can help these companies understand the behavior of their customers, support decision making, plan their business strategy, etc.

Companies and organizations are willing to share data for the sake of mutual benefit. The benefits derived from the sharing of such data are considerable

and they cannot be ignored. A typical example is a supermarket, which collects market basket data of its customers' purchases on a regular basis. These organizations might be willing to share their collected information with other parties, such as advisory organizations, for mutual benefit. For example, two stores, say A and B, cooperate in order to discover their customers' purchase behaviors.

Unfortunately, the sharing of such data does not come without risks, as problems with privacy may arise. Therefore, it must be done in such a way, that no sensitive information will be exposed to unauthorized parties. In our previous example, there might be some sensitive information that could reflect the business strategies and secrets of the participating companies that should not be revealed to their adversary competitors. For example, if data analysts of store A found out that its customers tend to purchase products x and y at the same time, they should regard this knowledge as sensitive information, and not disclose it to store B. With this knowledge, store B could offer sales with a lower price for customers who buy x and y together. Then, store A could possibly face the danger of losing some of its customers. Verykios et al. [33], Oliveira and Zaïane [25], and Evfimievski et al. [15] discuss other examples of situations where the sharing of operational databases could have serious adverse effects.

Privacy preserving data mining (PPDM) [3,22] is the research area that investigates techniques to preserve the privacy of data and patterns. Knowledge hiding [19], which is a subfield of PPDM, has as its goal to prevent the exposure of sensitive patterns included in the data to be published. Knowledge hiding can be achieved in several ways. The most commonly used is through the sanitization [4] of a number of transactions in the database, so that the sensitive information can no longer be extracted. Therefore, a hiding technique must be applied before making a database available for sharing. Many data mining tasks rely on frequent itemsets to be identified as a first step in their process. Thus, concealing the frequent patterns associated with the sensitive information would guarantee the preservation of the privacy of the sensitive relationships between patterns of the itemsets that may be discovered through any of these data mining tasks.

In this paper, we present the software architecture and implementation of a frequent itemset hiding (FIH) toolbox, which can be used to apply a suite of hiding techniques on real world datasets. The toolbox comes with a built-in library containing several implementations of FIH algorithms and a suite of performance metrics. Lastly, we present experimental results, to demonstrate the efficiency of the toolbox and the convenience it offers to data owners in comparing different frequent itemset hiding algorithms.

The rest of this paper is organized as follows. Section 2 provides an overview of the related work. In Sect. 3 we present the necessary background information and define the FIH problem. Section 4 describes the software architecture of the FIH toolbox and Sect. 5 presents its features. Section 6 summarizes the evaluation process and presents the experimental results. Finally, Sect. 7 concludes this work.

2 Related Work

Clifton et al. [12,13] are among the first to deal with privacy preservation in the field of data mining and propose data-obscuring techniques in order to avoid discovery of sensitive patterns. Atallah et al. [5] prove that optimally solving the frequent itemset hiding problem is NP-hard. The authors also present a greedy algorithm that turns ones into zeros in the database in order to hide sensitive frequent patterns. Various extensions of this work have been proposed over the years, including those by Verykios et al. [33] and Dasseni et al. [14].

Saygin et al. [30] and Verykios et al. [34] present another approach, according to which items may be added or removed from transactions, recording the possible participation of certain items to transactions with “?” (question marks), so that the hiding is achieved not by falsifying, but by fuzzifying the data. Oliveira and Zaïane [26,27] propose a technique for hiding multiple association rules simultaneously that requires only one pass over the whole database, regardless the size of the database. Pontikakis et al. [28,29] perform an exhaustive evaluation of distortion and blocking (using question marks for hiding) techniques. Bertino et al. [8] propose an evaluation framework that aims at measuring the performance of frequent itemset and association rule hiding techniques.

Menon et al. [23] were the first to introduce an integer linear program (ILP) formulation of the frequent itemset hiding problem. The solution of their ILP points out which transactions need to be sanitized in order to conceal the sensitive patterns. The sanitization process is addressed as a separate phase, independently of the linear programming solution, in a heuristic yet suboptimal way. Sun and Yu [32] introduce a greedy border-based approach for hiding sensitive frequent itemsets. They propose an algorithm that takes advantage of the interplay between preserving maximally non-sensitive and downsizing minimally sensitive frequent itemsets, and gives an accurate and efficient hiding solution.

In [20], Kagklis et al. formulate the FIH problem as an ILP and present a heuristic approach to calculate the coefficients of the objective function of the ILP, while at the same time minimizing the side effects introduced by the hiding process. They also propose a sanitization algorithm for the hiding process.

Stavropoulos et al. [31] relied on the enumeration on the minimal transversals of a hypergraph in order to induce the ideal border between frequent and sensitive itemsets. The ideal border is then utilized to formulate an ILP, the solution of which identifies the set of transactions that need to be sanitized so that the hiding can be achieved with maximum accuracy.

A large number of different approaches towards solving the frequent itemset hiding problem have been proposed. Additionally, several performance evaluation metrics or frameworks have been developed, so as to compare the quality of these techniques. Nevertheless, to our knowledge, there is yet no publicly available tool that offers implementations of such techniques and/or evaluation metrics, along with a common ground for performing experimental evaluation. The work proposed in this paper presents such a toolbox, which can be easily extended to host other hiding techniques, as well as additional evaluation metrics and visualization capabilities.

3 Problem Formulation

3.1 Definitions

Let $I = \{i_1, i_2, \dots, i_n\}$ be a set of items. Any non empty subset of I , $X \subseteq I$, is called an itemset. An itemset consisting of k elements is called a k -itemset. A transaction T is a pair $T = (tid, Z)$, where $Z \subseteq I$ is the itemset and tid is a unique identifier, used to distinguish among transactions that correspond to the same itemset. Let also D be a set of transactions. A transaction T supports an itemset X , if and only if $X \subseteq T$. Given a set of items I , we will denote as $P(I)$ the powerset of I , that is all possible combinations of items from I .

The support of an itemset X in a database D , denoted as $\sigma_D(X)$, is the number of transactions containing X . Support can also be expressed as the percentage of transactions in D . Support can also be expressed as the percentage of transactions in D . An itemset X is frequent in D , if and only if its support in D is at least equal to a support threshold σ_{min} . The set of all frequent itemsets, denoted as F , is given by $F = \{X \subseteq I \mid \sigma_D(X) \geq \sigma_{min}\}$.

Let S be the set of sensitive (frequent) itemsets that need to be hidden. We will denote as I^S the set of different items contained in S . Moreover, let all sensitive itemsets and their supersets in F be denoted as SS , where $SS = \{X \in F \mid \forall Y \in S, X \supseteq Y\}$ and $S \subseteq SS \subseteq F$. The revised set of frequent itemsets, denoted as \tilde{F} , is given by $\tilde{F} = F - SS$.

By its definition, the revised set of frequent itemsets \tilde{F} is the ideal set of itemsets that would still remain frequent after hiding the sensitive itemsets. The ideal case is when only the sensitive itemsets and their supersets are concealed. Sensitive itemsets is desirable to get concealed, while their supersets are inevitably concealed, due to the antimonotonicity property of support, that is $\forall X, Y \subseteq I: X \subseteq Y \Rightarrow \sigma_D(X) \geq \sigma_D(Y)$.

3.2 Problem Description

We are given a transactional database D , a support threshold σ_{min} and a set of sensitive frequent itemsets S . The frequent itemset hiding (FIH) problem involves sanitizing selected transactions in database D , so that itemsets in S cannot be mined from the sanitized database D' using a support threshold equal or above σ_{min} . A database is said to be sanitized, if it is altered in such a way that it no longer supports any sensitive itemset in it. Sanitizing a database involves the sanitization of one or more transactions. A transaction is said to be sanitized, if it is altered so that it no longer supports any sensitive itemset. In the best case scenario, the sensitive itemsets should be hidden with minimal damage to the database. In other words, we want to conceal the itemsets in S , whilst having the minimum impact on the utility of the database.

4 The Architecture of the FIH Toolbox

From an abstract point of view, the toolbox is divided into four layers (see Fig. 1). The top layer is the ‘‘Presentation Layer’’, which implements the Graphical User

Interface (GUI) of the toolbox and provides visualizations of the results received from the next layer below. Users interact with the toolbox through the presentation layer, depending on the options they make. The GUI is implemented in Python, using the Tkinter and ttk modules.

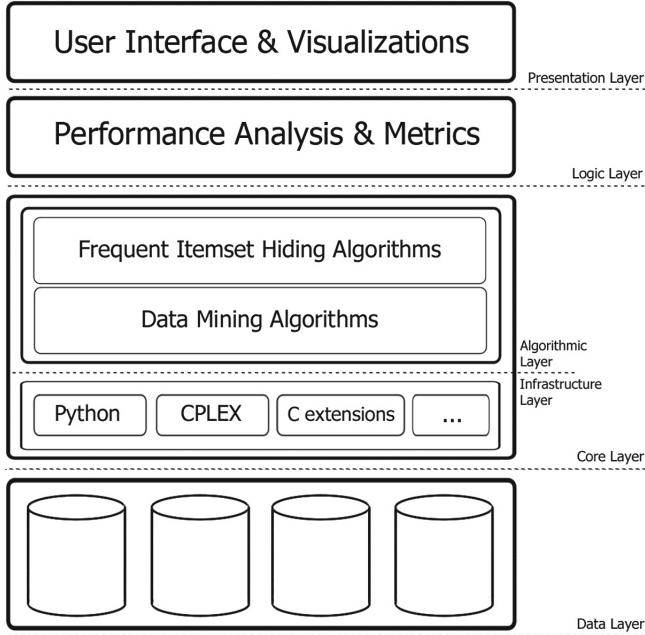


Fig. 1. The architecture of the FIH toolbox.

The next layer is the “Logic Layer”. In this layer, the performance analysis phase is implemented and the calculation of the metrics take place. Such metrics are the number of changes made in the raw data, the number of side effects, the execution time and the information loss incurred. A detailed description of the performance metrics is presented in Sect. 6. This layer is also responsible for moving and processing data between its two surrounding layers.

The third layer is the “Core Layer”, which consists of two sublayers: the “Algorithmic Sublayer” and the “Infrastructure Sublayer”. The “Algorithmic Sublayer” is basically the built-in library with the implemented FIH algorithms, described in Sect. 5.1, along with the implementation of some basic data mining methods. Currently, the only data mining method offered is the Apriori algorithm [3]. The “Infrastructure Sublayer” consists of all the tools used for the development of the toolbox. Most of the code is implemented in Python. The hiding algorithms are implemented in Cython [7]. The Apriori algorithm [3] is part of the PyFIM extension module. This module is a C extension for Python [10] to efficiently mine the set of frequent itemsets. The IBM ILOG CPLEX 12.6 [2] is used for solving any linear program.

Finally, the last layer is the “Data Layer”, which is related to the datasets that can be used and their corresponding supported formats. The toolbox supports both space-separated value (.ssv) and comma-separated value (.csv) formatted files. The delimiter is declared by the extension of the file. Therefore, a space-separated file must have the extension “.ssv”, while a comma-separated file must have the extension “.csv”. If the input file has a different extension, then by default the space delimiter is used.

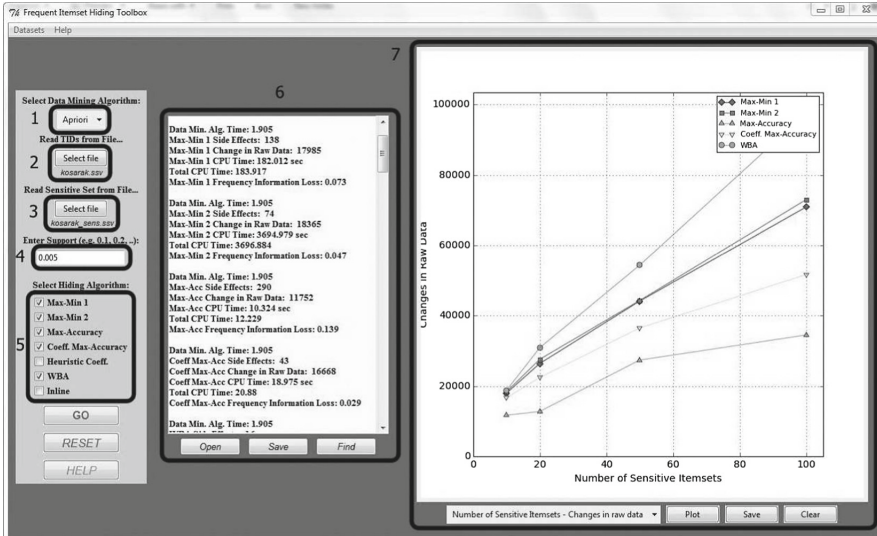


Fig. 2. The GUI of the FIH toolbox.

In Fig. 2, the GUI of the toolbox is presented. For the time being, there is a beta version of the toolbox available [1]. Linear programming techniques require a license for CPLEX, which can be obtained for free through IBM’s Academic Initiative program. The user can apply a hiding technique on a dataset by following a few easy steps, which are summarized in Fig. 3.

Firstly, a data mining algorithm from *field 1* must be selected. Then, a dataset must be supplied by the user, by using *field 2*. Respectively, the file with the sensitive itemsets must be given by using *field 3*. In *field 4*, the support threshold must be specified. *Field 5* is a group of checkboxes. By checking a checkbox, the user selects the corresponding algorithm to be executed. *Field 6* is a text editor, where the sanitized dataset and the calculated metrics are printed. The user can save these results by using the “Save” button below the text editor. Finally, *field 7* is a canvas that displays visualizations of the metrics. Below the canvas, there is drop-down list with options related to the axes of the figures. The “Plot” button should be used after an option from the drop-down list is selected, so as to plot the corresponding figure. Buttons “Save” and “Clear” can be used to save and clear the current figure respectively.

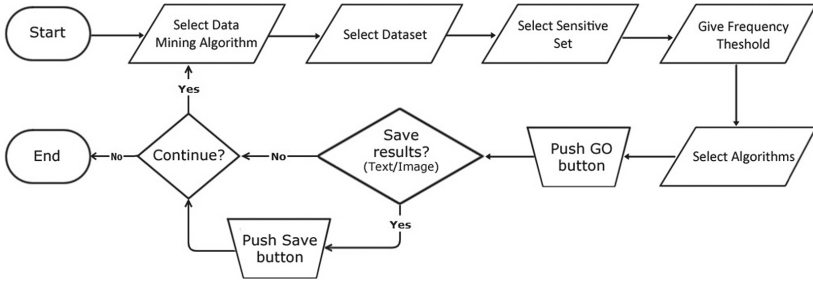


Fig. 3. “How-to-use” flow chart for the FIH toolbox.

5 Special Features of the FIH Toolbox

5.1 Built-In Library

The toolbox comes with a built-in library that consists of the following FIH algorithm implementations as independent modules: Max-Min Algorithms [24], Weight-Based Approach [32], Max-Accuracy Algorithm [23], Coefficient-Based Max-Accuracy Algorithm [21], Heuristic Coefficient Based Approach [20], and Inline Algorithm [16]. The Max-Min algorithms (Max-Min 1, Max-Min 2) and the Weight-Based Approach (WBA) use the border revision theory [32]. The Max-Accuracy algorithm, the Coefficient-Based Max-Accuracy algorithm and the Heuristic Coefficient Based Approach formulate the problem as an ILP. A heuristic algorithm for the sanitization is also used. The Inline algorithm combines border revision theory and linear programming. Future versions of the tool will include algorithms in [17, 18], as well as methods for hiding sensitive association rules [5, 14, 25, 26, 28–30, 33, 34].

5.2 Extensibility

The toolbox comes with a built-in library, which contains some implemented FIH algorithms. However, a non-extensible library would limit the utility of the Toolbox. An important feature of the toolbox is that it can be extended by its users. Users can implement and import new algorithms, and compare them with the existing algorithms in the built-in library.

User-implemented algorithms must be compatible with the toolbox. Therefore, any implementation must comply to the restrictions and guidelines, as described in the manual that can be found in [1]. Users can easily implement compatible source files by following the instructions given in the manual. After the implementation is completed, it can be used right away; create a folder with the name “Extensions” (without the quotes) in the same directory with the Toolbox and copy the source file in it.

5.3 Automatic Option Loading

In Sect. 4, we described how the user can manually load data and apply a hiding algorithm. The use of the toolbox can become even more convenient by defining option scenarios before using them. Instead of making the options manually, the Toolbox gives the capability to load automatically predefined option scenarios. Assume that we want to load the dataset “myDataset.dat” and the file of sensitive itemsets “HS1.dat”, and use a threshold equal to “0.05”. Firstly, we create a folder named “Datasets” in the same directory that the Toolbox is located. Then, we simply create the tree hierarchy of files and folders as shown in Fig. 4 (left). If we run the Toolbox, we can load the option scenario by clicking on *Datasets* → *myDataset* → 0.05 → *HS1.dat*, as shown in Fig. 4 (right). An option scenario can be also imported during runtime by following the same steps and by clicking *Datasets* → *Update Datasets*.

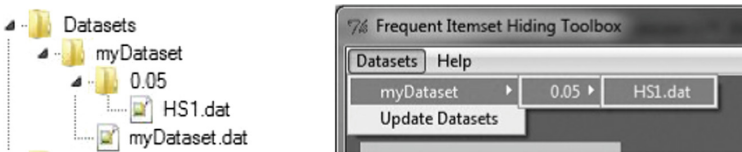


Fig. 4. Automatic loading predefined option scenario.

6 Experimental Evaluation

We evaluated some of the implemented algorithms on real datasets, by using different parameters such as the number of sensitive itemsets to be hidden and the support count threshold. In this section, we also present the datasets used with their special characteristics, the selected parameters and the experimental results. All experiments were conducted on a PC running Windows 7 with an Intel Core i5, 3.20 GHz processor. For the linear programming techniques, CPLEX [2] was used for solving the formulated linear programs.

6.1 Datasets

All datasets used for evaluation are publicly available in the FIMI repository (<http://fimi.ua.ac.be/data/>). These datasets have different characteristics in terms of transactions and items, and the average transaction length. The characteristics of the datasets used are presented in Table 1. The mushroom dataset was prepared by Roberto Bayardo (University of California, Irvine) [6]. The retail dataset is a market basket dataset from an anonymous Belgian store [11]. The kosarak dataset was provided by Ferenc Bodon [9] and contains anonymized click-stream data of a Hungarian online news portal.

Table 1. Characteristics of the datasets.

Dataset name	Number of transactions	Number of items	Avg. trans. length	σ_{min} used
Mushroom	8,124	119	23.00	1,625
Retail	88,162	16,470	10.30	22; 44; 66; 88
Kosarak	990,002	41,270	8.10	4,950

6.2 Evaluation Metrics and Framework

For the evaluation of the algorithms, we implemented and used several metrics along with the framework proposed by Bertino et al. [8]. The framework is based on several evaluation dimensions. We used the following metrics.

Efficiency. It is the ability of a PPDM algorithm to execute with good performance, in terms of all the resources consumed by the algorithm. Simply put, the efficiency of an algorithm quantifies how good is the relationship between its performance and the overall resources it uses. As in most cases, we assess efficiency in terms of time and space. In other words, efficiency is evaluated in terms of CPU time and the amount of memory that an algorithm requires.

Scalability. It is used to evaluate the behavior of the efficiency of a PPDM algorithm for a growing amount of input data, from which relevant information is mined while ensuring privacy. We conducted experiments with datasets of different size and density, so as to test the scalability of the algorithms implemented in the toolbox.

Data Quality. It refers to the quality of data after the hiding process. As mentioned earlier, attempting to hide sensitive information might have an impact on non-sensitive information as well. If data quality is too degraded, then the released database is useless for the purpose of knowledge extraction. According to Bertino et al. [8], the information loss can be measured in terms of the dissimilarity between the original dataset D and the sanitized D' . The information loss is defined as the ratio between the sum of the absolute errors made in computing the frequencies of items in the sanitized database and the sum of all the frequencies of items in the original database.

We also use two additional measures: (a) the number of raw changes that occurred in data, and (b) the number of side effects. The raw data changes is the total number of items that have been removed in order to sanitize the database. The number of side effects (SE) introduced by the application of the sanitization process can be measured by $SE(\tilde{F}, F') = |\tilde{F}| - |F'| \geq 0$, where $|\tilde{F}|$ is the number of itemsets in the revised set of frequent itemsets \tilde{F} , whilst $|F'|$ is the number of itemsets in the set of frequent itemsets F' mined from D' .

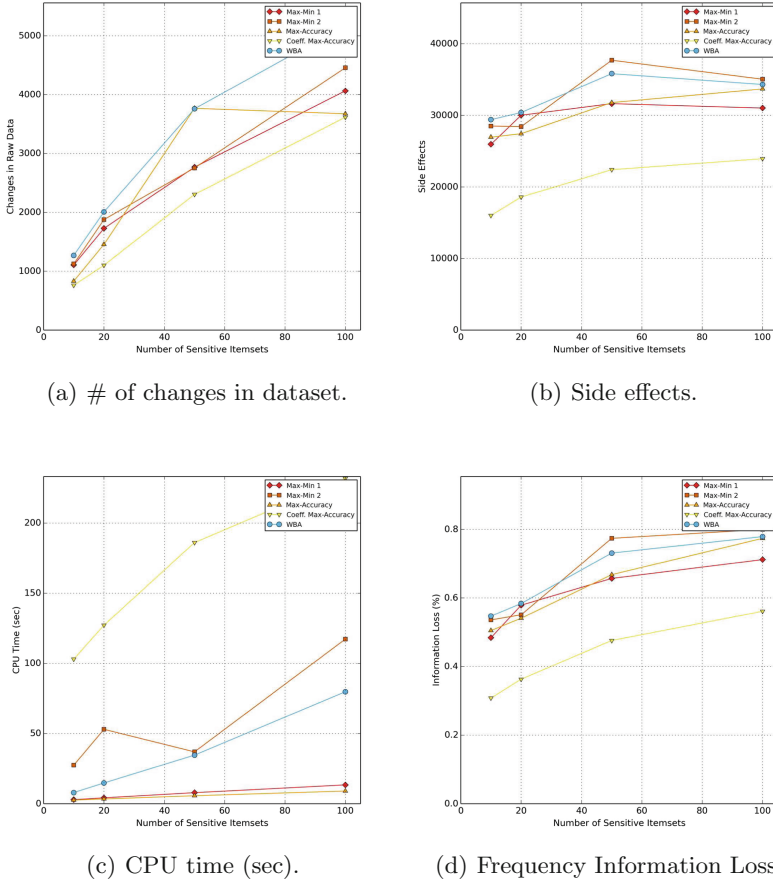


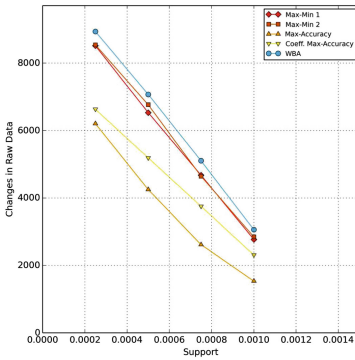
Fig. 5. Results for the mushroom dataset.

6.3 Experimental Results

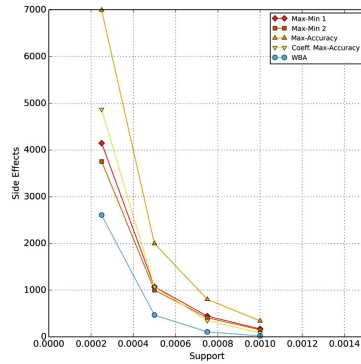
Figure 5 presents the results we obtained for the mushroom dataset. Figure 5(a) displays how many changes (item removals) each algorithm made in the original database. Figure 5(b) displays the number of side effects that occurred as a result of the concealing process. Figure 5(c) presents the times needed by each algorithm. Lastly, Fig. 5(d) presents the frequency information loss. For the evaluation with this dataset, we used 4 different hiding scenarios; hiding 10, 20, 50 and 100 sensitive itemsets of different, random length. The support threshold used is $\sigma_{min} = 1625$. We selected randomly the sensitive itemsets.

The mushroom dataset is a small, yet dense dataset. Thus, the number of frequent itemsets increases dramatically as we decrease the support threshold. Linear programming techniques (Max-Accuracy and Coefficient-Based Max-Accuracy) achieve better results than their heuristic-based counterparts

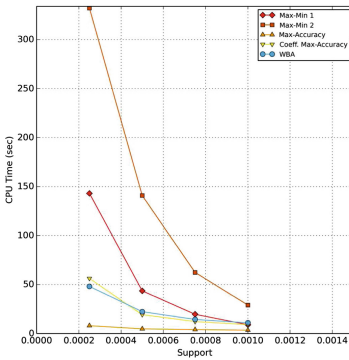
(Max-Min 1, Max-Min 2 and WBA). Concerning the time complexity, we notice that the simpler the algorithm, the less time is needed to run, as expected.



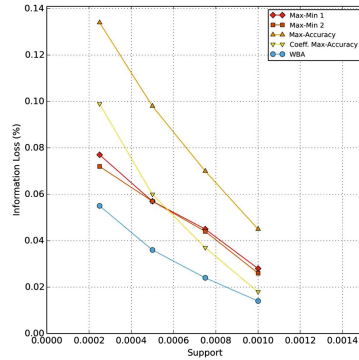
(a) # of changes in dataset.



(b) Side effects.



(c) CPU time (sec).



(d) Frequency Information Loss.

Fig. 6. Results for the retail dataset.

Figures 6(a)–(d) present the results for the retail dataset. We used a single hiding scenario of 100 sensitive itemsets. The set of sensitive itemsets consists of randomly selected itemsets. We performed experiments with this hiding scenario for different support thresholds, $\sigma_{min} = \{22, 44, 66, 88\}$.

The lower the mining threshold is, the larger the values of all metrics are. Notice that for this dataset, which is not as dense as the mushroom dataset, WBA achieves the best results, as far as side effects and information loss are concerned, with a fairly good time complexity. Although the results are printed in the text editor of the toolbox, the figures drawn in the canvas give a direct sense of which algorithm prevails.

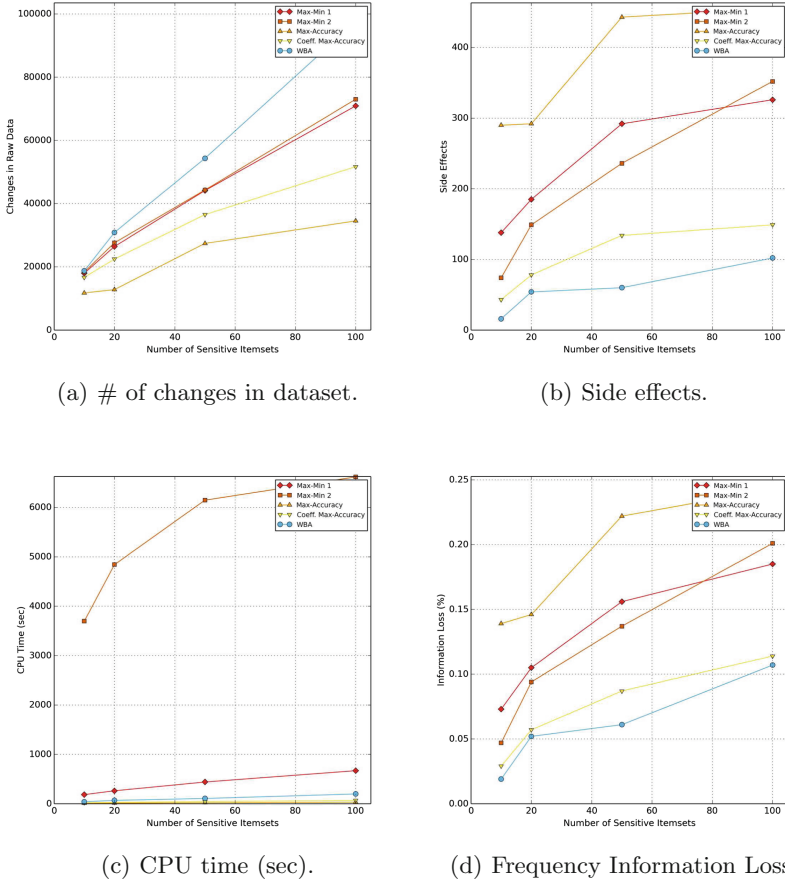


Fig. 7. Results for the kosarak dataset.

Finally, Figs. 7(a)–(d) present the results for the kosarak dataset. We used 4 different hiding scenarios for the evaluation with this dataset; hiding 10, 20, 50 and 100 sensitive itemsets of different, random length. The support threshold used is $\sigma_{min} = 4950$. The sensitive itemsets were picked randomly. Again WBA has the best results in terms of the number of side effects and the information loss. The time complexity is quite low for most of the algorithms and increases linearly with respect to the number of sensitive itemsets.

From the aforementioned experimental results, it is clear that the execution times of most of the techniques increase linearly as the number of sensitive itemsets increases, the size of the dataset increases, and the support threshold decreases. The density of the dataset has a great impact on the results. Linear programming techniques have good scalability. Then, the border-based techniques, such as Max-Min 1 and WBA, follow. Max-Min 2 appears to have a poor scalability compared to the rest of the heuristic algorithms.

7 Conclusions and Future Work

In this paper, we presented a FIH toolbox which can be used to apply a suite of hiding techniques on real world datasets. The toolbox comes with a built-in library containing several implementations of FIH algorithms and a suite of performance metrics. Currently the toolbox is in beta version and many improvements can be made, concerning both the GUI and the overall performance. Several features are currently under development, including a feature to recommend the appropriate FIH algorithm that, based on the characteristics of the input dataset, is expected to give the best results.

References

1. <https://github.com/kagklis/Frequent-Itemset-Hiding-Toolbox-x86>
2. IBM ILOG CPLEX User's Manual v12.6
3. Agrawal, R., Srikant, R.: Privacy-preserving data mining. In: SIGMOD Conference, pp. 439–450 (2000)
4. Askari, M., Safavi-Naini, R., Barker, K.: An information theoretic privacy and utility measure for data sanitization mechanisms. In: Proceedings of the 2nd ACM Conference on Data and Application Security and Privacy (CODASPY 12), pp. 283–294 (2012)
5. Atallah, M., Bertino, E., Elmagarmid, A., Ibrahim, M., Verykios, V.: Disclosure limitation of sensitive rules. In: Proceedings of the 1999 Workshop on Knowledge and Data Engineering Exchange (KDEX 99), pp. 45–52 (1999)
6. Bayardo Jr., R.J.: Efficiently mining long patterns from databases. In: Proceedings of the 1998 ACM SIGMOD International Conference on Management of Data (SIGMOD 98), pp. 85–93 (1998)
7. Behnel, S., Bradshaw, R., Citro, C., Dalcin, L., Seljebotn, D.S., Smith, K.: Cython: the best of both worlds. *Comput. Sci. Eng.* **13**(2), 31–39 (2011)
8. Bertino, E., Fovino, I.N., Provenza, L.P.: A framework for evaluating privacy preserving data mining algorithms. *Data Min. Knowl. Discov.* **11**(2), 121–154 (2005)
9. Bodon, F.: A fast APRIORI implementation. In: Proceedings of the IEEE ICDM Workshop on Frequent Itemset Mining Implementations (FIMI 03), vol. 90, pp. 56–65 (2003)
10. Borgelt, C.: Frequent item set mining. *Wiley Interdisc. Rev. Data Min. Knowl. Discov.* **2**(6), 437–456 (2012)
11. Brijs, T., Swinnen, G., Vanhoof, K., Wets, G.: Using association rules for product assortment decisions: a case study. In: Proceedings of the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 99), pp. 254–260 (1999)
12. Clifton, C., Marks, D.: Security and privacy implications of data mining. In: Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data, pp. 15–19 (1996)
13. Clifton, C., Kantarcioglu, M., Vaidya, J.: Defining privacy for data mining. In: National Science Foundation Workshop on Next Generation Data Mining (WNGDM), pp. 126–133 (2002)
14. Dasseni, E., Verykios, V.S., Elmagarmid, A.K., Bertino, E.: Hiding association rules by using confidence and support. In: Moskowitz, I.S. (ed.) IH 2001. LNCS, vol. 2137, pp. 369–383. Springer, Heidelberg (2001). <https://doi.org/10.1007/3-540-45496-9-27>

15. Evfimievski, A., Srikant, R., Agrawal, R., Gehrke, J.: Privacy preserving mining of association rules. *Inf. Syst.* **29**, 343–364 (2004)
16. Gkoulalas-Divanis, A., Verykios, V.S.: An integer programming approach for frequent itemset hiding. In: *CIKM*, pp. 748–757 (2006)
17. Gkoulalas-Divanis, A., Verykios, V.S.: Hiding sensitive knowledge without side effects. *Knowl. Inf. Syst.* **20**(3), 263–299 (2009)
18. Gkoulalas-Divanis, A., Verykios, V.S.: Exact knowledge hiding through database extension. *IEEE Trans. Knowl. Data Eng.* **21**(5), 699–713 (2009)
19. Johnsten, T., Raghavan, V.V.: A methodology for hiding Knowledge in databases. In: *Proceedings of the IEEE International Conference on Privacy, Security and Data Mining (CRPIT 14)*, pp. 9–17 (2002)
20. Kagklis, V., Verykios, V.S., Tzimas, G., Tsakalidis, A.K.: An integer linear programming scheme to sanitize sensitive frequent itemsets. In: *International Conference on Tools with Artificial Intelligence (ICTAI 14)*, pp. 771–775 (2014)
21. Leloglou, E., Ayav, T., Ergenc, B.: Coefficient-based exact approach for frequent itemset hiding. In: *eKNOW2014: The 6th International Conference on Information, Process, and Knowledge Management*, pp. 124–130 (2014)
22. Lindell, Y., Pinkas, B.: Privacy preserving data mining. In: *Proceedings of the 20th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO 00)*, pp. 36–54 (2000)
23. Menon, S., Sarkar, S., Mukherjee, S.: Maximizing accuracy of shared databases when concealing sensitive patterns. *INFORMS* **16**(3), 256–270 (2005)
24. Moustakides, G.V., Verykios, V.S.: A MaxMin approach for hiding frequent itemsets. *Data Knowl. Eng.* **65**(1), 75–89 (2008)
25. Oliveira, S.R., Zaïane, O.R.: Protecting sensitive knowledge by data sanitization. In: *Proceedings of the 3rd IEEE International Conference on Data Mining (ICDM 2003)*, pp. 99–106 (2003)
26. Oliveira, S.R., Zaïane, O.R.: Algorithms for balancing privacy and knowledge discovery in association rule mining. In: *IDEAS*, pp. 54–65 (2003)
27. Oliveira, S.R., Zaïane, O.R.: Protecting sensitive knowledge by data sanitization. In: *ICDM*, pp. 613–616 (2003)
28. Pontikakis, E.D., Tsitsonis, A.A., Verykios, V.S.: An experimental study of distortion-based techniques for association rule hiding. In: Farkas, C., Samarati, P. (eds.) *DBSec 2004. IIFIP*, vol. 144, pp. 325–339. Springer, Boston (2004). <https://doi.org/10.1007/1-4020-8128-6-22>
29. Pontikakis, E.D., Verykios, V.S., Theodoridis, Y.: On the comparison of association rule hiding heuristics. In: *Hellenic Database Management Symposium* (2004)
30. Saygin, Y., Verykios, V.S., Clifton, C.: Using unknowns to prevent discovery of association rules. *SIGMOD Rec.* **30**(4), 45–54 (2001)
31. Stavropoulos, E.C., Verykios, V.S., kagklis, V.: A transversal hypergraph approach for the frequent itemset hiding problem. *Knowl. Inf. Syst.* **47**(3), 625–645 (2016)
32. Sun, X., Yu, P.S.: Hiding sensitive frequent itemsets by a border-based approach. *JCSE* **1**(1), 74–94 (2007)
33. Verykios, V.S., Elmagarmi, A.K., Bertino, E., Saygin, Y., Dasseni, E.: Association rule hiding. *IEEE Trans. Knowl. Data Eng.* **16**(4), 434–447 (2004)
34. Verykios, V.S., Pontikakis, E.D., Theodoridis, Y., Chang, L.: Efficient algorithms for distortion and blocking techniques in association rule hiding. *Distrib. Parallel Databases* **22**(1), 85–104 (2007)