

Database Transposition for Constrained (Closed) Pattern Mining

Baptiste Jeudy, François Rioult

► **To cite this version:**

Baptiste Jeudy, François Rioult. Database Transposition for Constrained (Closed) Pattern Mining. Bart Goethals, Arno Siebes. Knowledge Discovery in Inductive Databases, Third International Workshop, KDID 2004, Pisa, Italy, Septembre 2004, Revised Selected and Invited Papers, Springer, pp.89-107, 2004, LNCS 3377. ujm-00359284

HAL Id: ujm-00359284

<https://hal-ujm.archives-ouvertes.fr/ujm-00359284>

Submitted on 6 Feb 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Database Transposition for Constrained (Closed) Pattern Mining

Baptiste Jeudy¹ and François Rioult²

¹ Équipe Universitaire de Recherche en Informatique de St-Etienne,
Université de St-Etienne, France.

`baptiste.jeudy@univ-st-etienne.fr`

² GREYC - CNRS UMR 6072,

Université de Caen Basse-Normandie, France.

`francois.rioult@info.unicaen.fr`

Abstract. Recently, different works proposed a new way to mine patterns in databases with pathological size. For example, experiments in genome biology usually provide databases with thousands of attributes (genes) but only tens of objects (experiments). In this case, mining the “transposed” database runs through a smaller search space, and the Galois connection allows to infer the closed patterns of the original database. We focus here on constrained pattern mining for those unusual databases and give a theoretical framework for database and constraint transposition. We discuss the properties of constraint transposition and look into classical constraints. We then address the problem of generating the closed patterns of the original database satisfying the constraint, starting from those mined in the “transposed” database. Finally, we show how to generate all the patterns satisfying the constraint from the closed ones.

1 Introduction

Frequent pattern mining is now well mastered, but these patterns, like association rules, reveal to be too numerous for the experts and very expensive to compute. They have to be filtered or constrained. However, mining and constraining have to be done jointly (pushing the constraint) in order to avoid combinatorial explosion [14]. Mining under complex constraint has become today a hot topic and the subject of numerous works (e.g., [14,7,16,20,10,8]). Moreover, new domains are interested in our applications, and data schemes vary consequently. In genome biology, biological experiments are very expensive and time consuming. Therefore, only a small number of these experiments can be processed. However, thanks to new devices (such as biochips), experiments can provide the measurements of the activity of thousands of genes. This leads to databases with lots of columns (the genes) and few rows (the experiments).

Numerous works present efficient algorithms which mine the patterns satisfying a user defined constraint in large databases. This constraint can combine minimum and maximum frequency threshold together with other syntactical

constraints. These algorithms are designed for databases with up to several millions of rows. However, their complexity is exponential in the number of columns and thus they are not suited for databases with too many columns, like those encountered in genome biology.

Recently, two propositions were done to solve this problem: instead of mining the original database, these algorithms work on the “transposed” database, i.e., columns of the original database become rows in the “transposed” database and rows becomes columns (this is indeed the same database but with a different representation). Therefore the “transposed” database has significantly less columns than the original one. The CARPENTER algorithm [18] is specifically designed for mining the frequent closed patterns, and our proposition [23,24] uses a classical algorithm for mining closed patterns with a monotonic (or anti-monotonic) constraint. Both approaches use the transposition principle, however the problem of mining under constraints is not fully studied, specially for complex constraints (i.e., conjunction and disjunction of simple constraints).

In this paper, we study this problem from a theoretical point of view. Our aim is to use classical algorithms (constrained pattern mining algorithms or closed patterns mining algorithms) in the “transposed” database and to use their output to regenerate patterns of the original database instead of directly mining in the original database.

There are several interesting questions which we will therefore try to answer:

1. What kind of information can be gathered in the “transposed” database on the patterns of the original database ?
2. Is it possible to “transpose” the constraints ? I.e., given a database and a constraint, is it possible to find a “transposed” constraint such that mining the “transposed” database with the “transposed” constraint gives information about the patterns which satisfy the original constraint in the original database ?
3. How can we regenerate the closed patterns in the original database from the patterns extracted in the “transposed” database ?
4. How can we generate *all* the itemsets satisfying a constraint using the extracted closed patterns.

These questions will be addressed respectively in Sec. 2, 3, 4 and 5.

The organization of the paper is as follows: we start Sec. 2 by recalling some usual definitions related to pattern mining and Galois connection. Then we show in Sec. 3 how to transpose usual and complex constraints. Section 4 is a complete discussion about mining constrained closed patterns using the “transposed” database and in Sec. 5 we show how to use this to compute all (i.e., not only closed) the patterns satisfying a constraint. Finally Sec. 6 is a short conclusion.

2 Definitions

To avoid confusion between rows (or columns) of the original database and rows (columns) of the “transposed” database, we define a database as a relation be-

tween two sets : a set of attributes and a set of objects. The set of **attributes** (or items) is denoted \mathcal{A} and the set of **objects** is \mathcal{O} . The **attribute space** $2^{\mathcal{A}}$ is the collection of the subsets of \mathcal{A} and the **object space** $2^{\mathcal{O}}$ is the collection of the subsets of \mathcal{O} . An **attribute set** (or **itemset** or attribute pattern) is a subset of \mathcal{A} . An **object set** (or **object pattern**) is a subset of \mathcal{O} . A **database** is a subset of $\mathcal{A} \times \mathcal{O}$.

In this paper we consider that the database has more attributes than objects and that we are interested in mining attributes sets. The database can be represented as an adjacency matrix where objects are rows and attributes are columns (original representation) or where objects are columns and attributes are rows (transposed representation).

Table 1. Original and transposed representations of a database. The attributes are $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$ and the objects are $\mathcal{O} = \{o_1, o_2, o_3\}$. We use a string notation for object sets or itemsets, e.g., $a_1a_3a_4$ denotes the itemset $\{a_1, a_3, a_4\}$ and o_2o_3 denotes the object set $\{o_2, o_3\}$. This dataset is used in all the examples.

object	attribute pattern
o_1	$a_1a_2a_3$
o_2	$a_1a_2a_3$
o_3	$a_2a_3a_4$

attribute	object pattern
a_1	o_1o_2
a_2	$o_1o_2o_3$
a_3	$o_1o_2o_3$
a_4	o_3

2.1 Constraints

Given a database, a **constraint** \mathcal{C} on an attribute set (resp. object set) is a boolean function on $2^{\mathcal{A}}$ (resp. on $2^{\mathcal{O}}$). Many constraints have been used in previous works. One of the most popular is the minimum frequency constraint which requires an itemset to be present in more than a fixed number of objects. But we can also be interested in the opposite, i.e., the maximum frequency constraint. Other constraints are related to Galois connection (see Sect. 2.2), such as closed [2] patterns, free [6], contextual free [7] or key [2] patterns, or even non-derivable [9] or emergent [25,11] patterns. There are also syntactical constraints, when one focuses only on itemsets containing a fixed pattern (superset constraint), contained in a fixed pattern (subset constraint), etc. Finally, when a numerical value (such as a price) is associated to items, aggregate functions such as sum, average, min, max, etc. can be used in constraints [16].

A constraint \mathcal{C} is **anti-monotonic** if $\forall A, B (A \subseteq B \wedge \mathcal{C}(B)) \implies \mathcal{C}(A)$. A constraint \mathcal{C} is **monotonic** if $\forall A, B (A \subseteq B \wedge \mathcal{C}(A)) \implies \mathcal{C}(B)$. In both definitions, A and B can be attribute sets or object sets. The frequency constraint is anti-monotonic, like the subset constraint. The anti-monotonicity property is important, because level-wise mining algorithms most of time use it to prune the search space. Indeed, when a pattern does not satisfy the constraint, its specialization neither and can be pruned [1].

Simple composition of constraints has good properties: the conjunction or the disjunction of two anti-monotonic (resp. monotonic) constraints is anti-monotonic (resp. monotonic). The negation of an anti-monotonic (resp. monotonic) constraints is monotonic (resp. anti-monotonic).

2.2 Galois Connection

The main idea underlying our work is to use the strong connection between the itemset lattice $2^{\mathcal{A}}$ and the object lattice $2^{\mathcal{O}}$ called the **Galois connection**. This connection was first used in pattern mining when closed itemset mining algorithms were proposed [19], while it relates to many works in concept learning [17,27].

Given a database db , the Galois operators f and g are defined as:

- f , called *intension*, is a function from $2^{\mathcal{O}}$ to $2^{\mathcal{A}}$ defined by

$$f(O) = \{a \in \mathcal{A} \mid \forall o \in O, (a, o) \in db\},$$

- g , called *extension*, is a function from $2^{\mathcal{A}}$ to $2^{\mathcal{O}}$ defined by

$$g(A) = \{o \in \mathcal{O} \mid \forall a \in A, (a, o) \in db\}.$$

Given an itemset A , $g(A)$ is also called the *support set* of A in db . It is also the set of objects for which all the attributes of A are true. The **frequency** of A is $|g(A)|$ and is denoted $\mathcal{F}(A)$.

Both functions enable us to link the attribute space to the object space. However, since both spaces have not the same cardinality, there is no one to one mapping between them³. This means that several itemsets can have the same image in the object space and conversely. We thus define two equivalence relations r_a and r_o on $2^{\mathcal{O}}$ and $2^{\mathcal{A}}$:

- if A and B are two itemsets, $A r_a B$ if $g(A) = g(B)$,
- if O and P are two sets of objects, $O r_o P$ if $f(O) = f(P)$.

In every equivalence class, there is a particular element: the largest (for inclusion) element of an equivalence class is unique and is called a **closed attribute set** (for r_a) or a **closed object set** (for r_o).

The Galois operators f and g lead by composition to two **closure** operators, namely $h = f \circ g$ and $h' = g \circ f$. They relate to lattice or hypergraph theory and have good properties [26]. The closed sets are then the fixed points of the closure operators and the closure of a set is the closed set of its equivalence class. In the following we will indifferently refer to h and h' with the notation cl . We denote $\mathcal{C}_{\text{close}}$ the constraint which is satisfied by the itemsets or the object sets which are closed.

If two itemsets are equivalent, their images are equal in the object space. There is therefore no mean to distinguish between them if the mining of the

³ This is fortunate since the whole point of transposition is to explore a smaller space.

closed patterns is performed in the object space. So, by using the Galois connection to perform the search in the object space instead of the attribute space, we will gather information about the equivalence classes of r_a (identified by their closed pattern), not about all individual itemsets. This answers the first question of the introduction, i.e. what kind of information can be gathered in the transposed database on the patterns of the original database. At best, we will only be able to discover closed patterns.

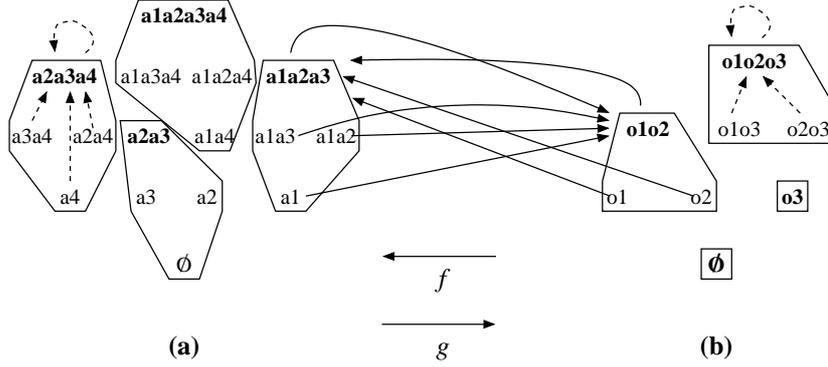


Fig. 1. The equivalence classes for r_a in the itemset lattice (a) and for r_o in the object set lattice (b) built on the database of Tab. 1. The closed sets are in bold face. The arrows represent the f and g operators between the $a_1a_2a_3$ and o_1o_2 equivalence classes. The dotted arrows represent the closure operators h and h'

Property 1. Some properties of f and g .

- f and g are decreasing w.r.t. the inclusion order: if $A \subseteq B$ then $g(B) \subseteq g(A)$ (resp. $f(B) \subseteq f(A)$)
- If A is an itemset and O an object set, then $g(A)$ is a closed object set and $f(O)$ a closed itemset
- fixed point: A is closed if and only if $f(g(A)) = \text{cl}(A) = A$ (resp. $g(f(O)) = \text{cl}(O) = O$)
- $f \circ g \circ f = f$ and $g \circ f \circ g = g$
- $A \subseteq \text{cl}(A)$

In the Galois connection framework, the association of a closed pattern of attributes and the corresponding closed pattern of objects is called a *concept*. Concept learning [17,27] has led to classification tasks and clustering processes. We use this connection in this article through the link it provides between the search spaces $2^{\mathcal{A}}$ and $2^{\mathcal{O}}$.

Example 1. In Fig. 1, the closed objects sets are \emptyset , o_3 , o_1o_2 , and $o_1o_2o_3$. The closed itemsets are a_2a_3 , $a_2a_3a_4$, $a_1a_2a_3$ and $a_1a_2a_3a_4$. Since $g(o_1o_2) = a_1a_2a_3$

and $f(a_1a_2a_3) = o_1o_2$, $(a_1a_2a_3, o_1o_2)$ is a concept. The others are $(a_2a_3, o_1o_2o_3)$, $(a_2a_3a_4, o_3)$, $(a_1a_2a_3a_4, \emptyset)$.

Closed sets of attributes are very useful for algorithms with support constraint, because they share, as maximal element of the equivalence class r_a , the same frequency with all patterns in the class. Closed set mining is now well known [12], and frequent closed patterns are known to be less numerous than frequent patterns [5,9]. Today's approaches relate to closed sets with constraints mining [3]. These patterns are good candidates for constituting relevant concepts, which associate at the same time the attributes and the objects. For example, biologists want to constraint their search to attribute patterns containing some specific genes, with a specified maximum length. They also will be interested in analyzing the other part of the concept. We specifically address here the problem of constrained closed mining in databases with more attributes than objects.

3 Constraint Transposition

Most algorithms extracting closed patterns are search algorithms. The size of the search space strongly determines their performance [12]. In our context, the object space $2^{\mathcal{O}}$ is smaller than the attribute space $2^{\mathcal{A}}$. We therefore choose to search the closed patterns in the smaller space ($2^{\mathcal{O}}$) by transposing the database. In order to mine under constraint, we study in this section how we can adapt constraints to the new transposed database, i.e., how we can transpose constraints. We will therefore answer question 2 of the introduction.

3.1 Definition and Properties

Given an itemset constraint \mathcal{C} , we want to extract the collection I of itemsets, $I = \{A \subseteq \mathcal{A} \mid \mathcal{C}(A)\}$. Therefore, we want to find in the transposed database a collection T of object sets (if it exists) such that the image by f of this collection is I , i.e., $\{f(O) \mid O \in T\} = I$. Since $f(O)$ is always a closed itemset, this is only possible if the collection I contains only closed itemsets (i.e., if the constraint \mathcal{C} includes the $\mathcal{C}_{\text{close}}$ constraint). In this case, a solution for T is the collection $\{O \subseteq \mathcal{O} \mid \mathcal{C}(f(O))\}$ which leads to the following definition of a transposed constraint:

Definition 1 (Transposed constraint). *Given a constraint \mathcal{C} , we define the transposed constraint ${}^t\mathcal{C}$ on a closed pattern O of objects as:*

$${}^t\mathcal{C}(O) = \mathcal{C}(f(O)).$$

Example 2. Consider the itemset constraint $\mathcal{C}(A) = (a_1 \in A)$. Its transposed constraint is (by definition) ${}^t\mathcal{C}(O) = (a_1 \in f(O))$. Using the dataset of Tab. 1, the object sets that satisfy ${}^t\mathcal{C}$ are $T = \{o_1, o_2, o_1o_2, o_1o_3, o_2o_3, o_1o_2o_3\}$. If we compute $\{f(O) \mid O \in T\}$, we get $\{a_1a_2a_3, a_1a_2a_3a_4\}$ which are exactly the closed itemsets that satisfy \mathcal{C} . Theorem 1 will show that this is always the case.

It is interesting to study the effect of transposition w.r.t. the monotonicity or anti-monotonicity of constraints, since many mining algorithms rely on them for efficient pruning:

Proposition 1. *If a constraint \mathcal{C} is monotonic (resp. anti-monotonic), the transposed constraint ${}^t\mathcal{C}$ is anti-monotonic (resp. monotonic).*

Proof: f and g are decreasing (Prop. 1), which inverts monotonicity and anti-monotonicity. \square

Since we also want to deal with complex constraints (i.e., constraints built with elementary constraints using boolean operators), we need the following:

Proposition 2. *If \mathcal{C} and \mathcal{C}' are two constraints then:*

$${}^t(\mathcal{C} \wedge \mathcal{C}') = {}^t\mathcal{C} \wedge {}^t\mathcal{C}'$$

$${}^t(\mathcal{C} \vee \mathcal{C}') = {}^t\mathcal{C} \vee {}^t\mathcal{C}'$$

$${}^t(\neg\mathcal{C}) = \neg{}^t\mathcal{C}$$

Proof: For the conjunction: ${}^t(\mathcal{C} \wedge \mathcal{C}')(O) = (\mathcal{C} \wedge \mathcal{C}')(f(O)) = \mathcal{C}(f(O)) \wedge \mathcal{C}'(f(O)) = ({}^t\mathcal{C} \wedge {}^t\mathcal{C}')(O)$. The proof is similar for the disjunction and the negation. \square

Many algorithms deal with conjunctions of anti-monotonic and monotonic constraints. The two last propositions mean that these algorithms can be used with the transposed constraints since the transposed constraint of the conjunction of a monotonic and an anti-monotonic constraint is the conjunction of a monotonic and an anti-monotonic constraint! The last proposition also helps in building the transposition of a composed constraint. It is useful for the algebraisation [22] of the constraint mining problem, where constraints are decomposed in disjunctions and conjunctions of elementary constraints.

3.2 Transposed Constraints of Some Classical Constraints

In the previous section, we gave the definition of the transposed constraint. In this definition (${}^t\mathcal{C}(O) = \mathcal{C}(f(O))$), in order to test the transposed constraint on an object set O , it is necessary to compute $f(O)$ (to come back in the attribute space) and then to test \mathcal{C} . This means that a mining algorithm using this constraint must maintain a dual context, i.e., it must maintain for each object set O the corresponding attribute set $f(O)$. Some algorithms already do this, for instance algorithms which use the so called *vertical representation* of the database (like CHARM [28]). For some classical constraints however, the transposed constraint can be rewritten in order to avoid the use of $f(O)$. In this section, we review several classical constraints and try to find a simple expression of their transposed constraint in the object space.

Let us first consider the minimum frequency constraint: the transposed constraint of $\mathcal{C}_{\gamma\text{-freq}}(A) = (\mathcal{F}(A) > \gamma)$ is, by definition 1, ${}^t\mathcal{C}_{\gamma\text{-freq}}(O) = (\mathcal{F}(f(O)) > \gamma)$. By definition of frequency, $\mathcal{F}(f(O)) = |g(f(O))| = |\text{cl}(O)|$ and if O is a closed object set, $\text{cl}(O) = O$ and therefore ${}^t\mathcal{C}_{\gamma\text{-freq}}(O) = (|O| > \gamma)$. Finally, the transposed constraint of the minimum frequency constraint is the “minimum size” constraint. The CARPENTER [18] algorithm uses this property and mines the closed patterns in a divide-and-conquer strategy, stopping when the length of the object set drops below the threshold.

The next two propositions give the transposed constraints of two other classical constraints : the subset and superset constraints:

Proposition 3 (subset constraint transposition). *Let $\mathcal{C}_{\subseteq E}$ be the constraint defined by: $\mathcal{C}_{\subseteq E}(A) = (A \subseteq E)$ where E is a constant itemset. Then if E is closed (O is an object set):*

$${}^t\mathcal{C}_{\subseteq E}(O) \Leftrightarrow g(E) \subseteq \text{cl}(O)$$

and if E is not closed

$${}^t\mathcal{C}_{\subseteq E}(O) \Rightarrow g(E) \subseteq \text{cl}(O).$$

Proof: ${}^t\mathcal{C}_{\subseteq E}(O) \Leftrightarrow \mathcal{C}_{\subseteq E}(f(O)) \Leftrightarrow (f(O) \subseteq E) \Rightarrow (g(E) \subseteq g(f(O))) \Leftrightarrow (g(E) \subseteq \text{cl}(O))$. Conversely (if E is closed): $(g(E) \subseteq g(f(O))) \Rightarrow (f(O) \subseteq \text{cl}(E)) \Rightarrow (f(O) \subseteq E)$. \square

Proposition 4 (superset constraint transposition). *Let $\mathcal{C}_{\supseteq E}$ be the constraint defined by: $\mathcal{C}_{\supseteq E}(A) = (A \supseteq E)$ where E is a constant itemset. Then:*

$${}^t\mathcal{C}_{\supseteq E}(O) \Leftrightarrow g(E) \supseteq \text{cl}(O).$$

Proof: ${}^t\mathcal{C}(O) \Leftrightarrow (E \subseteq f(O)) \Rightarrow (g(f(O)) \subseteq g(E)) \Leftrightarrow (\text{cl}(O) \subseteq g(E))$. Conversely, $(g(f(O)) \subseteq g(E)) \Rightarrow (fg(E) \subseteq fgf(O)) \Rightarrow fg(E) \subseteq f(O) \Rightarrow \text{cl}(E) \subseteq f(O) \Rightarrow E \subseteq f(O)$. \square

These two syntactical constraints are interesting because they can be used to construct many other kind of constraints. In fact, all syntactical constraints can be build on top of these using conjunctions, disjunctions and negations. With the proposition 2, it is then possible to compute the transposition of many syntactical constraints. Besides, these constraints have been identified in [13,4] to formalize dataset reduction techniques.

Table 2 gives the transposed constraints of several classical constraints if the object set O is closed (this is not an important restriction since we will use only closed itemsets extraction algorithms). These transposed constraints are easily obtained using the two previous propositions on the superset and the subset constraints and Prop. 2. For instance, if $\mathcal{C}(A) = (A \cap E \neq \emptyset)$, this can be rewritten $A \not\subseteq \overline{E}$ (\overline{E} denotes the complement of E , i.e. $\mathcal{A} \setminus E$) and then $\neg(A \subseteq \overline{E})$. The transposed constraint is therefore, using Prop. 2 and 3, $\neg(g(\overline{E}) \subseteq O)$ (if \overline{E} is closed) and finally $g(\overline{E}) \not\subseteq O$. If \overline{E} is not closed, then we write $E = \{e_1, \dots, e_n\}$ and we rewrite the constraint $\mathcal{C}(A) = (e_1 \in A \vee e_2 \in A \vee \dots \vee e_n \in A)$ and then, using Prop. 2 and 4, we obtain the transposed constraint ${}^t\mathcal{C}(O) = (O \subseteq$

Table 2. Transposed constraints of some classical constraints. A is a variable closed itemset, $E = \{e_1, e_2, \dots, e_n\}$ a constant itemset, O a variable closed object set and $\overline{E} = \mathcal{A} \setminus E = \{f_1, f_2, \dots, f_m\}$

Itemset constraint $\mathcal{C}(A)$	Transposed constraint ${}^t\mathcal{C}(O)$
$\mathcal{F}(A) \theta \alpha$	$ O \theta \alpha$
$A \subseteq E$	if E is closed: $g(E) \subseteq O$ else: $O \not\subseteq g(f_1) \wedge \dots \wedge O \not\subseteq g(f_m)$
$E \subseteq A$	$O \subseteq g(E)$
$A \not\subseteq E$	if E is closed: $g(E) \not\subseteq O$ else: $O \subseteq g(f_1) \vee \dots \vee O \subseteq g(f_m)$
$E \not\subseteq A$	$O \not\subseteq g(E)$
$A \cap E = \emptyset$	if \overline{E} is closed: $g(\overline{E}) \subseteq O$ else: $O \not\subseteq g(e_1) \wedge \dots \wedge O \not\subseteq g(e_n)$
$A \cap E \neq \emptyset$	if \overline{E} is closed: $g(\overline{E}) \not\subseteq O$ else: $O \subseteq g(e_1) \vee \dots \vee O \subseteq g(e_n)$
$\text{SUM}(A) \theta \alpha$	$\mathcal{F}_p(O) \theta \alpha$
$\text{MIN}(A) \theta \alpha$	see text
$\text{MAX}(A) \theta \alpha$	see text

$\theta \in \{<, >, \leq, \geq\}$

$g(e_1) \vee \dots \vee O \subseteq g(e_n)$). These expressions are interesting since they do not involve the computation of $f(O)$. Instead, there are $g(\overline{E})$ or $g(e_i) \dots$. However, since E is constant, these values need to be computed only once (during the first database pass, for instance).

Example 3. We show in this example how to compute the transposed constraints with the database of Tab. 1. Let the itemset constraint $\mathcal{C}(A) = (A \cap a_1 a_4 \neq \emptyset)$. In the database of Tab. 1, the itemset $\overline{a_1 a_4} = a_2 a_3$ is closed. Therefore, the transposed constraint is (Tab. 2) ${}^t\mathcal{C}(O) = (g(a_2 a_3) \not\subseteq O)$. Since $g(a_2 a_3) = o_1 o_2 o_3$, ${}^t\mathcal{C}(O) = (o_1 o_2 o_3 \not\subseteq O)$. The closed object sets that satisfy this constraint are $T = \{\emptyset, o_1 o_2, o_3\}$. If we apply f to go back to the itemset space: $\{f(O) \mid O \in T\} = \{a_1 a_2 a_3 a_4, a_1 a_2 a_3, a_2 a_3 a_4\}$ which are, as expected (and proved by Th. 1), the closed itemset which satisfy \mathcal{C} .

Consider now the constraint $\mathcal{C}(A) = (A \cap a_1 a_2 \neq \emptyset)$. In this case, $\overline{a_1 a_2} = a_3 a_4$ is not closed. Therefore, we use the second expression in Tab. 2 to compute its transposition. ${}^t\mathcal{C}(O) = (O \subseteq g(a_1) \vee O \subseteq g(a_2))$. Since $g(a_1) = o_1 o_2$ and $g(a_2) = o_1 o_2 o_3$, ${}^t\mathcal{C}(O) = (O \subseteq o_1 o_2 \vee O \subseteq o_1 o_2 o_3)$ which can be simplified in ${}^t\mathcal{C}(O) = (O \subseteq o_1 o_2 o_3)$. All the closed object sets satisfy this constraint ${}^t\mathcal{C}$, which is not surprising since all the closed itemsets satisfy \mathcal{C} .

Our last example is the constraint $\mathcal{C}(A) = (|A \cap a_1 a_2 a_4| \geq 2)$. It can be rewritten $\mathcal{C}(A) = ((a_1 a_2 \subseteq A) \vee (a_1 a_4 \subseteq A) \vee (a_2 a_4 \subseteq A))$. Using Prop. 2 and Tab. 2 we get ${}^t\mathcal{C}(O) = ((O \subseteq g(a_1 a_2)) \vee (O \subseteq g(a_1 a_4)) \vee (O \subseteq g(a_2 a_4)))$ which is ${}^t\mathcal{C}(O) = ((O \subseteq o_1 o_2) \vee (O \subseteq \emptyset) \vee (O \subseteq o_3))$. The closed object sets satisfying ${}^t\mathcal{C}$ are $T = \{\emptyset, o_1 o_2, o_3\}$ and $\{f(O) \mid O \in T\} = \{a_1 a_2 a_3 a_4, a_1 a_2 a_3, a_2 a_3 a_4\}$.

Other interesting constraints include aggregate constraints [16]. If a numerical value $a.v$ is associated to each attribute $a \in \mathcal{A}$, we can define constraints of the form $\text{SUM}(A) \theta \alpha$ for several aggregate operators such as SUM, MIN, MAX or AVG, where $\theta \in \{<, >, \leq, \geq\}$ and α is a numerical value. In this case, $\text{SUM}(A)$ denotes the sum of all $a.v$ for all attributes a in A .

The constraints $\text{MIN}(A) \theta \alpha$ and $\text{MAX}(A) \theta \alpha$ are special cases of the constraints of Tab. 2. For instance, if $\text{sup}_\alpha = \{a \in \mathcal{A} \mid a.v > \alpha\}$ then $\text{MIN}(A) > \alpha$ is exactly $A \subseteq \text{sup}_\alpha$ and $\text{MIN}(A) \leq \alpha$ is $A \not\subseteq \text{sup}_\alpha$. The same kind of relation holds for MAX operator: $\text{MAX}(A) > \alpha$ is equivalent to $A \cap \text{sup}_\alpha \neq \emptyset$ and $\text{MAX}(A) \leq \alpha$ is equivalent to $A \cap \text{sup}_\alpha = \emptyset$. In this case, since α is a constant, the set sup_α can be pre-computed.

The constraints $\text{AVG}(A) \theta \alpha$ and $\text{SUM}(A) \theta \alpha$ are more difficult. Indeed, we only found one expression (without $f(O)$) for the transposition of $\text{SUM}(A) \theta \alpha$. Its transposition is ${}^t\mathcal{C}(O) = (\text{SUM}(f(O)) \theta \alpha)$. In the database, $f(O)$ is a set of attributes, so in the transposed database, it is a set of rows and O is a set of columns. The values $a.v$ are attached to rows of the transposed database, and $\text{SUM}(f(O))$ is the sum of these values for the rows containing O . Therefore, $\text{SUM}(f(O))$ is a pondered frequency of O (in the transposed database) where each row a , containing O , contributes for $a.v$ in the total (we denote this pondered frequency by $\mathcal{F}_p(O)$). It is easy to adapt classical algorithms to count this pondered frequency. Its computation is the same as the classical frequency except that each row containing the counted itemset does contribute with a value different from 1 to the frequency.

4 Closed Itemsets Mining

In a previous work [23] we showed the complementarity of the set of concepts mined in the database, with constraining the attribute patterns, and the set of concepts mined in the transposed database with the *negation* of the transposed constraint, when the original constraint is anti-monotonic. The transposed constraint had to be negated in order to ensure the anti-monotonicity of the constraint used by the algorithm. This is important because we can keep usual mining algorithms which deal with anti-monotonic constraint and apply them in the transposed database with the negation of the transposed constraint. We also showed [24] a specific way of mining under monotonic constraint, by simply mining the transposed database with the transposed constraint (which is anti-monotonic). In this section, we generalize these results for more general constraints.

We define the **constrained closed itemset mining problem**: Given a database db and a constraint \mathcal{C} , we want to extract all the closed itemsets (and their frequencies) satisfying the constraint \mathcal{C} in the database db . More formally, we want to compute the collection:

$$\{(A, \mathcal{F}(A, db)) \mid \mathcal{C}(A, db) \wedge \mathcal{C}_{\text{close}}(A, db)\}.$$

The next theorem shows how to compute the above solution set using the closed object patterns extracted in the transposed database, with the help of the transposed constraint.

Theorem 1.

$$\{A \mid \mathcal{C}(A) \wedge \mathcal{C}_{\text{close}}(A)\} = \{f(O) \mid {}^t\mathcal{C}(O) \wedge \mathcal{C}_{\text{close}}(O)\}.$$

Proof: By def. 1, $\{f(O) \mid {}^t\mathcal{C}(O) \wedge \mathcal{C}_{\text{close}}(O)\} = \{f(O) \mid \mathcal{C}(f(O)) \wedge \mathcal{C}_{\text{close}}(O)\} = \{A \mid \exists O \text{ s.t. } \mathcal{C}(A) \wedge A = f(O)\} = \{A \mid \mathcal{C}(A) \wedge \mathcal{C}_{\text{close}}(A)\}.$ \square

This theorem means that if we extract the collection of all closed object patterns satisfying ${}^t\mathcal{C}$ in the transposed database, then we can get all the closed patterns satisfying \mathcal{C} by computing $f(O)$ for all the closed object patterns. The fact that we only need the *closed* object patterns and not all the object patterns is very interesting since the closed patterns are less numerous and can be extracted more efficiently (see CHARM [28], CARPENTER [18], CLOSET[21] or [7]). The strategy, which we propose for computing the solution of the constraint closed itemset mining problem, is therefore:

1. Compute the transposed constraint ${}^t\mathcal{C}$ using Tab. 2 and Prop. 2. This step can involve the computation of some constant object sets $g(E)$ used in the transposed constraint.
2. Use one of the known algorithms to extract the constrained closed sets of the transposed database. Most closed set extraction algorithms do not use constraints (like CLOSE, CLOSET or CARPENTER). However, it is not difficult to integrate them (by adding more pruning steps) for monotonic or anti-monotonic constraints. In [7], another algorithm to extract constrained closed sets is presented.
3. Compute $f(O)$ for each extracted closed object pattern. In fact, every algorithm already computes this when counting the frequency⁴ of O , which is $|f(O)|$. The frequency of $f(O)$ (in the original database) is simply the size of O and can therefore be provided without any access to the database.

The first and third steps can indeed be integrated in the core of the mining algorithm, as it is done in the CARPENTER algorithm (but only with the frequency constraint).

Finally, this strategy shows how to perform constrained closed itemset mining by processing all the computations in the transposed database, and using classical algorithms.

5 Itemsets Mining

In this section, we study how to extract *all* the itemsets that satisfy a user constraint (and not only the closed ones). We define the **constrained itemset**

⁴ This is the frequency in the *transposed database*

mining problem : Given a database db and a constraint \mathcal{C} , we want to extract all the itemsets (and their frequencies) satisfying the constraint \mathcal{C} in the database db . More formally, we want to compute the collection:

$$\{(A, \mathcal{F}(A, db)) \mid \mathcal{C}(A, db)\}.$$

In the previous section, we gave a strategy to compute the *closed* itemsets satisfying a constraint. We will of course make use of this strategy. Solving the constrained itemset mining problem will involve three steps : Given a database db and a constraint \mathcal{C} ,

1. find a constraint \mathcal{C}' ,
2. compute the collection $\{(A, \mathcal{F}(A, db)) \mid \mathcal{C}'(A, db) \wedge \mathcal{C}_{\text{close}}(A, db)\}$ of closed sets satisfying \mathcal{C}' using the strategy of Sec. 4,
3. compute the collection $\{(A, \mathcal{F}(A, db)) \mid \mathcal{C}(A, db)\}$ of all the itemsets satisfying \mathcal{C} from the closed ones satisfying \mathcal{C}' .

We will study the first step in the next subsection and the third one in Sec. 5.2, but first we will show why it is necessary to introduce a new constraint \mathcal{C}' . Indeed, it is not always possible to compute all the itemsets that satisfy \mathcal{C} from the closed sets that satisfy \mathcal{C} . Let us first recall how the third step is done in the classical case where \mathcal{C} is the frequency constraint [19]:

The main used property is that all the itemsets of an equivalence class have the same frequency than the closed itemset of the class. Therefore, if we know the frequency of the closed itemsets, it is possible to deduce the frequency of non-closed itemsets provided we are able to know in which class they belong. The regeneration algorithm of [19] use a top down approach. Starting from the largest frequent closed itemsets, it generates their subsets and assign them their frequencies, until all the itemsets have been generated.

Now, assume that the constraint \mathcal{C} is not the frequency constraint and that we have computed all the closed itemsets (and their frequencies) that satisfy \mathcal{C} . If an itemset satisfies \mathcal{C} , it is possible that its closure does not satisfies it. In this case, it is not possible to compute the frequency of this itemset from the collection of the closed itemsets that satisfy \mathcal{C} (this is illustrated in Fig. 2). Finally, the collection of the closed itemsets satisfying \mathcal{C} is not sufficient to generate the non-closed itemsets. In the next section, we show how the constraint \mathcal{C} can be relaxed to enable the generation all the non-closed itemsets satisfying it.

5.1 Relaxation of the Constraint

In order to be able to generate all the itemsets from the closed ones, it is necessary to have at least the collection of closed itemsets of all the equivalence classes that contain an itemset satisfying the constraint \mathcal{C} . This collection is also the collection of the closures of all itemsets satisfying \mathcal{C} : $\{\text{cl}(A) \mid \mathcal{C}(A, db)\}$.

We must therefore find a constraint \mathcal{C}' such that $\{\text{cl}(A) \mid \mathcal{C}(A, db)\}$ is included in $\{A \mid \mathcal{C}'(A, db) \wedge \mathcal{C}_{\text{close}}(A)\}$. We call such a \mathcal{C}' constraint a **good relaxation** of \mathcal{C} (see Fig. 3). If we have an equality instead of the inclusion, we call \mathcal{C}' an

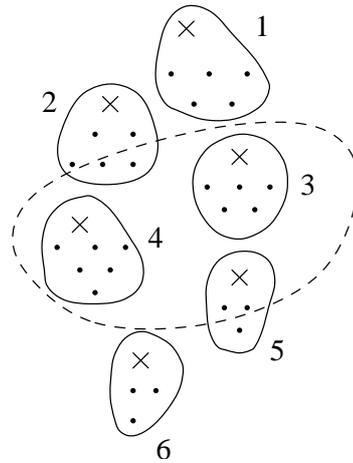


Fig. 2. The dots represent itemsets, the x are closed itemsets, the lines enclose the equivalence classes. The items inside the region delimited by the dashed line satisfy the constraint \mathcal{C} and the others do not. The closed sets satisfying \mathcal{C} are the closed sets of classes 3, 4 and 5. They will enable to generate the itemsets of these three classes. However, to get the two itemsets of class 2, we need the closed itemset of this class which does not satisfy \mathcal{C} . Therefore, in this case, having the closed itemsets satisfying \mathcal{C} is not enough to generate all itemsets satisfying \mathcal{C} .

optimal relaxation of \mathcal{C} . For example, the constant “true” constraint (which is true on all itemset) is a good relaxation of any constraint, however it is not very interesting since it will not provide any pruning opportunity during the extraction of step 2.

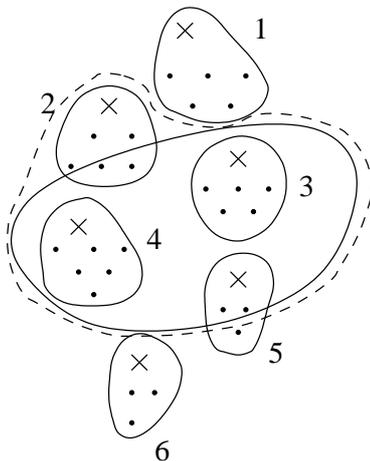


Fig. 3. An optimal relaxation of \mathcal{C} . The constraint \mathcal{C} is represented by the solid line and an optimal relaxation is represented by the dashed line.

If the closed itemsets (and their frequencies) satisfying an optimal relaxation of \mathcal{C} are computed in step 2, we will have enough information for regenerating all itemsets satisfying \mathcal{C} in step 3. However it is not always possible to find such an optimal relaxation. In this case, we can still use a good relaxation in step 2. In this case, some superfluous closed itemsets will be present in the collection and will have to be filtered out in step 3.

We will now give optimal relaxation for some classical constraints, and we start with two trivial cases :

Proposition 5. *The optimal relaxation of a monotonic constraint is the constraint itself and the optimal relaxation of the frequency constraint is the frequency constraint itself.*

Proof: Let \mathcal{C} be a monotonic constraint or a frequency constraint. We only have to prove that if an itemset A satisfy \mathcal{C} then $\text{cl}(A)$ also. If \mathcal{C} is monotonic, this is true since $S \subseteq \text{cl}(S)$ (Prop. 1. If \mathcal{C} is a minimum frequency constraint, it is true because A and $\text{cl}(A)$ have the same frequency. \square

The next proposition is used to compute the relaxation of a complex constraint from the relaxation of simple constraints.

Proposition 6. *If \mathcal{C}_1 and \mathcal{C}_2 are two constraints and \mathcal{C}'_1 and \mathcal{C}'_2 are optimal relaxation of them, then :*

- $\mathcal{C}'_1 \vee \mathcal{C}'_2$ is an optimal relaxation of $\mathcal{C}_1 \vee \mathcal{C}_2$ and
- $\mathcal{C}'_1 \wedge \mathcal{C}'_2$ is a good relaxation of $\mathcal{C}_1 \wedge \mathcal{C}_2$.

Proof: A constraint \mathcal{C}' is a good relaxation of a constraint \mathcal{C} if and only if $\forall A, \mathcal{C}(A) \Rightarrow \mathcal{C}'(\text{cl}(A))$. To prove that it is an *optimal* relaxation, we must also prove that if A is closed and satisfies \mathcal{C}' then there exists an itemset B satisfying \mathcal{C} such that $\text{cl}(B) = A$ (cf. definitions). We will use this two facts in our proofs.

Let A be an itemset satisfying $\mathcal{C}_1 \wedge \mathcal{C}_2$. This means that A satisfies \mathcal{C}_1 and \mathcal{C}_2 . Therefore, $\text{cl}(A)$ satisfies \mathcal{C}'_1 and \mathcal{C}'_2 , i.e., $\text{cl}(A)$ satisfies $\mathcal{C}'_1 \wedge \mathcal{C}'_2$. This means that $\mathcal{C}'_1 \wedge \mathcal{C}'_2$ is a good relaxation of $\mathcal{C}_1 \wedge \mathcal{C}_2$.

We can prove similarly that $\mathcal{C}'_1 \vee \mathcal{C}'_2$ is a good relaxation of $\mathcal{C}_1 \vee \mathcal{C}_2$. Let us now prove that it is optimal: Let A be a closed itemset satisfying $\mathcal{C}'_1 \vee \mathcal{C}'_2$. Then A satisfies \mathcal{C}'_1 or \mathcal{C}'_2 , suppose that it satisfies \mathcal{C}'_1 . Since \mathcal{C}'_1 is an optimal relaxation of \mathcal{C}_1 , there exists B satisfying \mathcal{C}_1 such that $\text{cl}(B) = A$. Therefore B satisfies $\mathcal{C}_1 \vee \mathcal{C}_2$ and $\text{cl}(B) = A$. \square

We found no relaxation for the negation of a constraint but this is not a problem. If the constraint is simple (i.e., in Tab. 2) its negation is also in the table and if it is complex, then we can “push” the negation into the constraint as shown in the next example.

Example 4. Let $\mathcal{C}(A) = (\neg(((\mathcal{F}(A) > 3) \wedge (A \not\subseteq E)) \vee (A \cap F = \emptyset)))$ where E and F are two constant itemsets. We can push the negation and we get: $\mathcal{C}(A) = ((\neg(\mathcal{F}(A) > 3) \vee \neg(A \not\subseteq E)) \wedge \neg(A \cap F = \emptyset))$, and finally :

$$\mathcal{C}(A) = (((\mathcal{F}(A) \leq 3) \vee (A \subseteq E)) \wedge (A \cap F \neq \emptyset)).$$

Then with Prop. 5, 6 and Tab. 3, we can compute a good relaxation \mathcal{C}' of \mathcal{C} :

$$\mathcal{C}'(A) = (((\mathcal{F}(A) \leq 3) \vee (A \subseteq \text{cl}(E))) \wedge (A \cap F \neq \emptyset)).$$

Table 3 gives good relaxation of the other constraints of Tab. 2 which are not covered by the previous proposition (i.e., which are not monotonic) except for the non-monotonic constraints involving SUM for which we did not find any interesting (i.e., other than the constant true constraint) good relaxation.

Proof: We prove here the results given in Tab. 3.

$\mathcal{C}(A) = (A \subseteq E)$, $\mathcal{C}'(A) = (A \subseteq \text{cl}(E))$: If $A \subseteq E$ then $\text{cl}(A) \subseteq \text{cl}(E)$. This means that $\mathcal{C}(A) \Rightarrow \mathcal{C}'(\text{cl}(A))$ therefore \mathcal{C}' is a good relaxation of \mathcal{C} .

$\mathcal{C}(A) = (A \cap E = \emptyset)$: \mathcal{C} can be rewritten $\mathcal{C}(A) = (A \subseteq \overline{E})$ and the previous case applies with \overline{E} instead of E .

Table 3. Good relaxation of some classical constraints. A is a variable closed itemset, $E = \{e_1, e_2, \dots, e_n\}$ a constant itemset.

Itemset constraint $\mathcal{C}(A)$	Good relaxation $\mathcal{C}'(A)$
$A \subseteq E$	$A \subseteq \text{cl}(E)$
$E \not\subseteq A$	$A \subseteq \text{cl}(\overline{e_1}) \vee A \subseteq \text{cl}(\overline{e_2}) \vee \dots \vee A \subseteq \text{cl}(\overline{e_n})$
$A \cap E = \emptyset$	$A \subseteq \text{cl}(\overline{E})$
$\text{MIN}(A) > \alpha$	$A \subseteq \text{cl}(sup_\alpha)$
$\text{MAX}(A) < \alpha$	$A \subseteq \text{cl}(supeq_\alpha)$

$\mathcal{C}(A) = (E \not\subseteq A)$: If $E = \{e_1, e_2, \dots, e_n\}$, this constraint can be rewritten $\{e_1\} \not\subseteq A \vee \{e_2\} \not\subseteq A \vee \dots \vee \{e_n\} \not\subseteq A$ which is also $A \subseteq \overline{\{e_1\}} \vee \dots \vee A \subseteq \overline{\{e_n\}}$. Then the first case and Prop 6 give the result.

$\mathcal{C}(A) = (\text{MIN}(A) > \alpha)$ and $\mathcal{C}(A) = (\text{MAX}(A) < \alpha)$: $\mathcal{C}(A) = (\text{MIN}(A) > \alpha)$ can be rewritten $A \subseteq sup_\alpha$ with $sup_\alpha = \{a \in \mathcal{A} \mid a.v > \alpha\}$ and we are in the first case. $\mathcal{C}(A) = (\text{MAX}(A) < \alpha)$ can be rewritten $A \cap supeq_\alpha = \emptyset$ with $supeq_\alpha = \{a \in \mathcal{A} \mid a.v \geq \alpha\}$ and we are in the second case. \square

5.2 Regeneration

Given a database db and a constraint \mathcal{C} , we suppose in this section that a collection $\{(A, \mathcal{F}(A, db)) \mid \mathcal{C}'(A, db) \wedge \mathcal{C}_{\text{close}}(A, db)\}$ of closed itemsets (and their frequencies) satisfying a good relaxation \mathcal{C}' of \mathcal{C} is available. The aim is to compute the collection $\{(A, \mathcal{F}(A, db)) \mid \mathcal{C}(A, db)\}$ of all itemset satisfying \mathcal{C} (and their frequencies).

If \mathcal{C} is a minimum frequency constraint, \mathcal{C} is an optimal relaxation of itself, therefore we take $\mathcal{C}' = \mathcal{C}$. The regeneration algorithm is then the classical algorithm 6 of [19]. We briefly recall this algorithm:

We suppose that the frequent closed itemsets (and their frequencies) of size i are stored in the list \mathcal{L}_i for $0 < i \leq k$ where k is the size of the longest frequent closed itemset. At the end of the algorithm, each \mathcal{L}_i contains all the frequent itemsets of size i and their frequencies.

```

1 for ( $i = k; i > 0; i --$ )
2   forall  $A \in \mathcal{L}_i$ 
3     forall subset  $B$  of size  $(i - 1)$  of  $A$ 
4       if  $B \notin \mathcal{L}_{i-1}$ 
5          $B.freq = A.freq$ 
6          $\mathcal{L}_{i-1} = \mathcal{L}_{i-1} \cup \{B\}$ 
7       endif
8   end
9 end
10 end

```

If \mathcal{C}' is not the frequency constraint, this algorithm generates all the subsets of the closed itemsets satisfying \mathcal{C}' and two problems arise:

1. Some of these itemsets do not satisfy \mathcal{C} . For instance, in Fig. 3, all the itemsets of classes 2, 3, 4, 5 and 6 are generated (because they are subsets of closed itemsets that satisfy \mathcal{C}') and only those of classes 3 and 4 and some of classes 2 and 5 satisfy \mathcal{C} .
2. The frequency computed in step 5 of the above algorithm for B is correct only if the closure of B is in the collection of the closed sets at the beginning of the algorithm. If it is not, then this computed frequency is smaller than the true frequency of B . In Fig. 3, this means that the computed frequency of the itemsets of class 6 are not correct.

However, the good news is that all the itemsets satisfying \mathcal{C} are generated (because \mathcal{C}' is a good relaxation of \mathcal{C}) and their computed frequencies are correct (because their closures belongs to the \mathcal{L}_i at the beginning).

A last filtering phase is therefore necessary to filter out all the generated itemsets that do not satisfy \mathcal{C} . This phase can be pushed inside the above generation algorithm if the constraint \mathcal{C} has good properties (particularly if it is a conjunction of a monotonic part and an anti-monotonic one). However, we will not detail this point here.

We are still facing a last problem: to test $\mathcal{C}(A)$, we can need $\mathcal{F}(A)$. However, if $\mathcal{C}(A)$ is false, it is possible that the computed frequency of A is not correct. To solve this problem, we propose the following strategy.

We assume that the constraint \mathcal{C} is a Boolean formula built using the atomic constraints listed in Tab. 2 and using the two operators \wedge and \vee (if the \neg operator appears, it can be pushed inside the formula as shown in Ex. 4). Then, we rewrite this constraint in disjunctive normal form (DNF), i.e., $\mathcal{C} = \mathcal{C}_1 \vee \mathcal{C}_2 \vee \dots \vee \mathcal{C}_n$ with $\mathcal{C}_i = \mathcal{A}_{m_{i-1}+1} \wedge \dots \wedge \mathcal{A}_{m_i}$ where each \mathcal{A}_i is a constraint listed in Tab. 2.

Now, consider an itemset A whose computed frequency is f (with $f \leq \mathcal{F}(A)$). First, we consider all the conjunction \mathcal{C}_i that we can compute, this include those where $\mathcal{F}(A)$ does not appear and those of the form $\mathcal{F}(A) > \alpha$ or $\mathcal{F}(A) < \alpha$ where $\alpha < f$ (in this two cases we can conclude since $\mathcal{F}(A) \geq f$). If one of them is true, then $\mathcal{C}(A)$ is true and A is not filtered out.

If all of them are false, we have to consider the remaining conjunctions of the form $\mathcal{A}_1 \wedge \dots \wedge (\mathcal{F}(A) > \alpha) \wedge \dots$ with $\alpha \geq f$. If one of the \mathcal{A}_i is false, then the conjunction is false. If all are true, we suppose that $\mathcal{F}(A) > \alpha$: in this case $\mathcal{C}(A)$ is true and therefore $\mathcal{F}(A) = f$ which contradict $\alpha \geq f$. Therefore, $\mathcal{F}(A) > \alpha$ is false and also the whole conjunction.

If it is still impossible to answer, it means that all the conjunctions are false, and that there are conjunction of the form $\mathcal{A}_1 \wedge \dots \wedge (\mathcal{F}(A) < \alpha) \wedge \dots$ with $\alpha \geq f$. In this case, it is not possible to know if $\mathcal{C}(A)$ is true without computing the frequency $\mathcal{F}(A)$.

Finally, all this means that if there is no constraints of the form $\mathcal{F}(A) < \alpha$ in the DNF of \mathcal{C} , we can do this last filtering phase efficiently. If it appears, then the filtering phase can involve access to the database to compute the frequency

of some itemsets. Of course, all these frequency computation should be made in one access to the database.

Example 5. In this example, we illustrate the complete process of the resolution of the constrained itemset mining problem on two constraints (we still use the dataset of Tab. 1):

$$\mathcal{C}(A) = ((\mathcal{F}(A) > 1) \vee (a_1 \in A)).$$

This constraint is its own optimal relaxation (cf. Prop. 5 and 6). According to Tab. 2 and Prop. 2, its transposed constraint is ${}^t\mathcal{C}(O) = ((|O| > 1) \vee (O \subseteq g(a_1)))$ and $g(a_1) = o_1o_2$. The closed objects sets that satisfy this constraints are $T = \{o_1o_2, o_1o_2o_3, \emptyset\}$. If we apply f to go back to the itemset space: $\{f(O) \mid O \in T\} = \{a_1a_2a_3a_4, a_1a_2a_3, a_2a_3\}$. Since this set contains $a_1a_2a_3a_4$, all the itemsets are generated. However, the generated frequency for the itemsets of the class of $a_2a_3a_4$ is 0. The other generated frequencies are correct. \mathcal{C} is in DNF with two simple constraints $(\mathcal{F}(A) > 1)$ and $(a_1 \in A)$. During the filtering step, when considering the itemsets of $a_2a_3a_4$'s class, the second constraint is always true. Since the generated frequency f is 0 and α is 1, $\alpha > f$ and therefore these itemsets must be filtered out. Finally, the remaining itemsets are exactly those that satisfy \mathcal{C} .

$$\mathcal{C}(A) = ((\mathcal{F}(A) > 1) \wedge (A \subseteq a_2a_4)).$$

A good relaxation of \mathcal{C} is $\mathcal{C}'(A) = ((\mathcal{F}(A) > 1) \wedge (A \subseteq \text{cl}(a_2a_4))) = ((\mathcal{F}(A) > 1) \wedge (A \subseteq a_2a_3a_4))$. The corresponding transposed constraint is ${}^t\mathcal{C}'(O) = ((|O| > 1) \wedge (g(a_2a_3a_4) \subseteq O)) = ((|O| > 1) \wedge (o_3 \subseteq O))$ since $a_2a_3a_4$ is closed. The closed objects sets that satisfy this constraints are $T = \{o_1o_2o_3\}$. If we apply f to go back to the itemset space: $\{f(O) \mid O \in T\} = \{a_2a_3\}$. Then all the subsets of a_2a_3 are generated and only \emptyset and a_2 remains after the filtering step.

6 Conclusion

In order to mine constrained closed patterns in databases with more columns than rows, we proposed a complete framework for the transposition: we gave the expression in the transposed database of the transposition of many classical constraints, and showed how to use existing closed set mining algorithms (with few modifications) to mine in the transposed database.

Then we gave a strategy to use this framework to mine all the itemset satisfying a constraint when a constrained closed itemset mining algorithm is available. This strategy consists of three steps: generation of a relaxation of the constraint, extraction of the closed itemset satisfying the relaxed constraint and, finally, generation of all the itemsets satisfying the original constraint.

We can therefore choose the smallest space between the object space and the attribute space depending on the number of rows/columns in the database. Our strategy gives new opportunities for the optimization of mining queries (also called inductive queries) in contexts having a pathological size. This transposition principle could also be used for the optimization of sequences of queries: the

closed object sets computed in the transposed database during the evaluation of previous queries can be stored in a cache and be re-used to speed up evaluation of new queries in a fashion similar to [15].

References

1. R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A. I. Verkamo. Fast discovery of association rules. In *Advances in Knowledge Discovery and Data Mining*, 1996.
2. Y. Bastide, N. Pasquier, R. Taouil, G. Stumme, and L. Lakhal. Mining minimal non-redundant association rules using frequent closed itemsets. In *Proc. Computational Logic*, volume 1861 of *LNAI*, pages 972–986, 2000.
3. J. Besson, C. Robardet, and J.-F. Boulicaut. Constraint-based mining of formal concepts in transactional data. In *Proc. PAKDD*, 2004. to appear.
4. F. Bonchi, F. Giannotti, A. Mazzanti, and D. Pedreschi. Exante: Anticipated data reduction in constrained pattern mining. In *PKDD'03*, 2003.
5. E. Boros, V. Gurvich, L. Khachiyan, and K. Makino. On the complexity of generating maximal frequent and minimal infrequent sets. In *Symposium on Theoretical Aspects of Computer Science*, pages 133–141, 2002.
6. J.-F. Boulicaut, A. Bykowski, and C. Rigotti. Free-sets : a condensed representation of boolean data for the approximation of frequency queries. *DMKD*, 7(1), 2003.
7. J.-F. Boulicaut and B. Jeudy. Mining free-sets under constraints. In *Proc. IDEAS*, pages 322–329, 2001.
8. C. Bucila, J. Gehrke, D. Kifer, and W. White. Dualminer: a dual-pruning algorithm for itemsets with constraints. In *Proc. SIGKDD*, pages 42–51, 2002.
9. T. Calders and B. Goethals. Mining all non-derivable frequent itemsets. In *Proc. PKDD*, volume 2431 of *LNAI*, pages 74–85, 2002.
10. L. de Raedt and S. Kramer. The levelwise version space algorithm and its application to molecular fragment finding. In *Proc. IJCAI*, pages 853–862, 2001.
11. G. Dong and J. Li. Efficient mining of emerging patterns : discovering trends and differences. In *Proc. SIGKDD*, pages 43–52, 1999.
12. H. Fu and E. M. Nguifo. How well do lattice algorithms on currently used machine learning testbeds ? In *1st Intl. Conf. on Formal Concept Analysis*, 2003.
13. B. Goethals and J. V. den Bussche. On supporting interactive association rule mining. In *DAWAK'00*, 2000.
14. B. Jeudy and J.-F. Boulicaut. Optimization of association rule mining queries. *Intelligent Data Analysis*, 6(4):341–357, 2002.
15. B. Jeudy and J.-F. Boulicaut. Using condensed representations for interactive association rule mining. In *Proc. PKDD*, volume 2431 of *LNAI*, 2002.
16. R. Ng, L. V. Lakshmanan, J. Han, and A. Pang. Exploratory mining and pruning optimizations of constrained associations rules. In *SIGMOD*, 1998.
17. E. M. Nguifo and P. Njiwoua. GLUE: a lattice-based constructive induction system. *Intelligent Data Analysis*, 4(4):1–49, 2000.
18. F. Pan, G. Cong, A. K. H. Tung, J. Yang, and M. J. Zaki. CARPENTER: Finding closed patterns in long biological datasets. In *Proc. SIGKDD*, 2003.
19. N. Pasquier, Y. Bastide, R. Taouil, and L. Lakhal. Efficient mining of association rules using closed itemset lattices. *Information Systems*, 24(1):25–46, Jan. 1999.
20. J. Pei, J. Han, and L. V. S. Lakshmanan. Mining frequent itemsets with convertible constraints. In *Proc. ICDE*, pages 433–442, 2001.

21. J. Pei, J. Han, and R. Mao. CLOSET an efficient algorithm for mining frequent closed itemsets. In *Proc. DMKD workshop*, 2000.
22. L. D. Raedt, M. Jaeger, S. Lee, and H. Mannila. A theory of inductive query answering (extended abstract). In *Proc. ICDM*, pages 123–130, 2002.
23. F. Rioult, J.-F. Boulicaut, B. Crémilleux, and J. Besson. Using transposition for pattern discovery from microarray data. In *DMKD workshop*, 2003.
24. F. Rioult and B. Crémilleux. Optimisation of pattern mining : a new method founded on database transposition. In *EIS'04*, 2004.
25. A. Soulet, B. Crémilleux, and F. Rioult. Condensed representation of emerging patterns. In *Proc. PAKDD*, 2004.
26. B. Stadler and P. Stadler. Basic properties of filter convergence spaces. *J. Chem. Inf. Comput. Sci.*, 42, 2002.
27. R. Wille. Concept lattices and conceptual knowledge systems. In *Computer mathematic applied*, 23(6-9):493-515, 1992.
28. M. J. Zaki and C.-J. Hsiao. CHARM: An efficient algorithm for closed itemset mining. In *Proc. SDM*, 2002.