



Free-Sets: A Condensed Representation of Boolean Data for the Approximation of Frequency Queries

JEAN-FRANÇOIS BOULICAUT
ARTUR BYKOWSKI
CHRISTOPHE RIGOTTI
*Laboratoire d'Ingénierie des Systèmes d'Information, INSA Lyon, Bâtiment 501, F-69621 Villeurbanne
Cedex, France*

Jean-François.Boulicaut@insa-lyon.fr
Artur.Bykowski@insa-lyon.fr
Christophe.Rigotti@insa-lyon.fr

Editors: Fayyad, Mannila, Ramakrishnan

Received August 10, 2000; Revised July 25, 2001

Abstract. Given a large collection of transactions containing items, a basic common data mining problem is to extract the so-called frequent itemsets (i.e., sets of items appearing in at least a given number of transactions). In this paper, we propose a structure called free-sets, from which we can approximate any itemset support (i.e., the number of transactions containing the itemset) and we formalize this notion in the framework of ϵ -adequate representations (H. Mannila and H. Toivonen, 1996. In Proc. of the Second International Conference on Knowledge Discovery and Data Mining (KDD'96), pp. 189–194). We show that frequent free-sets can be efficiently extracted using pruning strategies developed for frequent itemset discovery, and that they can be used to approximate the support of any frequent itemset. Experiments on real dense data sets show a significant reduction of the size of the output when compared with standard frequent itemset extraction. Furthermore, the experiments show that the extraction of frequent free-sets is still possible when the extraction of frequent itemsets becomes intractable, and that the supports of the frequent free-sets can be used to approximate very closely the supports of the frequent itemsets. Finally, we consider the effect of this approximation on association rules (a popular kind of patterns that can be derived from frequent itemsets) and show that the corresponding errors remain very low in practice.

Keywords: condensed representations, frequent pattern discovery, association rules

1. Introduction

Several data mining tasks (e.g., association rule mining (Agrawal et al., 1993)) are based on the evaluation of frequency queries to determine how often a particular pattern occurs in a large data set. We consider the problem of frequency query evaluation, when patterns are itemsets or conjunctions of properties, in dense data sets¹ like, for instance in the context of census data analysis (Bayardo, 1998) or log analysis (Bykowski and Gomez-Chantada, 2000). In these important but difficult cases, there is a combinatorial explosion of the number of frequent itemsets and computing the frequency of all of them turns out to be intractable. In this paper, we present an efficient technique to approximate closely the result of the frequency queries, and formalize it within the ϵ -adequate representation framework (Mannila and Toivonen, 1996). Intuitively, an ϵ -adequate representation is a representation of data that can be substituted to another representation to answer the same kind of queries,

Table 1. A relational table over four binary attributes.

<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>
1	1	0	0
1	0	0	1
0	1	1	1
1	1	1	0
1	0	0	0
1	0	0	0
0	1	0	0
0	1	0	0
0	0	1	0
0	0	1	0
0	0	0	1
0	0	0	1
1	1	1	1
1	1	1	1
1	1	1	1
1	1	1	1
1	1	1	1
1	1	1	1
1	1	1	1

but eventually with some lost of precision (bounded by the ϵ parameter). First evidences of the practical interest of such representations have been given in Boulicaut and Bykowski (2000) and Mannila and Toivonen (1996).

In this paper, we propose a new ϵ -adequate representation for the frequency queries. This representation, called *free-sets*, is more condensed than the ϵ -adequate representation based on itemsets Mannila and Toivonen (1996). The key intuition of the free-set representation is illustrated on the following example. Consider the binary attributes A, B, C, D in the relational table r depicted in Table 1 and suppose we are interested in the support of $\{A, B, C\}$ in r (i.e., the number of rows in r in which A, B and C are true). If we know that the rule $A, B \Rightarrow C$ nearly holds in r (i.e., when A and B are true in a row then, excepted in a few cases, C is also true) then we can approximate the support of itemset $\{A, B, C\}$ using the support of $\{A, B\}$. In Table 1 the rule $A, B \Rightarrow C$ has only one exception. So, we can use the support of $\{A, B\}$ as a value for the support of $\{A, B, C\}$. Moreover, we can approximate the support of any itemset X such that $\{A, B, C\} \subseteq X$ by the support of $X \setminus \{C\}$ because the rule $(X \setminus \{C\}) \Rightarrow C$ also holds with at most a few exceptions. For instance, the support of $\{A, B, C, D\}$ can be approximated by the support of $\{A, B, D\}$ since the rule $A, B, D \Rightarrow C$ can not have more exceptions than $A, B \Rightarrow C$. Furthermore, the support of $\{A, B, D\}$ does not need to be known directly, but can also be approximated itself. For example, the rule $A, D \Rightarrow B$ holds in Table 1 with one exception, so the support of $\{A, D\}$ can be used as an approximation of the support of $\{A, B, D\}$.

and then also of the support of $\{A, B, C, D\}$. It should be noticed that the framework presented in this paper can be restricted to rules with no exceptions. In this case, we still benefit of a significant condensation and speed-up when compared with frequent itemset extraction.

In the representation proposed in this paper, we call free-set an itemset Y such that the items in Y can not be used to form a nearly exact rule. For example, if we consider only rules having at most one exception, then the free-sets in Table 1 are $\{\emptyset, \{A\}, \{B\}, \{C\}, \{D\}, \{A, B\}, \{A, C\}, \{A, D\}, \{B, C\}, \{B, D\}, \{C, D\}\}$. All other subsets of $\{A, B, C, D\}$ contain items that can be used to form rules with zero or one exception (e.g., $A, B \Rightarrow C$ for $\{A, B, C\}$, $B, D \Rightarrow A$ for $\{A, B, D\}$, $A, C \Rightarrow D$ for $\{A, C, D\}$, $B, C \Rightarrow D$ for $\{B, C, D\}$, and $A, B, C \Rightarrow D$ for $\{A, B, C, D\}$) and thus are not free.

The freeness of itemsets is anti-monotonic in that sense that if a set is not a free-set then none of its supersets can be a free-set. The algorithm proposed to extract the free-sets takes advantage of this property. It first considers sets of size 0 (i.e., the empty itemset), then sets of size 1, and so on. When it determines that a set X is not free then it prunes the search space since there is no need to consider any of the supersets of X . For example, if the algorithm is executed on Table 1 and takes into account rules having at most one exception, then it will never consider the set $\{A, B, C, D\}$ because several sets among its subsets are not free (e.g., $\{A, B, C\}$).

The experiments show that *frequent* free-sets are an ϵ -adequate representation for frequency queries that can be extracted efficiently, even on dense data sets. They also show that the error made when approximating itemset support using frequent free-sets remains very low in practice.

Finally, we consider a popular application of frequent itemset discovery: the production of the so-called association rules (Agrawal et al., 1993). We determine bounds for the errors propagated on association rule characteristics when we use frequent free-sets to approximate the support of frequent itemsets, and we show that these bounds are very tight in practice.

This paper is a significant extension of a preliminary work presented in Boulicaut et al. (to appear). It includes proofs of the theorems, an in-depth error evaluation.

Organization of the paper. In the next section we introduce preliminary definitions used in this paper. In Section 3, we present the notion of free-set, and show that it can be used as an ϵ -adequate representation for the frequency queries. In Section 4, we present an algorithm to extract the frequent free-sets. In Section 5, we give practical evidences that frequent free-sets can be extracted efficiently and that the estimation of the supports of frequent itemsets using frequent free-sets leads in practice to very low errors. In Section 6, we show that these errors are propagated in a very favorable way when we produce association rules. We review related work in Section 7. Finally, we conclude with a summary and directions for future work.

2. Preliminary definitions

When applicable, we use the notational conventions and definitions from Mannila and Toivonen (1996, 1997).

2.1. Frequent sets

In this section, we recall standard definitions.

Definition 1 (binary database). Let R be a set of symbols called *items*. A *row* (also called *transaction*) is a subset of R . A *binary database* r over R is a multiset of transactions.

Definition 2 (support and frequency). We note $\mathcal{M}(r, X) = \{t \in r \mid X \subseteq t\}$ the multiset of rows matched by the itemset X and $Sup(r, X) = |\mathcal{M}(r, X)|$ the *support* of X in r , i.e., the number of rows matched by X . The *frequency* of X in r is $Sup(r, X)/|r|$. Let σ be a frequency threshold, $Freq(r, \sigma) = \{X \mid X \subseteq R \text{ and } Sup(r, X)/|r| \geq \sigma\}$ is the set of all σ -frequent itemsets in r .

For notational convenience, we also need the following specific definition.

Definition 3 (frequent sets). $FreqSup(r, \sigma)$ is the set of all pairs containing a frequent itemset and its support, i.e., $FreqSup(r, \sigma) = \{(X, Sup(r, X)) \mid X \subseteq R \text{ and } Sup(r, X)/|r| \geq \sigma\}$.

2.2. ϵ -adequate representation

Definition 4 (ϵ -adequate representation (Mannila and Toivonen, 1996)). Let \mathcal{S} be a class of structures. Let \mathcal{Q} be a class of queries for \mathcal{S} . The value of a query $Q \in \mathcal{Q}$ on a structure $s \in \mathcal{S}$ is assumed to be a real number in $[0, 1]$ and is denoted by $Q(s)$. An ϵ -adequate representation for \mathcal{S} w.r.t. a class of queries \mathcal{Q} , is a class of structures \mathcal{C} , a representation mapping $rep : \mathcal{S} \rightarrow \mathcal{C}$ and a query evaluation function $m : \mathcal{Q} \times \mathcal{C} \rightarrow [0, 1]$ such that $\forall Q \in \mathcal{Q}, \forall s \in \mathcal{S}, |Q(s) - m(Q, rep(s))| \leq \epsilon$.

Example 1. An example of a class of structures is the set noted \mathcal{DB}_R of all possible binary databases over a set of items R . An interesting query class is \mathcal{Q}_R , the set of all queries retrieving the frequency of an itemset $\subseteq R$. If we denote Q_X the query in \mathcal{Q}_R asking for the frequency of itemset X then $\mathcal{Q}_R = \{Q_X \mid X \subseteq R\}$ and the value of Q_X on a database instance $r \in \mathcal{DB}_R$ is defined by $Q_X(r) = Sup(r, X)/|r|$.

An example of ϵ -adequate representation for \mathcal{DB}_R w.r.t. \mathcal{Q}_R is the representation of $r \in \mathcal{DB}_R$ by means of $Freq(r, \epsilon)$. The corresponding rep, \mathcal{C} and m are as follows. $\forall r \in \mathcal{DB}_R, rep(r) = FreqSup(r, \epsilon), \mathcal{C} = \{rep(r) \mid r \in \mathcal{DB}_R\}, \forall Q_X \in \mathcal{Q}_R, \forall c \in \mathcal{C},$ if $\exists (X, \alpha) \in rep(r)$ then $m(Q_X, c) = \alpha/|r|$ else $m(Q_X, c) = 0$. It is straightforward to see that this is an ϵ -adequate representation for \mathcal{DB}_R w.r.t. \mathcal{Q}_R since $\forall Q_X \in \mathcal{Q}_R, \forall r \in \mathcal{DB}_R, |Q_X(r) - m(Q_X, rep(r))| \leq \epsilon$.

Interesting ϵ -adequate representations are *condensed representations*, i.e., ϵ -adequate representations where structures have a smaller size than the original structures.

3. The free-sets as a condensed representation

First, we recall the notion of *association rule*, and then define a class of rules called δ -strong rules in order to introduce the concept of free-set in a concise way.

Definition 5 (association rule). Let R be a set of items, an *association rule* based on R is an expression of the form $X \Rightarrow Y$, where $X, Y \subseteq R$, $Y \neq \emptyset$ and $X \cap Y = \emptyset$.

Definition 6 (δ -strong rule). A δ -strong rule² in a binary database r over R is an association rule $X \Rightarrow Y$ such that $Sup(r, X) - Sup(r, X \cup Y) \leq \delta$, i.e., the rule is violated in no more than δ rows.

In this definition, δ is supposed to have a small value, so a δ -strong rule is intended to be a rule with very few exceptions.

3.1. Free-sets

Definition 7 (δ -free-set). Let r be a binary database over R , $X \subseteq R$ is a δ -free-set w.r.t. r if and only if there is no δ -strong rule based on X in r . The set of all δ -free-sets w.r.t. r is noted $Free(r, \delta)$.

Since δ is supposed to be rather small, informally, a free-set is a set of items such that its subsets (seen as conjunction of properties) are not related by any very strong positive correlation.

One of the most interesting properties of *freeness* is its *anti-monotonicity* w.r.t. itemset inclusion.

Definition 8 (anti-monotonicity). A property ρ is *anti-monotone* if and only if for all itemsets X and Y , $\rho(X)$ and $Y \subseteq X$ implies $\rho(Y)$.

The anti-monotonicity has been identified as a key property for efficient pattern mining (Mannila and Toivonen, 1997; Ng et al., 1998), since it is the formal basis of a safe pruning criterion. Indeed, efficient frequent set mining algorithms like apriori (Agrawal et al., 1996) make use of the (anti-monotone) property “*is frequent*” for pruning.

The anti-monotonicity of freeness follows directly from the definition of free-set and is stated by the following theorem.

Theorem 1. *Let X be an itemset. For all $Y \subseteq X$ if $X \in Free(r, \delta)$ then $Y \in Free(r, \delta)$.*

3.2. Free-sets as an ϵ -adequate representation

We show now that δ -free-sets can be used to answer frequency queries with a bounded error. The following lemma states that the support of any itemset can be approximated using the support of one of the free-sets.

Lemma 1. *Let r be a binary database over a set of items R , $X \subseteq R$ and $\delta \in [0, |r|]$, then there exists $Y \subseteq X$ such that $Y \in \text{Free}(r, \delta)$ and $\text{Sup}(r, Y) \geq \text{Sup}(r, X) \geq \text{Sup}(r, Y) - \delta|X|$.*

Proof: We show this using a recurrence on $|X|$. The statement is true for $|X| = 0$ if we take $Y = \emptyset$. Suppose the statement is true for $|X| = i$. Let X be a subset of R such that $|X| = i + 1$. If $X \in \text{Free}(r, \delta)$ then we can simply choose $Y = X$. If $X \notin \text{Free}(r, \delta)$ then by definition of $\text{Free}(r, \delta)$ there exists a δ -strong rule $Z_1 \rightarrow Z_2$ based on X . Let A be an item in Z_2 and $Z_3 = X \setminus \{A\}$. As $|Z_3| = |X| - 1$ using the recurrence hypothesis we know that there exists $Y \subseteq Z_3$ such that $Y \in \text{Free}(r, \delta)$ and $\text{Sup}(r, Z_3) \geq \text{Sup}(r, Y) - \delta|Z_3|$. Since $Z_1 \rightarrow Z_2$ is a δ -strong rule, then $\text{Sup}(r, Z_1) - \text{Sup}(r, Z_1 \cup Z_2) \leq \delta$. $\text{Sup}(r, Z_1) - \text{Sup}(r, Z_1 \cup Z_2)$ is the number of rows not matched by Z_2 but matched by Z_1 , thus $\text{Sup}(r, Z_1) - \text{Sup}(r, Z_1 \cup Z_2)$ is greater or equal to $\text{Sup}(r, Z_1 \cup Z_3) - \text{Sup}(r, Z_1 \cup Z_2 \cup Z_3)$ (i.e., the number of rows not matched by Z_2 but matched by Z_1 and Z_3). So we have $\text{Sup}(r, Z_1 \cup Z_3) - \text{Sup}(r, Z_1 \cup Z_2 \cup Z_3) \leq \delta$ which simplifies to $\text{Sup}(r, Z_3) - \text{Sup}(r, X) \leq \delta$. Since $\text{Sup}(r, Z_3) \geq \text{Sup}(r, Y) - \delta|Z_3|$ and $|Z_3| = |X| - 1$ we deduce $\text{Sup}(r, X) \geq \text{Sup}(r, Y) - \delta|X|$. The other inequality $\text{Sup}(r, Y) \geq \text{Sup}(r, X)$ is straightforward because $Y \subseteq Z_3 \subseteq X$. \square

This lemma states that the support of an itemset X can be approximated using the support of one of the free-sets, but it does not determine which free-set to use. We now show that this can be done by simply choosing among the free-sets included in X any free-set with a minimal support value. This is stated more formally by the following theorem.

Theorem 2. *Let r be a binary database over a set of items R , $X \subseteq R$ and $\delta \in [0, |r|]$, then for any $Y \subseteq X$ such that $Y \in \text{Free}(r, \delta)$ and $\text{Sup}(r, Y) = \min(\{\text{Sup}(r, Z) \mid Z \subseteq X \text{ and } Z \in \text{Free}(r, \delta)\})$ we have $\text{Sup}(r, Y) \geq \text{Sup}(r, X) \geq \text{Sup}(r, Y) - \delta|X|$.*

Proof: Let Y be a subset of X such that $Y \in \text{Free}(r, \delta)$ and satisfying $\text{Sup}(r, Y) = \min(\{\text{Sup}(r, Z) \mid Z \subseteq X \text{ and } Z \in \text{Free}(r, \delta)\})$. Since $Y \subseteq X$ we have immediately that $\text{Sup}(r, Y) \geq \text{Sup}(r, X)$. By Lemma 1, there exists $Z \subseteq X$ such that $Z \in \text{Free}(r, \delta)$ and $\text{Sup}(r, Z) \geq \text{Sup}(r, X) \geq \text{Sup}(r, Z) - \delta|X|$. Since Y has the minimal support among all subsets of X in $\text{Free}(r, \delta)$, then $\text{Sup}(Z) \geq \text{Sup}(Y)$. Thus $\text{Sup}(Z) - \delta|X| \geq \text{Sup}(Y) - \delta|X|$. As $\text{Sup}(r, X) \geq \text{Sup}(r, Z) - \delta|X|$, we have $\text{Sup}(r, X) \geq \text{Sup}(Y) - \delta|X|$. \square

In practice, computing the whole collection of δ -free-sets is often intractable. We show now that such an exhaustive mining can be avoided since an ϵ -adequate representation to answer frequency queries can be obtained if we extract only *frequent* free-sets together with a subset of the corresponding negative border (Mannila and Toivonen, 1997).

Definition 9 (frequent free-set). Let r be a binary database over a set of items R , we denote $\text{FreqFree}(r, \sigma, \delta) = \text{Freq}(r, \sigma) \cap \text{Free}(r, \delta)$ the set of σ -frequent δ -free-sets w.r.t. r .

Let us adapt the concept of negative border from Mannila and Toivonen (1997) to our context.

Definition 10 (negative border of frequent free-sets). Let r be a binary database over a set of items R , the negative border of $\text{FreqFree}(r, \sigma, \delta)$ is denoted by $\text{Bd}^-(r, \sigma, \delta)$ and is

defined as follows: $\mathcal{B}d^-(r, \sigma, \delta) = \{X \mid X \subseteq R, X \notin \text{FreqFree}(r, \sigma, \delta) \wedge (\forall Y \subset X, Y \in \text{FreqFree}(r, \sigma, \delta))\}$.

Informally, the negative border $\mathcal{B}d^-(r, \sigma, \delta)$ consists of the smallest itemsets (w.r.t. set inclusion) that are not σ -frequent δ -free. Our approximation technique only needs a subset of the negative border $\mathcal{B}d^-(r, \sigma, \delta)$. This subset, denoted by $\text{FreeB}d^-(r, \sigma, \delta)$, is the set of all free-sets in $\mathcal{B}d^-(r, \sigma, \delta)$.

Definition 11. $\text{FreeB}d^-(r, \sigma, \delta) = \mathcal{B}d^-(r, \sigma, \delta) \cap \text{Free}(r, \delta)$

As in the case of an ϵ -adequate representation for \mathcal{DB}_R w.r.t. \mathcal{Q}_R using frequent itemsets (see Section 2.2), we need the free-sets and their supports.

Definition 12. $\text{FreqFreeSup}(r, \sigma, \delta)$ is the set of all pairs containing a frequent free-set and its support, i.e., $\text{FreqFreeSup}(r, \sigma, \delta) = \{\langle X, \text{Sup}(r, X) \rangle \mid X \in \text{FreqFree}(r, \sigma, \delta)\}$.

We can now define the ϵ -adequate representation w.r.t. the frequency queries.

Definition 13. The *frequent free-sets representation* w.r.t. σ, δ and a query class $\mathcal{Q} \subseteq \mathcal{Q}_R$, is defined by a class of structures \mathcal{C} , a representation mapping rep and a query evaluation function m , where $\forall r \in \mathcal{DB}_R, \text{rep}(r) = \langle \text{FreqFreeSup}(r, \sigma, \delta), \text{FreeB}d^-(r, \sigma, \delta) \rangle$, $\mathcal{C} = \{\text{rep}(r) \mid r \in \mathcal{DB}_R\}$, $\forall Q_X \in \mathcal{Q}, \forall c \in \mathcal{C}$, if $\exists Y \in \text{FreeB}d^-(r, \sigma, \delta), Y \subseteq X$ then $m(Q_X, c) = 0$ else $m(Q_X, c) = \min(\{\alpha \mid \exists Z \subseteq X, \langle Z, \alpha \rangle \in \text{FreqFreeSup}(r, \sigma, \delta)\})/|r|$.

Using this representation, the frequency of an itemset X is approximated as follows. If X has a subset Y which is free but not frequent then the frequency of X is considered to be 0. Otherwise we take the smallest support value among the supports of the subsets of X that are free and frequent.

We now establish that this representation is an ϵ -adequate representation for the following database class and query class.

Definition 14. $\mathcal{DB}_{R,s} = \{r \mid r \in \mathcal{DB}_R \text{ and } |r| \geq s\}$, i.e., the set of all binary databases having at least s rows. $\mathcal{Q}_{R,n} = \{Q_X \mid X \subseteq R \text{ and } |X| \leq n\}$, i.e., the set of frequency queries on itemsets having no more than n items.

Theorem 3. A frequent free-sets representation w.r.t. σ, δ and a query class $\mathcal{Q}_{R,n}$ is an ϵ -adequate representation for $\mathcal{DB}_{R,s}$ w.r.t. $\mathcal{Q}_{R,n}$ where $\epsilon = \max(\sigma, n\delta/s)$.

Proof: Let Q_X be a query in $\mathcal{Q}_{R,n}$ and r an database in $\mathcal{DB}_{R,s}$. If there exists $Y \in \text{FreeB}d^-(r, \sigma, \delta)$ such that $Y \subseteq X$ then X is not σ -frequent so $Q_X(r) \leq \sigma$. Since $m(Q_X, c) = 0$ we have $|Q_X(r) - m(Q_X, \text{rep}(r))| \leq \sigma$.

In the case where no $Y \in \text{FreeB}d^-(r, \sigma, \delta)$ is a subset of X , this means that all δ -free-set included in X are σ -frequent. Whence $\min(\{\text{Sup}(r, Z) \mid Z \subseteq X \text{ and } Z \in \text{Free}(r, \delta)\}) = \min(\{\alpha \mid \exists Z \subseteq X, \langle Z, \alpha \rangle \in \text{FreqFreeSup}(r, \sigma, \delta)\})$ which is equal to $m(Q_X, \text{rep}(r))$. Thus, by Theorem 2, $m(Q_X, \text{rep}(r)) \geq Q_X(r) \geq m(Q_X, \text{rep}(r)) - \delta|X|/|r|$. So we have $|Q_X(r) - m(Q_X, \text{rep}(r))| \leq n\delta/s$. \square

4. Discovering all frequent free-sets

In this section, we describe an algorithm, called MINEX, that generates all frequent free-sets. For clarity, we omit the fact that it outputs their supports as well. Implementation issues are presented in Section 4.2.

4.1. The algorithm—an abstract version

MINEX can be seen as an instance of the levelwise search algorithm presented in Mannila and Toivonen (1997). It explores the itemset lattice (w.r.t. set inclusion) levelwise, starting from the empty set and stopping at the level of the largest frequent free-sets. More precisely, the collection of candidates is initialized with the empty set as single member (the only set of size 0) and then the algorithm iterates on candidate evaluation and larger candidate generation. At each iteration of this loop, it scans the database to find out which candidates of size i are frequent free-sets. Then, it generates candidates for the next iteration, taking every set of size $i + 1$ such that all proper subsets are frequent free-sets. The algorithm finishes when there is no more candidate. The algorithm is given below as Algorithm 1.

Algorithm 1 (MINEX)

Input: r a binary database over a set of items R , σ and δ two thresholds.

Output: $\text{FreqFree}(r, \sigma, \delta)$

1. $C_0 := \{\emptyset\}$;
2. $i := 0$;
3. **while** $C_i \neq \emptyset$ **do**
4. $\text{FreqFree}_i := \{X \mid X \in C_i \text{ and } X \text{ is a } \sigma\text{-frequent } \delta\text{-free-set in } r\}$;
5. $C_{i+1} := \{X \mid X \subseteq R \text{ and } \forall Y \subset X, Y \in \bigcup_{j \leq i} \text{FreqFree}_j\} \setminus \bigcup_{j \leq i} C_j$;
6. $i := i + 1$;
7. **od**;
8. **output** $\bigcup_{j < i} \text{FreqFree}_j$;

Using the correctness result of the levelwise search algorithm given in Mannila and Toivonen (1997) the following theorem is straightforward.

Theorem 4 (Correctness). *Algorithm MINEX computes the sets of all σ -frequent δ -free-sets.*

4.2. Implementation issues

We used techniques similar to the ones described in Agrawal and Srikant (1994) for frequent itemset mining. The candidate generation is made using a join-based function, and the itemset support counters are updated w.r.t. a row of the database using a *prefix-tree* data structure.

The key point that needs a new specific technique is the freeness test in step 4 of the algorithm. An efficient computation of this test can be done, based on the following remark: Z is not a δ -free-set if and only if there exist $A \in Z$ and $X = Z \setminus \{A\}$ such that X is not δ -free or X is δ -free and $X \Rightarrow \{A\}$ is a δ -strong rule. Furthermore, the step 5 of the algorithm guarantees that if Z is a candidate then X must be δ -free since X is a subset of Z . Therefore, during the i th iteration, we might first compute the δ -strong rules of the form $X \Rightarrow \{A\}$, where $X \in \mathit{FreqFree}_i$ and $A \in R \setminus X$, and then use them to remove candidates in \mathcal{C}_{i+1} that are not δ -free. Thus, at the beginning of an iteration, only free-sets are candidates.

This is incorporated in the algorithm by replacing steps 4 and 5 with the following steps:

- 4.1 $\mathit{FreqFree}_i := \{X \mid X \in \mathcal{C}_i \text{ and } X \text{ is a } \sigma\text{-frequent}\};$
 4.2 $\mathit{NotFree}_{i+1} := \{Z \mid Z = X \cup \{A\} \text{ where } X \in \mathit{FreqFree}_i, A \in R \setminus X$
 and $X \Rightarrow \{A\}$ is a δ -strong rule};
 5.1 $\mathcal{C}_{i+1}^g := \{X \mid X \subseteq R \text{ and } \forall Y \subset X, Y \in \bigcup_{j \leq i} \mathit{FreqFree}_j\} \setminus \bigcup_{j \leq i} \mathcal{C}_j;$
 5.2 $\mathcal{C}_{i+1} := \mathcal{C}_{i+1}^g \setminus \mathit{NotFree}_{i+1};$

The steps 4.1 and 4.2 can be computed efficiently within the same database scan as follows. For each candidate X considered in step 4.1, we maintain a node n (in the prefix-tree) containing an integer denoted by $n.count$ to count the support of X and a set denoted by $n.rhs$ to determine the δ -strong rule having a left hand side equal to X . More precisely, $n.rhs$ is a set of pairs of the form $\langle A, e \rangle$. Such a pair $\langle A, e \rangle$ means that the rule $X \Rightarrow \{A\}$ has e exceptions.

Steps 4.1 and 4.2 are performed by first initializing for each candidate $X \in \mathcal{C}_i$ the corresponding node n in the prefix-tree with $n.count := 0$ and $n.rhs := \emptyset$. Then the database r is scanned, and for each row t the prefix-tree is used to find all candidates matching t . For each such candidate X , corresponding to a node n in the tree, we call $matched(t, X, n, \delta)$ to update $n.count$ and $n.rhs$. The description of $matched$ is given below as Algorithm 2.

Algorithm 2 (*matched*)

Input: a row t , a candidate X , a node n of the prefix-tree and the threshold δ .

Output: n updated.

```

if  $n.count \leq \delta$  then
  for all  $i \in t \setminus X$  do
    if  $\exists \langle j, e \rangle \in n.rhs$  with  $j = i$  then
       $n.rhs := n.rhs \cup \{i, n.count\};$ 
    fi
  od
fi
for all  $\langle j, e \rangle \in n.rhs$  do
  if  $j \notin t$  then
     $n.rhs := n.rhs \setminus \{\langle j, e \rangle\};$ 
  if  $e < \delta$  then

```

```

                                 $n.rhs := n.rhs \cup \{j, e + 1\};$ 
                                fi
                            fi
    od
     $n.count := n.count + 1;$ 
    output  $n$ ;

```

The key idea is that the set $n.rhs$ is created lazily, in the sense that a pair in $n.rhs$ is created for an item A only when the algorithm finds A in a transaction t matched by X . Moreover, when δ rows matched by X have been encountered, there is no need to create new entries for new items in $n.rhs$ since these items will lead obviously to rules with more than δ exceptions.

5. Experiments

The running prototype is implemented in C⁺⁺. We use a PC with 512 MB of memory and a 500 MHz Pentium III processor under Linux operating system.

For an experimental evaluation, we chose the PUMSB* data set, a PUMS census data set³ preprocessed by researchers from IBM Almaden Research Center. The particularity of PUMS data sets is that they are very dense and make the mining of all frequent itemsets together with their supports intractable for low frequency thresholds, because of the combinatorial explosion of the number of frequent itemsets (Bayardo, 1998).

5.1. Frequent free-set vs. frequent set condensation

Table 2 shows a comparison of the extraction of frequent sets and frequent free-sets for different frequency thresholds and different values of δ . The collections $FreqFree(r, \sigma, \delta)$ are significantly smaller than the corresponding $Freq(r, \sigma)$. For frequency thresholds of 15% and 20% $Freq(r, \sigma)$ is so large that it is clearly impossible to provide it on our platform, while the extraction of $FreqFree(r, \sigma, \delta)$ remains tractable. For these two frequency thresholds of 15% and 20%, we use lower-bound estimations of $|Freq(r, \sigma)|$. These lower-bounds are computed using the δ -strong rules collected by MINEX (see Section 4.2) to find the size of the largest frequent itemset. If this size is m then there are at least 2^m frequent itemsets. Figure 1 (left) emphasizes, using logarithmically scaled axes, the difference of the size of the various representations. We observe a brutal change between the size of $Freq(r, 0.25)$ and of $Freq(r, 0.20)$: 1000 times more frequent itemsets than expected by extrapolating the trend given by $Freq(r, 0.25)$ and $Freq(r, 0.30)$. If we look at the trend of the number of frequent free-sets it seems to be unchanged. The reason for this, is that between 0.25 and 0.20 we reach a support threshold where the number of strong rules increases significantly and then leads to the explosion of the number of frequent itemsets, but not to the explosion of the number of frequent free-sets.

Using also logarithmically scaled axes, figure 1 (right) shows that the extraction time for MINEX grows up exponentially when the frequency threshold is reduced. This is due to the

Table 2. Comparison of different representations at various frequency thresholds.

σ	15%			20%			25%			30%		
	0	10	20	0	10	20	0	10	20	0	10	20
δ	12	11	10	12	10	9	11	9	9	10	9	8
Max frequent free-set size (\approx MIN-EX DB scans)												
$ FreqFree(r, \sigma, \delta) $	909 806	324 743	232 887	253 107	105 615	76 413	78 220	36 310	27 137	26 972	14 631	11 079
$FreqFree(r, \sigma, \delta)$ extraction time in sec. (MIN-EX)	11 977	6 590	5 126	4 233	2 342	1 890	1 540	905	731	533	373	302
Max frequent set size (\approx apriori DB scans)		35			32			18			16	
$ Freq(r, \sigma) $		$>2^{35}$			$>2^{32}$			2 064 946			432 699	
$Freq(r, \sigma)$ extraction time in sec. (apriori)		N/A			N/A			14 559			3 469	




Figure 1. Extraction time and sizes of different representations.




Figure 2. Behavior of MINEX w.r.t. the number of rows and the number of items.

combinatorial explosion of the number of frequent free-sets. Apriori-based algorithms have a similar exponential evolution of the extraction time, due in this case to the combinatorial explosion of the number of frequent sets.

5.2. Scale-up experiment

On figure 2, we report the extraction time (for $\sigma = 20\%$) when changing the number of rows or the number of items in the data set. We observe an exponential complexity w.r.t. the number of items and a linear complexity w.r.t. number of rows in the data set if the value of δ follows the number of tuples (e.g., if we double the number of rows then we double the value of δ). This is emphasized by a superimposed straight line on figure 2 (left).

5.3. Approximation error in practice

In this section we report the practical error made on σ -frequent itemset supports when using the approximation based on σ -frequent δ -free-sets. This evaluation is made on the PUMSB* data set used in the previous experiments and also on a PUMS data set of Kansas in a less favorable case.

In the PUMSB* data set, for $\sigma = 0.3$, there are 432699 σ -frequent sets and the largest has $n = 16$ items. We computed the condensed representation $FreqFreeSup(r, 0.3, 20)$ which contains 11079 elements.

Table 3. Error observed on σ -frequent itemset supports by itemset size.

	Average abs. sup. error	Average rel. sup. error (%)	Maximal abs. sup. error	Maximal rel. sup. error (%)
Itemset size				
1	0	0	0	0
2	0.36	0.002	18	0.11
3	1.17	0.007	20	0.13
4	2.14	0.012	37	0.18
5	3.24	0.019	37	0.22
6	4.33	0.026	39	0.24
7	5.31	0.032	39	0.24
8	6.12	0.038	45	0.29
9	6.80	0.042	45	0.29
10	7.40	0.046	45	0.29
11	7.92	0.050	45	0.29
12	8.39	0.054	44	0.28
13	8.82	0.057	38	0.26
14	9.22	0.060	31	0.19
15	9.58	0.062	24	0.15
16	9.86	0.064	15	0.10

Theoretical error bounds for the frequent set support approximation can be determined using Theorem 2 as follows. In this experiment, the maximal absolute support error is $\delta * n = 20 * 16 = 320$ rows. The maximal relative support error can be obtained assuming that the maximal theoretical absolute error occurs on the σ -frequent set of minimal frequency (i.e., σ). The PUMSB* data set contains $N = 49046$ rows. So, the maximal relative support error is $\delta * n / (N * \sigma) = 2.18\%$.

The support of each of the 432699 σ -frequent itemsets is approximated using the collection $FreqFreeSup(r, 0.3, 20)$ and Theorem 2 and then compared to the exact support. The maximal observed absolute support error is 45 rows, and the maximal relative support error is 0.29%. The average absolute support error is 6.01 rows and the average relative support error is 0.037%. Tables 3 and 4 show that these errors remain very low even for frequent sets containing a lot of items and for low supports.

In this experiment the value of δ is small w.r.t. the minimal support required. The ratio is $20 / (0.3 * 49046) = 0.136\%$. We now report another experiment where the value of δ represents more than 1% of the minimal support required, and thus is likely to greatly increase the value of the error.

The data set used in this experiment is a PUMS data set of Kansas state.⁴ We use a version of this data set that has been preprocessed at the University of Clermont-Ferrand (France) in Prof. L. Lakhal's research group. We have reduced this data set to 10000 rows and 317 items to be able to extract all σ -frequent itemsets at a low frequency threshold. For

Table 4. Error observed on σ -frequent itemset supports by interval of support.

	Average abs. sup. error	Average rel. sup. error	Maximal abs. sup. error	Maximal rel. sup. error (%)
Itemset support (%)				
[30,40]	6.20	0.039%	45	0.29
(40,50]	3.29	0.016%	38	0.19
(50,60]	0.07	2.8×10^{-6}	9	0.03
(60,70]	0	0	0	0
(70,80]	0	0	0	0
(80,90]	0	0	0	0
(90,100]	0	0	0	0

$\sigma = 0.05$ (500 rows), there are 90755 σ -frequent sets and the largest has $n = 13$ items. We computed $\text{FreqFreeSup}(r, 0.05, 6)$ which contains 4174 elements.

In this experiment, the maximal absolute support error is $\delta * n = 6 * 13 = 78$ rows. The maximal relative support error is $\delta * n / (N * \sigma) = 15.6\%$ ($N = 10000$ rows in the experiment).

The supports of the σ -frequent itemsets are approximated using $\text{FreqFreeSup}(r, 0.05, 6)$ and compared to the exact supports. The maximal observed absolute support error is 18 rows, and the maximal relative support error is 3.1%. The average absolute support error is 2.12 rows and the average relative support error is 0.28%. A more detailed distribution of the error is given in Tables 5 and 6. These results show that the error remains low in practice even when the value of δ is high w.r.t. the minimal support.

Table 5. Error observed on σ -frequent itemset supports by itemset size.

	Average abs. sup. error	Average rel. sup. error (%)	Maximal abs. sup. error	Maximal rel. sup. error (%)
Itemset size				
1	0	0	0	0
2	0.24	0.03	6	1.1
3	0.65	0.07	10	1.3
4	1.10	0.13	12	2.1
5	1.53	0.18	14	2.7
6	1.92	0.24	18	3.1
7	2.31	0.31	18	3.1
8	2.75	0.38	18	3.1
9	3.28	0.47	18	3.1
10	3.90	0.58	18	2.9
11	4.58	0.71	18	2.9
12	5.20	0.83	15	2.9
13	5.50	0.88	11	2.0

Table 6. Error observed on σ -frequent itemset supports by interval of support.

	Average abs. sup. error	Average rel. sup. error (%)	Maximal abs. sup. error	Maximal rel. sup. error (%)
Itemset support (%)				
[5,10]	2.16	0.337	18	3.11
(10,20]	2.03	0.159	14	1.17
(20,30]	2.22	0.089	10	0.47
(30,40]	2.03	0.063	10	0.33
(40,50]	1.25	0.027	5	0.12
(50,60]	1.70	0.031	6	0.10
(60,70]	0.66	0.010	6	0.10
(70,80]	0	0	0	0
(80,90]	0	0	0	0
(90,100]	0	0	0	0

6. Effect of errors on association rules

A popular application of the extraction of frequent itemsets is the discovery of association rules (Agrawal et al., 1993). In this section, we give bounds for the error made on support and confidence of association rules when these rules are derived from frequent δ -free-sets instead of frequent itemsets. The notion of association rules has been recalled in Definition 5. Support and confidence are the two most widely used objective interestingness measures for association rules and are commonly defined as follows.

Definition 15 (support and confidence). Let $X \Rightarrow Y$ be an association rule based on the set of items R . The support and confidence of this rule in a database r over R are denoted by $Sup(r, X \Rightarrow Y)$ and $Conf(r, X \Rightarrow Y)$ and are defined respectively by $Sup(r, X \Rightarrow Y) = Sup(r, X \cup Y)$ and $Conf(r, X \Rightarrow Y) = Sup(r, X \cup Y) / Sup(r, X)$. The rule $X \Rightarrow Y$ is *frequent* in r w.r.t. a frequency threshold σ if $X \cup Y \in Freq(r, \sigma)$.

6.1. Error bounds for support approximation

The error on support of association rules is the same as the error on support of itemsets. For a frequent rule $X \Rightarrow Y$, if we use frequent δ -free-sets to determine its support, by Theorem 2 we always have an overestimate of its support with an error of at most $\delta|X \cup Y|$. In practice, we have the same approximation errors as those presented in Section 5.3.

6.2. Error bounds for confidence approximation

Let $X \Rightarrow Y$ be a frequent rule in r . Suppose we used frequent δ -free-sets to approximate $Sup(r, X \cup Y)$ and $Sup(r, X)$. These approximations are denoted respectively by $\overline{Sup}(r, X \cup Y)$ and $\overline{Sup}(r, X)$. Now, we can approximate $Conf(r, X \Rightarrow Y)$ by $\overline{Conf}(r, X \Rightarrow Y) =$

Table 7. Bounds for absolute error on rule confidence.

Confidence	Support						
	0.05	0.1	0.2	0.3	0.4	0.5	0.6
0.99	0.0308	0.0116	0.0047	0.0033	0.0012	0.0010	0.0010
0.95	0.0295	0.0111	0.0045	0.0031	0.0011	0.0010	0.0010
0.9	0.0280	0.0105	0.0042	0.0030	0.0011	0.0009	0.0009
0.85	0.0264	0.0099	0.0040	0.0028	0.0010	0.0008	0.0009

$\overline{Sup}(r, X \cup Y)/\overline{Sup}(r, X)$. By Theorem 2, and since we have overestimated the supports, then $Sup(r, X \cup Y)/(Sup(r, X) + \delta|X|) \leq \overline{Conf}(r, X \Rightarrow Y) \leq (Sup(r, X \cup Y) + \delta|X \cup Y|)/Sup(r, X)$.

Thus a bound for the absolute error made on the confidence when we use $\overline{Conf}(r, X \Rightarrow Y)$ instead of $Conf(r, X \Rightarrow Y)$ is $\max(Sup(r, X \cup Y)/(Sup(r, X) + \delta|X|) - Conf(r, X \Rightarrow Y), (Sup(r, X \cup Y) + \delta|X \cup Y|)/Sup(r, X) - Conf(r, X \Rightarrow Y))$.

Now, we derive values of this bound in practice, using the experiments reported in Section 5.3. We consider the PUMS data set of Kansas state, which is less favorable than the other (PUMSB*) since the error on the support was larger.

Let $ar(s, c)$ be the set of all association rules in this data set with support s and confidence c . For a given pair $\langle s, c \rangle$, we bound the error made on confidence for all rules in $ar(s, c)$ as follows. The support of the left hand side of any of these rules is $s' = s/c$. Using the experimental results of Section 5.3, we can find the maximal relative support error made on s and s' , denoted respectively by rse and rse' . Then we bound the absolute error made on the confidence by $\max(s/(s' + s' \times rse') - c, (s + s \times rse)/s' - c)$.

We consider support $s \in \{0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6\}$ and confidence $c \in \{0.99, 0.95, 0.9, 0.85\}$. For each pair $\langle s, c \rangle$, we used the maximal relative error on support given in Table 6 to bound the error made on confidence for the set of rule $ar(s, c)$. The corresponding values are presented in Table 7. For example, if we consider rules with confidence 0.99 and support 0.05, the maximal absolute error made on confidence is 0.0308. For higher rule supports the error decreases. This variation corresponds to the reduction of the maximal relative error for higher supports in Table 7. For lower confidence values the error also decreases. This is due to the fact that a lower confidence implies a higher support for the left hand side of the rule and thus a lower error on the left hand side support.

7. Related work

Using incomplete information about itemset frequencies for some mining task, e.g., Boolean rule mining, has been proposed in Mannila and Toivonen (1996), and formalized in the general framework of ϵ -adequate representations. Probabilistic approaches to the problem of frequency queries have also been investigated (see Pavlov et al., 2000).

Several search space reductions based on nearly exact (or exact) association rules have been proposed. The use of the nearly exact association rules to estimate the confidence of other rules and then to prune the search space has been suggested in Bayardo (1997) but

not investigated nor experimented. Efficient mining of nearly exact rules (more specifically rules with at most δ exceptions) with a single attribute in both the left and the right hand sides has been proposed in Fujiwara et al. (2000). Search space pruning using exact association rules has been experimented in Bayardo (1997) in the context of rule mining and developed independently in the context of frequent itemset mining in Pasquier et al. (1999). Pasquier et al. (1999) implicitly proposes a kind of condensed representation called *closed* itemsets which is strongly related to the notion of 0-free-sets (δ -free-sets with $\delta = 0$). Mining 0-free-sets or closed itemsets lead to similar gains, but mining δ -free-sets with $\delta \neq 0$ offers additional search space reductions (at the cost of an uncertainty on supports). It should also be noticed that by definition exact rules are very sensitive to noise. If we process a noisy data set (a very common case in practice) a few exceptions to the exact rules can appear easily. Then the pruning methods based on exact rules will be less effective, while the mining of δ -free-sets with $\delta \neq 0$ can still benefit of an important search space reduction.

The techniques mentioned in this section present important benefits on dense data sets, but if we consider very sparse data sets, we can hardly expect to have many exact or nearly exact rules that hold, and thus all these techniques are likely to be less interesting. Moreover, on very sparse data sets, these techniques may be a little bit slower than the direct extraction of frequent itemsets without pruning, since they can not take advantage of important search space reductions, but have to pay for a little overhead due to the tests performed to detect the rules.

8. Conclusion and future work

We proposed a structure called free-sets that can be extracted efficiently, even on dense data sets, and that can be used to approximate closely the support of frequent itemsets. We formalized this approximation in the framework of ϵ -adequate representations (Mannila and Toivonen, 1996) and gave a correct extraction algorithm formulated as an instance of the levelwise search algorithm presented in Mannila and Toivonen (1997).

We reported experiments showing that frequent free-sets can be extracted even when the extraction of frequent itemsets turns out to be intractable. The experiments also show that the error made when approximating the support of frequent itemsets using the support of frequent free-sets remains very low in practice. Finally, we considered the effect of this approximation on the support and confidence of association rules. We bounded the corresponding errors and the experiments show that these bounds are very tight in practice.

Interesting future work includes applications of the notion of δ -free-set to the the approximation of the support of general boolean formula as investigated in Mannila and Toivonen (1996).

Notes

1. E.g., data sets containing many strong correlations.
2. Stemming from the notion of *strong rule* of Piatetsky-Shapiro (1991).
3. http://www.almaden.ibm.com/cs/quest/data/long_11_patterns.bin.tar
4. <ftp://ftp2.cc.ukans.edu/pub/ippbr/census/pums/pums90ks.zip>

References

- Agrawal, R., Imielinski, T., and Swami, A. 1993. Mining association rules between sets of items in large databases. In Proc. of the 1993 ACM SIGMOD International Conference on Management of Data (SIGMOD'93), Washington, D.C., pp. 207–216.
- Agrawal, R., Mannila, H., Srikant, R., Toivonen, H., and Verkamo, A.I. 1996. Fast discovery of association rules. In Advances in Knowledge Discovery and Data Mining, AAAI Press, pp. 307–328.
- Agrawal, R. and Srikant, R. 1994. Fast algorithms for mining association rules in large databases. In Proc. of the 20th International Conference on Very Large Data Bases (VLDB'94), Santiago de Chile, Chile, pp. 487–499.
- Bayardo, R.J. 1997. Brute-force mining of high-confidence classification rules. In Proc. of the Third International Conference on Knowledge Discovery and Data Mining (KDD'97), Newport Beach, California, pp. 123–126.
- Bayardo, R.J. 1998. Efficiently mining long patterns from databases. In Proc. of the 1998 ACM SIGMOD International Conference on Management of Data (SIGMOD'98), Seattle, Washington, pp. 85–93.
- Boulicaut, J.-F. and Bykowski, A. 2000. Frequent closures as a concise representation for binary data mining. In Proc. of the Fourth Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'00), Vol. 1805 of LNAI, Kyoto, JP, Berlin: Springer-Verlag, pp. 62–73.
- Boulicaut, J.-F., Bykowski, A., and Rigotti, C. to appear. Approximation of frequency queries by mean of free-sets. In Proc. of the 4th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD'00), LNAI, Lyon, France, Berlin: Springer-Verlag.
- Bykowski, A. and Gomez-Chantada, L. 2000. Frequent itemset extraction in highly-correlated data: A web usage mining application. In Proc. of the 2000 International Workshop on Web Knowledge Discovery and Data Mining (WKDDM'00), Kyoto, Japan, pp. 27–42.
- Fujiwara, S., Ullman, J.D., and Motwani, R. 2000. Dynamic miss-counting algorithms: Finding implication and similarity rules with confidence pruning. In Proc. of the 16th International Conference on Data Engineering (ICDE'00), San Diego, California, pp. 501–511.
- Mannila, H. and Toivonen, H. 1996. Multiple uses of frequent sets and condensed representations. In Proc. of the Second International Conference on Knowledge Discovery and Data Mining (KDD'96), Portland, Oregon, pp. 189–194.
- Mannila, H. and Toivonen, H. 1997. Levelwise search and borders of theories in knowledge discovery. *Data Mining and Knowledge Discovery*, 1(3):241–258.
- Ng, R., Lakshmanan, L.V., Han, J., and Pang, A. 1998. Exploratory mining and pruning optimization of constrained association rules. In Proc. of the 1998 ACM SIGMOD International Conference on Management of Data (SIGMOD'98), Seattle, Washington, pp. 13–24.
- Pasquier, N., Bastide, Y., Taouil, R., and Lakhal, L. 1999. Efficient mining of association rules using closed itemset lattices. *Information Systems*, 24(1):25–46.
- Pavlov, D., Mannila, H., and Smyth, P. 2000. Probabilistic models for query approximation with large data sets. Technical Report 2000-07, Department of Information and Computer Science, University of California, Irvine, CA-92697-3425.
- Piatetsky-Shapiro, G. 1991. Discovery, analysis, and presentation of strong rules. In *Knowledge Discovery in Databases*, Menlo Park, CA: AAAI Press, pp. 229–248.

Jean-François Boulicaut is currently an associate professor at INSA Lyon (Ph.D. in 1992, Habilitation in 2001). His main research interests are databases and knowledge discovery from databases. He is the leader of the “Data Mining group” at INSA Lyon and coordinates the cInQ IST-2000-26469 project (consortium on knowledge discovery with Inductive Queries, 2001–2004) funded by the European Union.

Artur Bykowski earned his Ph.D. in Computer Science from INSA Lyon in 2002. His Ph.D. thesis was entitled “Condensed representations of frequent sets: application to descriptive pattern discovery”. His research interests include condensed representations of frequent patterns and rule mining from large databases.

Christophe Rigotti received his Ph.D. in Computer Science from INSA Lyon in 1996. His research interests include data mining, constraint programming, inductive programming and databases. He is currently an assistant professor at INSA Lyon and a visiting member of the Ludwig Maximilians University of Munich.