

# Frequent Pattern Outlier Detection without Exhaustive Mining

Arnaud Giacometti and Arnaud Soulet

Université François Rabelais Tours, LI EA 6300  
3 place Jean Jaurès, F-41029 Blois, France  
`firstname.lastname@univ-tours.fr`

**Abstract.** Outlier detection consists in detecting anomalous observations from data. During the past decade, pattern-based outlier detection methods have proposed to mine all frequent patterns in order to compute the outlier factor of each transaction. This approach remains too expensive despite recent progress in pattern mining field. In this paper, we provide exact and approximate methods for calculating the frequent pattern outlier factor (FPOF) without extracting any pattern or by extracting a small sample. We propose an algorithm that returns the exact FPOF without mining any pattern. Surprisingly, it works in polynomial time on the size of the dataset. We also present an approximate method where the end-user controls the maximum error on the estimated FPOF. Experiments show the interest of both methods for very large datasets where exhaustive mining fails to provide the exact solution. The accuracy of our approximate method outperforms the baseline approach for a same budget in time or number of patterns.

## 1 Introduction

Outlier detection consists in detecting anomalous observations from data [1]. The outlier detection problem has important applications, such as detection of credit card fraud or network intrusions. Recently, outlier detection methods were proposed for categorical data using the concept of frequent patterns [2,3,4]. The key idea of such approaches is to consider the number of frequent patterns supported by each data observation. A data observation is unlikely to be an outlier if it supports many frequent patterns since frequent patterns correspond to the “common features” of the dataset. Frequent pattern outlier detection methods first extract all frequent itemsets from the data and then assign an outlier score to each data observation based on the frequent itemsets it contains. These outlier detection methods follow the schema of pattern-based two-step methods.

Pattern-based two-step methods [5] aim at exhaustively mining all patterns (first step) in order to build models (second step) like pattern sets (e.g., classifier [6]) or pattern-based measures (e.g., FPOF [2] or CPCQ index [7]). The completeness of pattern mining is often considered as a crucial advantage for constructing accurate models or measures. However, the completeness of the first step requires to adjust thresholds which is recognized as being very difficult. If

the thresholds are too low, the extraction becomes unfeasible. If the thresholds are too high, some essential patterns are missed. Finally, completeness leads to huge pattern volumes without guaranteeing not missing important patterns. For a smaller budget (in time or number of patterns), we claim that non-exhaustive methods can produce collections of patterns better adapted to the task of the second step. Interestingly, a non-exhaustive method can even guarantee a certain quality on the second step.

This paper revisits the calculation of the Frequent Pattern Outlier Factor (FPOF) based on non-exhaustive methods i.e., by mining a small number of patterns (or zero!). We first propose a method for calculating the exact FPOF of each transaction. Surprisingly, our method is non-enumerative in the sense that no pattern is generated (a fortiori, this is also a non-exhaustive method). For this, we reformulate the FPOF by operating directly on transaction pairs. This method calculates the FPOF in polynomial time on the number of transactions and items of the dataset. Experiments show that this method arrives to calculate the exact FPOF where the usual approach fails. We also propose a non-exhaustive approximate method that exploits a pattern sample instead of the complete collection of frequent patterns. Using Bennett’s inequality, this method selects the sample size so as to guarantee a maximum error for a given confidence. Experiments show its efficiency with reasonable error.

The outline of this paper is as follows. Section 2 reviews some related work. Section 3 introduces the basic definitions about the FPOF and states the problem of its exact and approximate calculation. In Section 4, we propose our exact non-enumerative method for calculating FPOF. We introduce our approximate method based on sampling in Section 5. Section 6 provides experimental results.

## 2 Related Work

In this paper, we focus on the outlier detection methods based on frequent patterns [2,3,4]. A broader view of outlier detection is provided by surveys including [1]. Pattern-based methods benefit from the progress of pattern mining made over the past two decades. Such methods have a double interest. On the one hand, they are well suited to handle categorical data unlike most other methods dedicated to numerical data. In addition, they also remain efficient for high-dimensional spaces. The first approach [2] introduced the frequent pattern outlier factor that exploits the complete collection of frequent itemsets (while [3] uses an opposite approach by considering non-frequent itemsets). More recently, [4] replaces the collection of frequent itemsets by the condensed representation of Non-Derivable Itemsets (NDI) which is more compact and less expensive to mine. We would go further by showing that the frequent pattern outlier factor proposed in [2] can be calculated without extracting any pattern or by extracting a small sample.

Recently, there has been a resurgence in pattern mining for non-exhaustive methods through pattern sampling [8,9]. Pattern sampling aims at accessing the pattern space  $\mathcal{L}$  by an efficient sampling procedure simulating a distribution

$\pi : \mathcal{L} \rightarrow [0, 1]$  that is defined with respect to some interestingness measure  $m$ :  $\pi(\cdot) = m(\cdot)/Z$  where  $Z$  is a normalizing constant (formal framework and algorithms are detailed in [9]). In this way, the user has a fast and direct access to the entire pattern language and with no parameter (except possibly the sample size). Pattern sampling has been introduced to facilitate interactive data exploration [10]. In this paper, we investigate the use of pattern sampling for assigning an outlier score to each transaction. With a lower (pattern or time) budget than that of an exhaustive method, we obtain a higher quality with a bounded error.

### 3 Frequent Pattern Based Outlier Detection

#### 3.1 Basic definitions

Let  $\mathcal{I}$  be a set of distinct literals called *items*, an itemset (or a pattern) is a subset of  $\mathcal{I}$ . The language of itemsets corresponds to  $\mathcal{L} = 2^{\mathcal{I}}$ . A transactional dataset is a multi-set of itemsets of  $\mathcal{L}$ . Each itemset, usually called *transaction*, is a data observation. For instance, Table 1 gives three transactional datasets with 4 or 5 transactions  $t_i$  described by until 4 items  $A, B, C$  and  $D$ .

$\mathcal{D}$		$\mathcal{D}'$		$\mathcal{D}''$		
Trans.	Items	Trans.	Items	Trans.	Items	
$t_1$	$A \ B$	$t_1$	$A \ B$	$t_1$	$A \ B$	<b>D</b>
$t_2$	$A \ B$	$t_2$	$A \ B$	$t_2$	$A \ B$	<b>D</b>
$t_3$	$A \ B$	$t_3$	$A \ B$	$t_3$	$A \ B$	<b>D</b>
$t_4$	$C$	$t_4$	$C$	$t_4$	$C$	
		$t_5$	<b>A \ B</b>			

**Table 1.** Three toy datasets with slight variations

Pattern discovery takes advantage of interestingness measures to evaluate the relevancy of a pattern. The *support* of a pattern  $X$  in the dataset  $\mathcal{D}$  is the proportion of transactions covered by  $X$  [11]:  $supp(X, \mathcal{D}) = |\{t \in \mathcal{D} : X \subseteq t\}|/|\mathcal{D}|$ . A pattern is said to be *frequent* when its support exceeds a user-specified minimal threshold. The set of all frequent patterns for  $\sigma$  as minimal threshold in  $\mathcal{D}$  is denoted by  $\mathcal{F}_\sigma(\mathcal{D})$ :  $\mathcal{F}_\sigma(\mathcal{D}) = \{X \in \mathcal{L} : supp(X, \mathcal{D}) \geq \sigma\}$ .

In the following, we manipulate pattern multisets which are collections of patterns admitting several occurrences of the same pattern. The representativeness of a pattern multiset  $\mathcal{P}$ , denoted by  $Supp(\mathcal{P}, \mathcal{D})$ , is the sum of the support of each pattern in  $\mathcal{P}$ :  $Supp(\mathcal{P}, \mathcal{D}) = \sum_{X \in \mathcal{P}} supp(X, \mathcal{D})$ . The range of  $Supp(\mathcal{P}, \mathcal{D})$  is  $[0, |\mathcal{P}|]$ . Given a cardinality, high representativeness means the multiset contains very common patterns of the dataset. For comparing the content of two pattern multisets, we use the semi-join operator, denoted by  $\mathcal{P}_2 \triangleright \mathcal{P}_1$ , that returns all the patterns of  $\mathcal{P}_2$  occurring in  $\mathcal{P}_1$ :  $\mathcal{P}_2 \triangleright \mathcal{P}_1 = \{X \in \mathcal{P}_2 : X \in \mathcal{P}_1\}$ . For instance,  $\{A, AB, A, D\} \triangleright \{C, A, B\} = \{A, A\}$ .

### 3.2 Frequent Pattern Outlier Factor

Intuitively, a transaction is more representative when it contains many patterns which are very frequent within the dataset. In contrast, an outlier contains only few patterns and these patterns are not very frequent. The frequent pattern outlier factor [2] formalizes this intuition:

**Definition 1 (FPOF).** *The frequent pattern outlier factor of a transaction  $t$  in  $\mathcal{D}$  is defined as follows:*

$$fpof(t, \mathcal{D}) = \frac{Supp(2^t, \mathcal{D})}{\max_{u \in \mathcal{D}} Supp(2^u, \mathcal{D})}$$

The range of  $fpof$  is  $[0, 1]$  where 1 means that the transaction is the most representative transaction of the dataset while a value near 0 means that the transaction is an outlier. Other normalizations (denominator) are possible like  $Supp(\mathcal{L}, \mathcal{D})$  or  $\sum_{t \in \mathcal{D}} Supp(2^t, \mathcal{D})$ . Whatever the normalization method, two transactions remain ordered in the same way (so it does not affect the Kendall's tau that is used for evaluating our method). Under a certain Markov model, the score  $fpof(t, \mathcal{D})$  is also the proportion of time that an analyst would dedicate to study the transaction  $t$  considering the collection of frequent itemsets [12].

In the first dataset provided by Table 1,  $t_1$  is covered by  $\emptyset$  ( $supp(\emptyset, \mathcal{D}) = 1$ ) and,  $A$ ,  $B$  and  $AB$  whose support equals to 0.75 ( $Supp(\{\emptyset, A, B, AB\}, \mathcal{D}) = 3.25$ ) while  $t_4$  is only covered by  $\emptyset$  and  $C$  ( $Supp(\{\emptyset, C\}, \mathcal{D}) = 1.25$ ). Consequently,  $fpof(t_1, \mathcal{D}_1) = 3.25/3.25$  and  $fpof(t_4, \mathcal{D}_1) = 1.25/3.25$ . In this example,  $t_4$  appears to be an outlier. It is easy to see that increasing the frequency of the patterns covering the first transactions (e.g., dataset  $\mathcal{D}'$ ) decreases the FPOF of  $t_4$ . Similarly, increasing the number of patterns covering the first transactions also decreases the FPOF factor of  $t_4$  (e.g., dataset  $\mathcal{D}''$ ).

### 3.3 Problem Formulation

The outlier detection problem consists in computing the FPOF for each transaction:


*Problem 1 (Exact Problem).* Given a dataset  $\mathcal{D}$ , compute the frequent pattern outlier factor of each transaction  $t \in \mathcal{D}$ .

In practice, the exact calculation of frequent pattern outlier factor is performed by mining all patterns appearing at least once in the dataset (i.e., with  $\sigma = 1/|\mathcal{D}|$ ). Of course, this expensive task is not possible for very large datasets. Instead, FPOF is approximated with a collection of frequent patterns i.e., with a higher minimal support threshold:

**Definition 2 ( $\sigma$ -Exhaustive FPOF).** *Given a minimal support threshold  $\sigma$ , the  $\sigma$ -exhaustive FPOF of a transaction  $t$  in  $\mathcal{D}$  is defined as follows:*

$$fpof_{\sigma}(t, \mathcal{D}) = \frac{Supp(\mathcal{F}_{\sigma}(\mathcal{D}) \triangleright 2^t, \mathcal{D})}{\max_{u \in \mathcal{D}} Supp(\mathcal{F}_{\sigma}(\mathcal{D}) \triangleright 2^u, \mathcal{D})}$$

The approximation becomes accurate with low support thresholds. Figure 1 (left) plots the Kendall’s tau of  $fprof_\sigma$  in comparison with  $fprof$  for some benchmarks<sup>1</sup>. When the support threshold becomes very low, the number of patterns and the extraction time explode, see Figure 1 (right). Furthermore, the approximation error is not estimated. With a smaller budget, we claim that it is possible to approximate more precisely FPOF while having a bound on the error.



**Fig. 1.** Kendall’s tau and number of patterns with minimal support threshold

The Kendall’s tau varies significantly depending on the dataset for a same minimal support threshold. This threshold is not easy to set for obtaining a good compromise between efficiency and quality. Therefore, it seems interesting that the user sets the maximum error that he/she tolerates on the result rather than a threshold related to the method:

*Problem 2 (Approximate Problem).* Given a dataset  $\mathcal{D}$ , two reals  $\delta$  and  $\epsilon$ , find a function approximating the frequent pattern outlier factor such that for each transaction  $t \in \mathcal{D}$ ,  $|fprof(t, \mathcal{D}) - \widehat{fprof}(t, \mathcal{D})| \leq \epsilon$  with confidence  $1 - \delta$ .

This problem aims at assigning an approximate FPOF to each transaction with the probability  $1 - \delta$  that the error is less than  $\epsilon$ . Problem 1 is a particular case of Problem 2 by fixing  $\epsilon = 0$  and  $\delta = 0$ .

## 4 Exact Non-Enumerative Method

This section addresses Problem 1. To calculate the FPOF of a transaction  $t$ , Definition 1 formulates the problem in terms of frequent patterns appearing in  $t$ . The idea is to reformulate this factor by considering what each transaction  $u$  brings to the transaction  $t$ . For instance, in dataset  $\mathcal{D}$ , the FPOF of the first transaction relies on  $Supp(\{\emptyset, A, B, AB\}, \mathcal{D})$

<sup>1</sup> It is the proportion of pairs of transactions which would be ranked similarly with the approximate FPOF and with the true FPOF (see Section 6 for a formal definition).

---

**Algorithm 1** Exact Non-Enumerative Method

---

**Input:** A dataset  $\mathcal{D}$

**Output:** FPOF computed for each transaction

- 1: Initialize  $fprof[t] \leftarrow 0$  for  $t \in \mathcal{D}$
  - 2: **for all**  $(t, u) \in \mathcal{D} \times \mathcal{D}$  **do**
  - 3:      $fprof[t] \leftarrow fprof[t] + 2^{|t \cap u|}$
  - 4: **end for**
  - 5: Normalize  $fprof[t] \leftarrow fprof[t]/Z$  for  $t \in \mathcal{D}$  where  $Z = \max_{u \in \mathcal{D}} fprof[u]$
  - 6: **return**  $fprof$
- 

which is equal to  $|\{\emptyset, A, B, AB, \emptyset, A, B, AB, \emptyset, A, B, AB, \emptyset\}|/4$ . Each subset  $\{\emptyset, A, B, AB\}$  and  $\{\emptyset\}$  result from the intersection of patterns covering  $t_1$  with those covering another transaction  $u \in \mathcal{D}$ . Thereby,  $Supp(\{\emptyset, A, B, AB\}, \mathcal{D}) = |\{\bigcup_{u \in \mathcal{D}} 2^{t_1} \cap 2^u\}|/|\mathcal{D}| = |\{\bigcup_{u \in \mathcal{D}} 2^{t_1 \cap u}\}|/|\mathcal{D}|$ . This observation is at the core of the following property:

*Property 1 (Reformulation).* Given a dataset  $\mathcal{D}$ , the frequent pattern outlier factor can be reformulated as follows for all transaction  $t \in \mathcal{D}$ :

$$fprof(t, \mathcal{D}) = \frac{\sum_{u \in \mathcal{D}} 2^{|t \cap u|}}{\max_{v \in \mathcal{D}} \sum_{u \in \mathcal{D}} 2^{|v \cap u|}}$$

*Proof.* Let  $\mathcal{D}$  be a dataset. Given  $u \in \mathcal{D}$ , we define  $\kappa(X, u) = 1$  if  $X \subseteq u$  and 0 otherwise. For each transaction  $t \in \mathcal{D}$ , we obtain:

$$\begin{aligned} Supp(2^t, \mathcal{D}) &= \frac{1}{|\mathcal{D}|} \sum_{X \subseteq t} \left( \sum_{\{u \in \mathcal{D}: X \subseteq u\}} 1 \right) = \frac{1}{|\mathcal{D}|} \sum_{X \subseteq t} \sum_{u \in \mathcal{D}} \kappa(X, u) = \frac{1}{|\mathcal{D}|} \sum_{u \in \mathcal{D}} \sum_{X \subseteq t} \kappa(X, u) \\ &= \frac{1}{|\mathcal{D}|} \sum_{u \in \mathcal{D}} \left( \sum_{X \subseteq t \wedge X \subseteq u} 1 \right) = \frac{1}{|\mathcal{D}|} \sum_{u \in \mathcal{D}} \left( \sum_{X \subseteq t \cap u} 1 \right) = \frac{1}{|\mathcal{D}|} \sum_{u \in \mathcal{D}} 2^{|t \cap u|} \end{aligned}$$

Injecting this equation into Definition 1 proves that Property 1 is right.  $\square$

From a conceptual point of view, it is interesting to note that ultimately, the FPOF of a transaction is just the sum of its similarity with each of transactions (where similarity between  $t$  and  $u$  is  $2^{|t \cap u|}$ ). This measure is therefore very close to traditional methods relying on pair-wise distance among data observations.

Interestingly, FPOF formula of Property 1 can be simply calculated with a double loop (see Algorithm 1). Therefore, the first strong result of this paper is to prove that Problem 1 can be solved in polynomial time, contrary to what was envisaged in the literature:

*Property 2 (Complexity).* The frequent pattern outlier factor of all transactions can be calculated in time  $O(|\mathcal{D}|^2 \times |\mathcal{I}|)$ .

Due to lack of space, we omit certain proofs. To the best of our knowledge, our proposal is the first method to calculate FPOF in polynomial time. Nevertheless, for large datasets with a lot of transactions, this complexity remains high. Then it makes sense to consider an approximate solution obtained much faster.

## 5 $\epsilon$ -Approximate Sampling Method

This section addresses Problem 2 by using pattern sampling. First, we propose a method for approximating FPOF from a pattern sample drawn according to frequency. Then we show how to choose the sample size to control the error.

### 5.1 Pattern Sampling for FPOF

In Section 3.3, we showed that the use of the most frequent patterns is insufficient to approximate accurately FPOF. The most frequent patterns do not measure the singularity of each transaction that also relies on more specific patterns (whose frequency varies from small to average). Conversely, do not considering frequent patterns would also be a mistake because they contribute significantly to FPOF. A reasonable approach is to select patterns randomly with a probability proportional to their weight in the calculation of FPOF. Typically, in the dataset  $\mathcal{D}$  of Table 1, the itemset  $AB$  is 3 times more important than itemset  $C$  in the calculation of FPOF due to their frequency.

In recent years pattern sampling techniques have been proposed to randomly draw patterns in proportion to their frequency [9]. Such approaches are ideal to bring us a well-adapted collection of patterns. Of course, it remains the non-trivial task of approximating FPOF starting from this collection:

**Definition 3 ( $k$ -Sampling FPOF).** *Given an integer  $k > 0$ , a  $k$ -sampling frequent pattern outlier factor of a transaction  $t$  in  $\mathcal{D}$  is defined as follows:*

$$fpof_k(t, \mathcal{D}) = \frac{|\mathcal{S}_k(\mathcal{D})_{\triangleright 2^t}|}{\max_{u \in \mathcal{D}} |\mathcal{S}_k(\mathcal{D})_{\triangleright 2^u}|}$$

where  $\mathcal{S}_k(\mathcal{D})$  is a sample of  $k$  patterns drawn from  $\mathcal{D}$  according to support:  $\mathcal{S}_k(\mathcal{D}) \sim \text{supp}(\mathcal{L}, \mathcal{D})$ .

It is important to note that  $|\cdot|$  is used here instead of  $\text{Supp}(\cdot, \mathcal{D})$  as in Definition 1. As the sampling technique already takes into account the frequency when it draws patterns, it is not necessary to involve the support here. Indeed, the draw is with replacement for the correct approximation of FPOF (without this replacement the most frequent patterns would be disadvantaged). It induces that the same pattern can have multiple occurrences within the sample  $\mathcal{S}_k(\mathcal{D})$ .

For the same sample size  $k$  and for the same transaction  $t$ , it is possible to calculate different values of a  $k$ -sampling FPOF due to  $\mathcal{S}_k(\mathcal{D})$ . But, the higher the threshold  $k$ , the less the difference between values stemming from two samples is high. Furthermore, the greater the sample size  $k$ , the better the approximation:

*Property 3 (Convergence).* Given a dataset  $\mathcal{D}$ , a  $k$ -sampling FPOF converges to the FPOF for all transaction  $t \in \mathcal{D}$ .

*Proof.*  $\mathcal{S}_k(\mathcal{D}) \sim \text{supp}(\mathcal{L}, \mathcal{D})$  means that there exists a constant  $\alpha > 0$  such that  $\forall X \in \mathcal{L}, \lim_{k \rightarrow \infty} |\mathcal{S}_k(\mathcal{D})_{\triangleright \{X\}}| = \alpha \text{supp}(X, \mathcal{D})$ . Then, for each transaction  $t$ , we obtain that:  $\lim_{k \rightarrow \infty} |\mathcal{S}_k(\mathcal{D})_{\triangleright 2^t}| = \alpha \sum_{X \in 2^t} \text{supp}(X, \mathcal{D}) = \alpha \text{Supp}(2^t, \mathcal{D})$ . By injecting this result into Definition 3, we conclude that Property 3 is right.  $\square$

---

**Algorithm 2**  $\epsilon$ -Approximate Sampling Method

---

**Input:** A dataset  $\mathcal{D}$ , a confidence  $1 - \delta$ , a bound  $\epsilon$

**Output:** A  $k$ -sampling frequent pattern outlier factor of all transactions in  $\mathcal{D}$  with an error bounded by  $\epsilon$  for a confidence  $1 - \delta$

```
1:  $\tilde{\epsilon} \leftarrow 1$ ;  $\mathcal{S} \leftarrow \emptyset$ 
2: while  $\tilde{\epsilon} > \epsilon$  do
3:    $\mathcal{S} \leftarrow \mathcal{S} \cup \{X\}$  where  $X \sim \text{supp}(\mathcal{L}, \mathcal{D})$  // add a pattern in the sample
4:    $m \leftarrow \arg \max_{t \in \mathcal{D}} \text{cov}_{\mathcal{S}}(t)$  // select the most covered transaction
      // estimate the maximal error on  $\text{cov}_{\mathcal{S}}$ 
5:    $e_t \leftarrow \sqrt{2\bar{\sigma}_t \ln(1/\delta)/|\mathcal{S}|} + \ln(1/\delta)/(3|\mathcal{S}|)$  for each  $t \in \mathcal{D}$ 
      // estimate the maximal error on FPOF
6:    $\tilde{\epsilon} \leftarrow \max_{t \in \mathcal{D}} \{\min\{1, (\text{cov}_{\mathcal{S}}(t) + e_t)/(\text{cov}_{\mathcal{S}}(m) - e_m)\} - \text{cov}_{\mathcal{S}}(t)/\text{cov}_{\mathcal{S}}(m)\}$ 
7:    $\tilde{\epsilon} \leftarrow \max_{t \in \mathcal{D}} \{\text{cov}_{\mathcal{S}}(t)/\text{cov}_{\mathcal{S}}(m) - \max\{0, (\text{cov}_{\mathcal{S}_k}(t) - e_t)/(\text{cov}_{\mathcal{S}}(m) + e_m)\}; \tilde{\epsilon}\}$ 
8: od
9: return  $\langle \text{cov}_{\mathcal{S}}(t) / \max_{u \in \mathcal{D}} \text{cov}_{\mathcal{S}}(u) \rangle_{t \in \mathcal{D}}$ 
```

---

Beyond convergence, the interest of this approach is the speed of convergence far superior to that of the  $\sigma$ -exhaustive frequent pattern outlier factor as shown in the experimental study (see Section 6). This speed is accompanied by a good efficiency due to a reasonable complexity of pattern sampling:

*Property 4 (Complexity).* A  $k$ -sampling FPOF of all transactions can be calculated in time  $O(k \times |\mathcal{I}| \times |\mathcal{D}|)$ .

Given a number of patterns  $k$ , a  $k$ -sampling FPOF is therefore effective to calculate. However, the choice of the sample size  $k$  is both difficult and essential in order to achieve a desired approximation as suggested by Problem 2. The next section presents an iterative method for fixing this sample size.

## 5.2 Bounding the Error

This section shows how to determine the right sample size  $k$  for computing a  $k$ -sampling FPOF satisfying user specified parameters (a maximum error with a given confidence). The idea is to draw a sample and to bound the maximum error of FPOF using a statistical result known as Bennett's inequality. If this error is less than that allowed by the user, the algorithm returns a sampling FPOF based on the current sample. Otherwise, it increases the sample size by drawing more patterns and so on.

We use Bennett's inequality to estimate the current error because it is true irrespective of the probability distribution. After  $k$  independent observations of real-valued random variable  $r$  with range  $[0, 1]$ , Bennett's inequality ensures that, with confidence  $1 - \delta$ , the true mean of  $r$  is at least  $\bar{r} - \epsilon$  where  $\bar{r}$  and  $\bar{\sigma}$  are respectively the observed mean and variance of the samples and

$$\epsilon = \sqrt{\frac{2\bar{\sigma} \ln(1/\delta)}{k}} + \frac{\ln(1/\delta)}{3k}$$



In our case, the random variable is the average number of patterns within a sample  $\mathcal{S}_k \sim \text{supp}(\mathcal{L}, \mathcal{D})$  that cover the transaction  $t$ . It is denoted by  $\text{cov}_{\mathcal{S}_k}(t)$  and defined as follows:  $\text{cov}_{\mathcal{S}_k}(t) = |\mathcal{S}_k \triangleright 2^t|/k$ . It is easy to see that a  $k$ -sampling FPOF factor can be rewritten using  $\text{cov}_{\mathcal{S}_k}$ :  $\text{fprof}_k(t, \mathcal{D}) = \text{cov}_{\mathcal{S}_k}(t)/\max_{u \in \mathcal{D}} \text{cov}_{\mathcal{S}_k}(u)$ . Using Bennett’s inequality and the above definition enables us to bound FPOF:

*Property 5 (FPOF Bound).* Given a dataset  $\mathcal{D}$  and confidence  $1 - \delta$ , the FPOF of a transaction  $t$  is bounded as follows:

$$\max \left\{ 0, \frac{\text{cov}_{\mathcal{S}_k}(t) - \epsilon_t}{\text{cov}_{\mathcal{S}_k}(u) + \epsilon_u} \right\} \leq \text{fprof}(t, \mathcal{D}) \leq \min \left\{ \frac{\text{cov}_{\mathcal{S}_k}(t) + \epsilon_t}{\text{cov}_{\mathcal{S}_k}(u) - \epsilon_u}, 1 \right\}$$

where  $\mathcal{S}_k \sim \text{supp}(\mathcal{L}, \mathcal{D})$ ,  $u = \arg \max_{v \in \mathcal{D}} \text{cov}_{\mathcal{S}_k}(v)$  and  $\epsilon_v = \sqrt{2\sigma_v \ln(1/\delta)/k} + \ln(1/\delta)/(3k)$ .

Algorithm 2 returns the approximate FPOF of all transactions by guaranteeing a bounded error of  $\epsilon$  with confidence  $1 - \delta$ . Basically, the main loop is iterated until that the maximal error  $\tilde{\epsilon}$  is inferior to the expected bound  $\epsilon$ . Lines 4-7 calculate the maximal error  $\tilde{\epsilon}$  using Property 5. If the maximal error is less than  $\epsilon$ , Line 9 returns the  $k$ -sampling FPOF with the current sample  $\mathcal{S}$ . Otherwise, one more pattern is drawn (Line 3).

As desired by Problem 2, Algorithm 2 approximates the FPOF:

*Property 6 (Correctness).* Given a dataset  $\mathcal{D}$ , a confidence  $1 - \delta$ , a bound  $\epsilon$ , Algorithm 2 returns a  $k$ -sampling frequent pattern outlier factor of a transaction  $t$  in  $\mathcal{D}$  approximating FPOF with an error bounded by  $\epsilon$  with confidence  $1 - \delta$ .

## 6 Experimental Study

This experimental study aims to compare the speed of the non-enumerative method with that of the exact exhaustive method (i.e.,  $1/|\mathcal{D}|$ -exhaustive) and to estimate the error quality of the  $\epsilon$ -approximate sampling method faced to the  $\sigma$ -exhaustive method. Due to the lack of space, we do not provide new experiments showing the interest of FPOF for detecting outliers as this aspect is already detailed in literature [2,3,4]. Experiments are conducted on datasets coming from the UCI Machine Learning repository ([archive.ics.uci.edu/ml](http://archive.ics.uci.edu/ml)) and the FIMI repository ([fimi.ua.ac.be](http://fimi.ua.ac.be)). Table 2 gives the main features of datasets in first columns. All experiments are performed on a 2.5 GHz Xeon processor with the Linux operating system and 2 GB of RAM memory. Each reported evaluation measure is the arithmetic mean of 10 repeated measurements (interval confidence are narrow enough to be omitted).

### 6.1 Exact Non-Enumerative Method

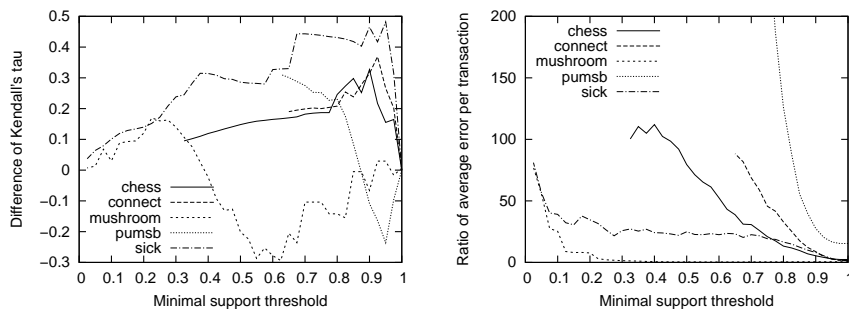
Table 2 reports the running time required for calculating the exact FPOF using the  $1/|\mathcal{D}|$ -exhaustive and the non-exhaustive methods (respectively the 4th and

$\mathcal{D}$	$ \mathcal{D} $	$ \mathcal{I} $	Exact methods		Approximate method
			Exh. Time (s)	Non-enum. Time (s)	0.1-Approx. Time (s)
chess	3,196	75	439.5	1.1	<b>0.3</b>
connect	67,557	129	748.5	<b>577.7</b>	<b>176.6</b>
mushroom	8,124	119	<b>0.4</b>	5.9	2.0
pumsb	49,096	7,117	time out	<b>1,970.5</b>	<b>175.0</b>
retail	88,162	16,470	<b>8.7</b>	5,969.9	<b>1.3</b>
sick	2,800	58	0.8	<b>0.5</b>	<b>0.5</b>

**Table 2.** Time comparison of the different methods

the 5th column). Note that the exact exhaustive method (as baseline) benefits from LCM which is one of the most recognized frequent itemset mining algorithm. The exact non-enumerative method is effective and rivals the exact exhaustive one. Its main advantage is to calculate the exact FPOF with datasets where the exact exhaustive method fails (e.g., pumsb where the execution was aborted after 5h).

## 6.2 Approximate Sampling Method



**Fig. 2.** Error of  $\sigma$ -exhaustive FPOF and  $k$ -sampling FPOF with  $\sigma$

This section compares our sampling method with the traditional heuristic based on frequent patterns as baseline. For this purpose, we use the Kendall's tau for comparing the ranking stemming from an approximate method  $f$  with that stemming from the FPOF (calculated with an exact method):

$$\tau(f, \mathcal{D}) = \frac{|\{(t, u) \in \mathcal{D}^2 : \text{sgn}(f(t, \mathcal{D}) - f(u, \mathcal{D})) = \text{sgn}(f_{\text{pof}}(t, \mathcal{D}) - f_{\text{pof}}(u, \mathcal{D}))\}|}{|\mathcal{D}|^2}$$

and in the same way, we also compute the average error per transaction:  $\epsilon(f, \mathcal{D}) = \sum_{t \in \mathcal{D}} |f(t, \mathcal{D}) - f_{\text{pof}}(t, \mathcal{D})| / |\mathcal{D}|$ .

The left chart of Figure 2 plots the difference between the Kendall’s tau of  $k$ -sampling FPOF and that of  $\sigma$ -exhaustive FPOF (when the curve is above 0, it means that the ranking of  $k$ -sampling FPOF is better than the ranking of  $\sigma$ -exhaustive FPOF). The right chart reports the average error of  $\sigma$ -sampling FPOF divided by that of  $k$ -sampling FPOF (when the curve is above 1, it means that the average error of  $k$ -sampling FPOF is smaller than that of  $\sigma$ -exhaustive FPOF). For each point, the minimal support threshold  $\sigma$  is used as parameter of the  $\sigma$ -exhaustive FPOF method. At the same time, the sample size  $k$  is fixed with the number of patterns mined with the minimal support threshold  $\sigma$ :  $k = |\mathcal{F}_\sigma(\mathcal{D})|$ . The  $k$ -sampling FPOF is clearly more accurate than the  $\sigma$ -exhaustive FPOF for a same pattern budget. For some datasets (e.g., `chess` or `sick`), the difference is always greater than 0. For other datasets, as soon as the number of patterns in the sample increases, the difference becomes positive. The average error ratio clearly shows that our method gives a better approximation of FPOF (especially, when the number of patterns is high).

Figure 3 plots the number of patterns and the average error per transaction of the  $\epsilon$ -approximate sampling method with bound  $\epsilon$  (for  $\delta = 0.1$ ). As expected, the smaller the error bound  $\epsilon$ , the greater the number of patterns in the sample. Therefore, the longer the execution time. It is interesting to note that the approximate sampling method (with  $\epsilon = 0.1$ ) is regularly faster than the exact methods (see Table 2). Finally, the real average error per transaction of the approximate sampling method is always much lower than the requested bound  $\epsilon$  (e.g.,  $\epsilon = 0.1$  actually gives an error less than 0.01). This difference results from the Bennett’s inequality that makes no assumption about the distribution.




Fig. 3. Number of patterns and average error per transaction with maximum error  $\epsilon$

## 7 Conclusion

We revisited the FPOF calculation in extracting the least possible patterns or no pattern. Despite this constraint, the proposed exact method has a complexity

better suited to certain datasets. Our approximate method using a sampling technique provides additional guarantees on the result with a maximum bound on the error. The experiments have shown the interest of these two approaches in terms of speed and accuracy compared to the usual exhaustive approach where all frequent patterns are mined.

Our proposal therefore combines the proven power of pattern-based methods by adding a guarantee on the quality of results without sacrificing speed thanks to sampling techniques. We think it can be generalized to other measures involving patterns or pattern-based models. We would also like to adapt this approach to build anytime algorithms. In the case of FPOF, it consists in extending the pattern sample indefinitely until the end-user wants to stop the process. Then, the algorithm returns the FPOF achieved with the current sample while estimating its error.

**Acknowledgements.** This work has been partially supported by the Prefute project, PEPS 2015, CNRS.

## References

1. Hawkins, D.M.: Identification of outliers. Volume 11. Springer (1980)
2. He, Z., Xu, X., Huang, Z.J., Deng, S.: FP-outlier: Frequent pattern based outlier detection. *Computer Science and Information Systems* **2**(1) (2005) 103–118
3. Otey, M.E., Ghoting, A., Parthasarathy, S.: Fast distributed outlier detection in mixed-attribute data sets. *Data Mining and Knowledge Discovery* **12**(2-3) (2006)
4. Koufakou, A., Secretan, J., Georgiopoulos, M.: Non-derivable itemsets for fast outlier detection in large high-dimensional categorical data. *Knowledge and information systems* **29**(3) (2011) 697–725
5. Knobbe, A., Crémilleux, B., Fürnkranz, J., Scholz, M.: From local patterns to global models: The lego approach to data mining. In: *From local patterns to global models: proceedings of the ECML PKDD 2008 Workshop*. (2008) 1–16
6. Liu, B., Hsu, W., Ma, Y.: Integrating classification and association rule mining. In: *International conference on Knowledge Discovery and Data mining*. (1998)
7. Liu, Q., Dong, G.: CPCQ: contrast pattern based clustering quality index for categorical data. *Pattern Recognition* **45**(4) (2012) 1739–1748
8. Chaoji, V., Hasan, M.A., Salem, S., Besson, J., Zaki, M.J.: ORIGAMI: A novel and effective approach for mining representative orthogonal graph patterns. *Statistical Analysis and Data Mining* **1**(2) (2008) 67–84
9. Boley, M., Lucchese, C., Paurat, D., Gärtner, T.: Direct local pattern sampling by efficient two-step random procedures. In: *ACM SIGKDD international conference on Knowledge discovery and data mining*. (2011) 582–590
10. van Leeuwen, M.: Interactive data exploration using pattern mining. In: *Interactive Knowledge Discovery and Data Mining in Biomedical Informatics*. Springer (2014) 169–182
11. Agrawal, R., Srikant, R., et al.: Fast algorithms for mining association rules. In: *International conference on Very Large Data Bases*. Volume 1215. (1994) 487–499
12. Giacometti, A., Li, D.H., Soulet, A.: Balancing the analysis of frequent patterns. In: *Advances in Knowledge Discovery and Data Mining*. Springer (2014) 53–64