Efficient Mining Under Rich Constraints Derived from Various Datasets

Arnaud Soulet[1], Jiří Kléma[1,2], and Bruno Crémilleux[1]

[1] GREYC, Université de Caen
Campus Côte de Nacre
F-14032 Caen Cédex France
`{Forename.Surname}@info.unicaen.fr`
[2] Department of Cybernetics
Czech Technical University, Prague
`klema@labe.felk.cvut.cz`

Abstract. Mining patterns under many kinds of constraints is a key point to successfully get new knowledge. In this paper, we propose an efficient new algorithm MUSIC-DFS which soundly and completely mines patterns with various constraints from large data and takes into account external data represented by several heterogeneous datasets. Constraints are freely built of a large set of primitives and enable to link the information scattered in various knowledge sources. Efficiency is achieved thanks to a new closure operator providing an interval pruning strategy applied during the depth-first search of a pattern space. A transcriptomic case study shows the effectiveness and scalability of our approach. It also demonstrates a way to employ background knowledge, such as free texts or gene ontologies, in the discovery of meaningful patterns.

Keywords: constraint-based mining, transcriptomic data.

1 Introduction

In current scientific, industrial or business data mining applications, the critical need is not to generate data, but to derive knowledge from huge and heterogeneous datasets produced at high throughput. In order to explore and discover new highly valuable knowledge it is necessary to develop environments and tools able to put all this data together. This involves different challenges, like designing efficient tools to tackle a large amount of data and the discovery of patterns of a potential user's interest through several datasets. There are various ways to interconnect the heterogeneous data sources and to express the mutual relations among the entities they address. Constraints provide a focus on the most promising knowledge by reducing the number of extracted patterns to those of a potential interest given by the user. Furthermore, when constraints can be pushed deep inside the mining algorithm, performance is improved, making the mining task computationally feasible and resulting in a human-workable output.

This paper addresses the issue of efficient pattern mining from large binary data under flexible constraints derived from additional heterogeneous datasets

synthetizing background knowledge (BK). Large datasets are characterized mainly by a large number of columns (i.e., items). This characteristic often encountered in a lot of domains (e.g., bioinformatics, text mining) represents a remarkable challenge. Usual algorithms show difficulties in running on this kind of data due to the exponential search space growth with the number of items. Known level-wise algorithms commonly fail in mining frequent or constrained patterns in such data [17]. On top of that, the user often would like to integrate BK in the mining process in order to focus on the most plausible patterns consistent with pieces of existing knowledge. BK is available in relational and literature databases, ontological trees and other sources. Nevertheless, mining in a heterogeneous environment allowing a large set of descriptions at various levels of detail is highly non-trivial. This paper solves the problem by pushing user-defined constraints that may stem both from the mined binary data and the BK summarized in similarity matrices or textual files.

The contribution of this paper is twofold. First we provide a new algorithm MUSIC-DFS which soundly and completely mines constrained patterns from large data while taking into account external data (i.e., several heterogeneous datasets). Except for specific constraints for which tricks like the transposition of data [14, 9] or the use of the extension [8] can be used, levelwise approaches cannot tackle large data due to the huge number of candidates. On the contrary, MUSIC-DFS is based on a depth first search strategy. The key idea is to use a new closure operator enabling an efficient interval pruning for various constraints (see Section 3). In [5], the authors also benefit from intervals to prune the search space, but their approach is restricted to the conjunction of one monotone constraint and one anti-monotone constraint. The output of MUSIC-DFS is an interval condensed representation: each pattern satisfying the given constraint appears once in the collection of intervals only. Second, we provide a generic framework to mine patterns with a large set of constraints based on several heterogeneous datasets like texts or similarity matrices. It is a way to take into account the BK. Section 4 depicts a transcriptomic case study. The biological demands require to mine the expression data with constraints concerning complex relations represented by free texts and gene ontologies. The discovered patterns are likely to encompass interesting and interpretable knowledge.

This paper differs from our work in [20] for a double reason. First, the framework is extended to external data. Second, MUSIC-DFS is deeply different from the prototype used in [20]: MUSIC-DFS integrates primitives to tackle external data and thanks to its strategy to prune the search space (new interval pruning based on prefix-free patterns, see Section 3), it is able to mine large data. Section 4 demonstrates the practical effectiveness of MUSIC-DFS in a transcriptomic case study and shows that other prototypes (including the prototype presented in [20]) fail. To the best of our knowledge, there is no other constraint-based tool to efficiently discover patterns from large data under a broad set of constraints linking the information distributed in various knowledge sources.

This paper is organized as follows. Section 2 introduces our framework to mine patterns satisfying constraints defined over several kinds of datasets. In Section 3,

we present the theoretical essentials that underlie the efficiency of Music-dfs and we provide its main features. Experiments showing the efficiency of Music-dfs and the cross-fertilization between several sources of genomic information are given in Section 4.

2   Defining Constraints on Several Datasets

2.1   Integrating Background Knowledge Within Constraints

Usual data-mining tasks rarely deal with a single dataset. Often it is necessary to connect knowledge scattered in several heterogeneous sources. In constraint-based mining, the constraints should effectively link different datasets and knowledge types. In the domain of genomics, there is a natural need to derive constraints both from expression data and descriptions of the genes and/or biological situations under consideration. Such constraints require to tackle various data types - transcriptome data and background knowledge may be stored in the boolean, numeric, symbolic or textual format.

Let us consider the transcriptomic mining context given in Figure 1. Firstly, the involved data include a transcriptome dataset also called internal data. The dataset is in the transactional format - the items correspond to genes and the transactions represent biological situations. The occurrence of an item in a transaction signifies over-expression of the corresponding gene in the corresponding biological situation (genes A, E and F are over-expressed in situation s_1). Secondly, external data – a similarity matrix and textual resources – are considered. They summarize background knowledge that contains various information on items (i.e., genes). This knowledge is transformed into a similarity matrix and a set of texts. Each field of the triangular matrix $s_{ij} \in [0, 1]$ gives a similarity measure between the items i and j. The textual dataset provides a description of genes. Each row of this dataset contains a list of phrases characterizing the given gene (details are given in Section 4.1). The mined patterns are composed of items of the internal data, the corresponding transactions are usually also noted (and possibly analyzed). The external data are used to further specify constraints in order to focus on meaningful patterns. In other words, the constraints may stem from all the datasets.

Table 1 provides the meaning of the primitive constraints applied in this text. The meaning of the primitives is also illustrated by their real values taken from the example in Figure 1. As primitives can address different datasets, the dataset makes another parameter of the primitive (for clarity not shown in Table 1).

A real example of the compound constraint $q(X)$ is given in Figure 1. The first part (a) of q addresses the internal data and means that the biologist is interested in patterns having a satisfactory size – a *minimal area*. Indeed, $area(X) = freq(X) \times length(X)$ is the product of the frequency of X and its length and means that the pattern must cover a minimum number of situations and contain a minimum number of genes. The other parts deal with the external data: (b) is used to discard ribosomal patterns (one gene exception per pattern is allowed), (c) avoids
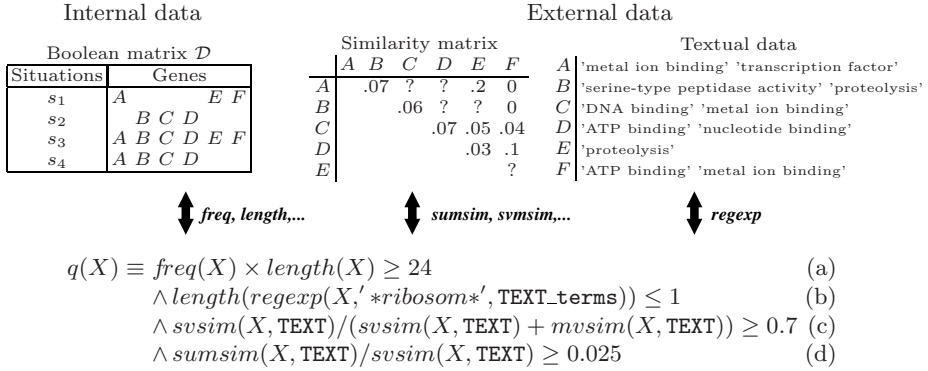
Internal data   External data

Boolean matrix $\mathcal{D}$

Situations	Genes
s_1	A E F
s_2	B C D
s_3	A B C D E F
s_4	A B C D

Similarity matrix

	A	B	C	D	E	F
A	.07	?	?	.2	0	
B		.06	?	?	0	
C			.07	.05	.04	
D				.03	.1	
E					?	

Textual data

A	'metal ion binding' 'transcription factor'
B	'serine-type peptidase activity' 'proteolysis'
C	'DNA binding' 'metal ion binding'
D	'ATP binding' 'nucleotide binding'
E	'proteolysis'
F	'ATP binding' 'metal ion binding'

$\updownarrow$ *freq, length,...* $\updownarrow$ *sumsim, svmsim,...* $\updownarrow$ *regexp*

$$q(X) \equiv freq(X) \times length(X) \geq 24 \qquad \text{(a)}$$
$$\wedge\, length(regexp(X,'*ribosom*', \texttt{TEXT\_terms})) \leq 1 \qquad \text{(b)}$$
$$\wedge\, svsim(X, \texttt{TEXT})/(svsim(X, \texttt{TEXT}) + mvsim(X, \texttt{TEXT})) \geq 0.7 \text{ (c)}$$
$$\wedge\, sumsim(X, \texttt{TEXT})/svsim(X, \texttt{TEXT}) \geq 0.025 \qquad \text{(d)}$$

Fig. 1. Example of a toy (transcriptomic) mining context and a constraint

Table 1. Examples of primitives and their values in the data mining context of Figure 1. Let us note that item pairs of the pattern ABC are (A,B), (A,C) and (B,C).

Primitives		Values
Boolean matrix		
$freq(X)$	frequency of X	$freq(ABC) = 2$
$length(X)$	length of X	$length(ABC) = 3$
Textual data		
$regexp(X, RE)$	items of X whose associated phrases match the regular expression RE	$regexp(ABC,'*ion*')$ $= AC$
Similarity matrix		
$sumsim(X)$	the similarity sum over the set of item pairs of X	$sumsim(ABC) = 0.13$
$svsim(X)$	the number of item pairs in X for which a similarity value is recorded	$svsim(ABC) = 2$
$mvsim(X)$	the number of item pairs in X for which a similarity value is missing	$mvsim(ABC) = 1$
$insim(X, min, max)$	the number of item pairs of X whose similarity lies between min and max	$insim(ABC, 0.07, 1) =$ 1

patterns with prevailing items of an unknown function and (d) is to ensure a minimal average gene similarity. Section 4 provides another constraint q'.

Let us generalize the previous informal description. Let $\mathcal{I}$ be a set of items. A pattern is a non-empty subset of $\mathcal{I}$. $\mathcal{D}$ is a transactional dataset (or boolean matrix) composed of rows usually called transactions. A pattern X is present in $\mathcal{D}$ whenever it is included in one transaction of $\mathcal{D}$ at least. The constraint-based mining task aims to discover all the patterns present in $\mathcal{D}$ and satisfying a constraint q. Unfortunately, real constraints adressing several datasets (the constraint q, for example) are difficult to mine because they have no suitable property as monotonicity [12] or convertibility [16].

2.2   Primitive-Based Constraints

This section presents our framework previously defined in [20] (and the declarative language) enabling the user to set compound and meaningful constraints. This framework naturally integrates primitives adressing external data (e.g., *sumsim* or *regexp*). Furthermore, in our framework constraints are freely built of a large set of primitives. Beyond the primitives mentioned earlier there are primitives such as $\{\wedge, \vee, \neg, <, \leq, \subset, \subseteq, +, -, \times, /, sum, max, min, \cup, \cap, \setminus\}$. The compound constraints of this framework are called *primitive-based constraints.* There are no formal properties required on the final constraints. The only property which is required on the primitives to belong to our framework is a property of monotonicity according to each variable of a primitive (when the others remain constant) [20]. We have already shown that the whole set of primitive-based constraints constitutes a super-class of monotone, anti-monotone, succinct and convertible constraints [19]. Consequently, the proposed framework provides a flexible and rich constraint (query) language. The user can iteratively develop complex constraints integrating various knowledge types.

Let us recall that the primitives and the constraints defined in [20] only address one boolean data set. Current constraints can consider properties taken from a wide scale of dataset types. In addition to the similarity and textual datasets, the framework also enables to access numerical datasets having items in rows and numerical attributes in columns. It implements the primitive X.val which gives the list of values of the attribute named *val* for the items contained in the pattern X.

We give below other examples of constraints belonging to primitive-based constraints and highlighting the generality of our framework:

$$\begin{cases} freq(X) \times length(X) \geq 6 & \text{minimal area (nothing)} \\ (min(X.val) + max(X.val))/2 \leq 50 & \text{maximal mean (loose anti-monotone [2])} \\ sum(X.val)/length(X) \geq 25 & \text{minimal average (convertible [16])} \\ AE \subseteq X & \text{having } AE \text{ (monotone [12])} \\ freq(X) \geq 2 & \text{minimal frequency (anti-monotone [1])} \end{cases}$$

A previous work [21] approximates primitive-based constraints by one anti-monotone and one monotone constraint which can be pushed by DUALMINER [5]. The next section describes an alternative solution in order to benefit from equivalence classes. This way is often more efficient because it avoids the enumeration of all the patterns which compose a particularly huge collection in the case of wide datasets. Besides, in context of wide datasets, previous algorithm MUSIC [20] is ineffective due to the breadth-first search approach (see experiments in Section 4.2). Then, Section 3 presents a new algorithm dedicated to primitive-based constraints in wide datasets.

3   MUSIC-DFS Tool

This section presents the MUSIC-DFS tool (Mining with a User-SpeciﬁIed Constraint, Depth-First Search approach) which benefits from the primitive-based

constraints presented in the previous section. Efficiency is achieved thanks to the exploitation of the primitive and constraint properties. We start by giving the key idea of the safe pruning process based on intervals.

3.1   Main Features of the Interval Pruning

The pruning process performed by MUSIC-DFS is based on the key idea to exploit properties of the monotonicity of the primitives (see Section 2) on the bounds of intervals to prune them. This new kind of pruning is called *interval pruning*. Given two patterns $X \subseteq Y$, the interval $[X, Y]$, also called sub-algebra or sub-lattice, corresponds to the set $\{Z \subseteq \mathcal{I} \mid X \subseteq Z \subseteq Y\}$. Figure 2 depicts an example with the interval $[AB, ABCD]$ and the values of the primitives *sumsim* and *svsim*.

$$AB \;\; 0.07/\underline{1} \longleftarrow sumsim(AB)/svsim(AB)$$
$$ABC \;_{?/?} \quad ABD \;_{?/?}$$
$$ABCD \;_{\underline{0.2}/3}$$

Fig. 2. Illustration of the interval pruning

Assume the constraint $sumsim(X)/svsim(X) \geq 0.25$. As the values associated to the similarities are positive, $sumsim(X)$ is an increasing function according X. Thus $sumsim(ABCD)$ is the highest *sumsim* value for the patterns in $[AB, ABCD]$. Similarly, all the patterns of this interval have a higher $svsim(X)$ value than $svsim(AB)$. Thereby, each pattern in $[AB, ABCD]$ has its average similarity lower or equal than $sumsim(ABCD)/svsim(AB) = 0.2/1$. As this fraction does not exceed 0.25, no pattern of $[AB, ABCD]$ can satisfy the constraint and this interval can be pruned. We say that this pruning is *negative* because no pattern satisfies the constraint. In the same way, if the values of proper combinations of the primitives on the bounds of an interval $[X, Y]$ show that all the patterns in $[X, Y]$ satisfy the constraint, then $[X, Y]$ is also pruned and this pruning is named *positive*. For instance, assuming that $sumsim(AB)/svsim(ABCD) \geq 0.02$, then all the patterns in $[AB, ABCD]$ satisfy the constraint.

In a more formal way, this approach is performed by two interval pruning operators $\lfloor . \rfloor$ and $\lceil . \rceil$ introduced in [20] (but only for primitives handling boolean data). The main idea of these operators is to recursively decompose the constraint to benefit from the monotone properties of the primitives and then to safely negatively or positively prune intervals as depicted above. This process is straightforwardly extended to all the primitives, no matter what kind of dataset they regard. This highlights the generic properties of our framework, as well as the feature of pushing all the parts of the constraint q into the mining step. Table 2 gives the description of the lower and upper bounding operators corresponding to the previous examples of primitives. In Table 2, the general notation

Table 2. The definitions of $\lfloor . \rfloor$ and $\lceil . \rceil$ with particular primitives

$e \in \mathcal{E}_i$	Primitive(s)	$\lfloor e \rfloor \langle X, Y \rangle$	$\lceil e \rceil \langle X, Y \rangle$
$e_1 \theta e_2$	$\theta \in \{\wedge, \vee, +, \times, \cup, \cap\}$	$\lfloor e_1 \rfloor \langle X, Y \rangle \theta \lfloor e_2 \rfloor \langle X, Y \rangle$	$\lceil e_1 \rceil \langle X, Y \rangle \theta \lceil e_2 \rceil \langle X, Y \rangle$
$e_1 \theta e_2$	$\theta \in \{>, \geq, \supset, \supseteq, -, /, \backslash\}$	$\lfloor e_1 \rfloor \langle X, Y \rangle \theta \lceil e_2 \rceil \langle X, Y \rangle$	$\lceil e_1 \rceil \langle X, Y \rangle \theta \lfloor e_2 \rfloor \langle X, Y \rangle$
θe_1	$\theta \in \{\neg, freq, \}$	$\theta \lceil e_1 \rceil \langle X, Y \rangle$	$\theta \lfloor e_1 \rfloor \langle X, Y \rangle$
$\theta(e_1.val)$	$\theta \in \{min\}$	$\theta(\lceil e_1 \rceil \langle X, Y \rangle.val)$	$\theta(\lfloor e_1 \rfloor \langle X, Y \rangle.val)$
$\theta(e_1)$	$\theta \in \{length\}$	$\theta \lfloor e_1 \rfloor \langle X, Y \rangle$	$\theta \lceil e_1 \rceil \langle X, Y \rangle$
$\theta(e_1.val)$	$\theta \in \{sum, max\}$	$\theta(\lfloor e_1 \rfloor \langle X, Y \rangle.val)$	$\theta(\lceil e_1 \rceil \langle X, Y \rangle.val)$
$\theta(e_1)$	$\theta \in \{sumsim, svsim,$ $mvsim\}$	$\theta(\lfloor e_1 \rfloor \langle X, Y \rangle)$	$\theta(\lceil e_1 \rceil \langle X, Y \rangle)$
$\theta(e_1, m, M)$	$\theta \in \{insim\}$	$\theta(\lfloor e_1 \rfloor \langle X, Y \rangle, m, M)$	$\theta(\lceil e_1 \rceil \langle X, Y \rangle, m, M)$
$\theta(e_1, RE)$	$\theta \in \{regexp\}$	$\theta(\lfloor e_1 \rfloor \langle X, Y \rangle, RE)$	$\theta(\lceil e_1 \rceil \langle X, Y \rangle, RE)$
$c \in E_i$	-	c	c
$X \in \mathcal{L}_\mathcal{I}$	-	X	Y

E_i designates one space among $\mathfrak{B}$, $\Re^+$ or $\mathcal{L}_\mathcal{I} = 2^\mathcal{I}$ and $\mathcal{E}_i$ the associated expressions (for instance, the set of constraints $\mathcal{Q}$ for the booleans $\mathfrak{B}$).

The next section indicates how the intervals are built.

3.2 Interval Condensed Representation

As indicated in Section 1, levelwise algorithms are not suitable to mine datasets with a large number of items due to the huge number of candidates growing exponentially according to the number of items. We adopt a depth-first search strategy instead of enumerating the candidate patterns and avoiding subsequent memory failures. We introduce a new and specific closure operator based on a prefix ordering relation $\preceq$. We show that this closure operator is central to the interval condensed representation (Theorem 1) and enables efficient pruning of the search space.

The prefix ordering relation $\preceq$ starts from an arbitrary order over items $A < B < C < \ldots$ as done in [16]. We say that an ordered pattern $X = x_1 x_2 \ldots x_n$ (i.e., $\forall i < j$, we have $x_i < x_j$) is a prefix of an ordered pattern $Y = y_1 y_2 \ldots y_m$ and note $X \preceq Y$ iff we have $n \leq m$ and $\forall i \in \{1, \ldots, n\}$, $x_i = y_i$. For instance, the prefixes of $ABCD$ are the patterns A, AB, ABC and $ABCD$. On the contrary, $AD \not\preceq ADC$ because the ordered form of ADC corresponds to ACD, and AD is not a prefix of ACD.

Definition 1 (Prefix-closure). *The prefix-closure of a pattern X, denoted* **cl**$_\preceq(X)$*, is the pattern* $\{a \in \mathcal{I} | \exists Y \subseteq X$ *such that* $Y \preceq Y \cup \{a\}$ *and* $freq(Ya) = freq(Y)\}$.

The pattern **cl**$_\preceq(X)$ gathers together the items occurring in all the transactions containing $Y \subseteq X$ such that Y is a prefix of $Y \cup \{a\}$. The fixed points of operator **cl**$_\preceq$ are named the *prefix-closed patterns*. Let us illustrate this definition on our running example (cf. Figure 1). The pattern ABC is not a prefix-closed pattern

because ABC is a prefix of $ABCD$ and $freq(ABCD) = freq(ABC)$. On the contrary, $ABCD$ is prefix-closed. We straightforwardly deduce that any pattern and its prefix-closure have the same frequency. For instance, as $\mathbf{cl}_\preceq(ABC) = ABCD$, $freq(ABC) = freq(ABCD) = 2$.

A closure operator is a function satisfying three main properties: extensivity, isotony, and idempotency [22]. Next property shows that $\mathbf{cl}_\preceq$ is a closure operator:

Property 1 (Closure operator). *The prefix-closure operator $\mathbf{cl}_\preceq$ is a closure operator.*

Proof. *Extensivity:* Let X be a pattern and $a \in X$. We have $\{a\} \subseteq X$ and obviously, $a \preceq a$ and $freq(a) = freq(a)$. Then, we obtain that $a \in \mathbf{cl}_\preceq(X)$ and $\mathbf{cl}_\preceq$ is extensive. *Isotony:* Let $X \subseteq Y$ and $a \in \mathbf{cl}_\preceq(X)$. There exists $Z \subseteq X$ such that $Z \preceq Za$ and $freq(Za) = freq(Z)$. As we also have $Z \subseteq Y$ (and $freq(Za) = freq(Z)$), we obtain that $a \in \mathbf{cl}_\preceq(Y)$ and conclude that $\mathbf{cl}_\preceq(X) \subseteq \mathbf{cl}_\preceq(Y)$. *Idempotency:* Let X be a pattern. Let $a \in \mathbf{cl}_\preceq(\mathbf{cl}_\preceq(X))$. There exists $Z \subseteq \mathbf{cl}_\preceq(X)$ such that $freq(Za) = freq(Z)$ with $Z \preceq Za$. As $Z \subseteq \mathbf{cl}_\preceq(X)$, for all $a_i \in Z$, there is $Z_i \subseteq X$ such that $freq(Z_i a_i) = freq(Z_i)$ with $Z_i \preceq Z_i a_i$. We have $\bigcup_i Z_i \preceq \bigcup_i Z_i a$ and $freq(\bigcup_i Z_i) = freq(\bigcup_i Z_i a)$ (because $freq(\bigcup_i Z_i) = freq(Z)$). As the pattern $\bigcup_i Z_i \subseteq X$, a belongs to $\mathbf{cl}_\preceq(X)$ and then, $\mathbf{cl}_\preceq$ is idempotent. $\qquad\square$

Property 1 is important because it enables to infer results requiring the properties of a closure operator. First, this new prefix-closure operator designs *equivalence classes* through the lattice of patterns. More precisely, two patterns X and Y are equivalent iff they have the same prefix-closure (i.e., $\mathbf{cl}_\preceq(X) = \mathbf{cl}_\preceq(Y)$). Of course, as $\mathbf{cl}_\preceq$ is idempotent, the maximal pattern (w.r.t. $\subseteq$) of a given equivalence class of X corresponds to the prefix-closed pattern $\mathbf{cl}_\preceq(X)$. Conversely, we call *prefix-free patterns* the minimal patterns (w.r.t. $\subseteq$) of equivalence classes. Second, closure properties enable to prove that the prefix-freeness is an anti-monotone constraint (see Property 2 in the next section).

Contrary to the equivalence classes defined by the Galois closure [4, 15], equivalence classes provided by $\mathbf{cl}_\preceq$ have a unique prefix-free pattern. This allows to prove that a pattern belongs to one interval only and provides the important result on the interval condensed representation (cf. Theorem 1). This result cannot be achieved without the new closure operator. Lemma 1 indicates that any equivalence class has a unique prefix-free pattern:

Lemma 1 (Prefix-freeness operator). *Let X be a pattern, there exists an unique minimal pattern (w.r.t. $\subseteq$), denoted $\mathbf{fr}_\preceq(X)$, in its equivalence class.*

Proof. Supposing that X and Y are two minimal patterns of the same equivalence class: we have $\mathbf{cl}_\preceq(X) = \mathbf{cl}_\preceq(Y)$. As X and Y are different, there exists $a \in X$ such that $a \notin Y$ and $a \leq min_\leq\{b \in Y\backslash X\}$ (or we swap X and Y). As X is minimal, no pattern $Z \subseteq X \cap Y$ satisfies that $Z \preceq Za$ and $freq(Za) = freq(Z)$. Besides, for all Z such that $Y \cap X \subset Z \subset Y$, we have $Z \not\preceq Za$ because a is smaller than any item of $Y\backslash X$. So, a does not belong to $\mathbf{cl}_\preceq(Y)$ and then, we

obtain that $\mathbf{cl}_{\preceq}(X) \neq \mathbf{cl}_{\preceq}(Y)$. Thus, we conclude that any equivalence class exactly contains one prefix-free pattern.   □

Lemma 1 means that the operator $\mathbf{fr}_{\preceq}$ links a pattern X to the minimal pattern of its equivalence class, i.e. $\mathbf{fr}_{\preceq}(X)$. X is prefix-free iff $\mathbf{fr}_{\preceq}(X) = X$. Any equivalence class corresponds to an interval delimited by one prefix-free pattern and its prefix-closed pattern (i.e., $[\mathbf{fr}_{\preceq}(X), \mathbf{cl}_{\preceq}(X)]$). For example, AB (resp. $ABCD$) is the prefix-free (resp. prefix-closed) pattern of the equivalence class $[AB, ABCD]$.

Now let us show that the whole collection of the intervals formed by all the prefix-free patterns and their prefix-closed patterns provides an *interval condensed representation* where each pattern X is present only once in the set of intervals.

Theorem 1 (Interval condensed representation). *Each pattern X present in the dataset is included in the interval $[\mathbf{fr}_{\preceq}(X), \mathbf{cl}_{\preceq}(X)]$. Besides, the number of these intervals is less than or equal to the number of patterns.*

Proof. Let X be a pattern and $R = \{[\mathbf{fr}_{\preceq}(X), \mathbf{cl}_{\preceq}(X)] | freq(X) \geq 1\}$. Lemma 1 proves that X is exactly contained in $[\mathbf{fr}_{\preceq}(X), \mathbf{cl}_{\preceq}(X)]$. The latter is unique. As X belongs to R by definition, we conclude that R is a representation of any pattern. Now, the extensivity and the idempotency of prefix-closure operator $\mathbf{cl}_{\preceq}$ ensure that $|R| \leq |\{X \subseteq \mathcal{I} \text{ such that } freq(X) \geq 1\}|$. This proves Theorem 1.   □

In the worst case the size of the condensed representation is the number of patterns (each pattern is its own prefix-free and its own prefix-closed pattern). But, in practice, the number of intervals is low compared to the number of patterns (in our running example, only 23 intervals sum up the 63 present patterns).

The condensed representation highlighted by Theorem 1 differs from the condensed representations of frequent patterns based on the Galois closure [4, 15]: in this last case, intervals are described by a free (or key) pattern and its Galois closure and a frequent pattern may appear in several intervals. We claim that the presence of a pattern in a single interval brings meaningful advantages: the mining is more efficient because each pattern is tested at most once. This property improves the synthesis of the output of the mining process and facilitates its analysis by the end-user. The next section shows that by combining this condensed representation and the interval pruning operators, we get an interval condensed representation of primitive-based constrained patterns.

3.3   Mining Primitive-Based Constraints in Large Datasets

When running, MUSIC-DFS enumerates all the intervals sorted in a lexicographic order and checks whether they can be pruned as proposed in Section 3.1. The enumeration benefits from the anti-monotonicity property of the prefix-freeness (cf. Property 2). The memory requirements grow only linearly with the number of items and the number of transactions.

Property 2. *The prefix-freeness is an anti-monotone constraint (w.r.t. $\subseteq$).*

The proof of Property 2 is very similar to those of the usual freeness [4, 15]:

Proof. Let X be a pattern which is not a prefix-free pattern. So, there is $Z \subset X$ such that $\mathbf{cl}_{\preceq}(Z) = \mathbf{cl}_{\preceq}(X)$. Let Y be a pattern with $X \subseteq Y$. First, we observe that $\mathbf{cl}_{\preceq}(Y) = \mathbf{cl}_{\preceq}(X \cup (Y \backslash X))$ and $\mathbf{cl}_{\preceq}(X \cup (Y \backslash X)) = \mathbf{cl}_{\preceq}(\mathbf{cl}_{\preceq}(X) \cup \mathbf{cl}_{\preceq}(Y \backslash X))$ (usual property of closure operators). As $\mathbf{cl}_{\preceq}(Z) = \mathbf{cl}_{\preceq}(X)$, we obtain that $\mathbf{cl}_{\preceq}(\mathbf{cl}_{\preceq}(X) \cup \mathbf{cl}_{\preceq}(Y \backslash X)) = \mathbf{cl}_{\preceq}(\mathbf{cl}_{\preceq}(Z) \cup \mathbf{cl}_{\preceq}(Y \backslash X))$ and then, $\mathbf{cl}_{\preceq}(\mathbf{cl}_{\preceq}(Z) \cup \mathbf{cl}_{\preceq}(Y \backslash X)) = \mathbf{cl}_{\preceq}(Z \cup (Y \backslash X))$. Finally, as Z is a proper subset of X, the pattern $Z \cup (Y \backslash X)$ is a proper subset of Y. Thus, we conclude that Y is not prefix-free. □

In other words, the anti-monotonicity ensures us that once we know that a pattern is not prefix-free, any superset of this pattern is not prefix-free anymore [1, 12]. Algorithms 1 and 2 give the sketch of MUSIC-DFS.

Algorithm 1. GLOBALSCAN

Input: A prefix-pattern X, a primitive based constraint q and a dataset $\mathcal{D}$
Output: Interval condensed representation of constrained patterns having X as prefix
1: **if** $\neg PrefixFree(X)$ **then return** $\emptyset$ // *anti-monotone pruning*
2: **return** LOCALSCAN$([X, \mathbf{cl}_{\preceq}(X)], q, \mathcal{D})$ // *local mining*
$\cup \bigcup\{$GLOBALSCAN$(Xa, q, \mathcal{D}) | a \in \mathcal{I} \wedge a \geq \max_{\leq} X\}$ // *recursive enumeration*

Algorithm 2. LOCALSCAN

Input: An interval $[X, Y]$, a primitive based constraint q and a dataset $\mathcal{D}$
Output: Interval condensed representation of constrained patterns of $[X, Y]$
1: **if** $\lfloor q \rfloor \langle X, Y \rangle$ **then return** $\{[X, Y]\}$ // *positive interval pruning*
2: **if** $\neg \lceil q \rceil \langle X, Y \rangle$ **then return** $\emptyset$ // *negative interval pruning*
3: **if** $q(X)$ **then return** $[X, X] \cup \bigcup\{$LOCALSCAN$([Xa, \mathbf{cl}_{\preceq}(Xa)], q, \mathcal{D}) | a \in Y \backslash X\}$
4: **return** $\bigcup\{$LOCALSCAN$([Xa, \mathbf{cl}_{\preceq}(Xa)], q, \mathcal{D}) | a \in Y \backslash X\}$ // *recursive division*

MUSIC-DFS scans the whole search space by running GLOBALSCAN on each item of $\mathcal{I}$. GLOBALSCAN recursively performs a depth-first search and stops whenever a pattern is not prefix-free (Line 1, GLOBALSCAN). For each prefix-free pattern X, it computes its prefix-closed pattern and builds $[X, \mathbf{cl}_{\preceq}(X)]$ (Line 2, GLOBALSCAN). Then, LOCALSCAN tests this interval by using the operators $\lfloor . \rfloor$ and $\lceil . \rceil$ informally presented in Section 3.1. If the interval pruning can be performed, the interval is selected (positive pruning, Line 1 from LOCALSCAN) or rejected (negative pruning, Line 2 from LOCALSCAN). Otherwise, the interval is explored by recursively dividing it (Line 3 or 4 from LOCALSCAN). The decomposition of the intervals is done so that each pattern is considered only once. The next theorem provides the correctness of MUSIC-DFS:

Theorem 2 (Correctness). MUSIC-DFS *mines soundly and completely all the patterns satisfying q by means of intervals.*

Proof. Property 2 ensures us that MUSIC-DFS enumerates all the interval condensed representation. Thereby, any pattern is considered (Theorem 1) individually or globally with the safe pruning stemmed from to the interval pruning (see Section 3.1).   □

An additional anti-monotone constraint can be pushed in conjunction of prefixfreeness (Line 1, GLOBALSCAN). This constraint (e.g., minimal frequency constraint) optimizes the extraction by reducing more the search space. Such anti-monotone constraint is automatically deduced from the original constraint q in [21].

4   Mining Constrained Patterns from Transcriptomic Data

This section depicts the effectiveness of our approach on a transcriptomic case study. We experimentally show two results. First, the usefulness of the interval pruning strategy of MUSIC-DFS (the other prototypes fail for such large data, cf. Section 4.2). Second, BK enables to automatically focus on the most plausible candidate patterns (cf. Section 4.3). This underlines the need to mine constrained patterns by taking into account external data. If not mentioned otherwise, the experiments are run on the genomic data described in Section 4.1.

4.1   Gene Expression Data and Background Knowledge

In this experiment we deal with the SAGE (Serial Analysis of Gene Expression) [24] human expression data downloaded from the NCBI website (`www.ncbi.nlm.nih.gov`). The final binary dataset contains 11082 genes tested in 207 biological situations, each gene can be either over-expressed in the given situation or not. The biological details regarding gene selection, mapping and binarization can be seen in [10].

BK available in literature databases, biological ontologies and other sources is used to help to focus automatically on the most plausible candidate patterns. We have experimented with the gene ontology (GO) and free-text data. First, the available gene databases were automatically searched and the records for each gene were built (around two thirds of genes have non-empty records, there is no information available for the rest of them). Then, various similarity metrics among the gene records were proposed and calculated. More precisely, the gene records were converted into the vector space model [18]. A single gene corresponds to a single vector, whose components correspond to a frequency of a single term from the vocabulary. The similarity between genes was defined as the cosine of the angle between the corresponding *term-frequency inverse-documentfrequency* (TFIDF) [18] vectors. TFIDF representation statistically measures how important a term is to a gene record. Moreover, the gene records were also simplified to get a condensed textual description. More details on text mining, gene ontologies and similarities are in [10].

4.2   Efficiency of Music-dfs

Dealing with large datasets Let us show the necessity of the depth-first search and usefulness of the interval pruning strategy of Music-dfs. All the experiments were conducted on a 2.2 GHz Xeon processor with 3GB RAM running Linux.

The first experiment highlights the importance of the depth-first search. We consider the constraint addressing patterns having an $area \geq 70$ (the minimal area constraint has been introduced in Section 2) and appearing at least 4 times in the dataset. Music-dfs only spends 7sec to extract 212 constrained patterns. In comparison, for the same binary dataset, the levelwise approach[1] presented in [20] fails after 963sec whenever the dataset contains more than 3500 genes. Indeed, the candidate patterns necessary to build the output do not fit in memory.

Comparison with prototypes coming from the FIMI repository (`fimi.cs.helsinki.fi`) shows that efficient implementations like KDCI [13], LCM (ver. 2) [23], COFI [25] or Borgelt's Apriori [3] fail with this binary dataset to mine frequent patterns occuring at least 4 times. Borgelt's Eclat [3] and Afopt [11] which are depth-first approaches, are able to mine with this frequency constraint. But they require a post-processing step for other constraints than the frequency (e.g., area, similarity-based constraints).

The power of Music-dfs can also be illustrated on any large benchmark dataset (i.e., containing many transactions). Let us consider the `mushroom` dataset taken from FIMI repository . Figure 3 presents the running times for the Music-dfs, Music, Apriori and Eclat algorithms with the constraints $freq(X) \times length(X) \geq \alpha$ (on the left) and $sum(X.val)/length(X) \geq \alpha$ (on the right). The latter is applied on item values (noted *val*) randomly generated within the range $[0, 100]$. An additional minimal frequency constraint $freq(X) \geq 100$ is used in order to make running of Apriori and Eclat feasible.

As Apriori and Eclat do not push the minimal area/average constraints into the mining, they require a post-processing step to select the right patterns
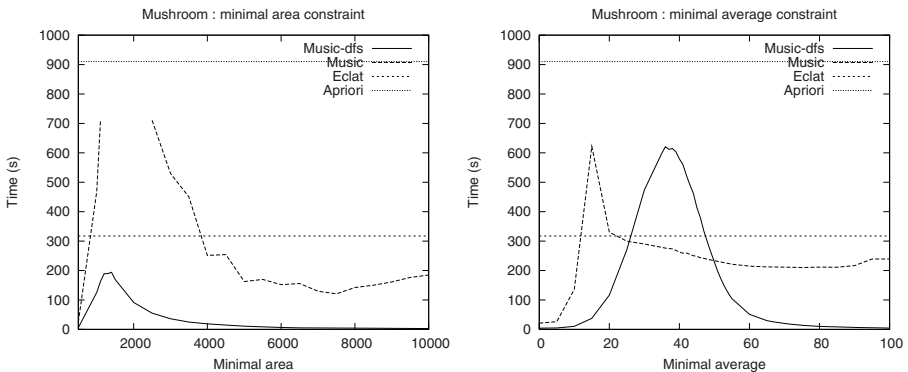

Fig. 3. Runtime performances with minimal area/average constraint on `mushroom`

---

[1] We do not use external data because this version does not deal with external data.

with respect to these constraints. Thus their curves (cf. Figure 3) do not depend on minimal area/average threshold α and are flat. Let us note that we neglect the time of the post-processing step therefore the total time spent by these methods is supposed to be even higher than shown. We observe that MUSIC-DFS clearly outperforms MUSIC and APRIORI. Moreover, MUSIC-DFS is often more efficient than ECLAT as it benefits from the constraint. The experimental study in [19] confirms that MUSIC-DFS is efficient with various constraints and various datasets.

Impact of interval pruning The next experiment shows the great role of the interval pruning strategy. For this purpose, we compare MUSIC-DFS with its modification that does not prune. The modification, denoted MUSIC-DFS-FILTER, mines all the patterns that satisfy the frequency threshold first, the other primitives are applied in the post-processing step. We use two typical constraints needed in the genomic domain and requiring the external data. These constraints and the time comparison between MUSIC-DFS and MUSIC-DFS-FILTER are given in Figure 4. The results show that post-processing is feasible until the frequency threshold generates reasonable pattern sets. For lower frequency thresholds, the number of patterns explodes and large intervals to be pruned appear. The interval pruning strategy decreases runtime and scales up much better than the comparative version without interval pruning and MUSIC-DFS becomes in the order of magnitude faster.


Fig. 4. Efficiency of interval pruning with decreasing frequency threshold. The left image deals with the constraint $freq(X) \geq thres \wedge lenght(X) \geq 4 \wedge sumsim(X)/svsim(X) \geq 0.9 \wedge svsim(X)/(svsim(X) + mvsim(X)) \geq 0.9$. The right image deals with the constraint $freq(X) \geq thres \wedge length(regexp(X,' *ribosom*',$ `GO_terms`$)) = 0$.

4.3   Use of Background Knowledge to Mine Plausible Patterns

This transcriptomic case study demonstrates that constraints coming from the BK can reduce the number of patterns, they can express various kinds of interest and the patterns that tend to reappear are likely to be recognized as interesting

by an expert. One of the goals of any pattern is to generalize the individual gene synexpressions observed in the individual situations. Although it seems that biologists focus on individual biological situations, they follow very similar generalization goals. The most valuable knowledge is extracted from the patterns that concern genes with interesting common features (e.g., process, function, location, disease) whose synexpression is observed in a homogeneous biological context (i.e., in a number of analogous biological situations). An example of this context is the cluster of medulloblastoma SAGE libraries discovered in one of the constrained patterns (see the end of this section). It is obvious that to get such patterns and to pursue the goals mentioned above, a tool dealing with external data is needed.

Let us consider all the patterns having a satisfactory size which is translated by the constraint $area \geq 20^2$. We get nearly half a million different patterns that are joined into 37852 intervals. Although the intervals prove to provide a good condensation, the manual search through this set is obviously infeasible as the interpretation of patterns is not trivial and asks for frequent consultations with medical databases. The biologists prefer sets with tens of patterns/intervals only.

Increasing the threshold of the area constraint to get a reasonable number of patterns is rather counter-productive. The constraint $area \geq 75$ led to a small but uniform set of 56 patterns that was flooded by the ribosomal proteins which generally represent the most frequent genes in the dataset. Biologists rated these patterns as valid but uninteresting.

The most valuable patterns expected by biologists – denoted as meaningful or plausible patterns – have non-trivial size containing genes and situations whose characteristics can be generalized, connected, interpreted and thus transformed into knowledge. To get such patterns, constraints based on the external data have to be added to the minimal area constraint just like in the constraint q given in Section 2. It joins the minimal area constraint with background constraints coming from the NCBI textual resources (gene summaries and adjoined PubMed abstracts). There are 46671 patterns satisfying the minimal area constraint (the part (a) of the constraint q), but only 9 satisfy q. This shows the efficiency of reduction of patterns brought by the BK.

A cross-fertilization with other external data is obviously favourable. So, we use the constraint q' which is similar to q, except that the functional Gene Ontology is used instead of NCBI textual resources and a similarity constraint is added (part (e) of q').

$$
\begin{aligned}
q'(X) \equiv\ & area(X) \geq 24 && \text{(a)}\\
& \wedge\, length(regexp(X,' *ribosom*', \texttt{GO\_terms})) \leq 1 && \text{(b)}\\
& \wedge\, svsim(X, \texttt{GO})/(svsim(X, \texttt{GO}) + mvsim(X, \texttt{GO})) \geq 0.7 && \text{(c)}\\
& \wedge\, sumsim(X, \texttt{GO})/svsim(X, \texttt{GO}) \geq 0.025 && \text{(d)}\\
& \wedge\, insim(X, 0.5, 1, \texttt{GO})/svsim(X, \texttt{GO}) \geq 0.6 && \text{(e)}
\end{aligned}
$$

---

[2] This threshold has been settled by statistical analysis of random datasets having the same properties as the original SAGE data. First spurious patterns start to appear for this threshold area.

Only 2 patterns satisfy q'. A very interesting observation is that the pattern[3] that was identified by the expert as one of the "nuggets" provided by q is also selected by q'. This pattern can be verbally characterized as follows: it consists of 4 genes that are over-expressed in 6 biological situations, it contains at most one ribosomal gene, the genes share a lot of common terms in their descriptions as well as they functionally overlap, at least 3 of the genes are known (have a non-empty record) and all of the biological situations are medulloblastomas which are very aggressive brain tumors in children. The constraints q and q' demonstrate two different ways to reach a compact and meaningful output that can be easily human surveyed.

5   Conclusion

Knowledge discovery from a large binary dataset supported by heterogeneous BK is an important task. We have proposed a generic framework to mine patterns with a large set of constraints linking the information scattered in various knowledge sources. We have presented an efficient new algorithm MUSIC-DFS which soundly and completely mines such constrained patterns. Effectiveness comes from an interval pruning strategy based on prefix free patterns. To the best of our knowledge, there is no other constraint-based tool able to solve such constraint-based tasks.

The transcriptomic case study demonstrates that our approach can handle large datasets. It also shows practical utility of the flexible framework integrating heterogeneous knowledge sources. The language of primitives applied to a wide spectrum of transcriptomic data results in constraints formalizing a viable notion of interestingness.

References

[1] Agrawal, R., Srikant, R.: Fast algorithms for mining association rules. In: Proc. 20th Int. Conf. Very Large Data Bases, VLDB, pp. 432–444 (1994)
[2] Bonchi, F., Lucchese, C.: Pushing tougher constraints in frequent pattern mining. In: Ho et al. [7] pp. 114–124
[3] Borgelt, C.: Efficient implementations of Apriori and Eclat. In: Goethals, Zaki [6]
[4] Boulicaut, J.F., Bykowski, A., Rigotti, C.: Free-sets: a condensed representation of boolean data for the approximation of frequency queries. Data Mining and Knowledge Discovery journal 7(1), 5–22 (2003)

---

[3] The pattern consists of 4 genes KHDRBS1 NONO TOP2B FMR1 over-expressed in 6 biological situations BM_P019 BM_P494 BM_P608 BM_P301 BM_H275 BM_H876. BM stands for brain medulloblastoma.

[5] Bucila, C., Gehrke, J., Kifer, D., White, W.M.: Dualminer: A dual-pruning algorithm for itemsets with constraints. Data Min. Knowl. Discov. 7(3), 241–272 (2003)

[6] Goethals, B., Zaki, M.J. (eds.): FIMI '03, Frequent Itemset Mining Implementations, Proceedings of the ICDM 2003 Workshop on Frequent Itemset Mining Implementations, 19 December 2003, Melbourne, Florida, USA, CEUR Workshop Proceedings, vol. 90 (2003) `CEUR-WS.org`

[7] Ho, T.-B., Cheung, D., Liu, H. (eds.): Advances in Knowledge Discovery and Data Mining, PAKDD 2005. LNCS (LNAI), vol. 3518. Springer, Heidelberg (2005)

[8] Hébert, C., Crémilleux, B.: Mining frequent δ-free patterns in large databases. In: Hoffmann, A., Motoda, H., Scheffer, T. (eds.) DS 2005. LNCS (LNAI), vol. 3735, pp. 124–136. Springer, Heidelberg (2005)

[9] Jeudy, B., Rioult, F.: Database transposition for constrained (closed) pattern mining. In: Goethals, B., Siebes, A. (eds.) KDID 2004. LNCS, vol. 3377, pp. 89–107. Springer, Heidelberg (2005)

[10] Kléma, J., Soulet, A., Crémilleux, B., Blachon, S., Gandrillon, O.: Mining plausible patterns from genomic data. In: Lee, D., Nutter, B., Antani, S., Mitra, S., Archibald, J. (eds.) CBMS 2006, the 19th IEEE International Symposium on Computer-Based Medical Systems, Salt Lake City, Utah, pp. 183–188. IEEE Computer Society Press, Los Alamitos (2006)

[11] Liu, G., Lu, H., Yu, J.X., Wei, W., Xiao, X.: AFOPT: An efficient implementation of pattern growth approach. In: Goethals, Zaki [6]

[12] Mannila, H., Toivonen, H.: Levelwise search and borders of theories in knowledge discovery. Data Mining and Knowledge Discovery 1(3), 241–258 (1997)

[13] Orlando, S., Lucchese, C., Palmerini, P., Perego, R., Silvestri, F.: kDCI: a multistrategy algorithm for mining frequent sets. In: Goethals, Zaki [6]

[14] Pan, F., Cong, G., Tung, A.K.H., Yang, Y., Zaki, M.J.: CARPENTER: finding closed patterns in long biological datasets. In: Proceedings of the 9th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD'03), Washington, DC, USA, pp. 637–642. ACM Press, New York (2003)

[15] Pasquier, N., Bastide, Y., Taouil, T., Lakhal, L.: Discovering frequent closed itemsets for association rules. In: Beeri, C., Bruneman, P. (eds.) ICDT 1999. LNCS, vol. 1540, pp. 398–416. Springer, Heidelberg (1998)

[16] Pei, J., Han, J., Lakshmanan, L.V.S.: Mining frequent item sets with convertible constraints. In: ICDE, pp. 433–442. IEEE Computer Society, Los Alamitos (2001)

[17] Rioult, F., Robardet, C., Blachon, S., Crémilleux, B., Gandrillon, O., Boulicaut, J.-F.: Mining concepts from large sage gene expression matrices. In: Boulicaut, J.-F., Dzeroski, S. (eds.) KDID, pp. 107–118. Rudjer Boskovic Institute, Zagreb, Croatia (2003)

[18] Salton, G., Buckley, C.: Term-weighting approaches in automatic text retrieval. Information Processing Management 24(5), 513–523 (1988)

[19] Soulet, A.: Un cadre générique de découverte de motifs sous contraintes fondées sur des primitives. PhD thesis, Université de Caen Basse-Normandie, France, 2006 (to appear)

[20] Soulet, A., Crémilleux, B.: An efficient framework for mining flexible constraints. In: Ho,, et al. (eds.), [7] pp. 661–671 (2005)

[21] Soulet, A., Crémilleux, B.: Exploiting Virtual Patterns for Automatically Pruning the Search Space. In: Bonchi, F., Boulicaut, J.-F. (eds.) KDID 2005. LNCS, vol. 3933, pp. 98–109. Springer, Heidelberg (2006)

[22] Stadler, B.M.R., Stadler, P.F.: Basic properties of filter convergence spaces (2002)

[23] Uno, T., Kiyomi, M., Arimura, H.: LCM ver. 2: Efficient mining algorithms for frequent/closed/maximal itemsets. In: Bayardo Jr., R.J., Goethals, B., Zaki, M.J. (eds.) FIMI. CEUR Workshop Proceedings, vol. 126 (2004), `CEUR-WS.org`
[24] Velculescu, V., Zhang, L., Vogelstein, B., Kinzler, K.: Serial analysis of gene expression. Science 270, 484–487 (1995)
[25] Zaïane, O.R., El-Hajj, M.: COFI-tree mining: A new approach to pattern growth with reduced candidacy generation. In: Goethals, Zaki [6]