

On When and How to use SAT to Mine Frequent Itemsets

Rui Henriques, Inês Lynce, and Vasco Manquinho

INESC-ID/IST, Technical University of Lisbon, Portugal

`rmch@ist.utl.pt, {ines,vmm}@sat.inesc-id.pt`

Abstract. A new stream of research was born in the last decade with the goal of mining itemsets of interest using Constraint Programming (CP). This has promoted a natural way to combine complex constraints in a highly flexible manner. Although CP state-of-the-art solutions formulate the task using Boolean variables, the few attempts to adopt propositional Satisfiability (SAT) provided an unsatisfactory performance. This work deepens the study on when and how to use SAT for the frequent itemset mining (FIM) problem by defining different encodings with multiple task-driven enumeration options and search strategies. Although for the majority of the scenarios SAT-based solutions appear to be non-competitive with CP peers, results show a variety of interesting cases where SAT encodings are the best option.

1 Introduction

Recent works [34, 24, 17, 9] show the cross-fertilization between *Pattern Mining* (PM) tasks, the discovery of patterns within large datasets, and *Constraint Programming* (CP), the programming paradigm wherein relations between variables are stated declaratively in the form of constraints. Traditional greedy approaches for PM contrast with optimal approaches developed within the artificial intelligence community. While traditional research aims at developing highly optimized and scalable implementations that are tailored towards specific tasks, CP employs a generic and declarative approach to model and mine patterns. This has motivated the adoption of high-level modeling languages or general solvers (that specify *what* the problem is, rather than outlining *how* a solution should be computed) for the flexible definition of constraints, which are critical for many PM applications and domains [22, 24, 21].

The core underlying task for every PM task is to count. Counting is required for every constraint: a specific pattern shape is only of interest above a minimum support threshold. However, here resides the efficiency bottleneck of CP solvers, which need to deal with large counting options to solve frequency-based inequalities. Even though state-of-the-art CP-based solutions are not yet as scalable as traditional PM solutions [34], they can be used for local scans [32], for the expressive definition of user-driven and non-trivial constraints [30], and their optimal search nature have led to significant performance improvements in a wide-diversity of problems [22].

In this work, we propose to study whether a specific class of CP solvers is tailored to solve this core task. The target class of solvers, propositional Satisfiability (SAT) solvers, aim to solve the Boolean SAT problem, which is the problem of finding an assignment for a set of Boolean variables that evaluates a target formula (usually restricted to a conjunctive normal form) to true. Although Boolean encodings are commonly adopted by CP solvers for PM tasks [17], to the best of our knowledge SAT-based solutions have only been proposed to a particular subclass of PM problems aiming at discovering a fixed number of patterns (k -PM) [30, 21].

Based on the critical need for an efficient CP solver for PM tasks, this work undertakes an extensive review to understand how SAT solutions compare to state-of-the-art CP-based alternatives. The target research question is: how SAT behaves in comparison to more general CP frameworks for PM underlying tasks? In section 2 the problem is defined and motivated. Section 3 introduces different SAT encodings with multiple enumeration and search options. Section 4 describes the properties of the adopted implementations and conducts an experimental analysis. The results are discussed and their implications synthesized. Section 5 reviews related research with relevant contributions to the target problem. Finally, concluding remarks and potential prospective research directions are presented.

2 Problem Definition

Definition 1. Let \mathcal{I} be a finite set, called the set of items, let \mathcal{T} be a finite set, called the set of transactions, and let I be an itemset, $I \subseteq \mathcal{I}$. A **transaction** $t \in \mathcal{T}$ over \mathcal{I} is a pair (t_{id}, I) , with t_{id} an identifier and $I \subseteq \mathcal{I}$.

Definition 2. An **itemset database** D over \mathcal{I} is a finite set of transactions. In a simplified way, a transactional dataset is a multi-set of itemsets (being the language of itemsets $L_{\mathcal{I}} = 2^{\mathcal{I}} \setminus \{\emptyset\}$). Equivalently, an itemset database D can be seen as a binary matrix of size $m \times n$, where $m = |\mathcal{T}|$ and $n = |\mathcal{I}|$, with $D_{ti} \in \{0, 1\}$, such that:

$$D = \{(t, I) | t \in \mathcal{T}, I \subseteq \mathcal{I}, \exists_{i \in I} D_{ti} = 1\} \quad (1)$$

Itemset	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	
t_1 A,E,G,K,N	t_1	1	0	0	0	1	0	1	0	0	0	1	0	0	1	0	0
t_2 C,E,H,L,N	t_2	0	0	1	0	1	0	0	1	0	0	0	1	0	1	0	0
t_3 A,D,H,J,O	t_3	1	0	0	1	0	0	0	1	0	1	0	0	0	0	1	0
t_4 B,D,H,J,N	t_4	0	1	0	1	0	0	0	1	0	1	0	0	0	1	0	0
t_5 A,D,H,J,N,P	t_5	1	0	0	1	0	0	0	1	0	1	0	0	0	1	0	1
t_6 A,E,G,K,N,P	t_6	1	0	0	0	1	0	1	0	0	0	1	0	0	1	0	1

Table 1: Illustrative itemset database: compact and Boolean views

A small example of an itemset database is given in Table 1. A traditional example of an itemset database is the supermarket shopping, where each transaction corresponds to a transacted basket and every item to a bought product.

However, common attribute-value tables can be easily converted into an itemset database. Both categorical data (where every attribute-value pair corresponds to an item) and numeric data (following an expressive discretization technique [32]) can be converted, with each row being mapped into a transaction.

Definition 3. The *coverage* $\varphi_{D(I)}$ of an itemset I is the set of all transactions in which the itemset occurs:

$$\varphi_{D(I)} = \{t \in \mathcal{T} \mid \forall_{i \in I} D_{ti} = 1\} \quad (2)$$

Definition 4. The *support* of an itemset I , denoted $sup_D(I)$, is its coverage size: $|\varphi_{D(I)}|$, and the *frequency* of an itemset, denoted $freq_D(I)$, is $sup_D(I)/n$.

Considering the itemset database from Table 1, we have $\varphi_D(\{J, N\}) = \{t_4, t_5\}$, $sup_D(\{J, N\}) = |\{t_4, t_5\}| = 2$ and $freq_D(\{J, N\}) = 0.3(3)$. Now the target PM problem can be formulated.

Definition 5. Given an itemset database D and a minimum support threshold θ , the *frequent itemset mining* problem consists of computing the set:

$$\{I \mid I \subseteq \mathcal{I}, sup_D(I) \geq \theta\} \quad (3)$$

Definition 6. Let a frequent itemset be an itemset with $sup_D(I) \geq \theta$, a *pattern* is a frequent itemset that satisfies any other placed constraints over D .

Considering the database from Table 1 and fixing $\theta=3$, $\{D, H, J\}$ and $\{E, N\}$ are examples of frequent itemsets. Finding frequent itemsets is the core underlying task of every pattern discovery problem and, additionally, form the basis for association-rule analysis, classification, regression, and clustering. FIM was initially proposed in 1993 by Agrawal et al. [1].

2.1 CP Mapping

Flexible constraint-based methods are key for PM as they:

- Focus on *what* the problem is, rather than outlining *how* a solution should be computed, is powerful enough to be used across a wide variety of applications and domains [22, 24, 21] as it suppresses the need of adapting the underlying traditional procedures in order to accommodate new types of constraints.
- Provide an easy method to adapt the search by changing the declarative specification to combine and add new constraints. This not only supports user-driven selection of which patterns are of interest, but FIM-based methods that may require iterative refinements as constraint-driven clustering and pattern-based classification.
- Can expressively capture background knowledge to prune the explosion of spurious and potentially non-interesting patterns [23]. These constraints may include properties from both closed pattern mining [10] and domain-driven pattern mining [3], which aim to incrementally improve results by refining the way patterns and domain-knowledge is represented.

- Support the introduction of a wide variety of expressive constraints, as pattern-set constraints (e.g. global patterns imposing overlapping relaxations over local patterns), with key implications on a wide-variety of problems ranging from web mining to bioinformatics [9, 17].

Although CP models like ConQueSt [10] or MusicDFS [37] already support a predefined number of constraints, they do not allow for the expressive definition and combination of constraints as CP approaches like FIMCP [17], PattCP [21], or GeMini [30]. The selected mapping is based on the constraint encodings of the later approaches. In these CP models for the FIM task, a Boolean variable is used for every individual item I_i and for every transaction T_t . One assignment of values to all I_i and T_t corresponds to one itemset and its corresponding transaction set.

Definition 7. Let T be a transaction set, $T \subseteq \mathcal{T}$. An **itemset** I can be defined by the true item variables: $I_i = 1$ if $i \in I$ and $I_i = 0$ if $i \notin I$. A **transaction set** T can be defined by the set of transactions that are covered by the itemset, $T = \varphi_D(I)$. Thus, $T_t = 1$ if $t \in \varphi_D(I)$.

Corollary 1. The **FIM** task can now be viewed as the computation of valid and frequent (I, T) combinations, i.e. on finding the set:

$$\{(I, T) \mid I \subseteq \mathcal{I}, T \subseteq \mathcal{T}, T = \varphi_D(I), |T| \geq \theta\} \quad (4)$$

We refer to $T = \varphi_D(I)$ as the **coverage** condition while $|T| \geq \theta$ expresses a **support** condition. These conditions restrict the valid variable assignments. Note that given that neither I nor T are fixed, there can be an arbitrary high number of valid attributions to I_i and T_t resulting in different (I, T) tuples that satisfy both constraints.

Property 1 (Coverage Constraint). Given a database D , an itemset I and a transaction set T , then

$$T = \varphi_D(I) \Leftrightarrow (\forall t \in \mathcal{T} : T_t = 1 \Leftrightarrow \sum_{i \in I} I_i (1 - D_{ti}) = 0) \quad (5)$$

where $I_i \in \{0, 1\}$, $T_t \in \{0, 1\}$ and $I_i = 1$ if $i \in I$ and $T_t = 1$ if $t \in T$ [34].

Property 2 (Frequency Constraint). Given a database D , a transaction set T and a threshold θ , then

$$|T| \geq \theta \Leftrightarrow \sum_{t \in \mathcal{T}} T_t \geq \theta \quad (6)$$

where $T_t \in \{0, 1\}$ and $T_t = 1$ if $t \in T$ [34].

We can now model the frequent itemset mining problem as a combination of the coverage constraint and the frequency constraint. To illustrate this, [17] provides an example of a potential implementation in Essence [20] (solver-independent modeling language):

- 1: given Freq : int, TDB : matrix[int(1..NrT),int(1..Nrl)] of int(0..1)
- 2: find **I** : matrix[int(1..Nrl)] of bool, **T** : matrix[int(1..NrT)] of bool
- 3: such that forall t: int(1..NrT)
- 4: **T**[t] $\Leftrightarrow ((\text{sum } i: \text{int}(1..Nrl). (1 - \text{TDB}[t,i]) \mathbf{I}[i])_i = 0)$ (*Coverage Constraint*)
- 5: $(\text{sum } t: \text{int}(1..NrT). \mathbf{T}[t]) \geq \text{Freq}$ (*Frequency Constraint*)

3 Mapping

Since current state-of-the-art CP formulations for FIM rely on Boolean variables, it is important to understand the impact of adopting approaches dedicated to solve Boolean formulae. In particular, efficient and scalable SAT solvers developed over the last decades have contributed to dramatic advances in the ability to solve problem instances involving thousands of variables and millions of clauses [30]. Additionally, a potential mapping of previous constraints into a conjunctive normal formula contains many binary clauses, which can be handled in a very effective manner by SAT solvers. For these reasons, this work encodes FIM as a SAT formula. Enumeration, search, encoding alternatives, and tuning options are covered.

3.1 Core Encoding

The previously introduced coverage and frequency constraints map the FIM problem into a high-level CP language. Next, an extended SAT encoding is proposed. SAT clauses and pseudo-Boolean constraints will be interchangeably adopted to facilitate their traceability.

Corollary 2 (Coverage Encoding). *Given a database D , an itemset I and a transaction set T , the SAT encoding for the coverage constraint is:*

$$\bigwedge_{t \in \mathcal{T}} (\bigwedge_{i \in I|D_{ti}} (\neg T_t \vee \neg I_i) \wedge (T_t \vee (\bigvee_{i \in I|D_{ti}} I_i))) \quad (7)$$

Proof. This formula is derived from equation (5) by:

1. rewriting the coverage sum:

$$\begin{aligned} \sum_{i \in I} I_i (1 - D_{ti}) &= 0 \\ \Leftrightarrow \sum_{i \in I|D_{ti}} I_i &= 0 \\ \Leftrightarrow \bigwedge_{i \in I|D_{ti}} \neg I_i \end{aligned}$$

2. decomposing the equivalence into CNF:

$$\begin{aligned} T_t = 1 &\Leftrightarrow \bigwedge_{i \in I|D_{ti}} \neg I_i \\ \Leftrightarrow (\neg T_t \vee (\bigwedge_{i \in I|D_{ti}} \neg I_i)) &\wedge (T_t \vee \neg(\bigwedge_{i \in I|D_{ti}} \neg I_i)) \\ \Leftrightarrow \bigwedge_{i \in I|D_{ti}} (\neg T_t \vee \neg I_i) &\wedge (T_t \vee (\bigvee_{i \in I|D_{ti}} I_i)) \end{aligned}$$

3. encoding the quantifier $\forall_{t \in \mathcal{T}}^m$ into m sets of clauses:

$$\begin{aligned} \forall_{t \in \mathcal{T}} (T_t = 1 &\Leftrightarrow \bigwedge_{i \in I|D_{ti}} \neg I_i) \\ \Leftrightarrow \bigwedge_{t \in \mathcal{T}} (\bigwedge_{i \in I|D_{ti}} (\neg T_t \vee \neg I_i) &\wedge (T_t \vee (\bigvee_{i \in I|D_{ti}} I_i))) \end{aligned}$$

Complexity. Considering n the size of \mathcal{I} and m the number of transactions in D , we have the following properties: an upper bound of $m \times n$ binary clauses, and m clauses with a maximum of $n + 1$ literals.

This SAT formula guarantees the consistency of T_t and I_i attributions.

To encode the *frequency constraints*, we need to extend the SAT notation to include *pseudo-Boolean (PB) constraints*, which are extensions of SAT clauses that support cardinality constraints and weighted literals. Additionally, we need to adapt equation (6) into a reified frequency constraint for a more focused search of space, as discussed in [22]. This model is equivalent to the original model.

Property 3 (Reified Frequency Constraint). Given a database D , a transaction set T and a threshold θ , then:

$$|T| \geq \theta \Leftrightarrow (\forall_{i \in \mathcal{I}} : I_i = 1 \rightarrow \Sigma_{T_t} D_{ti} \geq \theta) \quad (8)$$

Corollary 3 (Frequency Encoding). Given a database D , a transaction set T and a threshold θ , then the PB encoding for the frequency constraint is:

$$\bigwedge_{i \in \mathcal{I}} (\theta \neg I_i + \Sigma_{t \in T|D_{ti}} T_t \geq \theta) \quad (9)$$

Proof. This formula is derived from equation (8) by:

1. rewriting the frequency inequality into a cardinality constraint:

$$\begin{aligned} & \Sigma_{T_t} D_{ti} \geq \theta \\ \Leftrightarrow & \Sigma_{t \in T|D_{ti}} T_t \geq \theta \end{aligned}$$

2. decomposing the implication:

$$\begin{aligned} & I_i = 1 \Rightarrow \Sigma_{T_t} D_{ti} \geq \theta \\ \Leftrightarrow & I_i = 1 \vee (\Sigma_{t \in T|D_{ti}} T_t \geq \theta) \\ \Leftrightarrow & \neg I_i \vee (\wedge_{t \in T|D_{ti}} T_t \geq \theta) \end{aligned}$$

3. mapping the previous formula into conjunctive normal form or, as below, into pseudo-Boolean constraints:

$$\begin{aligned} & \neg I_i \vee (\Sigma_{t \in T|D_{ti}} T_t \geq \theta) \\ \Leftrightarrow & \neg I_i \vee (\frac{1}{\theta} \Sigma_{t \in T|D_{ti}} T_t \geq 1) \\ \Leftrightarrow & \theta \neg I_i + \Sigma_{t \in T|D_{ti}} T_t \geq \theta \end{aligned}$$

Potential SAT encodings for the $\neg I_i \vee (\Sigma_{t \in T|D_{ti}} T_t \geq \theta)$ constraint include the use of sequential counters [36], binary decision diagrams [19], sorting networks, or cardinality networks [4]. Note, however, that some of these encodings are not polynomial. Since solving these constraints is the core task of our problem, we benefit from solvers oriented to solve them. Moreover, a possible encoding using cardinality constraints would suffer from an additional complexity, since we would need to translate: $\neg I_i \vee \text{Card}_{Constraint}$. For these reasons, a pseudo-Boolean representation is the natural choice;

4. translating the quantifier $\forall_{i \in \mathcal{I}}$ into n sets of clauses:

$$\begin{aligned} & \forall_{i \in \mathcal{I}} (I_i = 1 \Rightarrow \Sigma_{T_t} D_{ti} \geq \theta) \\ \Leftrightarrow & \bigwedge_{i \in \mathcal{I}} (\theta \neg I_i + \Sigma_{t \in T|D_{ti}} T_t \geq \theta) \end{aligned}$$

Complexity. The incremental complexity added by this formula is n pseudo-Boolean constraints (of the form \geq) with a maximum of m unweighted literals and one weighted literal.

Concluding, equations (7) and (9) describe the resulting SAT encoding, which has: *i*) $m + n$ variables; *ii*) $\Theta(mn)$ clauses with mn being binary clauses and m clauses having $\Theta(n)$ variables; and *iii*) n pseudo-Boolean constraints with $\Theta(m)$ variables.

3.2 Enumeration Options

Using SAT solvers on the previously defined encoding would output either one frequent itemset (I_i literals assigned as true) or **unsat**, meaning that no itemset satisfies the constraints. Understandably, model enumeration needs to be present in order to solve the FIM problem, that is, to find the set of all frequent itemsets.

Note that an alternative strategy to define all frequent itemsets at the encoding level would mean an exponential growth of the search space, potentially behaving similarly to a simple enumeration strategy. However, since we can easily opt to adopt more expressive ways of performing enumerations, with significant space cuttings on each iteration, we only target enumeration strategies.

For this purpose, we need to introduce two properties of frequency that allow for pruning substantial parts of the search space, and dual formulation for FIM.

Definition 8. *Let I be a set of items. A transaction (t_{id}, J) contains I , denoted $I \subseteq (t_{id}, J)$, if $I \subseteq J$. FIM approaches rely heavily on:*

- **monotonicity** of frequency: if $I \subseteq J$, then the frequency of J is bounded from above by the frequency of I ;
- **anti-monotonicity** of frequency: if $I \subseteq J$ and I is not frequent, then J is also not frequent.

Definition 9. *Given an itemset database D and a minimum support threshold θ , the **dual-FIM** task centered on non-frequent itemsets is to compute:*

$$2^{\mathcal{I}} / \{I \mid I \subseteq \mathcal{I}, \text{sup}_D(I) < \theta\} \quad (10)$$

Table 2 rely on these notions to define different enumeration options. Explicit and compact negations based on the (anti-)monotonic property are proposed as well as further directions relying on alternative frequency-based properties.

3.3 Search Options

Previous enumeration strategies may not result in significant improvements if SAT iterations do not put any guarantee on the granularity of the itemsets found. For instance, if the subsets negation strategy is adopted and if the solver tends to output finer itemsets, the adoption of this strategy will not be relevant. The same is valid for the supersets negation strategy. To describe these search options, a new concept needs to be introduced.

Definition 10. *A **maximal frequent itemset** is a frequent itemset that also satisfies:*

$$\forall I' \supset I : |\varphi(I')| < \theta \quad (11)$$

All itemsets that are a superset of a maximal itemset are infrequent, while all itemsets that are subsets are frequent. Maximal frequent itemsets are the top border between itemsets that are frequent and not frequent.

In an itemset database where $ABCD$ is the only maximal frequent itemset, a SAT solution using subsets negation may either find all frequent itemsets within 1 iteration or across $\sum_{i=1}^{n=4} C_i^n = 15$ iterations. The following three search options guarantee an upper bound on the number of performed iterations.

Option	Description	Example $T=\{A, B, C, D\}$
Simple	<p>Observation. Adding a clause with the negation of the solution on every iteration hampers the search, as it may result in redundant searches with multiple assignments of transaction-based literals T_t for the same itemset.</p> <p>Method. The added clause must only include the negation of item-based literals I_i (i.e. exclude all transaction-based literals T_t).</p>	$sup(\{A, C\}) \geq \theta$ $(\neg A \vee B \vee \neg C \vee D)$
Subsets Negation	<p>Observation 1. If an itemset is frequent, then its subsets are also frequent (monotonicity).</p> <p>Method 1. Add the negation of the found itemset as well as the negation of its subsets, so the number of SAT iterations can be largely reduced.</p> <p>Observation 2. Negating every frequent itemset may result in an impractical growth of the number of clauses during the search. For a very simple itemset database with $mn = 100$, the number of added clauses can reach a thousand (significantly higher than the initial number of clauses used to encode the problem).</p> <p>Method 2. Compact the set of clauses obtained within one iteration in only one clause. That is, in next run the solver must be able to select, at least, one item that is not included in the previously found frequent itemsets:</p> $I \geq \theta \Rightarrow \bigvee_{i I_i} I_i$	$sup(\{A, C\}) \geq \theta$ $(\neg A \vee B \vee \neg C \vee D)$ $\wedge (\neg A \vee B \vee \neg C \vee D)$ $\wedge (A \vee B \vee \neg C \vee D)$ $sup(\{A, C\}) \geq \theta$ $(B \vee D)$
Supersets Negation	<p>Observations. For fixed I_i literal attributions satisfying the coverage constraints, if no combination of T_t satisfies the frequency constraints, new clauses can be directly learned corresponding to the supersets of the found non-frequent itemset. This property was found to be critical for a dual-FIM formulation.</p> <p>Methods. Adopt dual-FIM problem for medium-to-low frequency thresholds, and negate supersets negation in a similar fashion as subsets negation (items of non-frequent itemsets cannot jointly appear):</p> $I < \theta \Rightarrow \bigvee_{i I_i} \neg I_i$ <p>Note that the choice of when to adopt the FIM-dual should be dynamically made based on the dataset properties (mainly density, but also transactions-to-items ratio) and on the inputted frequency.</p>	$sup(\{A, C\}) < \theta$ $(\neg A \vee B \vee \neg C \vee D)$ $(\neg A \vee \neg B \vee \neg C \vee D)$ $(\neg A \vee B \vee \neg C \vee \neg D)$ $(\neg A \vee \neg B \vee \neg C \vee \neg D)$ $sup(\{A, C\}) < \theta$ $(\neg B \vee \neg D)$
Others	<p>Pointer 1. Adopt advanced enumeration strategy centered on implicit methods as, for instance, cube representations [29].</p> <p>Pointer 2. Use the monotonicity and anti-monotonicity properties to affect the exploitation of the structure of conflicts within a SAT solver.</p> <p>Pointer 3. Exploit important relationships between itemset frequencies beside monotonicity. These properties should not only affect the iterative insertions, but can be included as constraints in the initial encoding. This may result in significant improvements as the previous strategies. For example, in the MAXMINER algorithm [7], relations of the following form are exploited: $freq(\{a, b, c\}) = freq(\{a, b\}) + freq(\{a, c\}) - freq(\{a\})$. There are many more relations between the frequencies of itemsets. See [13] for extensions based on the inclusion-exclusion principle. For a generalization to other measures besides frequency, see [35].</p>	—

Table 2: Enumeration Options

Largest-to-Shortest Maximal (LSM) search

This search option guarantees that the number of iterations equals the number of maximal frequent itemsets by mapping the previous *decision* problem into an *optimization* problem. Each iteration returns a maximal frequent itemset, starting with the longest maximal frequent itemset until reaching the shortest maximal frequent itemset. This is done by defining the following goal function:

$$min : \sum_i^n \neg I_i, \quad (12)$$

and, understandably, by adopting one of the expressive enumeration strategies based on the (anti-)monotonicity property. Otherwise, the goal of finding maximal frequent itemsets would be no longer a valid effort.

Constrained Monotonic Growing (CMG) search

One of the undesirable consequences of the LSM search method is the need to fully exploit the search space within each iteration. That is, for an optimum value for the goal function, several (potentially maximal) frequent itemsets need to be incrementally compared. Additionally, most of the learned behavior needs to be restarted across iterations. This is particularly critical for itemset databases with either many maximal frequent itemsets (usually the case) or with multiple fine maximal frequent itemsets.

In order to overcome the referred problems, we propose CMG, a more relaxed search option formulated over decision tasks that do not require the larger maximal frequent itemsets to be found early. The CMG is based on two types of searches: α -search and β -search. An α -search is a simple SAT iteration (outputting one frequent itemset). As a result, not only a clause expressively negating its subsets is added, but also the itemset itself:

$$\text{temporaryClause} \leftarrow \bigvee_{i \in I_i} \{I_i\}$$

After an α -search, a set of β -searches are performed with the goal of finding larger frequent itemsets until a maximal frequent itemset is found (with `unsat` being returned). When this happens, the clauses related to the previous itemset's items are removed, the found frequent itemset is expressively negated and a new α -search is performed. This behavior is repeated until the α -search returns `unsat`. An illustrative instantiation segment of CMG behavior for an itemset database with $\mathcal{I} = \{A, B, C, D, E\}$ is presented next:

```
{A}  $\dashv$   $\alpha$ -search()
temporaryClause  $\leftarrow$  (B  $\vee$  C  $\vee$  D  $\vee$  E)
{A, B, E}  $\dashv$   $\beta$ -search({A, ?})
temporaryClause  $\leftarrow$  (C  $\vee$  D)
unsat  $\dashv$   $\beta$ -search({A, B, E, ?})
learnedClauses  $\leftarrow$  learnedClauses  $\wedge$  temporaryClause
temporaryClause  $\leftarrow$   $\emptyset$ 
{C, E}  $\dashv$   $\alpha$ -search()
```

Length Decreasing (LD) search

A third search option, LD search, benefits from a more focused search of space as it fixes the length of the itemsets to be found. To guarantee that only maximal frequent itemsets are selected, LD initially fixes this length to n and iteratively decrements it. Alternative length settings are possible if a separate initial scanning to the itemset database guarantees upper and lower bound restrictions on the length of maximal frequent itemsets.

LD accomplishes this behavior by adding and removing equalities of the form $\sum_i I_i = k$, with $k \in \{1, \dots, n\}$. However, since only few solvers support the addition-removal of pseudo-Boolean constraints, a new set of variables $A =$

$\{A_1, \dots, A_n\}$ is added into the following additional constraint:

$$\sum_i^n I_i + \sum_i^n A_i = k \quad (13)$$

In the simplest mode, all A_i variables are initially assigned to **false** until no maximal frequent itemset with length n is found (**unsat** output). Incrementally each A_i is reversed to **true**, so finer maximal frequent itemsets can be found. Understandably, either the subsets negation enumeration option or the encoding of equation (11) needs to be in place for an efficient search.

3.4 Mapping Restrictions

The need to test the proposed options against different pseudo-Boolean solvers may require adaptations over the initial encoding. For simplicity, this section only describes how to support the removal of clauses during enumeration. In appendix, additional adaptations to deal with non-negated variables are covered.

Encoding	Description
	<p>Method. Recur to: <i>i</i>) n additional variables ($N = \{N_1, \dots, N_n\}$), to <i>ii</i>) n additional clauses:</p> $\bigwedge_{i \in I} (I_i \vee N_i) \quad (14)$
<i>Incremental Clauses</i>	<p>and to <i>iii</i>) manipulations over the vector of assumptions that SAT solvers usually disclose. In α-searches the N Boolean vector of assumptions is set to true, so these new n clauses are directly satisfied. In β-searches, the index i of the I_i items belonging to the target itemset are fixed, and the respective N_i variables are set to false. Recurring to N_i assumptions, the SAT solver behavior is similar to solvers that allow for clause removal and the level of performance is closely maintained. This is referred as incremental clauses encoding, because although no clauses need to be deleted, still new clauses need to be added.</p>
<i>Fixed Clauses</i>	<p>Observation. As the number of inserted clauses grows significantly with the number of iterations, an encoding with an increasing number of clauses penalizes the performance. Method. The strategy non-prone to insertions requires part of the reasoning to be done outside of the solver. The challenge is that all new clauses are relevant, and, thus, need to be maintained in memory. For instance, in a $n=5$ itemset database, if the solver finds in initial iterations $\{I_1 I_2 I_3\}$ and $\{I_1 I_3 I_5\}$ (i.e. learned clauses are, respectively, $I_4 \vee I_5$ and $I_2 \vee I_4$), the new solutions need to satisfy all the learned clauses. In this strategy, the storage and reasoning is done separately to affect the values of the N Boolean vector. Following the introduced example, either N_4 or N_2 and N_5 will iteratively assume the value false. So the binary clauses to be satisfied, described in the previous option, will guarantee that the respective items will be true. This method of defining assumptions, needs however to be complemented with a control variable y. This control variable is required for CMG search option to distinguish between the α- and β-searches. Additionally, the following constraint needs to be satisfied:</p> $y + \sum I_i + \sum N_i \geq n + 1, \quad (15)$ <p>to guarantee that in β-searches (y assumed to be false) an additional item is selected.</p>

Table 3: Encoding Options for Clause Removals

The majority of available pseudo-Boolean solvers do not support the dynamic insertion or removal of clauses. Since the insertion of clauses is critical, when a solver of interest does not allow for insertion, its implementation needs to be adapted (usually by turning visible invocations to SAT solver methods at the level of the pseudo-Boolean solver interface). Although the removal of clauses

is required within CMG and LD search options, it is rarely allowed and even non-easily disclosed recurring to the interface of SAT solvers. Two encoding adaptations to deal with removals are depicted in Table 3.

3.5 Tuning Options

Although multiple options for an efficient FIM were introduced, additional improvements can be performed either by adapting the initial encoding or the solver behavior.

Constraints Reduction

From the multiple encoding adaptations that were studied, only one resulted in a significant performance improvement (for medium-to-high support thresholds). This adaptation is centered on a FIM-property that can lead to the reduction of the initial $\Theta(mn)$ constraints or binary clauses (if the solver is able to clausify all of these constraints) into only $\Theta(m + n)$ constraints:

$$\begin{aligned} & \wedge_{t \in T} (\neg T_t \vee (\wedge_{i \in I|D_{ti}} \neg I_i)) \\ \leftrightarrow & \wedge_{t \in T} (\wedge_{i \in I|D_{ti}} \overline{(\neg T_t \vee \neg I_i)}) \\ \leftrightarrow & \wedge_{t \in T} (|D_{ti}| \neg T_t + \sum_{i \in I|D_{ti}} \neg I_i \geq 0) \end{aligned}$$

Polarity Suggestions and Parameters

A simple and effective way to adapt the solver behavior is to change the polarity suggestions. Since we are interested in the early finding of larger itemsets, polarity suggestions for I_i variables should be set to positive (this only degrades performance when maximal frequent itemsets have a very fine length). Orthogonally, the polarity suggestions for T_t variables depend on a wide variety of factors (as the given frequency, density, iteration step and search strategy), and, therefore, can be dynamically attributed in a scope-sensitive manner.

Additionally, solver parameters as the decay factors for variable activity and clause activity can be dynamically tuned on the basis of sensitive analysis.

Finally, the solver resolution can be adapted to be, for instance, sensitive to the difference between the transaction and item literals in a way that promotes a more focused search. This can result in significant performance improvements. Note, however, that the solver functionally must be ensured to support the addition of new flexible constraints.

4 Results

This section details the undertaken evaluation of the previous options against state-of-the-art CP solutions. First, we visit the properties of our implementation, then we describe the most significant observations, and, finally, we discuss the results to retrieve a set of implications.

Covered Options

The supported encodings are the target FIM encoding and its dual formulation. The simple and expressive subsets and supersets negation are covered enumeration options. The modeling of (anti-)monotonicity properties at the encoding level are not supported since they imply an exponential growth on the number of variables and constraints. All the advanced search options (simple, LSM, CMG and LD) are supported as well. Additionally, all the restrictions introduced were found when linking some solvers, and, therefore, they are addressed in our experiments. Finally, an extensive set of tuning options were implemented, with the most significant being the ones described in the previous section.

Codification Alternatives

Different codifications were defined with the goal of supporting an efficient and flexible interface with multiple pseudo-Boolean solvers. For instance, a codification in Java can only interface efficiently with solvers in Java through direct invocation. Otherwise, interaction needs to be done between executables (the developed layer and the solver), leading to an additional latency as a result of the required synchronization between them. Additionally, the exchanged information requires parsing. This hampers the performance as information is extensively exchanged in every iteration (note that the number of iterations is usually greater than n^2 for low frequencies).

Two classes of SAT solvers were adopted. The first class comprises the solvers with open-source for whom all the covered options were implemented. The second class includes the solvers with undisclosed-source with whom a simple goal was assessed: see how they performed for specific single-iterations against the alternatives. This was justified by the fact that since most of them do not support neither the use of assumptions nor the insertion of new clauses required to perform enumerations. Although the full-feeding of these solvers with the encoding for every iteration was tried, the fact that they do not keep the learned clauses in memory turned their performance impractical. These solvers are PBS [2], BSOLO [28] and WBO [27].

The adopted solvers belonging to the first class are SAT4J [26] and MiniSat+ [19]. Note that, since SAT4J is implemented in Java and MiniSat+ is implemented in C++, two codifications for the target solution were supported: under Java and C++¹.

Datasets

The adopted datasets were taken from the UCI repository². The density of the dataset is defined by the average number of items per transaction divided by the size of the items' alphabet. Although the selected datasets are not large (note that optimal approaches suffer from scalability problems), they are dense by nature and, therefore, their use within traditional approaches is still largely computationally expensive.

¹ web.ist.utl.pt/rmch/research/software

An additional third codification is available in C# upon request.

² <http://archive.ics.uci.edu/ml/>

4.1 Observations

The computer used to run the experiments was an Intel Core i5 2.80GHz with 6GB of RAM. The algorithms were implemented using Java (JVM version 1.6.0-24) and C++ (GCC 4.4.5) in 64-bit Linux (Debian 2.30.2) operating system.

Comparative analysis

Table 4 synthesizes the main results over UCI datasets for tuned MiniSat+ and SAT4J implementations under CMG search option and for the state-of-the-art CP performer, FIMCP. The proposed SAT-based solutions have a phase transition that relaxes for very low and medium-to-high frequency thresholds, as illustrated in Fig.1. Contrasting, FIMCP behavior increasingly deteriorates with the decrease of the frequency threshold.

Dataset	#Items	#Trans	Density	$\theta=0,02$			$\theta=0,05$		
				MiniSat+	SAT4J	FIMCP	MiniSat+	SAT4J	FIMCP
<i>Tic-tac-toe</i>	27	958	0,33	366,7	—*	0,2	550,0	—*	0,1
<i>Primary-tumor</i>	31	336	0,48	17,8	—*	1,0	300,0	—*	0,5
<i>Zoo-1</i>	36	101	0,44	0,3	4,3	1,5	0,8	7,2	0,6
<i>Vote</i>	48	435	0,33	528,5	—*	2,0	—*	—*	0,7
<i>Soybean</i>	50	630	0,32	—*	—*	1,9	—*	—*	0,5
<i>Hepatitis</i>	68	137	0,5	60,5	—*	—*	229,8	—*	—*
<i>Lymph</i>	68	148	0,4	—*	—*	—*	—*	—*	—*
<i>Kr-vs-kp</i>	73	3169	0,49	380,5	—*	—*	—*	—*	—*
<i>Hypothyroid</i>	88	3247	0,49	433,3	—*	—*	—*	—*	—*
<i>Heart-cleveland</i>	95	296	0,47	—*	—*	—*	—*	—*	—*
<i>German-credit</i>	112	1000	0,34	—*	—*	—*	—*	—*	—*
<i>Mushroom</i>	119	8124	0,18	—*	—*	—*	—*	—*	49126
<i>Aust-credit</i>	125	653	0,41	—*	—*	—*	—*	—*	—*
<i>Audiology</i>	148	216	0,45	—*	—*	—*	—*	—*	—*

Dataset	$\theta=0,1$			$\theta=0,2$			$\theta=0,4$		
	MiniSat+	SAT4J	FIMCP	MiniSat+	SAT4J	FIMCP	MiniSat+	SAT4J	FIMCP
<i>Tic-tac-toe</i>	33,8	36,0	0,1	2,3	1,9	0,1	0,2	0,0	0,0
<i>Primary-tumor</i>	330,0	—*	0,2	126,2	124,0	0,1	38,0	9,5	0,0
<i>Zoo-1</i>	1,2	9,0	0,2	1,4	17,0	0,0	0,8	0,8	0,0
<i>Vote</i>	—*	—*	0,3	149,3	397,5	0,1	5,4	2,1	0,0
<i>Soybean</i>	—*	166,9	0,2	554,0	15,7	0,0	12,9	2,0	0,0
<i>Hepatitis</i>	—*	—*	—*	—*	—*	2,5	382,1	—†	0,2
<i>Lymph</i>	—*	—*	19447	—*	—*	0,1	14,2	9,1	0,0
<i>Kr-vs-kp</i>	—*	—*	—*	—*	—*	—*	—*	—†	—*
<i>Hypothyroid</i>	—*	—*	—*	—*	—*	—*	—*	—*	—*
<i>Heart-cleveland</i>	—*	—*	—*	—*	—*	—*	508,5	—†	0,8
<i>German-credit</i>	—*	—*	—*	—*	—*	15,9	—*	—*	0,5
<i>Mushroom</i>	—*	—*	11,9	—*	—*	2,9	99,9	—†	0,5
<i>Aust-credit</i>	—*	—*	—*	—*	—*	—*	—*	—*	5,5
<i>Audiology</i>	—*	—*	—*	—*	—*	—*	—*	—*	—*

Table 4: Overall efficiency of the proposed solvers against FIMCP (seconds)

*timeout; †memory out;

FIMCP is the best option when targeting either medium-to-high frequency thresholds or normal-to-low dense datasets. However, when the target problem

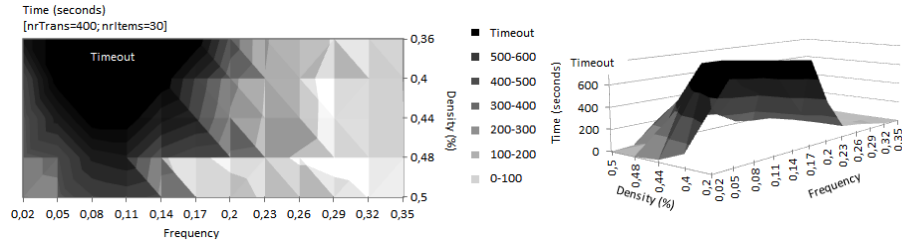


Fig.1: Phase transitions for varying densities of generated datasets (using MiniSat+)

is based on low frequency thresholds over dense or large datasets, our adapted MiniSat+ solution is the choice. Note that these cases are the most common scenario in FIM problems, where the required frequency range to perform association rules falls between 1 and 2%. An illustrative example of this advantageous behavior can be observed over the hepatitis dataset. The performance of FIMCP for this dataset is only scalable until frequency thresholds near and above 20%. For $\theta < 10\%$, FIMCP performance exponentially deteriorates with θ decrements. Interestingly, under this same θ range, MiniSat+ is able to answer to the FIM problem in useful time as shown in Fig.2.

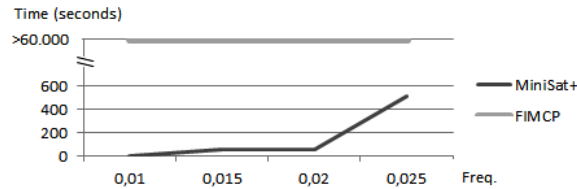


Fig.2: MiniSat+ behavior for Hepatitis dataset under very low frequencies

Since the performance of SAT4J is hampered by its resolution properties (not tuned to deal with low frequency thresholds and not able to classify key constraints) and by a bad memory management (dependent on a garbage collector), the adoption of MiniSat+ or FIMCP is overall preferred.

Selecting SAT-based solvers

The selection of best performer selections based on the inputted frequency and dataset properties implies an extensive analysis of the behavior of the solvers across different axes of choice. Table 5 synthesizes the main results of the undertaken analysis, which are further detailed in Appendix C.

4.2 Deepening the analysis behind SAT vs. CSP

Up to now, the most efficient CP approaches map FIM as a Constraint Satisfaction Problem (CSP) [34]. A CSP problem is specified by a finite set of variables V , an initial domain D (which maps every variable $v \in V$ to a finite set of values $D(v)$), and a finite set of constraints C in first-order logic. The goal is to output the variable domains which satisfy all constraints. Thus, the solution to a FIM problem can be directly retrieved from the codification of equations (5) and (8).

Axis	Observations
Enumeration options	<ul style="list-style-type: none"> ◦ Although the negation of subsets or supersets within iterations lead to significant performance improvements, the level of impact depends on the search ability to early discovery the largest itemsets (as in LSM, CMG and LD); ◦ Interestingly, the explicit negation of each subset/superset is preferred over the expressive one-clause-only negations when there is a high number of iterations as solvers are able to remove duplicated negated subsets, while in the later option there is an increasing redundancy that may hamper the search performance.
Search options	<ul style="list-style-type: none"> ◦ CMG search is overall the best choice (β-searches are efficient); ◦ Simple search methods are only able to perform searches with acceptable efficiency on high frequency thresholds over very sparse datasets, and should only be adopted when resolution promotes an early discovery of larger itemsets; ◦ LSM searches are only competitive when a few number of iterations is performed since each search is a lot heavier than a full CMG search (including one α-search and multiple β-searches); ◦ LD is competitive with CMG for medium frequency thresholds ($0.05 < \theta < 0.2$) on small datasets. LD performance quickly degrades with their growing size. Although <code>sat</code> searches are light (focused on discovering frequent itemsets with a given length) a significant overhead is added by the fixed number of heavier <code>unsat</code> searches, particularly for a low-to-medium number of maximal frequent itemsets.
Encoding options (restrictions)	<ul style="list-style-type: none"> ◦ Clause-oriented encodings seems to be preferred over minimal encodings for lower-to-medium frequency thresholds. Although minimal encodings have significantly fewer constraints, they are not easy to classify; ◦ Restricted encodings penalize the performance significantly (5-25%).
Tuning options	<ul style="list-style-type: none"> ◦ Positive suggestions are adequate for low-to-medium frequencies (indicate a preference towards larger itemsets), while negative suggestions are indicated for higher frequencies (since large itemsets are not frequent, more conflicts are found increasing the number of backtracks and potential restarts);
Implementation options	<ul style="list-style-type: none"> ◦ SAT4J solver is the option for high frequency thresholds and for very sparse datasets since it accepts non-constrained encodings and its resolution is more tuned to find finer itemsets; ◦ MiniSat+ solver is the natural choice for the rest of the options due to, among others, the tuned parameterizations as polarity suggestions and decaying factors.

Table 5: Key observations for the target SAT-based solvers across six dimensions

The first key concept used to speed up the search is constraint propagation to reduce the domains of variables such that the domain remains locally consistent. To maintain local consistencies, propagators are used to remove values from a domain that can never satisfy a constraint.

Besides this property, constraint-based solvers are well-prepared to deal with certain types of constraints. Two examples are: the summation constraint and the reified summation constraint [22]. Flexible solvers as Gecode are well-prepared to deal with them. The FIM problem heavily relies on these constraints.

Let $x \in V \subseteq \mathcal{V}$ be a variable with an associated weight w_x , a summation constraint as the following form [22]:

$$\sum_{x \in V} w_x x \geq \theta \quad (16)$$

The propagator task is to discover as early as possible whether the constraint is violated (i.e. whether the upper-bound of the sum is still above the threshold).

In a reified summation constraint, the evaluation of a summation constraint depends on a Boolean variable b (as adopted in the target frequency constraints):

$$b \rightarrow C' \text{ (usually } C' = \sum_{x \in V} w_x x \geq \theta) \quad (17)$$

The most important propagation that occurs for this constraint is the one that updates the domain of b [22]. When an item variable is fixed, the following is possible for the coverage constraint:

- if for some t : $\sum_{i \in I} (1 - D_{ti}) I_i^{min} > 0$ then remove 1 from $D(T_t)$
- if for some t : $\sum_{i \in I} (1 - D_{ti}) I_i^{max} = 0$ then remove 0 from $D(T_t)$

Once the domain of a variable T_t changes, the support constraint is activated. The support constraint is simply a summation constraint, which checks whether: $\sum_{t \in A} T_t^{max} \geq \theta$. If this constraint fails, CSP solvers do not need to branch further and, therefore, can backtrack.

Contrasting, the mapping of reified frequency constraint into SAT is not concisely handled, leading to the generation of nm clauses. Jointly these observations and the fact that SAT solvers are not able to expressively deal with enumerations (in particular, when an arbitrary number of clauses are added between iterations) justify the poor performance of the developed solutions across several datasets and frequencies.

4.3 Discussion

From the previous observations, several implications with impact on when to use and how to tune SAT-based solutions can be retrieved.

When to use SAT-based solutions:

- for low frequency thresholds, with most promising results on dense datasets (for instance, datasets typically adopted for classification tasks with nominal attributes with few labels, or with numeric attributes that are binarized using thresholds). The level of frequency used to opt for a SAT-based solution depends on the density: if density is near 45-50% the frequency can reach 10%, while for other cases the frequency should not exceed 4%;
- when the problem is not defined as a complete enumeration, but aims to find a fixed number of patterns of interest, or to verify its satisfiability;
- in specific cases for higher frequencies (mainly between 10% and 25%) when the FIM problem is formulated as its dual;

How to tune SAT-based solutions:

- use an expressive enumeration strategy such as the compact subsets negation, change polarity suggestions (item-related variables to positive and transaction-related variables to false) and prefer clause-oriented encodings;
- adopt the search strategy according to the target instances and problem:
 - by default and for the majority of cases, CMG search is the most efficient;
 - *simple* search should only be adopted when the polarity suggestions can be set according to the proposed guidelines for datasets with a good distribution of frequent itemsets among transactions (otherwise the use polarity suggestions do not guarantee that, within each iteration, large itemsets are discovered);

- LSM search should be selected when we are interested in a subset of itemsets with major interest (i.e. when there is the requirement of obtaining the k largest maximal frequent itemsets). In particular, the value of k should be significantly lower than the total number of maximal frequent itemsets (phase transition), otherwise a CMG search should be preferred (although there is the need to generate all maximal frequent itemsets since it does not guarantee if a maximal frequent itemset is one of the k largest);
- be able to translate the flexible user-defined constraints into a SAT formula, by defining generic methods to translate equivalences and implications. Note that the distinguishing feature of CP is that it provides general principles for solving problems with any type of constraints. Although some streams claim that this observation sets it apart from SAT solving [9], we show that this is a simple step as it is illustrated by the following constraints' encodings:
 - $I^p = I^q \rightarrow \wedge_{i \in \mathcal{I}} (I_i^p \Leftrightarrow I_i^q)$
 - $i \in I^p \rightarrow I_i^p$
 - $I^p \cap I^q = I^r \rightarrow \wedge_{i \in \mathcal{I}} (I_i^r \Leftrightarrow I_i^p \wedge \neg I_i^q)$
 - $I^p \cap I^q = I^r \rightarrow \wedge_{i \in \mathcal{I}} (I_i^r \Leftrightarrow I_i^p \wedge I_i^q)$
 - $I^p \cup I^q = I^r \rightarrow \wedge_{i \in \mathcal{I}} (I_i^r \Leftrightarrow I_i^p \vee I_i^q)$
 - $coverItems(I^1, \dots, I^k) \rightarrow \wedge_{i \in \mathcal{I}} (\vee_{j \in 1..k} I_i^j)$
 - $coverTrans(I^1, \dots, I^k) \rightarrow \wedge_{t \in \mathcal{T}} (\vee_{j \in 1..k} T_t^j)$

Unfortunately, both the proposed SAT-based solution and any other CP solutions are featured by high computational complexity and their straightforward implementations are not applicable to large data sets. This advocates the need for local learning where transactions are partitioned and multiple criteria can be used for the integration of the frequent patterns found within each fragment (for instance, through voting techniques [32]). Alternatively, the FIM constraints can be used to compute information gain metrics [14] as the entropy-measure or its dual formulation as a basis for pruning techniques.

5 Related research

The main research streams approaching PM within a CP framework can be classified according to: *i*) the extent to which an user can define novel constraints and combine them, and *ii*) according to the type of supported constraints. Within the first axis, although approaches as Patternist [10], Molfea [18] and MusicDFS [37] support a predefined number of constraints, they do not allow for the expressive definition of novel constraints as FIMCP [17], PattCP [21] and GeMini [30]. In the following section the constraints covered by existing approaches and variations to the FIM problem are briefly presented. Finally, other related work with potential relevant contributions is covered.

Extending Constraints.

Expressive CP models [30, 24, 21] enable the flexible definition of constraints using: *constants* including numerical values, items as A , specific itemsets as $\{A, B\}$,

and specific transactions as t_i ; *variables* ($\bigvee_{i \in C} I_i$ and $\bigvee_{t \in C} T_t$); set and numerical *operators* ($\cup, \cap, \setminus, \times, +, -$); and *function symbols* involving one or several terms, which can be built-in as $overlapTrans(I, J) = |cover(I) \cap cover(J)|$ or defined by the user as $area(I) = freq(I) \times size(I)$ or $coverage(I, J) = freq(I \cup J) \times size(I \cap J)$.

These constraints can be used to filter the patterns of interest. However they can be used with a different purpose: closed itemset mining, discriminative pattern mining, pattern-based clustering and pattern-based classification. A recent direction is taking into account the relationships between local patterns to produce global patterns or pattern sets [24]. Despite their importance, there are very few attempts to mine patterns involving several local patterns and the existing methods tackle just particular cases by using devoted techniques [22]. In [30, 9], the importance of adopting a declarative CP-based approach to mine global patterns was highlighted by several examples coming from clustering tasks based on associations. Due to the complexity of this task and its easy modeling through constraints, CP-based approaches as pattern teams [25] have been providing encouraging results when compared against heuristic-based approaches that consider the added value of a new global pattern given extensive combinations of selected patterns [11].

Problem Variations.

The introduced FIM approach can be extended and adapted in different ways. One of the most common is the k -pattern set mining [21, 30]. Unlike FIM, the problem of k -pattern set mining is concerned with finding a set of k related patterns under constraints. The discovery of k representative patterns often uses probabilistic models for summarizing frequent patterns [31] and other condensed representations of patterns as the dataset compression using Minimum Description Length Principle [38]. These approaches mainly aim at reducing the redundancy between patterns and, like our SAT approach, often focus on maximal frequent patterns. The k -pattern set mining problem is a very general problem that can be instantiated to a wide variety of mining tasks including concept-learning, rule-learning, re-description mining, conceptual clustering and tiling [21, 30].

Other Relevant Work.

Instead of mining rules that rely on frequent itemsets, some approaches encode the problem within the CP framework using (compact) reducts [5], i.e., subsets of most informative attributes. In this stream, a SAT representation is often formulated as an Integer Programming (IP) model to solve the minimal reduct constraints. In [32], an extensive research is done over Boolean reasoning methodologies for Rough Set theory. Two problems are encoded: the search for reducts and the search for decision rules which are building units of many rule-based classification methods.

Another important direction is the verification whether PM constraints are satisfiable. This is a one-iteration-only decision problem that, according to our previous analysis, can be handled using SAT solvers. An example of this task with multiple applications in privacy preserving data mining, condensed representations and the FIM problem, is the called FREQSAT problem [12]: given

some itemset frequency-interval pairs, does there exist a database such that for every pair the frequency of the itemset falls into the interval? That is, given a set of frequency constraints $C = \{freq(I_j) \in [l_j, u_j], j = 1 \dots m\}$, verify if there exists a database D over $\cup_{j=1}^m I_j$ that satisfies C ³. The problem can be further extended to include arbitrary Boolean expressions over items and conditional frequency expressions in the form of association rules. Additionally, FREQSAT is equivalent to probabilistic satisfiability (pSAT) [14].

6 Concluding Remarks

The use of SAT within the CP framework to address constrained-pattern mining problems was demonstrated to be valid: not only the adopted variables are Boolean, but expressive constraints can be easily mapped into a Boolean formulae. Multiple search options, enumeration strategies, encoding alternatives and parameterizations were studied in order to improve its performance. Understandably, these adaptations aim to orient SAT reasoning to the main properties of frequent itemset mining, without losing the ability to flexibly support novel constraints.

Experimental results show that SAT-based solutions are competitive with the state-of-the-art CSP solutions for an important range of frequencies (range commonly adopted to discover association rules or to perform classification tasks). The efficiency problems found for higher frequencies mainly result from the fact that SAT was not developed with the intent of perform enumerations under the addition of new (and potentially conflicting) clauses. Finally, a set of guidelines were introduced to understand when to use and how to tune SAT-based solutions.

Future Work. In the next steps one should expect:

- the extension of this approach to evaluate its impact not only in terms of efficiency but also in terms of accuracy for FIM-based problems as the discriminative-FIM problem;
- the exploitation of SAT limits and performance in comparison to CP approaches under an intensive use of constraints;
- the development of hybrid approaches with clear rules to select and tune the best performer under certain conditions: *i*) the selected frequency, *ii*) the nature of the problem constraints and relaxations, and *iii*) the dataset properties – by order of relevance: the density, the number of items, the item-to-transaction ratio and the number of transactions;
- the exploitation of potential improvements from the codification of the adopted enumeration and search strategies at the encoding level;

³ suppose that the following set C of frequency constraints C is given: $\{freq(\{a, b\}) \in [3/4, 1]; freq(\{a, c\}) \in [3/4, 1], freq(\{b, c\}) \in [3/4, 1], freq(\{d, e\}) \in [3/4, 1], freq(\{d, f\}) \in [1/2, 1], freq(\{e, f\}) \in [1/2, 1], freq(\{a, b, c, d, e, f\}) = 0\}$; C is in FREQSAT, because it is satisfiable by the following database: $D = \{(1, \{a, b, c, d, e\}), (2, \{a, b, c, d, e\}), (3, \{a, b, c, d, e\}), (4, \{a, b, c, d, f\}), (5, \{a, b, c, e, f\}), (6, \{a, b, d, e, f\}), (7, \{a, c, d, e, f\}), (8, \{b, c, d, e, f\})\}$

- the adoption of more expressive constraints besides (anti-)monotonicity, including extensions based on the inclusion-exclusion principle [13] and other frequency-based relations [7, 35];
- the finding of patterns in continuous data (as, for instance, required in many bioinformatic datasets), which may require discretization techniques beyond support-vectors, rough set and Boolean reasoning theories [32];
- the mining of frequent itemsets in structured data as sequences, trees and graphs. New formulations are required to represent these problems under CP, which may not be trivial to encode when recurring to a fixed number of features or variables;
- the assessment of SAT approaches to perform constraint-based clustering and constraint-based classifier induction (not necessarily relying on frequent itemsets). In constraint-based clustering the challenge is to cluster examples when additional knowledge is available about these examples, for instance, prohibiting certain examples from being clustered together (so-called cannot-link constraints) [30, 34]. Similarly, in constraint-based classifier induction, one may wish to find a decision tree that satisfies size and cost-constraints [8]. In traditional data mining, the relationship between itemset mining and constraint-based decision tree learning was studied in [33], however such relation was not yet exploited in a CP setting;
- the adoption of SAT verification (as the previously introduced FREQSAT problem [14]) using deduction rules to prune/constrain the search of frequent itemsets. The monotonicity rule is a very simple example of deduction. More advanced rules, as the partial frequency available for some itemsets, bound on the frequencies of itemsets yet to be counted. Examples of deduction rules to improve pruning and speed-up FIM approaches are given in [6, 16];
- the study of potential techniques to turn SAT and other CP-based approaches scalable. Options may include the local scanning on dataset partitions [32] or the use of data stream mining approaches [15].

References

1. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. *SIGMOD Rec.* 22(2), 207–216 (Jun 1993)
2. Aloul, F.A., Ramani, A., Sakallah, K.A., Markov, I.L.: Solution and optimization of systems of pseudo-boolean constraints. *IEEE Trans. Comput.* 56(10), 1415–1424 (Oct 2007)
3. Antunes, C.: An ontology-based framework for mining patterns in the presence of background knowledge. In: 1st Int. Conf. on Advanced Intelligence (ICAI’08). pp. 163–168. Chinese Association for AI, Post and Telecom Press, Beijing, China (2008)
4. Asín, R., Nieuwenhuis, R., Oliveras, A., Rodríguez-Carbonell, E.: Cardinality networks: a theoretical and empirical study. *Constraints* 16(2), 195–221 (2011)
5. Bakar, A.A., Sulaiman, M.N., Othman, M., Selamat, M.H.: Propositional satisfiability algorithm to find minimal reducts for data mining. *Int. J. Comput. Math.* 79(4), 379–389 (2002)
6. Bastide, Y., Taouil, R., Pasquier, N., Stumme, G., Lakhal, L.: Mining frequent patterns with counting inference. *SIGKDD Explor. Newsl.* 2(2), 66–75 (Dec 2000)

7. Bayardo, Jr., R.J.: Efficiently mining long patterns from databases. *SIGMOD Rec.* 27(2), 85–93 (Jun 1998)
8. Bessiere, C., Hebrard, E., O’Sullivan, B.: Minimising decision tree size as combinatorial optimisation. In: *Proceedings of the 15th international conference on Principles and practice of constraint programming*. pp. 173–187. CP’09, Springer-Verlag, Berlin, Heidelberg (2009)
9. Boizumault, P., Crémilleux, B., Khiari, M., Loudni, S., Métivier, J.P.: Discovering knowledge using a constraint-based language. *CoRR abs/1107.3407* (2011)
10. Bonchi, F., Lucchese, C.: Extending the state-of-the-art of constraint-based pattern discovery. *Data Knowl. Eng.* 60(2), 377–399 (Feb 2007)
11. Bringmann, B., Zimmermann, A.: The chosen few: On identifying valuable patterns. In: *Proceedings of the 2007 Seventh IEEE International Conference on Data Mining*. pp. 63–72. ICDM ’07, IEEE Computer Society, Washington, DC, USA (2007)
12. Calders, T.: Computational complexity of itemset frequency satisfiability. In: *Proceedings of the twenty-third ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. pp. 143–154. PODS ’04, ACM, New York, NY, USA (2004)
13. Calders, T.: Deducing bounds on the support of itemsets. In: *In Database Technologies for Data Mining - Discovering Knowledge with Inductive Queries*, volume 2682 of LNCS. pp. 214–233. Springer (2004)
14. Calders, T.: Itemset frequency satisfiability: Complexity and axiomatization. *Theor. Comput. Sci.* 394(1-2), 84–111 (Mar 2008)
15. Calders, T., Dexters, N., Goethals, B.: Mining frequent itemsets in a stream. In: *Proceedings of the 2007 Seventh IEEE International Conference on Data Mining*. pp. 83–92. ICDM ’07, IEEE Computer Society, Washington, DC, USA (2007)
16. Calders, T., Goethals, B.: Mining all non-derivable frequent itemsets. In: *Proceedings of the 6th European Conference on Principles of Data Mining and Knowledge Discovery*. pp. 74–85. PKDD ’02, Springer-Verlag, London, UK, UK (2002)
17. De Raedt, L., Guns, T., Nijssen, S.: Constraint programming for itemset mining. In: *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 204–212. KDD ’08, ACM, New York, NY, USA (2008)
18. De Raedt, L., Kramer, S.: The levelwise version space algorithm and its application to molecular fragment finding. In: *Proceedings of the 17th international joint conference on Artificial intelligence - Volume 2*. pp. 853–859. IJCAI’01, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2001)
19. Eén, N., Sörensson, N.: Translating pseudo-boolean constraints into sat. *JSAT* 2(1-4), 1–26 (2006)
20. Gent, I.P., Miguel, I., Rendl, A.: Tailoring solver-independent constraint models: a case study with essence’ and minion. In: *Proceedings of the 7th International conference on Abstraction, reformulation, and approximation*. pp. 184–199. SARA’07, Springer-Verlag, Berlin, Heidelberg (2007)
21. Guns, T., Nijssen, S., de Raedt, L.: k-pattern set mining under constraints. *IEEE Transactions on Knowledge and Data Engineering* 99(PrePrints) (2011)
22. Guns, T., Nijssen, S., Raedt, L.D.: Itemset mining: A constraint programming perspective. *Artif. Intell.* 175(12-13), 1951–1983 (2011)
23. Henriques, R., Antunes, C.: An integrated approach for healthcare planning over dimensional data using long-term prediction. In: *1st Proc. in Healthcare Information Systems*. Springer-Verlag, Beijing, China (2012)
24. Khiari, M., Boizumault, P., Crémilleux, B.: Constraint programming for mining n-ary patterns. In: *Proceedings of the 16th international conference on Principles and*

- practice of constraint programming. pp. 552–567. CP’10, Springer-Verlag, Berlin, Heidelberg (2010)
25. Knobbe, A., Ho, E.: Pattern teams. In: Fürnkranz, J., Scheffer, T., Spiliopoulou, M. (eds.) Knowledge Discovery in Databases: PKDD 2006, Lecture Notes in Computer Science, vol. 4213, pp. 577–584. Springer Berlin Heidelberg (2006)
 26. Le Berre, D., Parrain, A.: The Sat4j library, release 2.2. *Journal on Satisfiability, Boolean Modeling and Computation* 7, 59–64 (2010)
 27. Manquinho, V., Marques-Silva, J., Planes, J.: Algorithms for weighted boolean optimization. In: Proceedings of the 12th International Conference on Theory and Applications of Satisfiability Testing. pp. 495–508. SAT ’09, Springer-Verlag, Berlin, Heidelberg (2009)
 28. Manquinho, V.M., Marques-Silva, J.P.: On using cutting planes in pseudo-boolean optimization. *JSAT* 2(1-4), 209–219 (2006)
 29. McMillan, K.L.: Applying sat methods in unbounded symbolic model checking. In: Proceedings of the 14th International Conference on Computer Aided Verification. pp. 250–264. CAV ’02, Springer-Verlag, London, UK, UK (2002)
 30. Métivier, J.P., Boizumault, P., Crémilleux, B., Khiari, M., Loudni, S.: A constraint language for declarative pattern discovery. In: Proceedings of the 27th Annual ACM Symposium on Applied Computing. pp. 119–125. SAC ’12, ACM, New York, NY, USA (2012)
 31. Mielikäinen, T., Mannila, H.: The pattern ordering problem. In: Proceedings of the 7th European Conference on Principles of Data Mining and Knowledge Discovery, Lecture Notes in Artificial Intelligence. pp. 327–338. Springer-Verlag (2003)
 32. Nguyen, H.S.: Approximate Boolean Reasoning: Foundations and Applications in Data Mining. *Transactions on Rough Sets V* 4100, 334–506 (2006)
 33. Nijssen, S., Fromont, E.: Optimal constraint-based decision tree induction from itemset lattices. *Data Min. Knowl. Discov.* 21(1), 9–51 (Jul 2010)
 34. Raedt, L.D., Guns, T., Nijssen, S.: Constraint programming for data mining and machine learning. In: Fox, M., Poole, D. (eds.) AAAI. AAAI Press (2010)
 35. Sayrafi, B., Van Gucht, D.: Differential constraints. In: Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems. pp. 348–357. PODS ’05, ACM, New York, NY, USA (2005)
 36. Sinz, C.: Towards an optimal cnf encoding of boolean cardinality constraints. In: van Beek, P. (ed.) CP. Lecture Notes in Computer Science, vol. 3709, pp. 827–831. Springer (2005)
 37. Soulet, A., Crémilleux, B.: An efficient framework for mining flexible constraints. In: Proceedings of the 9th Pacific-Asia conference on Advances in Knowledge Discovery and Data Mining. pp. 661–671. PAKDD’05, Springer-Verlag, Berlin, Heidelberg (2005)
 38. Vreeken, J., Leeuwen, M., Siebes, A.: Krimp: mining itemsets that compress. *Data Min. Knowl. Discov.* 23(1), 169–214 (Jul 2011)

A Positive Literals

Some solvers only admit a simplified and constrained pseudo-Boolean notation as input. An example is the exclusion of negated literals, which may not be trivial to handle. The translation of *coverage* constraints into constraints with non-negated literals is trivial:

$$\begin{aligned} & \bigwedge_{t \in T} (\bigwedge_{i \in I|D_{ti}} (\neg T_t \vee \neg I_i) \wedge (T_t \vee (\bigvee_{i \in I|D_{ti}} I_i))) \\ \leftrightarrow & \bigwedge_{t \in T} (\bigwedge_{i \in I|D_{ti}} (T_t \vee I_i \leq 1) \wedge (T_t \vee (\bigvee_{i \in I|D_{ti}} I_i) \geq 1)) \end{aligned}$$

The *frequency* constraints were translated with the goal of maintaining the advantageous \geq operator:

$$\begin{aligned} & \bigwedge_{i \in I} (\theta - I_i + \sum_{t \in T|D_{ti}} T_t \geq \theta) \\ \leftrightarrow & \bigwedge_{i \in I} (-\theta - I_i + \sum_{t \in T|D_{ti}} T_t \geq 0) \end{aligned}$$

In this fashion, solvers as `Minisat+` [19] can be tested without significant overhead, as most of them internally are able to classify both of the pseudo-Boolean constraints defined for the coverage restrictions.

B Results

The following appendix sections detail the observations made in Table 3. The adopted datasets for these experiments are either: real UCI datasets (distributed in Fig.3 according to their properties) or generated datasets (with customized number of items, number of transactions and density) whose generation depend on a biasing parameter γ for the emerging of patterns according to a distribution similar as real datasets. The small but representative `zoo` dataset is often used to compare options.

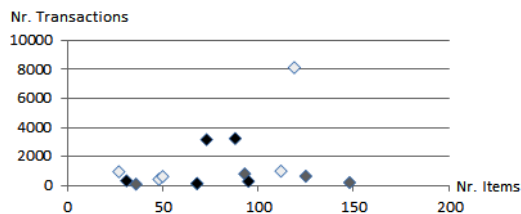


Fig.3: Characterization of the target datasets (high density in black, low in gray)

B.1 Enumeration options

As expected, the adoption of the monotonicity principle within each enumeration have significantly impacted the performance of SAT solvers. According to an extract of the results in Table 6, two main observations can be drawn. *First*, the level of impact by (either simply or expressively) negating subsets depends on the search option and on the target frequency.

Understandably, such impact depends on the ability to early discover the largest itemsets. This is the reason why subsets negation is key to LSM and CMG searches (its development relies on this principle), and important, but not so significant, within simple search methods. Additionally, cutting space through subsets negation is more critical for low frequencies, as an increasing length and number of itemsets is observed.

Second, the choice of when to explicitly insert a negation of each subset or to expressively insert one unique clause requiring the selection of an item not observed in

Search	Enumeration	$\theta=0,2$	$\theta=0,4$	$\theta=0,6$
Optimization (SAT4J)	Simple	1514,0	230,0	1,0
	Expressive Simple	943,7	89,2	0,1
	Subsets Negation	31,3	2,6	0,1
	Expr. Subsets Neg.	18,2	1,8	0,1
Decision (SAT4J)	Simple	1009,4	12,6	0,5
	Expressive Simple	589,8	5,4	0,2
	Subsets Negation	208,8	4,8	0,1
	Expr. Subsets Neg.	183,0	6,5	0,2

Table 6: Comparison of enumeration strategies for the zoo dataset (seconds)

the found itemset is not simple. This choice depends on two factors: the total number of similar itemsets and the relative length of the found itemset. In the first case, when we have multiple similar frequent itemsets (sharing the majority of items), the explicit insertion of repeated negated sub-itemsets is detected by the solver, and the repeated clauses are removed, while in the expressive insertion every new clause is inserted as a new problem constraint. In the second case, when the length of itemsets is small in comparison with the item alphabet this means that, although the negated sub-itemsets inserted by an explicit negation strategy generate multiple clauses, the number of clauses is not high and the number of literals per clause is low. Contrasting, although expressive negation strategies only add a clause per iteration, the number of literals can be significantly high and may hamper the resolution performance. These observations claim for an increased attention on the strategy selection based on the inputted frequency and dataset properties.

B.2 Search options

The first observation coming from the search option results is that, although CMG search is overall the best choice, the performance of the searches highly vary according to the dataset density and input frequency. Maximal-oriented searches (as LSM, CMG and LD) perform better for low frequency thresholds in dense datasets and for high frequency thresholds in sparse datasets. Table 7 performs a two-axes evaluation – over generated datasets with varying densities and over a fixed dataset with varying frequency thresholds.

Search	Dense Gen.	Normal Gen.	Sparse Gen		
Simple	—*	—*	—*		
LSM	3491,4	176,2	75,4		
CMG	268,6	97,5	62,9		
LD	1296,4	343,9	291,2		
	$\theta=0,02$	$\theta=0,05$	$\theta=0,1$	$\theta=0,2$	$\theta=0,4$
Simple	1020,1	978,9	792,3	183,0	6,5
LSM	5,9	14,1	26,8	17,0	1,8
CMG	4,9	7,4	9,6	7,6	1,0
LD	5,4	9,8	9,5	7,4	3,4

Table 7: Comparison of search options on generated datasets with different densities ($\theta=0,05$) and on zoo dataset with different frequencies using SAT4J (seconds)

		$\theta=0,02$	$\theta=0,05$	$\theta=0,1$	$\theta=0,2$	$\theta=0,4$
Simple	Number of searches	62032	30369	9354	4087	367
	Average time per search	2,3	2,5	3,7	7,6	49
LSM	Number of searches	119	221	298	267	90
	Average time per search	8	23	43	78	195
CMG	Number of α -searches	119	221	298	267	90
	Average time per main search	2,2	2,2	2,8	4,3	38
	Number of β -searches	244	783	1075	622	134
	Average time per constrained search	1,02	1,02	1,11	1,22	9,1
LD	Number of unsat searches	32	32	32	32	32
	Average time per main search	7,2	8,4	15,3	27,4	82
	Number of sat searches	119	221	298	267	90
	Average time per constrained search	2,7	3,2	3,9	6,3	28,3

Table 8: Search options in MiniSat+: number of searches and avg. time (milliseconds)

Simple search methods are only able to perform searches with acceptable efficiency on high frequency thresholds over very sparse datasets. In fact, even when adopting the monotonicity principle, simple searches do not scale as it is visible based on the increasing number of iterations (Table 8) as frequency thresholds decrease. Simple searches should only be adopted for an implementation that is able to promote the early selection of larger itemsets by, for instance, adjusting the polarity of the item variables. If this is the case where maximal frequent itemsets are guaranteed, simple search should be able to outperform constrained CMG.

LSM searches are only competitive when a few number of iterations is performed since each search is a lot heavier than a full α -search (include multiple β -searches) as it has a larger space to exploit. This is usually the case where multiple similar itemsets (e.g. ACD, ACE, ADE, BCD, BDE) collapse into unique maximal frequent itemsets (e.g. ABCDE) as a result of the threshold frequency decreasing.

CMG searches smooths the heavy computational cost of discovering multiple similar itemsets as each β -search is very light as depicted in Table 8. The average time per β -search is less than a half of an α -search (and this relation even decreases under more larger datasets).

Finally, LD is as competitive as CMG for medium frequency thresholds ($0.05 < \theta < 0.2$) on small datasets (LD performance quickly degrades with their growing size). Instead of performing α - and β -searches, it performs one unique type of search for a fixed length of itemsets, which is a very focused search. In order to cover all possible lengths, n of these searches are **unsat**. The number of **sat** searches is equal to the number of maximal frequent itemsets. A search returning **sat** has a similar performance as an α -search, being the additional overhead added by a fixed number of **unsat** searches (understandably, not varying with the input frequency). This overhead can be critical as a search returning **unsat** is largely heavier than a search returning **sat** or a β -search (see Table 8). Therefore, the adoption of this search option essentially depends on how the number of items (defining the number of **unsat** searches in LD) compares to the number of maximal frequent itemsets (influencing the number of β -searches in CMG).

B.3 Encoding options

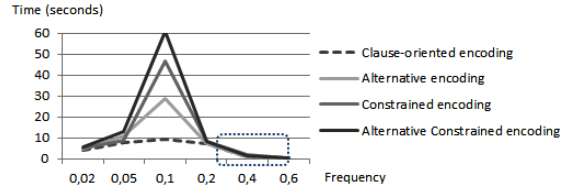
Two main observations can be derived from the experimental tests over different encodings (see Table 9). *First* a clause-oriented encoding seems to be preferred over a constrained encoding. Among other aspects, the adopted constrained encoding requires

the use of non-negated variables and additional clauses and variables to support later clause-removals. This is particularly true for SAT4J as this solver is not able to clausefy some of these constraints.

Encoding	$\theta=0,02$	$\theta=0,05$	$\theta=0,1$	$\theta=0,2$	$\theta=0,4$	$\theta=0,6$
Clause-oriented	4,3	7,9	9,1	7,1	1,0	0,3
Alternative	4,8	10,9	29,0	7,0	1,0	0,2
Constrained	5,3	8,8	46,8	8,4	2,0	0,4
Alternative Constrained	5,9	12,9	60,7	8,3	1,3	0,2

Table 9: Comparison of encodings for the zoo dataset using SAT4J (seconds)

Second, alternative encodings that aim to reduce the number of clauses from $\Theta(mn)$ to $\Theta(m+n)$ (see section 3.5), may not result in significant improvements as the solver instead of having to deal with nm simple binary clauses has to deal with m complex constraints. The choice of whether to adopt or not this encoding mainly depends on the inputted frequency threshold (adopt for $\theta > 20\%$ and avoid its use under low frequency thresholds).



B.4 Tuning options

As depicted in Fig.4, two simple variable polarity suggestions were undertaken. The positive suggestion is the choice for low frequency thresholds. This results from the fact that since item variables are set to true, larger itemsets tend to be initially identified.

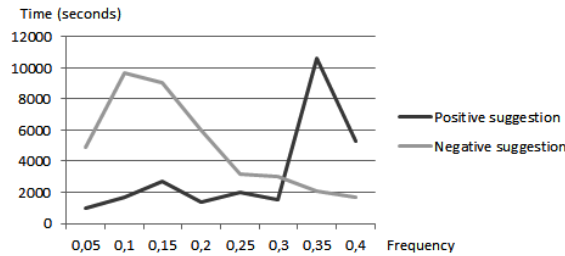


Fig.4: Positive vs. negative polarity suggestion for the zoo.txt dataset (milliseconds)

However, the negative suggestion becomes the option for frequencies below 20% in many datasets ($\sim 30\%$ in Fig.2). This derives from the fact that in the positive suggestion many of the initial large itemset options will not be verified with this higher thresholds, so significantly more conflicts are found within each iteration leading to an additional inefficiency related to the number of backtracks.

Advanced polarity suggestions should not only take into account an overall suggestion for all the variables, but also be able to: *i*) differentiate the polarity suggestion

between item variables and transaction variables, and *ii*) be able to locally adapt suggestions based on task-driven heuristics (e.g. a low number of occurrences of an item variable relative to others may result in a negative polarity suggestion).

B.5 Implementation options

Interestingly, Fig.5 shows that each adopted SAT solver has a unique behavior when answering to the FIM problem. SAT4J is the best option when we are targeting high frequency thresholds and when mining very sparse datasets. Beyond its resolution specificities, this is also a result of accepting non-constrained encodings (including, among others, negated variables, differentiated insertion of clauses and pseudo-Boolean constraints, and clause removal). SAT4J main problems are related to memory inefficiency when dealing with large datasets and with the fact that its algorithm is more tuned to find smaller itemsets, which hampers the behavior of CMG searches since it exponentially increases the number of β -searches.

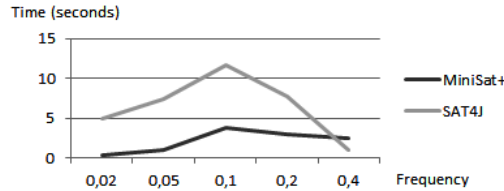
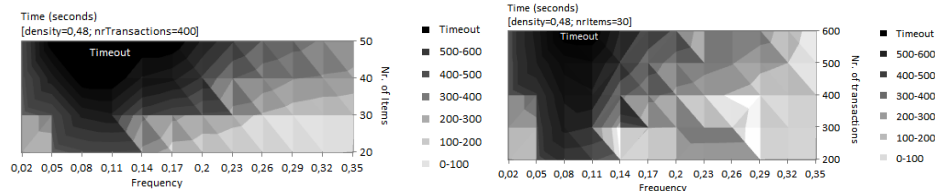


Fig.5: Implementation options for the zoo dataset under the best method

MiniSat+ is the natural choice for the rest of the options – dense datasets and low frequency thresholds. This is not only a consequence of the resolution methods or of C++ additional efficiency, but also a result of multiple improvements related to the solver parameterizations, with positive polarity suggestion being the most significant.

The discussion of the behavior of other adopted solvers as PBS [2], BSOLO [28] and WBO [27], is out of the scope as a result of an excessive latency caused by the need to call them as executables. The successive memory refreshes among iterations and recurrent need to parse and classify formulas distorts any potential analysis.

B.6 Fixing phase transitions



(a) Phase transitions for varying #items

(b) Phase transitions for varying #transactions

Fig.6: Phase transitions for varying size of generated datasets (using MiniSat+)

The study of phase transitions on generated datasets with varying density (Fig.1), number of items (Fig.6a) and number of transactions (Fig.6b) led to two main observations. First, density is the property of datasets with highest impact on the performance of the proposed solution. The variation of a few percentage points can exponentially hamper the performance of our solution. Second, and understandably, the size of the dataset is also a decisive criterion affecting the behavior of our solution. A variation on the number of items is more critical than on the number of transactions. Although our solution hardly scales for a number of items above 100, for low frequencies in dense datasets or high frequencies in sparse datasets it can handle up to 10.000 transactions.

C Software Capabilities

- flexible selection and combination of the search options, enumeration strategies, target datasets, frequencies of interest and optimization parameters (a fragment of the testing code in Java is depicted below);
- extensible codification that gives a basis to model real-problems using simple and enumeration-centered SAT or PB:
 - Utils package contains a general set of encoding functionalities as the generation of `.opb` and `.cnf` files adapted to the restrictions of a particular solver;
 - Solver package provides the interface to SAT solvers (supporting both a direct interface through methods invocation or via executables) and the ability to select multiple search and enumeration options in a task-independent manner;
- parameterizable generation of datasets and their expressive and usable adoption to test limits of performance;
- flexible addition of pattern mining constraints (by extending the `SatPM` class);

```

1: List<Dataset> datasets = DatasetGeneration.getDatasets();
2: List<SATHandler> handlers = SolverOptions.getSolvers();
3: List<Strategy> strategies = StrategyOptions.getStrategies();
4: for(Dataset dataset : datasets)
5:     dataset.encodingOption(encodingID);
6:     for(SATHandler handler : handlers)
7:         handler.setPolarity(polarityID);
8:         for(Strategy strategy : strategies)
9:             for(double freq=0.01; freq<=0.8; freq+=0.01)
10:                 results.add(new StandardFIM(dataset,handler,strategy).run());

```

The software is available in web.ist.utl.pt/rmch/research/software.