

Efficient Discovery of Association Rules and Frequent Itemsets through Sampling with Tight Performance Guarantees*

Matteo Riondato and Eli Upfal

Department of Computer Science, Brown University, Providence, RI, USA
{matteo,eli}@cs.brown.edu

Abstract. The tasks of extracting (top- K) Frequent Itemsets (FI's) and Association Rules (AR's) are fundamental primitives in data mining and database applications. Exact algorithms for these problems exist and are widely used, but their running time is hindered by the need of scanning the entire dataset, possibly multiple times. High quality approximations of FI's and AR's are sufficient for most practical uses, and a number of recent works explored the application of sampling for fast discovery of approximate solutions to the problems. However, these works do not provide satisfactory performance guarantees on the quality of the approximation, due to the difficulty of bounding the probability of under- or over-sampling any one of an unknown number of frequent itemsets. In this work we circumvent this issue by applying the statistical concept of *Vapnik-Chervonenkis (VC) dimension* to develop a novel technique for providing tight bounds on the sample size that guarantees approximation within user-specified parameters. Our technique applies both to absolute and to relative approximations of (top- K) FI's and AR's. The resulting sample size is linearly dependent on the VC-dimension of a range space associated with the dataset to be mined. The main theoretical contribution of this work is a characterization of the VC-dimension of this range space and a proof that it is upper bounded by an easy-to-compute characteristic quantity of the dataset which we call *d-index*, namely the maximum integer d such that the dataset contains at least d transactions of length at least d . We show that this bound is strict for a large class of datasets. The resulting sample size for an absolute (resp. relative) (ϵ, δ) -approximation of the collection of FI's is $O(\frac{1}{\epsilon^2}(d + \log \frac{1}{\delta}))$ (resp. $O(\frac{2+\epsilon}{\epsilon^2(2-\epsilon)\theta}(d \log \frac{2+\epsilon}{(2-\epsilon)\theta} + \log \frac{1}{\delta}))$) transactions, which is a significant improvement over previous known results. We present an extensive experimental evaluation of our technique on real and artificial datasets, demonstrating the practicality of our methods, and showing that they achieve even higher quality approximations than what is guaranteed by the analysis.

1 Introduction

Discovery of frequent itemsets and association rules is a fundamental computational primitive with application in data mining (market basket analysis), databases (histogram construction), networking (heavy hitters) and more [15, Sect. 5]. Depending on the particular application, one is interested in finding all itemsets with frequency greater or

* Work was supported in part by NSF award IIS-0905553.

equal to a user defined threshold (FIs), identifying the K most frequent itemsets (top- K), or computing all association rules (ARs) with user defined minimum support and confidence level. Exact solutions to these problems require scanning the entire dataset, possibly multiple times. For large datasets that do not fit in main memory, this can be prohibitively expensive. Furthermore, such extensive computation is often unnecessary, since high quality approximation are sufficient for most practical applications. Indeed, a number of recent papers [4, 6, 7, 9, 10, 12, 13, 17–22, 25, 27–30, 32, 33, 36–41] explored the application of sampling for approximate solutions to these problems. However, the efficiency and practicality of the sampling approach depends on a tight relation between the size of the sample and the quality of the resulting approximation. Previous works do not provide satisfactory solutions to this problem.

The technical difficulty in analyzing any sampling technique for frequent itemsets discovery problems is that a-priori any subset of items can be among the most frequent ones, and the number of subsets is exponential in the number of distinct items appearing in the dataset. A standard analysis begins with a bound on the probability that a given itemset is either over or under represented in the sample. Such bound is easy to obtain using a Chernoff-like bound or the Central Limit theorem. The difficulty is in combining the bounds for individual itemsets into a global bound that holds simultaneously for all the itemsets. A simple application of the union bound vastly overestimates the error probability because of the large number of possible itemsets, a large fraction of which may not be present in the dataset and therefore should not be considered. More sophisticated techniques, developed in recent works [6, 12, 29], give better bounds only in limited cases. A loose bound on the required sample size for achieving the user defined performance guarantees, decreases the gain obtained from the use of sampling.

In this work we circumvent this problem through a novel application of the *Vapnik-Chervonenkis (VC)* dimension concept, a fundamental tool in statistical learning theory. Roughly speaking, the VC-dimension of a collection of indicator functions (a range space) is a measure of its complexity or expressiveness (see Sect. 2.2 for formal definitions). A major result [35] relates the VC-dimension of a range space to the sufficient size for a random sample to simultaneously approximate all the indicator functions within predefined parameters. The main obstacle in applying the VC-dimension theory to particular computation problems is computing the VC-dimension of the range spaces associated with these problems.

We apply the VC-dimension theory to frequent itemsets problems by viewing the presence of an itemset in a transaction as the outcome of an indicator function associated with the itemset. The major theoretical contributions of our work are a complete characterization of the VC-dimension of the range space associated with a dataset, and a tight bound to this quantity. We prove that the VC-dimension is upper bounded by an easy-to-compute characteristic quantity of the dataset which we call *d-index*, namely, the maximum integer d such that the dataset contains at least d transactions of length at least d . We show that this bound is tight by demonstrating a large class of datasets with a VC-dimension that matches the bound.

The VC-dimension approach provides a unified tool for analyzing the various frequent itemsets and association rules problems (i.e., the market basket analysis tasks). We use it to prove tight bounds on the required sample size for extracting FI's with a

minimum frequency threshold, for mining the top- K FI's, and for computing the collection of AR's with minimum frequency and confidence thresholds. Furthermore, we compute bounds for both absolute and relative approximations (see Sec 2.1 for definitions). We show that high quality approximations can be obtained by mining a very small random sample of the dataset. For example, the required sample size for an absolute (ϵ, δ) -approximation of the collection of FI's is $O(\frac{1}{\epsilon^2}(d + \log \frac{1}{\delta}))$ transactions, which is a significant improvement over previous known results, as it is smaller and, more importantly, less dependent on parameters such as the minimum frequency threshold and the dataset size. Similar results are proven for the top- K FI's and AR's tasks.

We present an extensive experimental evaluation of our method using real and artificial datasets, to assess the practicality of our approach. The experimental results show that indeed our method achieves, and even exceeds, the analytically proven guarantees for the quality of the approximations.

1.1 Previous Work

Agrawal et al. [1] introduced the problem of mining association rules in the basket data model, formalizing a fundamental task of information extraction in large datasets. Almost any known algorithm for the problem starts by solving a FI's problem and then generate the association rules implied by these frequent itemsets. Agrawal and Srikant [2] presented *Apriori*, the most well-known algorithm for mining FI's, and *Fast-GenRules* for computing association rules from a set of itemsets. Various ideas for improving the efficiency of FI's and AR's algorithms have been studied, and we refer the reader to the survey by Ceglar and Roddick [5] for a good presentation of recent contributions. However, the running times of all known algorithms heavily depend on the size of the dataset.

Mannila et al. [27] first suggested the idea that sampling can be used to efficiently obtain the collection of FI's, presenting some empirical results to validate the intuition. Toivonen [33] presents an algorithm that, by mining a random sample of the dataset, builds a candidate set of frequent itemsets which contains all the frequent itemsets with a probability that depends on the sample size. There are no guarantees that that all itemsets in the candidate set are frequent, but the set of candidates can be used to efficiently identify the set of frequent itemsets with at most two passes over the entire dataset. The work also suggests a bound on the sample size sufficient to ensure that the frequencies of itemsets in the sample are close to their real one. The analysis uses Chernoff bounds and the union bound. The major drawback of this sample size is that it depends linearly on the number of individual items appearing in the dataset.

Zaki et al. [39] show that static sampling is an efficient way to mine a dataset, but choosing the sample size using Chernoff bounds is too conservative, in the sense that it is possible to obtain the same accuracy and confidence in the approximate results at smaller sizes than what the theoretical analysis suggested.

Other works tried to improve the bound to the sample size by using different techniques from statistic and probability theory like the central limit theorem [19, 22, 40] or hybrid Chernoff bounds [41].

Since theoretically-derived bounds to the sample size were too loose to be useful, a corpus of works applied progressive sampling to extract FI's [4, 7, 9, 10, 12, 17, 18, 20,

21, 25, 28, 38]. Progressive sampling algorithms work by selecting a random sample and then trimming or enriching it by removing or adding new sampled transactions according to a heuristic or a self-similarity measure that is fast to evaluate, until a suitable stopping condition is satisfied. The major downside of this approach is that it offers no guarantees on the quality of the obtained results.

Another approach to estimating the required sample size is presented in [13]. The authors give an algorithm that studies the distribution of frequencies of the itemsets and uses this information to fix a sample size for mining frequent itemsets, but without offering any theoretical guarantee.

A recent work by Chakaravarthy et al. [6] gives the first analytical bound on a sample size that is linear in the length of the longest transaction, rather than in the number of items in the dataset. This work is also the first to present an algorithm that uses a random sample of the dataset to mine approximated solutions to the AR’s problem with quality guarantees. No experimental evaluation of their methods is presented, and they do not address the top- K FI’s problem. Our approach gives better bounds for the problems studied in [6] and applies to related problems such as the discovery of top- K FI’s and absolute approximations.

Extracting the collection of top- K frequent itemsets is a more difficult task since the corresponding minimum frequency threshold is not known in advance [11, 14]. Some works solved the problem by looking at *closed* top- K frequent itemsets, a concise representation of the collection [30, 37], but they suffers from the same scalability problems as the algorithms for exactly mining FI’s with a fixed minimum frequency threshold.

Previous works that used sampling to approximate the collection of top- K FI’s [29, 32] used progressive sampling. Both works provide (similar) theoretical guarantees on the quality of the approximation. What is more interesting to us, both works present a theoretical upper bound to the sample size needed to compute such an approximation. The size depended linearly on the number of items. In contrast, our results give a sample size that only in the worst case is linear in the number of items but can be (and is, in practical cases) much less than that, depending on the dataset, a flexibility not provided by previous contributions. Sampling is used by Vasudevan and Vojonović [36] to extract an approximation of the top- K frequent individual *items* from a sequence of items, which contains no item whose actual frequency is less than $f_K - \varepsilon$ for a fixed $0 < \varepsilon < 1$, where f_K is the *actual* frequency of the K -th most frequent item. They derive a sample size sufficient to achieve this result, but they assume the knowledge of f_K , which is rarely the case. An empirical sequential method can be used to estimate the right sample size. Moreover, the results cannot be directly extended to the mining of top- K frequent item(set)s from datasets of transactions with length greater than one.

1.2 Our Contributions

By applying tools from statistical learning theory, we develop a general technique for bounding the sample size required for generating high quality approximations to frequent itemsets and association rules tasks. Table 1 compares our technique to the best previously known results for the various problems (see Sect. 2.1 for definitions). Our bounds, which are linear in the VC-dimension associated with the dataset, are consistently smaller and less dependent on other parameters of the problem than previous

results. An extensive experimental evaluation demonstrates the advantage of our technique in practice.

Task	Approx.	This work	Best previous work
FI's	absolute	$\frac{4c}{\varepsilon^2} (d + \log \frac{1}{\delta})$	$O(\frac{1}{\varepsilon^2} (\mathcal{I} + \log \frac{1}{\delta}))$ [19, 22, 33, 40]
	relative	$\frac{4(2+\varepsilon)c}{\varepsilon^2(2-\varepsilon)\theta} \left(d \log \frac{2+\varepsilon}{\theta(2-\varepsilon)} + \log \frac{1}{\delta} \right)$	$\frac{24}{\varepsilon^2(1-\varepsilon)\theta} \left(\Delta + 5 + \log \frac{4}{(1-\varepsilon)\theta\delta} \right)$ [6]
top- K	absolute	$\frac{16c}{\varepsilon^2} (d + \log \frac{1}{\delta})$	$O(\frac{1}{\varepsilon^2} (\mathcal{I} + \log \frac{1}{\delta}))$ [29, 32]
	relative	$\frac{4(2+\varepsilon)c}{\varepsilon^2(2-\varepsilon)\theta} \left(d \log \frac{2+\varepsilon}{\theta(2-\varepsilon)} + \log \frac{1}{\delta} \right)$	not available
AR's	absolute	$O\left(\frac{(1+\varepsilon)}{\varepsilon^2(1-\varepsilon)\theta} \left(d \log \frac{1+\varepsilon}{\theta(1-\varepsilon)} + \log \frac{1}{\delta} \right)\right)$	not available
	relative	$\frac{16c(4+\varepsilon)}{\varepsilon^2(4-\varepsilon)\theta} \left(d \log \frac{4+\varepsilon}{\theta(4-\varepsilon)} + \log \frac{1}{\delta} \right)$	$\frac{48}{\varepsilon^2(1-\varepsilon)\theta} \left(\Delta + 5 + \log \frac{4}{(1-\varepsilon)\theta\delta} \right)$ [6]

Table 1: Required sample sizes (as number of transactions) as a function of the VC-dimension d , the maximum transaction size Δ , the number of items $|\mathcal{I}|$, the accuracy ε , the failure probability δ , the minimum frequency θ , and the minimum confidence γ . Note that $d \leq \Delta \leq |\mathcal{I}|$.

To the best of our knowledge, our work is the first to provide a characterization and an explicit bound for the VC-dimension of the range space associated to a dataset and to apply the result to the extraction of FI's and AR's from random sample of the dataset. We believe that this connection with statistical learning theory can be further exploited in other data mining problems.

We also believe that our approach can be applied not just to mining collections of frequent itemsets and association rules, which can be massive, but also to the mining of small collections of itemsets/association rules that describe the dataset with the minimal number of itemsets/association rules possible, as presented in [26].

Outline. In Sect. 2 we formally define the problem and our goals, and introduce definitions and lemmas used in the analysis. The main part of the analysis with derivation of a strict bound to the VC-dimension of association rules is presented in Sect. 3, while our algorithms and sample sizes for mining FI's, top- K FI's, and association rules through sampling are in Sect. 4. Section 5 contains an extensive experimental evaluation of our techniques. Due to space constraints, the proofs of our theorems and lemmas are not presented in this paper. We refer the interested reader to the full version [31].

2 Preliminaries

2.1 Datasets, Itemsets, and Association Rules

A *dataset* \mathcal{D} is a collection of *transactions*, where each transaction τ is a subset of a ground set \mathcal{I} . There can be multiple identical transactions in \mathcal{D} . Members of \mathcal{I} are called *items* and members of $2^{\mathcal{I}}$ are called *itemsets*. Let $|\tau|$ denote the number of items

in transaction τ . Given an itemset $A \in 2^{\mathcal{I}}$, let $T_{\mathcal{D}}(A)$ denote the set of transactions in \mathcal{D} that contain A . The *support* of A , $\sigma_{\mathcal{D}}(A) = |T_{\mathcal{D}}(A)|$, is the number of transaction in \mathcal{D} that contains A , and the *frequency* of A , $f_{\mathcal{D}}(A) = \frac{|T_{\mathcal{D}}(A)|}{|\mathcal{D}|}$, is the fraction of transactions in \mathcal{D} that contain A .

Definition 1. Given a minimum frequency threshold θ , $0 < \theta \leq 1$, the FI's mining task with respect to θ is finding all itemsets with frequency $\geq \theta$, i.e., the set

$$\text{FI}(\mathcal{D}, \mathcal{I}, \theta) = \{(A, f_{\mathcal{D}}(A)) : A \in 2^{\mathcal{I}} \text{ and } f_{\mathcal{D}}(A) \geq \theta\}.$$

To define the collection of top- K FI's, we assume a fixed *canonical ordering* of the itemsets in $2^{\mathcal{I}}$ by decreasing frequency in \mathcal{D} , with ties broken arbitrarily, and label the itemsets A_1, A_2, \dots, A_m according to this ordering. For a given K , with $1 \leq K \leq m$, we denote with $f_{\mathcal{D}}^{(K)}$ the frequency $f_{\mathcal{D}}(A_K)$ of the K -th most frequent itemset A_K , and define the set of top- K FI's (with their respective frequencies) as

$$\text{TOPK}(\mathcal{D}, \mathcal{I}, K) = \text{FI}(\mathcal{D}, \mathcal{I}, f_{\mathcal{D}}^{(K)}).$$

One of the main uses of frequent itemsets is in the discovery of *association rules*.

Definition 2. An association rule W is an expression " $A \Rightarrow B$ " where A and B are itemsets such that $A \cap B = \emptyset$. The *support* $\sigma_{\mathcal{D}}(W)$ (resp. *frequency* $f_{\mathcal{D}}(W)$) of the association rule W is the support (resp. frequency) of the itemset $A \cup B$. The *confidence* $c_{\mathcal{D}}(W)$ of W is the ratio $\frac{f_{\mathcal{D}}(A \cup B)}{f_{\mathcal{D}}(A)}$ of the frequency of $A \cup B$ to the frequency of A .

Intuitively, an association rule " $A \Rightarrow B$ " expresses, through its support and confidence, how likely it is for the itemset B to appear in the same transactions as itemset A , so that when A is found in a transaction it is then possible to infer that B will be present in the same transaction with a probability equal to the confidence of the association rule.

Definition 3. Given a dataset \mathcal{D} with transactions built on a ground set \mathcal{I} , and given a minimum frequency threshold θ and a minimum confidence threshold γ , the AR's task with respect to θ and γ consist in finding the set

$$\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma) = \{(W, f_{\mathcal{D}}(W), c_{\mathcal{D}}(W)) \mid \text{Assoc. Rule } W, f_{\mathcal{D}}(W) \geq \theta, c_{\mathcal{D}}(W) \geq \gamma\}.$$

Often, with an abuse of the notation, we will say that an itemset A (resp. an association rule W) is in $\text{FI}(\mathcal{D}, \mathcal{I}, \theta)$ or in $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$ (resp. in $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$) and denote this fact with $A \in \text{FI}(\mathcal{D}, \mathcal{I}, \theta)$ or $A \in \text{TOPK}(\mathcal{D}, \mathcal{I}, K)$ (resp. $W \in \text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$), meaning that there is a pair $(A, f) \in \text{FI}(\mathcal{D}, \mathcal{I}, \theta)$ or $(A, f) \in \text{TOPK}(\mathcal{D}, \mathcal{I}, K)$ (resp. a triplet $(W, f_w, c_w) \in \text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$).

In this work we are interested in extracting absolute and relative approximations of the sets $\text{FI}(\mathcal{D}, \mathcal{I}, \theta)$, $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$ and $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$.

Definition 4. Given a parameter ε_{abs} (resp. ε_{rel}), an absolute ε_{abs} -close approximation (resp. a relative ε_{rel} -close approximation) of $\text{FI}(\mathcal{D}, \mathcal{I}, \theta)$ is a set $\mathcal{C} = \{(A, f_A) : A \in 2^{\mathcal{I}}, f_A \in [0, 1]\}$ of pairs (A, f_A) where f_A approximates $f_{\mathcal{D}}(A)$. \mathcal{C} is such that:

1. \mathcal{C} contains all itemsets appearing in $\text{FI}(\mathcal{D}, \mathcal{I}, \theta)$;
2. \mathcal{C} contains no itemset A with frequency $f_{\mathcal{D}}(A) < \theta - \varepsilon_{\text{abs}}$ (resp. $f_{\mathcal{D}}(A) < (1 - \varepsilon_{\text{rel}})\theta$);
3. For every pair $(A, f_A) \in \mathcal{C}$, it holds $|f_{\mathcal{D}}(A) - f_A| \leq \varepsilon_{\text{abs}}$ (resp. $|f_{\mathcal{D}}(A) - f_A| \leq \varepsilon_{\text{rel}}f_{\mathcal{D}}(A)$).

This definition extends easily to the case of top- K frequent itemsets mining using the equivalence $\text{TOPK}(\mathcal{D}, \mathcal{I}, K) = \text{FI}(\mathcal{D}, \mathcal{I}, f_{\mathcal{D}}^{(K)})$: an absolute (resp. relative) ε -close approximation to $\text{FI}(\mathcal{D}, \mathcal{I}, f_{\mathcal{D}}^{(K)})$ is an absolute (resp. relative) ε -close approximation to $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$.

For the case of association rules, we have the following definition.

Definition 5. Given a parameter ε_{abs} (resp. ε_{rel}), an absolute ε_{abs} -close approximation (resp. a relative ε_{rel} -close approximation) of $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$ is a set

$$\mathcal{C} = \{(W, f_W, c_W) : \text{association rule } W, f_W \in [0, 1], c_W \in [0, 1]\}$$

of triplets (W, f_W, c_W) where f_W and c_W approximate $f_{\mathcal{D}}(W)$ and $c_{\mathcal{D}}(W)$ respectively. \mathcal{C} is such that:

1. \mathcal{C} contains all association rules appearing in $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$;
2. \mathcal{C} contains no association rule W with frequency $f_{\mathcal{D}}(W) < \theta - \varepsilon_{\text{abs}}$ (resp. $f_{\mathcal{D}}(W) < (1 - \varepsilon_{\text{rel}})\theta$);
3. For every triplet $(W, f_W, c_W) \in \mathcal{C}$, it holds $|f_{\mathcal{D}}(W) - f_W| \leq \varepsilon_{\text{abs}}$ (resp. $|f_{\mathcal{D}}(W) - f_W| \leq \varepsilon_{\text{rel}}\theta$).
4. \mathcal{C} contains no association rule W with confidence $c_{\mathcal{D}}(W) < \gamma - \varepsilon_{\text{abs}}$ (resp. $c_{\mathcal{D}}(W) < (1 - \varepsilon_{\text{rel}})\gamma$);
5. For every triplet $(W, f_W, c_W) \in \mathcal{C}$, it holds $|c_{\mathcal{D}}(W) - c_W| \leq \varepsilon_{\text{abs}}$ (resp. $|c_{\mathcal{D}}(W) - c_W| \leq \varepsilon_{\text{rel}}c_{\mathcal{D}}(W)$).

Note that the definition of relative ε -close approximation to $\text{FI}(\mathcal{D}, \mathcal{I}, \theta)$ (resp. to $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$) is more stringent than the definition of ε -close solution to frequent itemset mining (resp. association rule mining) in [6, Sect. 3]. Specifically, we require an approximation of the frequencies (and confidences) in addition to the approximation of the collection of itemsets or association rules (property 3 in Def. 4 and properties 3 and 5 in Def. 5).

2.2 VC-Dimension

The Vapnik-Chernovenkis (VC) Dimension of a space of points is a measure of the complexity or expressiveness of a family of indicator functions (or equivalently a family of subsets) defined on that space [35]. A finite bound on the VC-dimension of a structure implies a bound on the number of random samples required for approximately learning that structure. We outline here some basic definitions and results and refer the reader to the works of Alon and Spencer [3, Sect. 14.4], Chazelle [8, Chap. 4], and Vapnik [34] for more details on VC-dimension.

VC-dimension is defined on *range spaces*:

Definition 6. A range space is a pair (X, R) where X is a (finite or infinite) set and R is a (finite or infinite) family of subsets of X . The members of X are called points and those of R are called ranges.

To define the VC-dimension of a range space we consider the projection of the ranges into a set of points:

Definition 7. Let (X, R) be a range space and $A \subset X$. The projection of R on A is defined as $P_R(A) = \{r \cap A : r \in R\}$.

The definition of *shattered* set will be heavily used in our proofs:

Definition 8. Let (X, R) be a range space and $A \subset X$. If $P_R(A) = 2^A$, then A is said to be shattered by R .

The VC-dimension of a range space is the cardinality of the largest set shattered by the space:

Definition 9. Let $S = (X, R)$ be a range space. The Vapnik-Chervonenkis dimension (or VC-dimension) of S , denoted as $\text{VC}(S)$ is the maximum cardinality of a shattered subset of X . If there are arbitrary large shattered subsets, then $\text{VC}(S) = \infty$.

The main application of VC-dimension in statistics and learning theory is its relation to the size of the sample needed to approximate learning the ranges, in the following sense.

Definition 10. Let (X, R) be a range space and let A be a finite subset of X .

1. For $0 < \varepsilon < 1$, a subset $B \subset A$ is an ε -approximation for A if $\forall r \in R$, we have

$$\left| \frac{|A \cap r|}{|A|} - \frac{|B \cap r|}{|B|} \right| \leq \varepsilon. \quad (1)$$

2. For $0 < p, \varepsilon < 1$, a subset $B \subset A$ is a relative (p, ε) -approximation for A if for any range $r \in R$ such that $\frac{|A \cap r|}{|A|} \geq p$ we have $\left| \frac{|A \cap r|}{|A|} - \frac{|B \cap r|}{|B|} \right| \leq \varepsilon \frac{|A \cap r|}{|A|}$ and for any range $r \in R$ such that $\frac{|A \cap r|}{|A|} < p$ we have $\frac{|B \cap r|}{|B|} \leq (1 + \varepsilon)p$.

An ε -approximation (resp. a relative (p, ε) -approximation) can be constructed by random sampling points of the point space [16, Thm. 2.12 (resp. 2.11)].

Theorem 1. There is an absolute positive constant c (resp. c') such that if (X, R) is a range-space of VC-dimension at most d , $A \subset X$ is a finite subset and $0 < \varepsilon, \delta < 1$ (resp. and $0 < p < 1$), then a random subset $B \subset A$ of cardinality m , where

$$m \geq \min \left\{ |A|, \frac{c}{\varepsilon^2} \left(d + \log \frac{1}{\delta} \right) \right\}, \quad (2)$$

(resp. $m \geq \min \left\{ |A|, \frac{c'}{\varepsilon^2 p} \left(d \log \frac{1}{p} + \log \frac{1}{\delta} \right) \right\}$) is an ε -approximation (resp. a relative (p, ε) -approximation) for A with probability at least $1 - \delta$.

Note that throughout the work we assume the sample to be drawn *with* replacement if $m < |A|$ (otherwise the sample is exactly the set A). Löffler and Phillips [24] showed experimentally that the absolute constant c is approximately 0.5. It is also interesting to note that an ε -approximation of size $O(\frac{d}{\varepsilon^2} \log \frac{d}{\varepsilon})$ can be built *deterministically* in time $O(d^{3d}(\frac{1}{\varepsilon^2} \log \frac{d}{\varepsilon})^d |X|)$ [8].

3 The Dataset's Range Space and its VC-dimension

We define one range space that is used in the derivation of the sample sizes needed to approximate the solutions to the tasks of market basket analysis.

Definition 11. Let \mathcal{D} be a dataset of transactions that are subsets of a ground set \mathcal{I} . We define $S = (X, R)$ to be a range space associated with \mathcal{D} such that:

1. $X = \mathcal{D}$ is the set of transactions in the dataset.
2. $R = \{T_{\mathcal{D}}(W) \mid W \subseteq \mathcal{I}\}$ is a family of sets of transactions such that for each itemset $W \subseteq \mathcal{I}$, the set $T_{\mathcal{D}}(W) = \{\tau \in \mathcal{D} \mid W \subseteq \tau\}$ of all transactions containing W is an element of R .

Theorem 2. Let \mathcal{D} be a dataset and let $S = (X, R)$ be the associated range space. Let $d \in \mathbb{N}$. Then $\text{VC}(S) \geq d$ if and only if there exists a set $\mathcal{A} \subseteq \mathcal{D}$ of d transactions from \mathcal{D} such that for each subset $\mathcal{B} \subseteq \mathcal{A}$ of \mathcal{A} , there exists an itemset $I_{\mathcal{B}}$ such that 1) all transactions in \mathcal{B} contain $I_{\mathcal{B}}$ and 2) no transaction $\rho \in \mathcal{A} \setminus \mathcal{B}$ contains $I_{\mathcal{B}}$.

Corollary 1. Let \mathcal{D} be a dataset and $S = (\mathcal{D}, R)$ be the corresponding range space. Then, the VC-Dimension $\text{VC}(S)$ of S , is the maximum integer d such that there is a set $\mathcal{A} \subseteq \mathcal{D}$ of d transactions from \mathcal{D} such that for each subset $\mathcal{B} \subseteq \mathcal{A}$ of \mathcal{A} , there exists an itemset $I_{\mathcal{B}}$ such that 1) all transactions in \mathcal{B} contain $I_{\mathcal{B}}$ and 2) no transaction $\rho \in \mathcal{A} \setminus \mathcal{B}$ contains $I_{\mathcal{B}}$.

Computing the exact VC-dimension of a dataset is extremely expensive from a computational point of view. This does not come as a surprise, as it is known that computing the VC-dimension of a range space (X, R) can take time $O(|R||X|^{\log |R|})$ [23, Thm. 4.1]. It is instead possible to give an upper bound to the VC-dimension of a dataset, and a procedure to efficiently compute the bound.

Definition 12. Let \mathcal{D} be a dataset. The d -index of a dataset is defined as the maximum integer d such that \mathcal{D} contains at least d transactions of length at least d .

A note of folklore: if the dataset represents the scientific publications of a given scientist, with transactions corresponding to articles and items in a transaction corresponding to the citations received by the paper, then the d -index of the dataset corresponds to the h -index of the scientist.

The d -index d of a dataset \mathcal{D} can be computed in one scan of the dataset and with total memory $O(d)$. The scan starts with $d^* = 1$ and it stores the length of the first transaction. At any given step the procedure stores d^* , the current estimate of d , computed as the maximum d' such that the scan up to this step found at least d' transactions

with length at least d' , and keeps a list of the sizes of the transactions longer than d' found so far. There can be no more than d' such transactions. As the scan proceeds, the procedure updates d^* and the list of transactions sizes greater than d^* .

The d -index is an upper bound to the VC-dimension of a dataset.

Theorem 3. *Let \mathcal{D} be a dataset with d -index d . Then the range space $S = (X, R)$ corresponding to \mathcal{D} has VC-dimension at most d .*

This bound is strict, i.e., there are indeed datasets with VC-dimension exactly d , as formalized by the following Theorem.

Theorem 4. *There exists a dataset \mathcal{D} with d -index d and such the corresponding range space has VC-dimension exactly d .*

The datasets built in the proof of Thm. 4 are extremely artificial. Our experiments suggest that the VC-dimension of real datasets is usually much smaller than the upper bound presented in Thm. 3.

4 Mining (top- K) Frequent Itemsets and Association Rules

We apply the VC-dimension results to constructing efficient sampling algorithms with performance guarantees for approximating the collections of FI's, top- K FI's and AR's.

4.1 Mining Frequent Itemsets

We construct bounds for the sample size needed to obtain relative/absolute ε -close approximations to the collection of FI's. The algorithms to compute the approximations use a standard exact FI's mining algorithm on the sample, with an appropriately adjusted minimum frequency threshold, as formalized in the following lemma.

Lemma 1. *Let \mathcal{D} be a dataset with transactions built on a ground set \mathcal{I} , and let d be the d -index of \mathcal{D} . Let $0 < \varepsilon, \delta < 1$. Let \mathcal{S} be a random sample of \mathcal{D} with size $|\mathcal{S}| = \min\{|\mathcal{D}|, \frac{4c}{\varepsilon^2} (d + \log \frac{1}{\delta})\}$, for some constant c . Then $\text{FI}(\mathcal{S}, \mathcal{I}, \theta - \frac{\varepsilon}{2})$ is an absolute ε -close approximation to $\text{FI}(\mathcal{D}, \mathcal{I}, \theta)$ with probability at least $1 - \delta$.*

One very interesting consequence of this result is that we do not need to know the minimum frequency threshold θ in advance to build the sample: the properties of the ε -approximation allow to use the same sample for any threshold and for different thresholds, i.e., the sample does not need to be rebuilt if we want to mine it with a threshold θ first and with another threshold θ' later.

It is important to note that the VC-dimension of a dataset, and therefore the sample size from (2) needed to probabilistically obtain an ε -approximation, is independent from the size (number of transactions) in \mathcal{D} and also of the size of $\text{FI}(\mathcal{S}, \mathcal{I}, \theta)$. It only depends on the quantity d , which is always less or equal to the length of the longest transaction in the dataset, which in turn is less or equal to the number of different items $|\mathcal{I}|$.

To obtain a relative ε -close approximation, we need to add a dependency on θ as shown in the following Lemma.

Lemma 2. Let \mathcal{D} , d , ε , and δ as in Lemma 1. Let \mathcal{S} be a random sample of \mathcal{D} with size $|\mathcal{S}| = \min\{|\mathcal{D}|, \frac{4(2+\varepsilon)c}{\varepsilon^2\theta(2-\varepsilon)} \left(d \log \frac{2+\varepsilon}{\theta(2-\varepsilon)} + \log \frac{1}{\delta} \right)\}$, for some constant c . Then $\text{Fl}(\mathcal{S}, \mathcal{I}, (1 - \frac{\varepsilon}{2})\theta)$ is a relative ε -close approximation to $\text{Fl}(\mathcal{D}, \mathcal{I}, \theta)$ with probability at least $1 - \delta$.

4.2 Mining Top- K Frequent Itemsets

Given the equivalence $\text{TOPK}(\mathcal{D}, \mathcal{I}, K) = \text{Fl}(\mathcal{D}, \mathcal{I}, f_{\mathcal{D}}^{(K)})$, we could use the above FI's sampling algorithms if we had a good approximation of $f_{\mathcal{D}}^{(K)}$, the threshold frequency of the top- K FI's.

For the absolute ε -close approximation we first execute a standard top- K FI's mining algorithm on the sample to estimate $f_{\mathcal{D}}^{(K)}$ and then run a standard FI's mining algorithm on the same sample using a minimum frequency threshold depending on our estimate of $f_{\mathcal{S}}^{(K)}$. Lemma 3 formalizes this intuition.

Lemma 3. Let \mathcal{D} , d , ε , and δ be as in Lemma 1. Let K be a positive integer. Let \mathcal{S} be a random sample of \mathcal{D} with size $|\mathcal{S}| = \min\{|\mathcal{D}|, \frac{16c}{\varepsilon^2} (d + \log \frac{1}{\delta})\}$, for some constant c , then $\text{Fl}(\mathcal{S}, \mathcal{I}, f_{\mathcal{S}}^{(K)} - \frac{\varepsilon}{2})$ is an absolute ε -close approximation to $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$ with probability at least $1 - \delta$.

Note that as in the case of the sample size required for an absolute ε -close approximation to $\text{Fl}(\mathcal{D}, \mathcal{I}, \theta)$, we do not need to know K in advance to compute the sample size for obtaining an absolute ε -close approximation to $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$.

Two different samples are needed for computing a relative ε -close approximation to $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$, the first one to compute a lower bound to $f_{\mathcal{D}}^{(K)}$, the second to extract the approximation. Details for this case are presented in Lemma 4.

Lemma 4. Let \mathcal{D} , d , ε , and δ be as in Lemma 1. Let K be a positive integer. Let δ_1, δ_2 be two reals such that $(1 - \delta_1)(1 - \delta_2) \geq (1 - \delta)$. Let \mathcal{S}_1 be a random sample of \mathcal{D} with some size $|\mathcal{S}_1| = \frac{\phi c}{\varepsilon^2} \left(d + \log \frac{1}{\delta_1} \right)$ for some $\phi > 2\sqrt{2}/\varepsilon$ and some constant c . If $f_{\mathcal{S}_1}^{(K)} \geq \frac{2\sqrt{2}}{\varepsilon\phi}$, then let $p = \frac{2-\varepsilon}{2+\varepsilon}\theta$ and let \mathcal{S}_2 be a random sample of \mathcal{D} of size $\min\{|\mathcal{D}|, \frac{4c}{\varepsilon^2 p} (d \log \frac{1}{p} + \log \frac{1}{\delta})\}$ for some constant c . Then $\text{Fl}(\mathcal{S}_2, \mathcal{I}, (1 - \varepsilon/2)(f_{\mathcal{S}_1}^{(K)} - \varepsilon/\sqrt{2\phi}))$ is a relative ε -close approximation to $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$ with probability at least $1 - \delta$.

4.3 Mining Association Rules

Our final theoretical contribution concerns the discovery of relative/absolute approximations to $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \eta)$ from a sample. Lemma 5 builds on a result from [6, Sect. 5] and covers the *relative* case, while Lemma 6 deals with the *absolute* one.

Lemma 5. Let $0 < \delta, \varepsilon, \theta, \gamma < 1$, $\phi = \max\{3, 2 - \varepsilon + 2\sqrt{1 - \varepsilon}\}$, $\eta = \frac{\varepsilon}{\phi}$, and $p = \frac{1-\eta}{1+\eta}\theta$. Let \mathcal{D} be a dataset with d -index d . Let \mathcal{S} be a random sample of \mathcal{D} of size $\min\{|\mathcal{D}|, \frac{c}{\eta^2 p} (d \log \frac{1}{p} + \log \frac{1}{\delta})\}$ for some constant c . Then $\text{AR}(\mathcal{S}, \mathcal{I}, (1 - \eta)\theta, \frac{1-\eta}{1+\eta}\gamma)$ is a relative ε -close approximation to $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$ with probability at least $1 - \delta$.

Lemma 6. Let \mathcal{D} , d , θ , γ , ε , and δ be as in Lemma 5 and let $\varepsilon_{\text{rel}} = \frac{\varepsilon}{\max\{\theta, \gamma\}}$.

Fix $\phi = \max\{3, 2 - \varepsilon_{\text{rel}} + 2\sqrt{1 - \varepsilon_{\text{rel}}}\}$, $\eta = \frac{\varepsilon_{\text{rel}}}{\phi}$, and $p = \frac{1-\eta}{1+\eta}\theta$. Let \mathcal{S} be a random sample of \mathcal{D} of size $\min\{|\mathcal{D}|, \frac{c}{\eta^2 p}(d \log \frac{1}{p} + \log \frac{1}{\delta})\}$ for some constant c . Then $\text{AR}(\mathcal{S}, \mathcal{I}, (1 - \eta)\theta, \frac{1-\eta}{1+\eta}\gamma)$ is an absolute ε -close approximation to $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$.

Note that the sample size needed for absolute ε -close approximations to $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$ depends on θ and γ , which was not the case for absolute ε -close approximations to $\text{FI}(\mathcal{D}, \mathcal{I}, \theta)$ and $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$.

5 Experimental Evaluation


In this section we present an extensive experimental evaluation of our methods to extract approximations of $\text{FI}(\mathcal{D}, \mathcal{I}, \theta)$, $\text{TOPK}(\mathcal{D}, \mathcal{I}, K)$, and $\text{AR}(\mathcal{D}, \mathcal{I}, \theta, \gamma)$. Due to space constraints, we focus on a subset of the results.

Our first goal is to evaluate the *quality* of the approximations obtained using our methods, by comparing the experimental results to the analytical bounds. We also evaluate how strict the bounds are by testing whether the same quality of results can be achieved at sample sizes smaller than those suggested by the theoretical analysis. We then show that our methods can significantly speed-up the mining process, fulfilling the motivating promises of the use of sampling in the market basket analysis tasks. Lastly, we compare the sample sizes from our results to the best previous work [6].


We tested our methods on both real and artificial datasets. The real datasets come from the FIMI'04 repository (<http://fimi.ua.ac.be/data/>). Since most of them have a moderately small size, we replicated their transactions a number of times, with the only effect of increasing the size of the dataset but no change in the distribution of the frequencies of the itemsets. The artificial datasets were built such that their corresponding range spaces have VC-dimension equal to the maximum transaction length d , which is the maximum possible as shown in Thm. 3. To create these datasets, we followed the proof of Thm. 4 and used the generator included in ARtool (<http://www.cs.umb.edu/~laur/ARtool/>), which is similar to the one presented in [2]. We used the FP-Growth and Apriori implementations in ARtool to extract frequent itemsets and association rules. In all our experiments we fixed $\delta = 0.1$. In the experiments involving absolute (resp. relative) ε -close approximations we set $\varepsilon = 0.01$ (resp. $\varepsilon = 0.05$). The constant c was fixed to 0.5 as suggested by [24]. For each dataset we selected a range of minimum frequency thresholds and a set of values for K when extracting the top- K frequent itemsets. For association rules discovery we set the minimum confidence threshold $\gamma \in \{0.5, 0.75, 0.9\}$. For each dataset and each combination of parameters we created random samples with size as suggested by our theorems and with smaller sizes to evaluate the strictness of the bounds. We measured, for each set of parameters, the *absolute frequency error* and the *absolute confidence error*, defined as the error $|f_{\mathcal{D}}(X) - f_{\mathcal{S}}(X)|$ (resp. $|c_{\mathcal{D}}(Y) - c_{\mathcal{S}}(Y)|$) for an itemset X (resp. an association rule Y) in the approximate collection extracted from sample \mathcal{S} . When dealing with the problem of extracting *relative* ε -close approximations, we defined the *relative frequency error* to be the absolute frequency error divided by the

real frequency of the itemset and analogously for the relative confidence error (dividing by the real confidence). In the figures we plot the maximum and the average for these quantities, taken over all itemsets or association rules in the output collection. In order to limit the influence of a single sample, we computed and plot in the figures the maximum (resp. the average) of these quantities in three runs of our methods on three different samples for each size.

The first important result of our experiments is that, for all problems, for every combination of parameters and every run, the collection of itemsets or of association rules obtained using our methods always satisfied the requirements to be an absolute or relative ε -close approximation to the real collection. Thus in practice our methods indeed achieve or exceed the theoretical guarantees for approximations of the collections $FI(\mathcal{D}, \mathcal{I}, \theta)$, $TOPK(\mathcal{D}, \mathcal{I}, \theta)$, and $AR(\mathcal{D}, \mathcal{I}, \theta, \gamma)$.




(a) Absolute Frequency Error



(b) Relative Frequency Error

Fig. 1: Absolute / Relative ε -close Approximation to $FI(\mathcal{D}, \mathcal{I}, \theta)$

Evaluating the strictness of the bounds to the sample size was the second goal of our experiments. In Figure 1a we show the behaviour of the maximum frequency error as function of the sample size in the itemsets obtained from samples using the method presented in Lemma 1 (i.e., we are looking for an *absolute* ε -close approximation to $FI(\mathcal{D}, \mathcal{I}, \theta)$). The rightmost plotted point corresponds to the sample size suggested by the theoretical analysis. We are showing the results for the dataset BMS-POS replicated 40 times (d-index $d = 134$), mined with $\theta = 0.02$. It is clear from the picture that the guaranteed error bounds are achieved even at sample sizes smaller than what suggested by the analysis and that the error at the sample size derived from the theory (rightmost plotted point for each line) is one to two orders of magnitude smaller than the maximum tolerable error $\varepsilon = 0.01$. This fact seems to suggest that there is still room for improvement in the bounds to the sample size needed to achieve an absolute ε -close approximation to $FI(\mathcal{D}, \mathcal{I}, \theta)$. In Fig. 1b we report similar results for the problem of computing a *relative* ε -close approximation to $FI(\mathcal{D}, \mathcal{I}, \theta)$ for an artificial dataset whose range space has VC-dimension d equal to the length of the longest transaction in the dataset, in this case 33. The dataset contained 100 million transactions. The sample size, suggested by Lemma 2, was computed using $\theta = 0.01$, $\varepsilon = 0.05$, and $\delta = 0.1$. The conclusions we can draw from the results for the behaviour of the relative frequency error are similar to those we got for the absolute case. For the case of absolute and relative ε -close approximation to $TOPK(\mathcal{D}, \mathcal{I}, K)$, we observed similar results, which we do not report here because of space constraints.



(a) Relative Frequency Error (b) Relative Confidence Error
 Fig. 2: Relative ε -close approximation to $AR(\mathcal{D}, \mathcal{I}, \theta, \gamma)$

The results of the experiments to evaluate our method to extract a relative ε -close approximation to $AR(\mathcal{D}, \mathcal{I}, \theta, \gamma)$ are presented in Fig. 2a and 2b. The same observations as before hold for the relative frequency error, while it is interesting to note that the relative confidence error is even smaller than the frequency error, most possibly because the confidence of an association rule is the ratio between the frequencies of two itemsets that appear in the same transactions and their sample frequencies will therefore have similar errors that cancel out when the ratio is computed. Similar conclusions can be made for the absolute ε -close case (not reported due to space constraints).

The major motivating intuition for the use of sampling in market basket analysis tasks is that mining a sample of the dataset is faster than mining the entire dataset. Nevertheless, the mining time does not only depend on the number of transactions, but also on the number of frequent itemsets. Given that our methods suggest to mine the sample at a lowered minimum frequency threshold, this may cause an increase in running time that would make our method not useful in practice, because there may be many more frequent itemsets than at the original frequency threshold. We performed a number of experiments to evaluate whether this was the case and present the results in Fig. 3. We mined the artificial dataset introduced before for different values of θ , and created samples of size sufficient to obtain a relative ε -close approximation to $FI(\mathcal{D}, \mathcal{I}, \theta)$, for $\varepsilon = 0.05$ and $\delta = 0.1$. Figure 3 shows the time needed to mine the large dataset and the time needed to create and mine the samples. It is possible to appreciate that, even considering the sampling time, the speed up achieved by our method is relevant, proving the usefulness of sampling.





Fig. 3: Runtime Comparison. The sample line includes the sampling time.

Comparing our results to previous work we note that the bounds generated by our technique are always linear in the VC-dimension d associated with the dataset. As reported in Table 1, the best previous work [6] presented bounds that are linear in the maximum transaction size Δ for two of the six problems studied here. Figures 4a and 4b shows a comparison of the actual sample sizes for relative ε -close approximations to $FI(\mathcal{D}, \mathcal{I}, \theta)$ for as function of θ and ε . To compute the points for these figures, we set

$\Delta = d = 50$, corresponding to the worst possible case for our method, i.e., when the VC-dimension of the range space associated to the dataset is exactly equal to the maximum transaction length. We also fixed $\delta = 0.05$ (the two methods behave equally as δ changes). For Fig. 4a, we fixed $\varepsilon = 0.05$, while for Fig. 4b we fixed $\theta = 0.05$. From the Figures we can appreciate that both bounds have similar, but not equal, dependencies on θ and ε . More precisely the bound presented in this work is less dependent on ε and only slightly more dependent on θ . It also evident that the sample sizes suggested by the bound presented in this work are always much smaller than those presented in [6] (the vertical axis has logarithmic scale). In this comparison we used $\Delta = d$, but almost all real datasets we encountered have $d \ll \Delta$ as shown in Table 2 which would result in a larger gap between the sample sizes provided by the two methods.



(a) Sample size as function of θ

(b) Sample size as function of ε

Fig. 4: Sample sizes for relative ε -close approximations to $F1(\mathcal{D}, \mathcal{I}, \theta)$.

	accidents	BMS-POS	BMS-Webview-1	kosarak	mushroom	pumsb*	retail	webdocs
Δ	51	164	267	2497	23	63	76	71472
d	46	81	57	443	23	59	58	2452

Table 2: Values for maximum transaction length Δ and d-index d for real datasets

References

- [1] Agrawal, R., Imieliński, T., and Swami, A. Mining association rules between sets of items in large databases. *SIGMOD Rec.* 22, 207–216, 1993.
- [2] Agrawal, R. and Srikant, R. Fast algorithms for mining association rules in large databases. VLDB '94.
- [3] Alon, N. and Spencer, J. H. *The Probabilistic Method* 3rd Ed. Wiley, 2008.
- [4] Brönnimann, H., Chen, B., Dash, M., Haas, P., and Scheuermann, P. Efficient data reduction with ease. KDD '03.
- [5] Ceglar, A. and Roddick, J. F. Association mining. *ACM Comput. Surv.* 38, 5, 2006.
- [6] Chakaravarthy, V. T., Pandit, V., and Sabharwal, Y. Analysis of sampling techniques for association rule mining. ICDT '09.
- [7] Chandra, B. and Bhaskar, S. A new approach for generating efficient sample from market basket data. *Expert Sys. with Appl.* 38, 3, 1321–1325, 2011.
- [8] Chazelle, B. *The discrepancy method: randomness and complexity*. Cambridge, 2000.
- [9] Chen, B., Haas, P., and Scheuermann, P. A new two-phase sampling based algorithm for discovering association rules. KDD '02.
- [10] Chen, C., Horng, S.-J., and Huang, C.-P. Locality sensitive hashing for sampling-based algorithms in association rule mining. *Expert Sys. with Appl.* 38, 10, 12388–12397, 2011.
- [11] Cheung, Y.-L. and Fu, A. W.-C. Mining frequent itemsets without support threshold: With and without item constraints. *IEEE Trans. on Knowl. and Data Eng.* 16, 1052–1069, 2004.

- [12] Chuang, K.-T., Chen, M.-S., and Yang, W.-C. Progressive sampling for association rules based on sampling error estimation. In *Adv. in Knowl. Disc. and Data Mining*. Springer, 2005.
- [13] Chuang, K.-T., Huang, J.-L., and Chen, M.-S. Power-law relationship and self-similarity in the itemset support distribution: analysis and applications. *The VLDB Journal* 17, 5, 2008.
- [14] Fu, A. W.-C., Kwong, R. W.-w., and Tang, J. Mining n-most interesting itemsets. ISMIS'00.
- [15] Han, J., Cheng, H., Xin, D., and Yan, X. Frequent pattern mining: current status and future directions. *Data Min. Knowl. Discov.* 15, 55–86, 2007.
- [16] Har-Peled, S. and Sharir, M. Relative (p, ϵ) -approximations in geometry. *Discr. & Comput. Geometry* 45, 3, 462–496, 2011.
- [17] Hu, X. and Yu, H. The research of sampling for mining frequent itemsets. In *Rough Sets and Knowledge Technology*. Springer, 2006.
- [18] Hwang, W. and Kim, D. Improved association rule mining by modified trimming. CIT '06.
- [19] Jia, C. and Lu, R. Sampling ensembles for frequent patterns. In *Fuzzy Systems and Knowledge Discovery*. Springer, 2005.
- [20] Jia, C.-Y. and Gao, X.-P. Multi-scaling sampling: An adaptive sampling method for discovering approximate association rules. *J. of Comp. Sci. and Tech.* 20, 309–318, 2005
- [21] John, G. H. and Langley, P. Static versus dynamic sampling for data mining. KDD '96.
- [22] Li, Y. and Gopalan, R. Effective sampling for mining association rules. In *AI 2004: Advances in Artificial Intelligence*. Springer, 2005
- [23] Linial, N., Mansour, Y., and Rivest, R. L. Results on learnability and the Vapnik-Chervonenkis dimension. *Information and Computation* 90, 1, 33–49, 1991.
- [24] Löffler, M., Phillips, J. M. Shape fitting on point sets with probability distributions. ESA'09.
- [25] Mahafzah, B. A., Al-Badarneh, A. F., and Zakaria, M. Z. A new sampling technique for association rule mining. *J. of Information Science* 35, 3, 358–376, 2009.
- [26] Mampaey, M., Tatti, N., and Vreeken, J. Tell me what i need to know: succinctly summarizing data with itemsets. KDD '11.
- [27] Mannila, H., Toivonen, H., and Verkamo, I. Efficient algorithms for discovering association rules. KDD'94
- [28] Parthasarathy, S. Efficient progressive sampling for association rules. ICDM '02.
- [29] Pietracaprina, A., Riondato, M., Upfal, E., and Vandin, F. Mining top-K frequent itemsets through progressive sampling. *Data Min. Knowl. Discov.* 21, 310–326, 2010.
- [30] Pietracaprina, A. and Vandin, F. Efficient incremental mining of top-K frequent closed itemsets. DS '07.
- [31] Riondato, M. and Upfal, E. Efficient discovery of association rules and frequent itemsets through sampling with tight performance guarantees. *CoRR abs/1111.6937*, 2011.
- [32] Scheffer, T. and Wrobel, S. Finding the most interesting patterns in a database quickly by using sequential sampling. *J. Mach. Learn. Res.* 3, 833–862, 2002.
- [33] Toivonen, H. Sampling large databases for association rules. VLDB '96.
- [34] Vapnik, V. N. *The Nature of Statistical Learning Theory*. Springer-Verlag, 1999.
- [35] Vapnik, V. N. and Chervonenkis, A. J. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Prob. and its Appl.* 16, 2, 264–280, 1971.
- [36] Vasudevan, D. and Vojnović, M. Ranking through random sampling. MSR-TR-2009-8 8, Microsoft Research, 2009.
- [37] Wang, J., Han, J., Lu, Y., and Tzvetkov, P. TFP: An efficient algorithm for mining top-k frequent closed itemsets. *IEEE Trans. on Knowl. and Data Eng.* 17, 652–664, 2005.
- [38] Wang, S., Dash, M., Chia, L.-T. Efficient sampling: Application to image data. PAKDD'05.
- [39] Zaki, M., Parthasarathy, S., Li, W., and Ogihara, M. Evaluation of sampling for data mining of association rules. RIDE '97.
- [40] Zhang, C., Zhang, S., and Webb, G. I. Identifying approximate itemsets of interest in large databases. *Applied Intelligence* 18, 91–104, 2003.
- [41] Zhao, Y., Zhang, C., Zhang, S. Efficient frequent itemsets mining by sampling. AMT '06.