



HAL
open science

Efficiency Analysis of ASP Encodings for Sequential Pattern Mining Tasks

Thomas Guyet, Yves Moinard, René Quiniou, Torsten Schaub

► **To cite this version:**

Thomas Guyet, Yves Moinard, René Quiniou, Torsten Schaub. Efficiency Analysis of ASP Encodings for Sequential Pattern Mining Tasks. Bruno Pinaud; Fabrice Guillet; Bruno Cremilleux; Cyril de Runz. *Advances in Knowledge Discovery and Management*, 7, Springer, pp.41–81, 2017, 978-3-319-65405-8. hal-01631879

HAL Id: hal-01631879

<https://inria.hal.science/hal-01631879v1>

Submitted on 13 Nov 2017

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

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

Efficiency Analysis of ASP Encodings for Sequential Pattern Mining Tasks

Thomas Guyet and Yves Moinard and René Quiniou and Torsten Schaub

Abstract This article presents the use of Answer Set Programming (ASP) to mine sequential patterns. ASP is a high-level declarative logic programming paradigm for high level encoding combinatorial and optimization problem solving as well as knowledge representation and reasoning. Thus, ASP is a good candidate for implementing pattern mining with background knowledge, which has been a data mining issue for a long time. We propose encodings of the classical sequential pattern mining tasks within two representations of embeddings (*fill-gaps* vs *skip-gaps*) and for various kinds of patterns: frequent, constrained and condensed. We compare the computational performance of these encodings with each other to get a good insight into the efficiency of ASP encodings. The results show that the *fill-gaps* strategy is better on real problems due to lower memory consumption. Finally, compared to a constraint programming approach (CPSM), another declarative programming paradigm, our proposal showed comparable performance.

1 Introduction

Pattern mining is a data analysis task aiming at identifying “meaningful” patterns in a database of structured data (e.g. itemsets, sequences, graphs). Sequential pattern mining consists in discovering subsequences as patterns in a sequence database

Thomas Guyet
AGROCAMPUS-OUEST/IRISA-UMR 6074/Inria, e-mail: thomas.guyet@agrocampus-ouest.fr

Yves Moinard
Inria – Centre de Rennes e-mail: yves.moinard@inria.fr

René Quiniou
Inria – Centre de Rennes e-mail: rene.quiniou@inria.fr

Torsten Schaub
Potsdam University/Inria e-mail: torsten@cs.uni-potsdam.de

[Shen et al., 2014]. This problem has been introduced at the early beginning of the pattern mining field [Agrawal and Srikant, 1995] with the itemsets mining task [Agrawal et al., 1993]. Sequential pattern mining is known to have a higher complexity than itemsets mining, but it has broad applications [Gupta and Han, 2013]. It includes – but is not limited to – the analysis of patient care pathways, education traces, digital logs (web access for client profiling, intrusion detection from network logs), customer purchase (rules for purchases recommendations), text and bioinformatic sequences.

In most cases, “interesting” patterns are the frequent ones. A pattern is said to be frequent if it appears at least f_{min} times in the database, where f_{min} is a frequency threshold given by the data analyst. This interestingness criterion reveals some important behaviours in the datasets and, above all, it benefits from an interesting property (anti-monotonicity) that make algorithms very efficient, even on large databases. Two decades of research on the specific task of frequent sequential pattern mining have led to very efficient mining methods and strategies to extract the complete set of frequent patterns or condensed representation of frequent patterns [Wang and Han, 2004]. These methods can currently process millions of sequences with very low frequency threshold.

The challenge of mining a deluge of data is about to be solved, but is also about to be replaced by another issue: the deluge of patterns. In fact, the size of the complete set of frequent patterns explodes with the database size and density [Lhote, 2010]. The data analyst cannot handle such volumes of results. A broad range of research, from data visualization [Perer and Wang, 2014] to database sampling [Low-Kam et al., 2013] is currently attempting to tackle this issue. The data-mining community has mainly focused on the addition of expert constraints on sequential patterns [Pei et al., 2004].

Recent approaches have renewed the field of Inductive Logic Programming [Muggleton and De Raedt, 1994] by exploring declarative pattern mining. Similarly, some works have tackled the itemset mining task [Guns et al., 2015, Järvisalo, 2011]. Recently, some propositions have extended the approach to sequence mining [Negrevergne and Guns, 2015, Coquery et al., 2012, Métivier et al., 2013]. Their practical use depends on the efficiency of their encoding to process real datasets. Thanks the improvements on satisfiability (SAT) or constraints programming (CP) solving techniques and solvers, such approaches become realistic alternatives for highly constrained mining tasks. Their computational performances closely reach those of dedicated algorithms.

The long term objective is to benefit from the genericity of solvers to let a user specify a potentially infinite range of constraints on the patterns. Thus, we expect to go from specific algorithm constraints to a rich query language for pattern mining.

The approach we propose in this paper uses the formalism of Answer Set Programming (ASP) and the solver *clingo*. ASP is a logic programming language, as Prolog. Its first order syntax makes ASP programs easy to understand. Furthermore,

ASP benefits from efficient solvers to compute efficiently the solution answer sets [Lifschitz, 2008].

The contributions of this article are twofold. 1) The article presents a declarative approach which provides a high-level specification of a broad range of sequential pattern mining tasks in a unique framework. It demonstrates that this mining task and its related problems – mining closed, maximal and constrain patterns – can easily be encoded with pure declarative ASP. 2) The article extensively evaluates the proposed encodings to draw the computational strengths and limits of ASP for declarative pattern mining. It gives also experimental results about time/memory computing efficiency of the solving process and provides alternative encodings to improve the computing efficiency. The proposed encodings are compared to the results of the CPSM software, based on CP programming [Negrevergne and Guns, 2015].

The article is organized as follows: Sect. 2 introduces ASP programming, its principles and the solver *clingo*. Then in Sect. 3, we introduce sequential pattern mining. In Sect. 4, we give several ASP encodings of the basic sequential pattern mining task. Sect. 5 presents encodings for alternative sequential pattern mining tasks, including the use of constraints and the extraction of condensed representations. After presenting some related works in Sect. 6, we present our experiments in Sect. 7.

2 ASP – Answer Set Programming

In this section we introduce the Answer Set Programming (ASP) paradigm, syntax and tools. Sect. 2.1 introduces the main principles and notations of ASP. Sect. 2.2 illustrates them on the well-known graph coloring problem.

2.1 Principles of Answer Set Programming

ASP is a declarative programming paradigm. From a general point of view, declarative programming gives a description of what is a problem instead of specifying how to solve it. Several declarative paradigms have been proposed, differing in the modelling formalism they use. For instance, logic programming [Lallouet et al., 2013] specifies the problem using a logic formalism, the SAT paradigm encodes the problem with boolean expressions [Biere et al., 2009], the CP (constraint programming) paradigm specifies the problem using constraints [Rossi et al., 2006]. ASP belongs to the class of logic programming paradigms, such as Prolog. The high-level syntax of logic formalisms makes generally the program easier to understand than other declarative programming paradigms.

An *ASP program* is a set of rules of the form

$$a_0 :- a_1, \dots, a_m, \text{not } a_{m+1}, \dots, \text{not } a_n. \quad (1)$$

where each a_i is a propositional atom for $0 \leq i \leq n$ and **not** stands for *default negation*. In the body of the rule, commas denote conjunctions between atoms. Contrary to Prolog, the order of atoms is meaningless. In ASP, rule (1) may be interpreted as “if a_1, \dots, a_m are all true and if none of a_{n+1}, \dots, a_n can be proved to be true, then a_0 is true.”

If $n = 0$, i.e. the rule body is empty, (1) is called a *fact* and the symbol “:-” may be omitted. Such a rule states that the atom a_0 has to be true. If a_0 is omitted, i.e. the rule head is empty, (1) represents an integrity constraint.

Semantically, a logic program induces a collection of so-called *answer sets*, which are distinguished models of the program determined by answer sets semantics; see [Gelfond and Lifschitz, 1991] for details. For short, a model assigns a truth value to each propositional atoms of the program and this set of assignments is valid. An answer set is a minimal set of true propositional atoms that satisfies all the program rules. Answer sets are said to be minimal in the way that only atoms that have to be true are actually true.

To facilitate the use of ASP in practice, several extensions have been developed. First of all, rules with *variables* are viewed as shorthands for the set of their ground instances. This allows for writing logic programs using a first order syntax. Such kind of syntax makes program shorter, but it hides the grounding step and its specific encoding issues, especially from the memory management point of view.

Further language constructs include *conditional literals* and *cardinality constraints* [Simons et al., 2002]. The former are of the form

$$a : b_1, \dots, b_m$$

the latter can be written as

$$s \{c_1; \dots; c_n\} t$$

where a and b_i are possibly default negated literals for $0 \leq i \leq m$, and each c_j is a conditional literal for $1 \leq j \leq n$. The purpose of conditional literals is to govern the instantiation of a literal a through the literals b_1, \dots, b_m . In a cardinality constraint, s (resp. t) provides the lower (resp. upper) bound on the number of literals from $c_1; \dots; c_n$ that must be satisfied in the model.

A cardinality constraint in the head of the rule defines a specific rule called a *choice rule*:

$$s \{c_1; \dots; c_n\} t :- a.$$

If a is true then all atoms of a subset $\mathcal{S} \subset \{c_1, \dots, c_n\}$ of size between s and t have to be true. All such subsets are admissible according to this unique rule, but not in the same model. All such subsets contribute to alternative answer sets. It should be noted that alternative models are solved independently. It is not possible to specify constraints that involve several models.

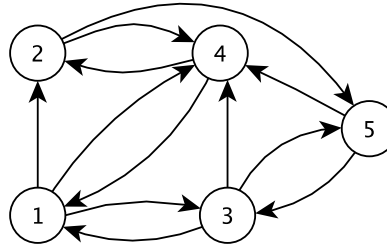


Fig. 1 An example graph for the graph coloring problem.

ASP problem solving is ensured by efficient solvers [Lifschitz, 2008] which are based on the same technologies as constraint programming solvers or satisfiability checking (SAT) solvers. *smodels* [Syrjänen and Niemelä, 2001], *dlv* [Leone et al., 2006], *ASPeRiX* [Lefèvre and Nicolas, 2009] or *clingo* [Gebser et al., 2011] are well-known ASP solvers. Due to the computational efficiency it has demonstrated and its broad application to real problems, we use *clingo* as a basic tool for designing our encodings.

The basic method for programming in ASP is to follow a *generate-and-test* methodology. Choice rules generate solution candidates, while integrity constraints are tested to eliminate those candidates that violate the constraints. The programmer should not have any concern about how solutions are generated. He/she just has to know that all possible solutions will be actually evaluated. From this point of view, the ASP programming principle is closer to CP programming than to Prolog programming. Similarly to these declarative programming approaches, the difficulty of programming in ASP lies in the choices for the best way to encode problem constraints: it must be seen as the definition of the search space (*generate* part) or as an additional constraint on solutions within this search space (*test* part). This choices may have a large impact on the efficiency of the problem encoding.

2.2 A simple example of ASP program

The following example illustrates the ASP syntax on encoding the graph coloring problem. Lines 9-10 specify the problem as general rules that solutions must satisfy while lines 1-6 give the input data that defines the problem instance related to the graph in Fig. 1.

The problem instance is a set of colors, encoded with predicates `col/1` and a graph, encoded with predicates `vertex/1` and `edge/2`. The input graph has 5 vertices numbered from 1 to 5 and 12 edges. The 5 fact-rules describing the vertex are listed in the same line (Line 2). It should be noted that, generally speaking, `edge`

```

1 % instance of the problem: the graph and colors
2 vertex(1). vertex(2). vertex(3). vertex(4). vertex(5).
3 edge(1,2). edge(1,3). edge(1,4). edge(2,4). edge(2,5).
4 edge(3,1). edge(3,4). edge(3,5). edge(4,1). edge(4,2).
5 edge(5,3). edge(5,4).
6 col(r). col(b). col(g).
7
8 % graph coloring problem specification
9 1 { color(X, C) : col(C) } 1 :- vertex(X).
10 :- edge(X, Y), color(X, C), color(Y, C).

```

Listing 1 Encoding of graph coloring – ASP syntax and encoding example

(1,2) is different from `edge(2,1)`, but considering that integrity constraints for the graph coloring problem are symmetric, it is sufficient to encode directed edge in only one direction. Line 6 encodes the 3 colors that can be used: `r`, `g` and `b`. Lower case letters represent values, internally encoded as integers, while strings beginning with upper case letters represent variables (see line 9 for instance).

Lines 9 and 10 specify the graph coloring problem. The predicate `color/2` encodes the color of a vertex: `color(X,C)` expresses that vertex `X` has color `C`. Line 10 is an integrity constraint. It forbids neighbor vertices `X` and `Y` to have the same color `C`¹. The ease of expressing such integrity constraints is a major feature of ASP.

Line 9 is a choice rule indicating that for a given vertex `X`, an answer set must contain exactly one atom of the form `color(X,C)` where `C` is a color. The grounded version of this rule is the following:

```

1 { color(1, r), color(1, b), color(1, g) } 1.
1 { color(2, r), color(2, b), color(2, g) } 1.
1 { color(3, r), color(3, b), color(3, g) } 1.
1 { color(4, r), color(4, b), color(4, g) } 1.
1 { color(5, r), color(5, b), color(5, g) } 1.

```

The variable `X` is expanded according to the facts in line 2 and for each vertex, a specific choice rule is defined. Within the brackets, the variable `C` is expanded according to the conditional expression in the rule head of line 9: the only admissible values for `C` are color values. For each line of the grounded version of the program, one and only one atom within brackets can be chosen. This corresponds to a unique mapping of a color to a vertex. Line 9 can be seen as a search space generator for the graph coloring problem.

The set `color(1,b) color(2,r) color(3,r) color(4,g) color(5,b)` is an answer set for the above program (among several others).

For more detailed presentation of ASP programming paradigm, we refer the reader to recent article of Janhunnen and Niemelä [Janhunnen and Niemelä, 2016].

¹ It is important to notice that the scope of a variable is the rule and each occurrence of a variable in a rule represents the same value.

2.3 The *Potassco* collection of ASP tools

The *Potassco* collection is a set of tools for ASP developed at the University of Potsdam. The main tool of the collection is the ASP solver *clingo* [Gebser et al., 2011]. This solver offers both a rich syntax to facilitate encodings² and a good solving efficiency. It is worth-noting that the ASP system *clingo* introduced many facilities to accelerate the encoding of ASP programs. For the sake of simplicity, we do not use them in the presented programs. A complete description of the *clingo* syntax can be found in [Gebser et al., 2014].

The *clingo* solving process follows two consecutive main steps:

1. *grounding* transforms the initial ASP program into a set of propositional clauses, cardinality constraints and optimisation clauses. Note that grounding is not simply a systematic problem transformation. It also simplifies the rules to generate the as short as possible equivalent grounded program.
2. *solving* consists in finding from one to all solutions of the grounded program. This step is performed by *clasp* which is a conflict-driven ASP solver. The primary *clasp* algorithm relies on conflict-driven *nogood* learning. It is further optimized using sophisticated reasoning and implementation techniques, some specific to ASP, others borrowed from CDCL-based SAT solvers.

The overall process may be controlled using procedural languages, e.g. *Python* or *lua* [Gebser et al., 2014]. These facilities are very useful to automate processes and to collect statistics on solved problems. Despite this procedural control which enables to interact with the grounder or the solver, it is important to note that once a program has been grounded, it cannot be changed.

3 Sequential pattern mining: definition and notations

Briefly, the sequential pattern mining problem consists in retrieving from a sequence database \mathcal{D} every frequent non empty sequence P , so-called a sequential pattern. Sequences, either in \mathcal{D} or sequential patterns P , are multiset sequences of itemsets over a fixed alphabet of symbols (also called items). A pattern is frequent if it is a subsequence of at least f_{min} sequences of \mathcal{D} , where f_{min} is an arbitrary given threshold. In this section, we introduce classical definitions and notations for frequent sequential pattern mining which will be useful to formulate the problem in an ASP setting. In the sequel, if not specified otherwise, a pattern is a sequential pattern.

² *clingo* is fully compliant with the recent ASP standard:
<https://www.mat.unical.it/aspcomp2013/ASPStandardization>

3.1 Sequences

We introduce here the basic definitions of sequences of itemsets. $[n] = \{1, \dots, n\}$ denotes the set of the first n strictly positive integers.

Let $(\mathcal{I}, =, <)$ be the set of items (alphabet) with a total order (e.g. lexicographic order). An *itemset* $A = (a_1, a_2, \dots, a_n)$ is a subset of distinct increasingly ordered items from \mathcal{I} : $\forall i \in [n], a_i \in \mathcal{I} \wedge \forall i \in [n-1], a_i < a_{i+1} \wedge \forall i, j \in [n], i \neq j \implies a_i \neq a_j$. An itemset $\beta = (b_i)_{i \in [m]}$ is a sub-itemset of $\alpha = (a_i)_{i \in [n]}$, denoted $\beta \sqsubseteq \alpha$, iff there exists a sequence of m integers $1 \leq i_1 < i_2 < \dots < i_m \leq n$ such that $\forall k \in [m], b_k = a_{i_k}$. A *sequence* S is an ordered set of itemsets $S = \langle s_1, s_2, \dots, s_n \rangle$: $\forall i, j \in [n], i < j$ means that s_i occurs before s_j . The length of sequence S , denoted $|S|$, is equal to its number of itemsets. Two sequences $S = \langle s_1, s_2, \dots, s_n \rangle$ and $T = \langle t_1, t_2, \dots, t_m \rangle$ are equal iff $n = m$ and $\forall i \in [n], s_i = t_i$.

$T = \langle t_1, t_2, \dots, t_m \rangle$ is a *sub-sequence* of $S = \langle s_1, s_2, \dots, s_n \rangle$, denoted $T \preceq S$, iff there exists a sequence of integers $1 \leq i_1 < i_2 < \dots < i_m \leq n$ such that $\forall k \in [m], t_k \sqsubseteq s_{i_k}$. In other words, $(i_k)_{1 \leq k \leq m}$ defines a mapping from $[m]$, the set of indexes of T , to $[n]$, the set of indexes of S . We denote by $T \prec S$ the strict sub-sequence relation such that $T \preceq S$ and $T \neq S$.

$T = \langle t_1, t_2, \dots, t_m \rangle$ is a *prefix* of $S = \langle s_1, s_2, \dots, s_n \rangle$, denoted $T \preceq_b S$, iff $\forall i \in [k-1], t_i = s_i$ and $t_m \sqsubseteq s_m$. Thus, we have $T \preceq_b S \implies T \preceq S$.

A sequence T is *supported* by a sequence S if T is a sub-sequence of S , i.e. $T \preceq S$.

Example 1 (Sequences, subsequences and prefixes). Let $\mathcal{I} = \{a, b, c\}$ with a lexicographic order ($a < b, b < c$) and the sequence $S = \langle a(bc)(abc)cb \rangle$. To simplify the notation, parentheses are omitted around itemsets containing only one item. The length of S is 5. $\langle abb \rangle$, $\langle (bc)(ac) \rangle$ or $\langle a(bc)(abc)cb \rangle$ are sub-sequences of S . a , $\langle a(bc) \rangle$ and $\langle a(bc)a \rangle$ are prefixes of S .

Proposition 1 \prec and \prec_b induces two partial orders on the sequence set. For all sequences (s, s') , $s \prec_b s' \implies s \prec s'$.

3.2 Sequential pattern mining

Let $\mathcal{D} = \{S_1, S_2, \dots, S_N\}$ be a set of sequences. \mathcal{D} is called a *sequence database*. The *support of a sequence S* in \mathcal{D} , denoted by $\text{supp}_{\mathcal{D}}(S)$, is the number of sequences of \mathcal{D} that support S :

$$\text{supp}_{\mathcal{D}}(S) = |\{S_i \in \mathcal{D} \mid S \prec S_i\}|$$

Proposition 2 $\text{supp}_{\mathcal{D}}(\cdot)$ is an anti-monotonic measure on the set of subsequences of a sequence database \mathcal{D} structured by \prec or \prec_b .

This proposition implies that for all pairs of sequences P, Q :

$$P \prec Q \implies \text{supp}_{\mathcal{D}}(P) \geq \text{supp}_{\mathcal{D}}(Q), \text{ and}$$

$$P \prec_b Q \implies \text{supp}_{\mathcal{D}}(P) \geq \text{supp}_{\mathcal{D}}(Q).$$

Let f_{min} be a *frequency threshold* defined by the analyst. For any sequence S , if $\text{supp}_{\mathcal{D}}(S) \geq f_{min}$, we say that S is a *frequent sequence* or a (*frequent*) *sequential pattern* of \mathcal{D} . *Mining sequential patterns* consists in extracting all frequent subsequences in a sequence database \mathcal{D} .

Every pattern mining algorithm [Agrawal and Srikant, 1995, Wang and Han, 2004, Pei et al., 2007] uses the anti-monotonicity property to browse efficiently the pattern search space. In fact, this property ensures that a sequence P including a sequence Q which is not frequent, cannot be frequent itself. So, the main idea of classical algorithms is to extend the patterns until identifying a non frequent pattern.

Example 2 (Sequential pattern mining). To illustrate the concepts introduced above, we consider the following sequence database \mathcal{D} containing sequences built on items in $\mathcal{I} = \{a, b, c, d\}$ such that $a < b$, $b < c$ and $a < c$. In this running example, and in the rest of this article, we focus on *simple* sequences of items instead of sequences of itemsets.

SeqId	Sequence
1	$\langle ac \rangle$
2	$\langle dabc \rangle$
3	$\langle b \rangle$
4	$\langle abc \rangle$
5	$\langle ab \rangle$
6	$\langle acbc \rangle$
7	$\langle abc \rangle$

Given a threshold value $f_{min} = 3$, the frequent sequential patterns are: $\langle a \rangle$, $\langle b \rangle$, $\langle c \rangle$, $\langle ac \rangle$, $\langle bc \rangle$, $\langle ab \rangle$ and $\langle abc \rangle$.

It is interesting to relate the sequential pattern mining task with the presentation of ASP principles. The sequential pattern mining task rises two issues: 1) exploring a large search space, i.e. the potentially infinite set of sequences and 2) assessing the frequency constraint (with respect to the given database). Thus, sequential pattern mining can be considered as a *generate and test* process which makes it straightforward to encode the mining task using ASP principles: 1) choice rules will define the search space and 2) the frequency assessment will be encoded using integrity constraints.

4 Frequent sequential pattern mining with ASP

In this section, we present several ASP encodings for sequential pattern mining. We assume that the database contains sequences of itemsets. But for the sake of simplicity, we will restrict patterns to sequences of items (each itemset is a singleton).

Listing 10 in Appendices gives an encoding for the full general case of sequential pattern mining.

Our proposal is borrowed from Järvisalo’s [Järvisalo, 2011]: the solution of the sequential pattern mining ASP program is all the answer sets (AS), each of which contains the atoms describing a single frequent pattern as well as its occurrences in database sequences. The solution relies on the “generate and test principle”: generate combinatorially all the possible patterns and their related occurrences in the database sequences and test whether they satisfy the specified constraints.

4.1 Modelling database, patterns and parameters

A sequence database \mathcal{D} is modelled by the predicate $\text{seq}(\text{T}, \text{I}_s, \text{I})$ which holds if sequence T contains item I at index I_s .

Example 3 (A sequence database encoded in ASP). The following facts encode the database of Example 2:

```
seq(1, 1, a) . seq(1, 2, c) .
seq(2, 1, d) . seq(2, 2, a) . seq(2, 3, b) . seq(2, 4, c) .
seq(3, 1, b) .
seq(4, 1, a) . seq(4, 2, b) . seq(4, 3, c) .
seq(5, 1, a) . seq(5, 2, b) .
seq(6, 1, a) . seq(6, 2, c) . seq(6, 3, b) . seq(6, 4, c) .
seq(7, 1, a) . seq(7, 2, b) . seq(7, 3, c) .
```

Similarly, the current pattern is modelled by the predicate $\text{pat}(\text{I}_p, \text{I})$ which holds if the current pattern contains item I at index I_p .

For example, the pattern $\langle abc \rangle$ is modelled by the following atoms:

```
pat(1, a) . pat(2, b) . pat(3, c) .
```

In addition, we define two program constants:

- `#const th=23`. represents f_{min} , the minimal frequency threshold, i.e. the requested minimal number of supporting sequences
- `#const maxlen=10`. represents the maximal pattern length.

Let S be a set of ground atoms and $P \subset S$ the set of $\text{pat}(\text{I}_p, \text{I})$ atoms in S , according to the Järvisalo’s encoding principle we would like an ASP program π such that S is an answer set of π iff the pattern defined by P is a frequent sequential pattern in the database \mathcal{D} .

4.2 Two encodings for sequential pattern mining

The main difficulty in declarative sequential pattern mining is to decide whether a pattern $P = \langle p_1, \dots, p_n \rangle$ supports a sequence $S = \langle s_1, \dots, s_m \rangle$ of the database. According to Def. 1, it means that it exists a mapping $e = (e_i)_{1 \leq i \leq n}$ such that $p_i = s_{e_i}$.

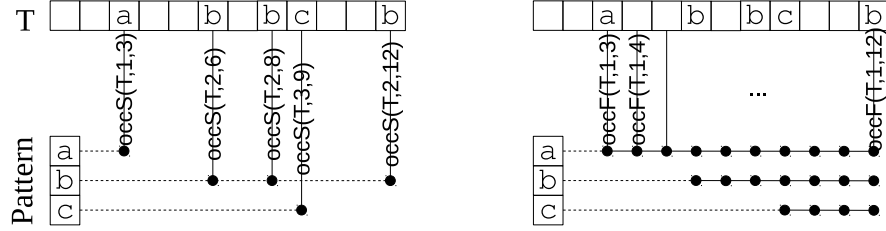


Fig. 2 Illustration of embeddings strategies. On the left, *skip-gaps*, on the right, *fill-gaps*.

Unfortunately, this definition is not usable in practice to implement an efficient ASP encodings. The difficulty comes from the possible multiple mappings of a pattern in a single sequence. On the other hand, the detailed mapping description is not required here, we simply have to defined embeddings that exists iff a pattern supports a sequence. An embedding of a pattern in a sequence is given by the description of a relation between pattern item indexes to sequence item indexes.

This section presents two encodings of sequential pattern mining. These two encodings differ in their representation of embeddings, as illustrated in Fig. 2. Two embedding strategies have been defined and compared in our results: *skip-gaps* and *fill-gaps*.

More formally, let $P = \langle p_1, \dots, p_n \rangle$ be a pattern sequence and $S_T = \langle s_1, \dots, s_m \rangle$ be the T -th sequence of \mathcal{D} . In the *skip-gaps* strategy, an embedding \mathcal{E} is a relation over $[1, m] \times [1, n]$ such that $\forall (i, j) \in \mathcal{E}, i \leq j \wedge p_i = s_j$ and $\forall (i, j), (i', j') \in \mathcal{E}, i < i' \implies j < j'$. In the *fill-gaps* strategy, an embedding \mathcal{E}' is the same relation as \mathcal{E} (i.e. $(i, j) \in \mathcal{E} \implies (i, j) \in \mathcal{E}'$) with the additional specification: $\forall i \in [1, m], \forall j, j' \in [1, n], (i, j) \in \mathcal{E}' \wedge j < j' \implies (i, j') \in \mathcal{E}'$. This additional specification expresses that once a pattern item has been mapped to the leftmost (having the lowest index, let it be j), the knowledge of this mapping is maintained on remaining sequences items with indexes $j' > j$. So, a *fill-gaps* embedding makes only explicit the “leftmost admissible matches” of P items in sequence S_T .

Relations \mathcal{E} and \mathcal{E}' are interesting because (i) that can be computed in a constructive way (i.e. without encoding guesses) and (ii) they contains the information required to decide whether the pattern supports the sequence.

The two following sections detail the ASP programs for extracting patterns under each embedding strategy.

4.2.1 The skip-gaps approach

In the first ASP encoding, an embedding of the current pattern $P = \langle p_i \rangle_{1 \leq i \leq n}$ in sequence $T = \langle s_i \rangle_{1 \leq i \leq m}$ is described by a set of atoms $\text{occS}(T, I_p, I_s)$ which holds if I_p -th the pattern item (occurring at index I_p) is identical to the I_s -th item in sequence T (formally, $p_{I_p} = s_{I_s} \wedge \langle p_i \rangle_{1 \leq i \leq I_p} \prec \langle s_i \rangle_{1 \leq i \leq I_s}$). The set of valid atoms $\text{occS}(T, _, _)$ encodes the relation \mathcal{E} above and is illustrated in Fig. 2 (on the left).

```

1 item(I) :- seq(_, _, I).
2
3 %sequential pattern generation
4 patpos(1).
5 0 { patpos(Ip+1) } 1 :- patpos(Ip), Ip<maxlen.
6 patlen(L) :- patpos(L), not patpos(L+1).
7
8 1 { pat(Ip,I): item(I) } 1 :- patpos(Ip).
9
10 %pattern embeddings
11 occS(T,1,_,Is):- seq(T,Is,I), pat(1,_,I).
12 occS(T,Ip+1,Is):- seq(T,Is,I), pat(Ip+1,I), occS(T,Ip,Js), Js<Is.
13
14 %frequency constraint
15 support(T) :- occS(T,L,_), patlen(L).
16 :- { support(T) } < th.

```

Listing 2 Encoding of frequent sequential pattern mining – skip-gaps strategy

Example 4 (Illustration of skip-gaps embedding approach). Let $P = \langle ac \rangle$ be a pattern represented by $\text{pat}(1, a) . \text{pat}(2, c)$. Here follows, the embeddings of pattern P in the sequences of example 2:

```

occS(1,1,1) occS(1,2,2)
occS(2,1,2) occS(2,2,4)
occS(4,1,1) occS(4,2,3)
occS(5,1,1)
occS(6,1,1) occS(6,2,2) occS(6,2,4)
occS(7,1,1) occS(7,2,3)

```

The pattern could not be fully identified in the fifth sequence. There are two possible embeddings in the sixth sequence. Atom $\text{occS}(6,1,1)$ is used for both. Nonetheless, this sequence must be counted only once in the support.

Listing 2 gives the ASP program for sequential pattern mining. The first line of the program is called a *projection*. It defines a new predicate that provides all items from the database. The symbol “_” denotes an anonymous (or don’t care) variable.

Lines 4 to 8 of the program encode the pattern generation. Predicate $\text{patpos}/1$ defines the allowed sequential pattern indexes, beginning at index 1 (line 4). Line 5 is a choice rule that generates the successive pattern positions up to an ending index iterating from 2 to maxlen : $\text{patpos}(Ip+1)$ is true if there is a pattern position at index Ip and Ip is lower than maxlen . Line 6 defines the length of a pattern: $\text{patlen}(L)$ holds if L is the index of the last pattern item (there is no pattern item with a greater index). This predicate is used to decide whether an embedding has been completed or not. Finally, line 8 is a choice rule that associates exactly one item with each position x . We can note that each possible sequence is generated once and only once. So, there is no redundancy in the search space exploration.

Lines 11 to 12 encode pattern embedding search. Line 11 guesses a sequence index for the first pattern item: $\text{occS}(T,1,Is)$ holds if the first pattern item is iden-

tical to the I_S -th of sequence T (i.e. $p_1 = s_{I_S}$). Line 12 guesses sequence indexes for pattern items at indexes strictly greater than 1. $\text{occS}(T, I_P, I_S)$ holds if the I_P -th pattern item is equal to the I_S -th sequence item (i.e. $p_{I_P} = s_{I_S}$) and the preceding pattern item is mapped to a sequence item at an index strictly lower than I_S . Formally, this rule expresses the following implication $(J_P, I_S - 1) \in \mathcal{E} \wedge p_{I_P} = s_{I_S} \wedge I_P > J_P \implies (I_P, I_S) \in \mathcal{E}$ and recursively, we have $\langle p_i \rangle_{1 \leq i \leq I_P} \prec \langle s_i \rangle_{1 \leq i \leq I_S}$. It should be noted that this encoding generates all the possible embeddings of some pattern.

Finally, lines 15 to 16 are dedicated to assess the pattern frequency constraint. $\text{support}(T)$ holds if the database sequence T supports the pattern, i.e. if an atom occS holds for the last pattern position. The last line of the program is an integrity constraint ensuring that the number of supported sequences is not lower than the threshold th or, in other words, that the support of the current pattern is greater than or equal to the threshold.

4.2.2 The fill-gaps approach

In the *fill-gap* approach, an embedding of the current pattern P is described by a set of atoms $\text{occF}(T, I_P, I_S)$ having a slightly different semantics than in the *skip-gap* approach. $\text{occF}(T, I_P, I_S)$ holds if at sequence index I_S it is true that the I_P -th pattern item has been mapped (to some sequence index equal to I_S or lower than I_S if $\text{occF}(T, I_P, I_S - 1)$ holds). More formally, we have $\langle p_i \rangle_{1 \leq i \leq I_P} \prec \langle s_i \rangle_{1 \leq i \leq I_S}$. The set of atoms $\text{occF}(T, _, _)$ encodes the relation \mathcal{E}^f above and is illustrated in Fig. 2 (on the right).

Example 5 (Fill-gaps approach embedding example). Pattern $P = \langle a, c \rangle$ has the following fill-gaps embeddings (represented by occF atoms) in the sequences of the database of example 2:

```

occF(1, 1, 1)  occF(1, 1, 2)  occF(1, 2, 2)
occF(2, 1, 2)  occF(2, 1, 3)  occF(2, 1, 4)  occF(2, 2, 4)
occF(4, 1, 1)  occF(4, 1, 2)  occF(4, 1, 3)  occF(4, 2, 3)
occF(5, 1, 1)  occF(5, 1, 2)
occF(6, 1, 1)  occF(6, 1, 2)  occF(6, 1, 3)  occF(6, 1, 4)  occF(6, 2, 2)  occF
(6, 2, 3)  occF(6, 2, 4)
occF(7, 1, 1)  occF(7, 1, 2)  occF(7, 1, 3)  occF(7, 2, 3)

```

Contrary to the skip-gap approach example (see Example 4), the set of $\text{occF}(T, I_P, I_S)$ atoms alone is not sufficient to deduce all occurrences. For instance, occurrence with indexes (3, 8, 9) is masked.

Listing 3 gives the ASP program for sequential pattern mining with the fill-gaps strategy. The rules are quite similar to those encoding the skip-gaps method. The main difference comes from the computation of embeddings (lines 11-13). As in listing 2, line 11 guesses a sequence index for the first pattern item: $\text{occF}(T, 1, I_S)$ holds if the first pattern item is identical to the I_S -th of sequence T (i.e. $p_{I_P} = s_{I_S}$).

Line 12 guesses sequence indexes for pattern items at indexes strictly greater than 1. $\text{occS}(T, I_P, I_S)$ holds if the I_P -th pattern item is equal to the I_S -th sequence

```

10 %pattern embeddings
11 occF(T, I, Is) :- seq(T, Is, I), pat(I, I).
12 occF(T, Ip, Is) :- occF(T, Ip-1, Is-1), seq(T, Is, I), pat(I, I).
13 occF(T, Ip, Is) :- occF(T, Ip, Is-1), seq(T, Is, _).
14
15 %frequent constraint
16 seqLen(T, L) :- seq(T, L, _), not seq(T, L+1, _).
17 support(T) :- occF(T, L, LS), patLen(L), seqLen(T, LS).
18 :- { support(T) } < th.

```

Listing 3 Encoding of frequent sequential pattern mining – fill-gaps strategy (see Listing 2 for first lines).

item and the preceding pattern item is mapped to some sequence item at some index strictly lower than I_s . More formally, we have that $p_{I_p} = s_{I_s} \wedge (I_p - 1, I_s - 1) \in \mathcal{E}' \implies (I_p, I_s) \in \mathcal{E}'$.

Line 13 simply maintains the knowledge that the I_p -th pattern item has been mapped all along the further sequence indexes, i.e. $\text{occF}(T, I_p, I_s)$ holds if $\text{occF}(T, I_p, I_s - 1)$ holds. More formally, $(I_p - 1, I_s - 1) \in \mathcal{E}' \implies (I_p, I_s) \in \mathcal{E}'$. In combination with previous rules, we thus have recursively that $\text{occF}(T, I_p, I_s)$ is equivalent to $\langle p_i \rangle_{1 \leq i \leq I_p} \prec \langle s_i \rangle_{1 \leq i \leq I_s}$.

Line 17 a sequence is supported by the pattern an occF atoms exists at the last position LS of the sequence, computed line 16. The remaining rules for testing whether it is greater than the threshold th are identical to those in the skip-gaps approach.

4.3 Sequential pattern mining improvements

The main objective of this subsection is to present alternative encodings of the sequential pattern mining task. These encodings attempt to take advantage of known properties of the sequential pattern mining task to support the solver to mine datasets more efficiently or with less memory requirements. The efficiency of these improvements will be compared in the experimental part.

4.3.1 Filter out unfrequent items

The first improvement consists in generating patterns from only frequent items. According to the anti-monotonicity property, all items in a pattern have to be frequent. The rules in listing 4 may replace the projection rule previously defining the available items. Instead, an explicit aggregate argument is introduced to evaluate the frequency of each item I and to prune it if it is unfrequent.

In the new encoding, the predicate $\text{item}/1$ defines the set of items that occurs in the database and $\text{item}/1$ defines the frequent items that can generate patterns.

```

1 sitem(I) :- seq(_, _, I).
2 item(I)  :- sitem(I), #count{ T:seq(T,_, I) } >= th.

```

Listing 4 Restriction of the pattern traversal to sequences made of frequent items only

4.3.2 Using projected databases

The idea of this alternative encoding is to use the principle of *projected databases* introduced by algorithm PrefixSpan [Pei et al., 2004]. Let $P = \langle p_1, \dots, p_n \rangle$ be a pattern, the projected database of $\mathcal{D} = \{S_1, \dots, S_n\}$ is $\{S'_1, \dots, S'_n\}$ where S'_i is the projected sequence of S_i with respect to P . Let $S_i = \langle s_1, \dots, s_m \rangle$ be a sequence. Then the projected sequence of S_i is $S'_i = \langle s_{k+1}, \dots, s_m \rangle$ where k is the position of the last item of the first occurrence of P in S_i . If P does not occur in S_i then S'_i is empty.

A projected database is smaller than the whole database and the set of its frequent items is consequently much smaller than the original set of frequent items. The idea is to improve the candidate generation part of the algorithm by making use of items from projected databases. Instead of generating a candidate (a sequential pattern) by extending a frequent pattern with an item that is frequent in the whole database, the pattern extension operation uses only the items that are frequent in the database projected along this pattern.

```

8 item(I, I) :- sitem(I),
9             #count{ T: seq(T,_, I) } >= th.
10 item(Ip+1, I) :- item(Ip, I),
11                #count{ T: seq(T, Js, I), occS(T, Ip, Is), Js>Is } >= th.
12 1 { pat(Ip, I) : item(Ip, I) } 1 :- patpos(Ip).

```

Listing 5 Pattern generation with prefix-projection principle – skip-gaps strategy

```

8 item(I, I) :- sitem(I),
9             #count{ T: seq(T,_, I) } >= th.
10 item(Ip+1, I) :- item(Ip, I),
11                #count{ T: seq(T, Is, I), occF(T, Ip, Is) } >= th.
12 1 { pat(Ip, I) : item(Ip, I) } 1 :- patpos(Ip).

```

Listing 6 Pattern generation with prefix-projection principle – fill-gaps strategy

The ASP encoding of the prefix-projection principle is given in Listing 5 for the skip-gaps strategy and in Listing 6 for the fill-gaps strategy. The programs of Listings 2 and 3 remain the same except for the generation of patterns defined by `patpos/1` and the new predicate `item/2`. `item(Ip, I)` defines an item I that is frequent in sequence suffixes remaining after removing the prefix of the sequence containing the first occurrence of the $X-1$ -pattern prefix (consisting of the I_p-1 first positions of the pattern). Lines 8-9 are similar to those in Listing 4. `item(I, I)` defines the frequent items, i.e. those that are admissible as first item of a frequent pattern. Lines 10-11 generates the admissible items for pattern position I_p+1 . Such an item must be admissible for position I_p and be frequent in sequence suffixes (sub-sequence after at least one (prefix) pattern embedding). For skip-gaps, the sequence suffix is defined by `seq(T, Js, I)`, `occS(T, Ip, Is)`, `Js>Is` (the items at


```

1 sitem(I) :- seq(_, _, I).
2
3 % extract frequent itemsets
4 db(T, I) :- seq(T, _, I).
5 0 { in_itemset(I) } 1 :- th { in_support(T):db(T, I) }, sitem(I).
6 in_support(T) :- 0 { conflict_at(T, I) : sitem(I) } 0, db(T, _).
7 conflict_at(T, I) :- not db(T, I), in_itemset(I), db(T, _).
8
9 %sequential pattern generation from frequent itemsets
10 patpos(1).
11 0 { patpos(Ip+1) } 1 :- patpos(Ip), Ip < maxlen.
12 patlen(L) :- patpos(L), not patpos(L+1).
13
14 1 { pat(Ip, I) : in_itemset(I) } 1 :- patpos(Ip).
15
16 :- 0 { pat(Ip, I) : patpos(Ip) } 0, in_itemset(I).
17
18 %pattern occurrences
19 occS(T, l, Is) :- in_support(T), seq(T, Is, I), pat(l, I).
20 occS(T, Ip+1, Is) :- occS(T, Ip, Js), seq(T, Is, I), pat(Ip+1, I), Js < Is.
21
22 support(T) :- occS(T, L, _), patlen(L).
23 :- { support(T) } < th.

```

Listing 7 Mining frequent sequences from frequent itemsets – skip-gaps strategy

sequence positions farther away than the last position that matches the last (partial) pattern item at position I_p). For fill-gaps, $\text{seq}(T, J_s, I)$, $\text{occF}(T, I_p, I_s)$ is sufficient because $\text{occF}(T, I_p, I_s)$ atoms represent the sequence suffix beginning at the sequence position that matches the last (partial) pattern item (at position I_p).

4.3.3 Mixing itemsets and sequences mining

In [Järvisalo, 2011], Järvisalo showed that ASP can be efficient for itemset pattern mining. The main idea of this last alternative approach is to mine frequent itemsets and to derive sequential patterns from them.

This time, the itemset mining step extracts a frequent itemset pattern $I = (e_i)_{i \in [n]}$, $e_i \in \mathcal{I}$. A sequential pattern $S = (s_i)_{i \in [m]}$ is generated using the items of the itemset, i.e. $\forall i \in [m], \exists j \in [n], s_i = e_j$ taking into account that items may be repeated within a sequential pattern and that every item from I must appear in S . If not, there would exist a subset $J \subset I$ that would generate the same sequence s . This would lead to numerous redundant answer sets for similar frequent sequences and would cause a performance drop.

Listing 7 gives the entire encoding of this alternative for the skip-gaps strategy³. Rules in Lines 4-7 extract frequent itemsets, represented by the predicate

³ A similar encoding can be done for the fill-gaps strategy applying the same changes as above.

`in_itemset/1`, borrowed from Järvisalo’s encoding [Järvisalo, 2011]. Next, the generation of sequential patterns in line 14 uses only items from such a frequent itemset. Line 16 defines a constraint required to avoid answer set redundancies. The remaining part of the program is left unchanged.

5 Alternative sequential pattern mining tasks

In this section, we illustrate how the previous encodings can be modified to solve more complex mining tasks. Our objective is to show the impressive expressiveness of ASP which let us encode a very wide range of mining tasks. We focus our attention on the most classical alternative sequential pattern mining tasks: constrained sequential patterns and condensed representation of sequential patterns.

In [Negrevergne and Guns, 2015], the authors organize the constraints on sequential patterns in three categories: 1) constraints on patterns, 2) constraints on patterns embeddings, 3) constraints on pattern sets. These constraints are provided by the user and capture his background knowledge.

The following subsection shows that our ASP approach enables to add constraints on individual patterns (constraints of categories 1 and 2). But, as ASP cannot compare models with each others, the third category of constraints can not be encoded directly.

In Sect. 5.2, we transform the classical definition of the most known constraints of the third category – the condensed representations – to encode them in pure ASP. Condensed representations (maximal and closed patterns) have been widely studied due to their monotonicity property, and to their representativeness with respect to frequent patterns. Concerning more generic constraints on pattern sets, such as the extraction of skypatterns [Ugarte et al., 2015], we have proposed in [Gebser et al., 2016] an ASP-based approach for mining sequential skypatterns using *asprin* for expressing preferences on answer sets. *asprin* [Brewka et al., 2015] provides a generic framework for implementing a broad range of preferences relations on ASP models and can easily manage them. This approach is out of the scope of this article.

5.1 Constraints on patterns and embeddings

Pei et al. [Pei et al., 2007] defined seven types of constraints on patterns and embeddings. In this subsection, we describe each of these constraints keeping their original numbering. Constraints 1, 2, 3 and 5 are pattern constraints, while constraints 4, 6 and 7 are embedding constraints. If not stated otherwise, the base encoding is the skip-gaps strategy and line numbers refers to Listing 2.

In a first approach, constraints on patterns and on embeddings may be trivially encoded by adding integrity constraints. But these integrity constraints acts a posteriori, during the test stage, for invalidating candidate models. A more efficient method

consists in introducing constraints in the generate stage, specifically in choice rules, for pruning the search space early.

Constraint 1 – item constraint. An item constraint specifies what are the particular individual or groups of items that should or should not be present in the patterns. For instance, the constraint “patterns must contain item 1 but not item 2 nor item 3” can be encoded using `must_have/1` and `cannot_have/1` predicates: `must_have(1). cannot_have(2). cannot_have(3).`

To cope with this kind of constraint, Line 8 of Listing 2 is modified as:

```
8 1 { pat(X,I): item(I), not cannot_have(I) } 1 :- patpos(X).
9 :- { pat(X,I) : must_have(I) } < 1.
```

The encoding of Line 8 modifies the choice rule to avoid the generation of known invalid patterns, i.e. patterns with forbidden items. Line 9 is a new constraint that imposes to have at least one of the required items.

Constraint 2 – length constraint. A length constraint specifies a prerequisite on pattern length. The maximal length constraint is anti-monotonic while the minimal length is not anti-monotonic. The maximal length constraint is already encoded using the program constant `maxlen` in our encodings. A new constant `minlen` is defined to encode the minimal length constraint and a new rule is added to predicate `patpos/1` to impose at least `minlen` positions in patterns instead of only one.

```
#const minlen = 2.
patpos(1).
patpos(X+1) :- patpos(X), X<=minlen.
0 { patpos(X+1) } 1 :- patpos(X), X<maxlen.
```

Constraint 3 – super-pattern constraint. A super-pattern constraint enforces the extraction of patterns that contain one or more given sub-patterns. Mandatory sub-patterns are defined by means of the new predicate `subpat(SP,P,I)` expressing that sub-pattern `SP` contains item `I` at position `P`.

Predicate `issubpat(SP)` verifies that the sub-pattern `SP` is included in the pattern. An approach similar to embedding computation may be used:

```
issubpat(SP,1,P) :- pat(P,I), subpat(SP,1,I).
issubpat(SP,Pos+1,P) :- issubpat(SP,Pos,Q), pat(P,I),
                        subpat(SP,Pos+1,I), Q<P.
issubpat(SP) :- issubpat(SP,L,_), subpatlen(SP,L).
```

`issubpat(SP)` is true if the sub-pattern `SP` is a sub-pattern of the current pattern. This predicate is used to define the final integrity constraint:

```
:- #count{ SP : issubpat(SP), subpat(SP,_,_) } = 0.
```

Constraint 4 – aggregate constraint. An aggregate constraint is a constraint on an aggregation of items in a pattern, where the aggregate function can be *sum*, *avg*, *max*, *min*, *standard deviation*, etc. The only aggregates that are provided by *clingo* are `#sum`, `#max` and `#min`. For example, let us assume that to each item `I` is assigned a cost `C`, which is given by predicate `cost(I,C)`. The following constraint enforces the selection of patterns having a total cost of at least 1000.

```
:- #sum{ C,X : cost(I,C), pat(X,I) } < 1000.
```

As an integrity constraint, this rule means that it is not possible to have a total amount lower than 1000 for pattern. It should be noted that c values are summed for each pair (C, X) . Thus, item repetitions are taken into account.

Constraint 5 – Regular expression. Such a constraint is satisfied if the pattern is an accepted regular expression as stated by the user. A regular expression can be encoded in ASP as its equivalent deterministic finite automata. Expressing such a constraint is mainly technical and is not detailed here. SPIRIT [Garofalakis et al., 1999] is one of the rare algorithms that considers complex pattern constraints expressed as regular expressions.

Constraint 6 – Duration constraints. The duration (or span) of some pattern is the difference between its last item timestamp and its first item timestamp. A duration constraint requires that the pattern duration should be longer or shorter than a given time period. In the database encoding introduced Sect. 4.1, predicate $\text{seq}(T, P, I)$ defines the timestamp of I in sequence T as the integer position P . A global constraint such as *max-span* cannot be expressed through simple local constraints on successive pattern item occurrences, as gap constraints described in the next paragraph. In fact, the predicate $\text{occS}/3$ does not describe the embeddings precisely enough to express the *max-span* constraint: for some pattern embedding, there is no explicit link between its first item occurrence and its last item occurrence. The proposed solution is to add an argument to $\text{occS}/3$ to denote the position of the occurrence of the first pattern item:

```

11 %pattern embeddings (skip-gaps strategy)
12 occS(T, 1, P, P)      :- seq(T, P, I), pat(1, I).
13 occS(T, Pos+1, P, IP) :- occS(T, Pos, Q, IP), seq(T, P, I), pat(Pos+1, I),
14                        P-IP+1<=maxspan, P-IP+1>=minspan.

```

Constraint 7 – Gap constraints. A gap constraint specifies the maximal/minimal number of positions (or timestamp difference) between two successive itemsets in an embedding. The maximal gap constraint is anti-monotonic while the minimal gap is not anti-monotonic. Contrary to pattern constraints, embedding constraints cannot be encoded simply by integrity constraints. In fact, an integrity constraint imposes a constraint on all embeddings. If an embedding does not satisfy the constraint, the whole interpretation – i.e. the pattern – is unsatisfied.

In the following we give an encoding for the *max-gap* and *min-gap* constraints. For such local constraint, the solution consists in modifying the embedding generation (lines 11-12 in Listing 2) for yielding only embeddings that satisfy gap constraints:

```

11 occS(T, 1, P)      :- seq(T, P, I), pat(1, I).
12 occS(T, Pos+1, P) :- seq(T, P, I), pat(Pos+1, I), occS(T, Pos, Q),
13                    P-Q-1>=mingap, P-Q-1<=maxgap.

```

This encoding assumes that the value of constants *mingap* and *maxgap* have been provided by the user (using `#const` statements).

Constraints of type 6 and 7 can be mixed by merging the two encodings of occS above:

```

11 occS(T, I, P, P)      :- seq(T, P, I), pat(I, I).
12 occS(T, Pos+1, P, IP) :- seq(T, P, I), pat(Pos+1, I), occS(T, Pos, Q, IP),
13                               P-Q-1>=mingap, P-Q-1<=maxgap,
14                               P-IP+1<=maxspan, P-IP+1>=minspan.

```

5.2 Condensed representation of patterns: closed and maximal sequences

In this section, we study the encodings for two well-studied pattern types, closed and maximal patterns. A closed pattern is such that none of its frequent super-patterns has the same support. A maximal pattern is such that none of its super-patterns is frequent. Thus, it is necessary to compare the supports of several distinct patterns. Since a solution pattern is encoded through an answer set, a simple solution would be to put constraints on sets of answer sets. However, such a facility is not provided by basic ASP language⁴. So, these constraints have been encoded without any comparison of answer sets but as additional constraints on the requested patterns. The next section introduces the definitions of these alternative mining tasks and the properties that were used to transform the pattern set constraints as constraints on individual patterns. Sect. 5.2.2 gives encodings for closed and maximal patterns extraction.

5.2.1 Definitions and properties

A frequent pattern S is *maximal* (resp. *backward-maximal*) with respect to the relation \prec (resp. \prec_b) iff there is no other frequent pattern S' such that $S \prec S'$ (resp. $S \prec_b S'$).

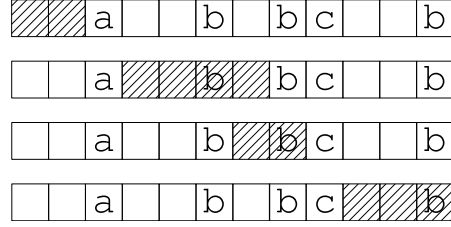
A frequent pattern S is *closed* (resp. *backward-closed*) with respect to the relation \prec (resp. \prec_b) iff there is no proper superpattern S' such that $S \prec S'$ (resp. $S \prec_b S'$) and $\text{supp}(S) = \text{supp}(S')$. Mining the closed patterns significantly reduces the number of patterns without loss of information for the analyst. Having the closed patterns and their support, the support of any pattern can be computed. This is not the case for maximal patterns.

Example 6 (Maximal and closed-sequential pattern mining). Considering the database of Example 2, among the frequent patterns with $f_{min} = 3$, the only maximal pattern is $\langle abc \rangle$. The set of backward-maximal is $\{\langle c \rangle, \langle bc \rangle, \langle ac \rangle, \langle abc \rangle\}$.

The set of closed patterns is $\{\langle a \rangle, \langle b \rangle, \langle ab \rangle, \langle ac \rangle, \langle abc \rangle\}$. $\langle bc \rangle$ is not closed because in any sequence it occurs, it is preceded by an a . Thus $\text{supp}(\langle bc \rangle) = \text{supp}(\langle abc \rangle) = 4$.

⁴ *asprin* [Brewka et al., 2015] is a *clingo* extension that allows for this kind of comparison. For more details about the use of *asprin* to extract skypatterns, see [Gebser et al., 2016].

Fig. 3 Illustration of the notion of insertable region on the example of Fig. 2 for pattern $\langle abc \rangle$. Each line shows an insertable region, from top to bottom: insertion in the prefix, insertion between b and b , insertion in b and c , insertion in the suffix.



The set of backward-closed patterns is $\{\langle a \rangle, \langle b \rangle, \langle c \rangle, \langle bc \rangle, \langle ac \rangle, \langle abc \rangle\}$. $\langle bc \rangle$ is backward-closed because any pattern $\langle bc? \rangle$ is frequent.

Now, we introduce alternative maximality/closure conditions. The objective of these equivalent conditions is to define maximality/closure without comparing patterns. Such conditions can be used to encode the mining of condensed pattern representations. The main idea is to say that a sequence S is maximal (resp. closed) if and only if for every sequence S' s.t. S is a subsequence of S' with $|S'| = |S| + 1$, then S' is not frequent (resp. S' has not the same support as S).

More precisely, a frequent pattern S is maximal iff any sequence S_a^j , obtained by adding to S any item a at any position j , is non frequent. Such an a will be called an *insertable* item.

Proposition 3 (Maximality condition) *A frequent sequence $S = \langle t_1, \dots, t_n \rangle$ is maximal iff $\forall a \in \mathcal{I}, \forall j \in [0, n], |\{T \in \mathcal{D} \mid S \prec T \wedge S_a^j \prec T\}| < f_{min}$, where $S_a^0 = \langle a, t_1, \dots, t_n \rangle$, $S_a^j = \langle t_1, \dots, t_j, a, t_{j+1}, \dots, t_n \rangle$ and $S_a^n = \langle t_1, \dots, t_n, a \rangle$.*

A frequent pattern S is closed iff for any frequent sequence S_a^j , obtained by adding any item a at any position j in S , any sequence T that supports S supports also S_a^j .

Proposition 4 (Closure condition) *A frequent sequence $S = \langle t_1, \dots, t_n \rangle$ is closed iff $\forall a \in \mathcal{I}, \forall j \in [0, n], \text{supp}(S_a^j) \geq f_{min} \implies (\forall T \in \mathcal{D}, S \prec T \implies S_a^j \prec T)$, where $S_a^0 = \langle a, t_1, \dots, t_n \rangle$, $S_a^j = \langle t_1, \dots, t_j, a, t_{j+1}, \dots, t_n \rangle$ and $S_a^n = \langle t_1, \dots, t_n, a \rangle$.*

A consequence (the contraposition) of these properties is that if an item may be inserted between items of an embedding for at least f_{min} sequences (resp. for all supported sequences) then the current pattern is not maximal (resp. not closed). The main idea of our encodings is grounded on this observation.

The main difficulty is to construct the set of insertable items for each in-between position of a pattern, so-called insertable regions. Fig. 3 illustrates the insertable regions of a sequence for the pattern $\langle abc \rangle$.

Definition 1 (Insertable item/insertable region). Let $P = \langle p_i \rangle_{i \in [l]}$ be an l -pattern, $S = \langle s_i \rangle_{i \in [n]}$ be a sequence and $\mathcal{E}^j = (e_i^j)_{i \in [l]}$, $j \in [k]$ be the k embeddings of P in S , $k > 0$. An *insertable region* $R_i = [l_i + 1, u_i - 1] \subset [n]$, $i \in [l + 1]$ is a set of positions in S where $l_i \stackrel{\text{def}}{=} \min_{j \in [k]} e_{i-1}^j$, $i \in [2, l + 1]$, $u_i \stackrel{\text{def}}{=} \max_{j \in [k]} e_i^j$, $i \in [1, l]$, $l_1 \stackrel{\text{def}}{=} 0$, $u_{l+1} \stackrel{\text{def}}{=} n + 1$.

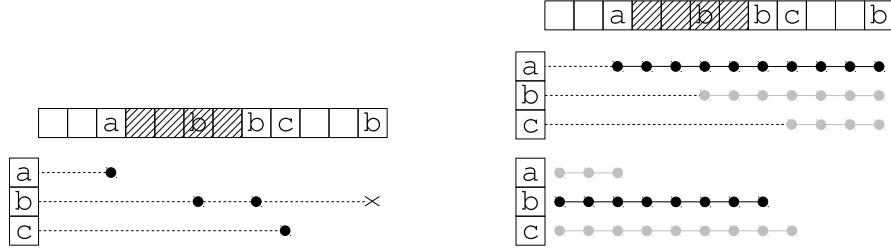


Fig. 4 Illustration of the computation of an insertable region (hatched area) where insertable items are located between a and b and related to the first and second element of pattern $\langle abc \rangle$. On the left, valid $\text{occS}/3$ atoms in the skip-gaps strategy. In the figures, the leftmost occurrences and the rightmost occurrences are the same. Concerning the fill-gaps strategy, $\text{occF}/3$ and $\text{roccF}/3$ atoms are illustrated on the right. Black occurrences are used to compute the hatched region.

Any item $a \in s_p$, $p \in [l_i, u_i]$, $i \in [l+1]$ is called an *insertable item* and is such that S supports the pattern P' obtained by inserting a in P at position i as follows:

- $P' = \langle a, p_1, \dots, p_l \rangle$ if $i = 1$,
- $P' = \langle p_1, \dots, p_l, a \rangle$ if $i = l + 1$,
- $P' = \langle p_1, \dots, p_{i-1}, a, p_i, \dots, p_l \rangle$ otherwise.

In the sequel, we present encodings for closed and maximal patterns which are based on the notations introduced in Definition 1. These encodings cope with the most general case of condensed patterns. It should be noted that, for efficiency reasons, most of procedural algorithms for condensed sequential pattern mining process backward-condensed patterns only. Specific ASP encodings for backward-condensed pattern mining can be found in [Guyet et al., 2016]. These encodings are known to be more efficient but are less generic. In Sect. 7, the performance of the encodings introduced here will be compared with other existing approaches that often implement only backward closure/maximality constraints.

5.2.2 Encoding maximal and closed patterns constraints

The encoding below describes how is defined the set of items that can be inserted between successive items of an embedding. These itemsets are encoded by the atoms of predicate $\text{ins}(\mathbb{T}, \mathbb{X}, \mathbb{I})$ where \mathbb{I} is an item which can be inserted in an embedding of the current pattern in sequence \mathbb{T} between items at position \mathbb{X} and $\mathbb{X}+1$ in the pattern. We give the encodings for the two strategies skip-gaps and fill-gaps: Listing 8 (resp. Listing 9) has to be added to the encoding of skip-gaps strategy (Listing 2), resp. fill-gaps strategy (Listing 3). We illustrate the way they proceed in Fig. 4.

Listing 8 gives the encoding for computing insertable items using the skip-gaps strategy. This encoding is based on the idea that the insertable region i is roughly defined by the first occurrence of the $(i-1)$ -th pattern item and the last occurrence of the i -th pattern item. However, not all occurrences of an item I represented by

```

18 % leftmost "valid" embeddings
19 mlocc(T,1,P) :- occS(T,1,P), 0 { occS(T,1,Q) : Q<P } 0,
20     support(T).
21 mlocc(T,X,P) :- occS(T,X,P), mlocc(T,X-1,Q), Q<P, X>1,
22     support(T).
23
24 % rightmost "valid" embeddings
25 mrocc(T,L,P) :- occS(T,L,P), 0 { occS(T,L,R) : R>P } 0, patlen(L).
26 mrocc(T,X,P) :- occS(T,X,P), mrocc(T,X+1,R), R>P, X<L, patlen(L).
27
28 %insertable items
29 ins(T,1,I) :- seq(T,P,I), P<Q, mrocc(T,1,Q).
30 ins(T,X,I) :- seq(T,P,I), P<Q, mrocc(T,X,Q),
31     P>R, mlocc(T,X-1,R), X>1, patpos(X).
32 ins(T,L+1,I) :- seq(T,P,I), P>R, mlocc(T,L,R), patlen(L).

```

Listing 8 Computation of insertable items – skip-gaps strategy

$\text{occS}/3$ atoms are valid. For instance, in Fig. 4, on the left, the last occurrence of b is not valid because it can not be used to define an occurrence of $\langle abc \rangle$. The valid occurrences are those which have both a preceding and a following valid occurrence. Thus, this validity property is recursive. The encoding of Listing 8 selects two types of occurrences: the leftmost occurrences (resp. rightmost occurrences) corresponding to the earlier (resp. the later) embeddings.

Lines 19 and 25 are boundary cases. A leftmost occurrence is valid if it is the first occurrence in the sequence. Lines 21-22 expresses that an occurrence of the X -th item is a valid leftmost occurrence if it follows a valid leftmost occurrence of the $(X-1)$ -th item. Note that it is not required to compute a unique leftmost occurrence here. Lines 25-26 do the same operation starting from the end of the sequence, precisely, the rightmost occurrence.

Lines 29-32 define insertable items. There are three cases. Lines 29 and 32 are specific boundary cases, i.e. insertion respectively in the prefix and in the suffix. The rule in lines 30-31 specifies that insertable items \mathbb{I} are the items of a sequence \mathbb{T} at position \mathbb{P} such that \mathbb{P} is strictly between a leftmost position of the $(X-1)$ -th item and a rightmost position of the X -th item. In Fig. 4 left, the hatched segment defines the second insertable region for pattern $\langle abc \rangle$ (strictly between a and b).

The encoding of Listing 9 achieves the same task using the alternative semantics for predicate $\text{occF}/3$ defining the fill-gaps strategy. As noted for the previous encoding, only the positions of the last and the first valid occurrences are required for any pattern item. It can be noted that the fill-gaps strategy provides the first valid occurrence of an item X as the first atom of the $\text{occF}(T, X, _)$ sequence. Then, computing the last occurrence for each pattern item can be done in the same manner considering an embedding represented in reverse order. The right part of Fig. 4 illustrates $\text{occF}/3$ and $\text{roccF}/3$ (reverse order) occurrences (see Listing 9, lines 21-23). We can notice that the hatched insertable region is the intersection of occurrences re-


```

20 %embeddings in a reverse order
21 roccF(T,L,P) :- seq(T,P,I), pat(L,I), patlen(L).
22 roccF(T,L,P) :- roccF(T, L, P+1), seq(T,P,_).
23 roccF(T,L,P) :- roccF(T, L+1, P+1), seq(T,P,C), pat(L,C).
24
25 %insertable items
26 ins(T,l ,I) :- seq(T,P,I), roccF(T,l, P+1).
27 ins(T,L+1,I) :- seq(T,P,I), occF(T,L,P-1), patlen(L).
28 ins(T,X, I) :- seq(T,P,I), roccF(T,X,P+1),
29                occF(T,X-1,P-1), patpos(X), X>1.

```

Listing 9 Computation of insertable items – fill-gaps strategy

lated to a and reverse occurrences related to b , after having removed intersection bounds.

The computation of insertable items, Listing 9 lines 26-29, exploits the above remark. Line 26 defines the insertable region in a prefix using `roccF(T, l, P)`. Since items are insertable if they are strictly before the first position, we consider the value of `roccF(T, l, P+1)`. Line 27 uses `occF(T, L, P)` to identifies the suffix region. Line 28-29 combines both constraints for in-between cases.

We can now define the (integrity) constraints for closed and maximal patterns. These constraints are the same for the two embedding strategies.

To extract only maximal patterns, the following constraint denies patterns for which it is possible to insert an item which will be frequent within sequences that support the current pattern.

```
:- item(I), X = 1..maxlen+1, { ins(T,X,I) : support(T) } >= th.
```

The following constraint concerns the extraction of closed-patterns. It specifies that for each insertion position (from 1, in the prefix, to `maxlen+1`, in the suffix), it not possible to have a frequent insertable item I for each supported transaction.

```
:- item(I), X = 1..maxlen+1, { ins(T,X,I) } >=th,
    ins(T,X,I) : support(T).
```

Though interesting from a theoretical point of view, these encodings leads to more complex programs and should be more difficult to ground and to solve, especially the encoding in Listing 8. Backward-closure/maximality constraints are more realistic from a practical point of view.

Finally, it is important to notice that condensed constraints have to be carefully combined with other patterns/embedding constraints. As noted by Negrevergne et al. [Negrevergne et al., 2013], in such cases the problem is not clearly specified. For instance, with our database of Example 2, extracting closed patterns amongst the patterns of length at most 2 will not yield the same results as extracting closed patterns of length at most 2. In the first case, $\langle bc \rangle$ is closed because there is no extended pattern (of length at most 2) with the same support. In the second case, this pattern is not closed (see Example 6), even if its length is at most 2.

6 Related works

Sequential pattern mining in a sequence database have been addressed by numerous algorithms inspired by algorithms for mining frequent itemsets. The most known algorithms are GSP [Srikant and Agrawal, 1996], SPIRIT [Garofalakis et al., 1999], SPADE [Zaki, 2001], PrefixSpan [Pei et al., 2004], and CloSpan [Yan et al., 2003] or BIDE [Wang and Han, 2004] for closed sequential patterns. It is worth-noting that all these algorithms are based on the anti-monotonicity property which is essential to obtain good algorithmic performances. The anti-monotonicity property states that if some pattern is frequent then all its sub-patterns are also frequent. And reciprocally, if some pattern is not-frequent then all its super-patterns are non-frequent. This property enables the algorithm to prune efficiently the search space and thus reduces its exploration. These algorithms differ by their strategy for browsing the search space. GSP [Srikant and Agrawal, 1996] is based on a breadth-first strategy, while PrefixSpan [Pei et al., 2004] combines a depth-first strategy with a database projection that consists in reducing the database size after each pattern extension. LCM_seq [Uno, 2004] is also based on the PrefixSpan principle, but it uses the data structures and processing method of LCM, which is the state of the art algorithm for frequent itemsets mining. Finally, SPADE [Zaki, 2001] introduces a vertical representation of database to propose an alternative to the two previous type of algorithms. For more details about these algorithms, we refer the reader to the survey of Mooney and Roddick [Mooney and Roddick, 2013].

Many algorithms extend the principles of these algorithms to extract alternative forms of sequential patterns. Constraints and condensed patterns are among the most studied alternative patterns due to their relevance to a wide range of applications or to their concise representation of frequent patterns. Integrating constraints in sequential pattern mining is often limited to the use of anti-monotonic temporal constraints such as `maxgap` constraints. When constraints are not anti-monotonic, the previous pruning technique cannot be applied and the computation may require an exhaustive search, which is not reasonable. The usual technique consists in defining an anti-monotonic upper-bound of the measure such that a large part of the search space can be prune (e.g. high occupancy patterns [Zhang et al., 2015]). The tighter the upper-bound is, the better the computing performances are. However, any new type of constraint requires a long effort before being integrated in an efficient algorithm. Integrating flexible and generic constraints in a pattern mining algorithm remains a challenge.

The design of a generic framework for data mining is not a new problem. It has been especially studied within the field of inductive databases as proposed by Imielinski and Mannila [Imielinski and Mannila, 1996]. In an inductive database, knowledge discovery is viewed as a querying process. The idea is that queries would return patterns and models. This framework is based on a parallel between database and data mining theory and has as ultimate goal the discovery of a relational algebra for supporting data mining.

In the specific field of pattern mining, designing such query languages has recently attracted interest in the literature [De Raedt, 2015, Guns et al., 2015, Negrevergne et al., 2013, Bonchi et al., 2006, Boulicaut and Jeudy, 2005, Vautier et al., 2007]. For instance, Vautier et al. [Vautier et al., 2007] proposed a framework which is based on an algebraic specification of pattern mining operators. Bonchi et al. [Bonchi et al., 2006] proposed the Conquest system which is an algorithmic framework that accepts constraints with different properties (anti-monotonic, convertible, loose anti-monotonic, etc.). Boulicaut and Jeudy [Boulicaut and Jeudy, 2005] survey the field of constraint-based data mining. Negrevergne et al. [Negrevergne et al., 2013] recently proposed an algebra for programming pattern mining problems. This algebra allows for the generic combination of constraints on individual patterns with dominance relations between patterns.

More recently, the declarative approaches have shown a strong potential to be relevant frameworks for implementing the principles of inductive databases [De Raedt, 2015]. Many data mining problems can be formalized as combinatorial problems in a declarative way. For instance, tasks such as the discovery of patterns in data, or finding clusters of similar examples in data [Dao et al., 2015], often require constraints to be satisfied and require solutions that are optimal with respect to a given scoring function. The aim of these declarative approaches is to obtain a declarative constraint-based language even at the cost of degraded runtime performance compared to a specialized algorithm. Three types of state-of-the-art solvers have been used: SAT solvers [Coquery et al., 2012], CP solvers [Guns et al., 2015] and ASP solvers [Järvisalo, 2011].

MiningZinc [Guns et al., 2015] is a CP-based approach providing a specific language built upon MiniZinc, a medium-level constraint modelling language [Nethercote et al., 2007]. A similar declarative language has been proposed by Bruynooghe et al. [Bruynooghe et al., 2015] using the IDP3 system. IDP3 is a Knowledge Base System (KBS) that intends to offer the user a range of inference methods and to make use of different state of the art technologies including SAT, SAT Modulo Theories, Constraint Programming and various technologies from Logic Programming. One example of application of their system concerns the problem of learning a minimal automaton consistent with a given set of strings. In ASP, Järvisalo [Järvisalo, 2011] has proposed the first attempt of encoding pattern mining in ASP. Järvisalo addressed this problem as a new challenge for the ASP solver, but he did not highlight the potential benefit of this approach to improve the expressiveness of pattern mining tools. Nonetheless, the first order expressions of ASP encodings can easily be understood by users without higher abstracted languages. Following Guns et al.'s proposal [Guns et al., 2011], Järvisalo designed an ASP program to extract frequent itemsets in a transaction database. A major feature of Järvisalo's proposal is that each answer set (AS) contains only one frequent itemset associated with the identifiers of the transactions where it occurs. To the best of our knowledge, there is no comprehensive language provided for SAT-based data mining approaches.

All these approaches were conducted on itemset mining in transaction databases, which is much simpler than sequential pattern mining in a sequence database. Some

recent works have proposed to explore declarative programming for sequential pattern mining. In fact, dealing with expressive constraints is especially interesting for sequential pattern mining. The range of constraints on sequential patterns is wider than on itemsets and are meaningful for various concrete data analysis issues.

Negrevergne et Guns [Negrevergne and Guns, 2015] proposed the CPSM approach which can be considered as the state of the art of declarative sequential pattern mining. Their contribution is twofold: i) the first declarative encodings of the standard sequential pattern mining task, ii) an efficient CP-based approach based on dedicated propagators that remains compatible with sequential pattern constraints. By combining efficiency and declarativity, CPSM is a proof of concept that a declarative approach can be efficient to solve pattern mining tasks.

Métivier et al. [Métivier et al., 2013] have developed a constraint programming method for mining sequential patterns with constraints in a sequence database. The constraints are based on *amongst* and *regular expression* constraints and expressed by automata. Coquery et al. [Coquery et al., 2012] have proposed a SAT based approach for sequential pattern mining. The patterns are of the form $ab?c$ and an occurrence corresponds to an exact substring (without gap) with joker (the character $?$ replaces exactly one item different from b and c). Coletta and Negrevergne [Coletta and Negrevergne, 2016] have proposed a purely boolean SAT formulation of sequential pattern mining (including closed and maximal patterns) that can be easily extended with additional constraints.

ASP has also been used for sequential pattern mining [Gebser et al., 2016, Guyet et al., 2014]. Gebser et al. [Gebser et al., 2016] have proposed, firstly, an efficient encoding for sequential pattern mining. Secondly, they have proposed to use the *asprin* system for the management of pattern set constraints using preferences. In [Guyet et al., 2014], the mining task is the extraction of serial episodes in a unique long sequence of itemsets where occurrences are the minimal occurrences with constraints. Counting the number of occurrences of a pattern, or of a set of patterns, in a long sequence introduces additional complexity compared to mining sequential patterns from a sequence database since two pattern occurrences can overlap. The main contribution is a method for enumerating pattern occurrences that ensures the anti-monotonicity property.

7 Experiments

Having demonstrated that modelling in ASP is powerful yet simple, it is now interesting to examine the computational behavior of ASP-based encodings.

The first experiments compare the performance, in runtime and memory requirements, of the various ASP programs presented before. The objective is to better understand the advantages and drawbacks of each encoding. The questions we would like to answer are: which of the two embedding strategies is the best? does the encoding improvement really reduce computing resources needs? what is the behaviour of our encoding with added pattern constraints?

Next, we compare our results with the CP-based ones of CPSM [Negrevergne and Guns, 2015]. CPSM constitutes a natural reference since it aims at solving a mining task similar to the present one and since CPSM adopts a semi-declarative approach, in particular, occurrence search is performed by a dedicated constraint propagator.

In all presented experiments, we use the version 4.5 of *clingo*⁵, with default solving parameters. For benchmarking on synthetic data, the ASP programs were run on a computing server with 8Go RAM without using the multi-threading mode of *clingo*. Multi-threading reduces the mean runtime but introduces variance due to the random allocation of tasks. Such variance is inconvenient for interpreting results with repeated executions. For real datasets, we used the multi-threading mode with 4 threads and 20Go shared RAM. This large amount of memory is required for large datasets.

7.1 Encodings comparisons on synthetic datasets

The first experiments were conducted on synthetic databases to control the most important features of data. It allows for an easier and more reliable analysis of time and memory requirements with respect to these parameters. We designed a sequential database simulator to generate datasets with controlled characteristics. The generator⁶ is based on a “retro-engineering” process: 1) a set of random patterns is generated, 2) occurrences of patterns are assigned to a given percentage of database sequences, and 3) each sequence of items is randomly generated according to the patterns it must contain and a mean length.

The parameters of the generator and their default values are sum up in Table 1. Default values are those used when not explicitly specified.

The task to be solved is the extraction of the complete set of frequent patterns (see Sect. 3). It should be noted that every encoding extracts exactly the same set of patterns. Resource requirements are thus fairly comparable. The computation runtime is the time needed to extract all the patterns. It includes both grounding and solving of the ASP programs using the quiet *clingo* mode (no printed output). The memory consumption is evaluated from the size of the grounded program, i.e. the number of grounded atoms and rules. This approximation is accurate to compare ASP encodings. The solving process may require additional memory. This memory requirement is negligible compared to grounding.

We start with an overall comparison of the different encodings and their refinements with respect to parameters th_D and l . Fig. 5 compares the runtimes for different encodings and the two embedding strategies, fill-gaps and skip-gaps. For each setting, 6 databases with the same characteristics were generated. Figure curves

⁵ <http://potassco.sourceforge.net/>

⁶ The generator and databases used in our experiments are available at <https://sites.google.com/site/aspseqmining>.

Table 1 Sequence generator parameters.

Parameter	Default value	Description
D	500	number of sequences in the database
l	20	sequence mean length (sequence length follows a normal law)
n	20	number of different patterns
lp	5	pattern mean length
th_D	10%	minimum number of occurrences generated for each pattern
k	50	alphabet size. The distribution of item occurrences follows a normal law ($\mu = .5$ and $\sigma = 0.05$). Some items occur more often than others.

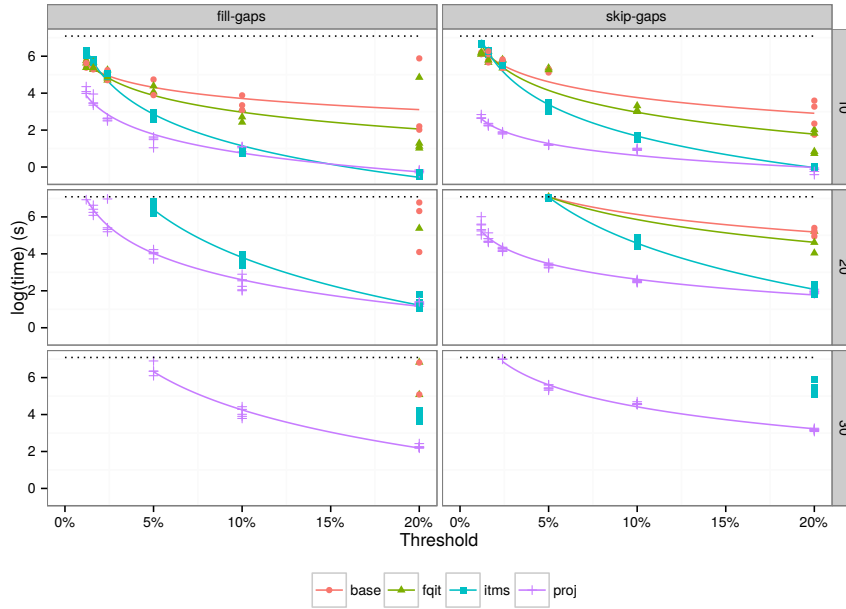


Fig. 5 Mean computation time for synthetic databases with sequences length from 10 to 30 (in rows). Each curve represents an improvement to the basic encoding: none (base), frequent items (fqit), itemsets (itms), projection (proj). Each dot represents the mean of results on 6 datasets. The left-hand (resp. right-hand) column gives the results for the fill-gaps (resp. skip-gaps) strategy. The dashed horizontal line denotes the timeout of 20 minutes.

show the mean runtime of the successful executions, i.e. those that extract the complete set of frequent pattern within the timeout period. The timeout was set to 20 minutes.

The exponential growth of the runtime when the threshold decreases is a classical result in pattern mining considering that the number of patterns grows exponentially. Every approach conforms to this behaviour. In more details:

- the longer the sequences, the greater the runtime. Most problem instances related to databases with $l = 10$ can be solved by any approach. When the mean

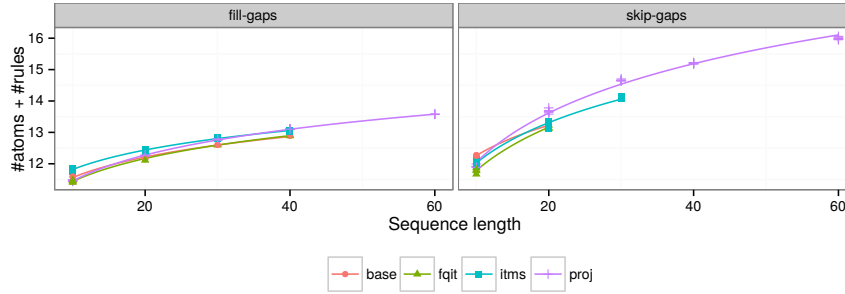


Fig. 6 Memory requirement with respect to sequence length. Problem size (estimated memory requirement) for synthetic databases of sequences of length from 10 to 30 (in rows), under frequency threshold of 20%.

length of sequences increases, the computation time increases also and the number of instances solved within the timeout period decreases. This can be easily explained by the combinatorics of computing embeddings which increases with the sequence length.

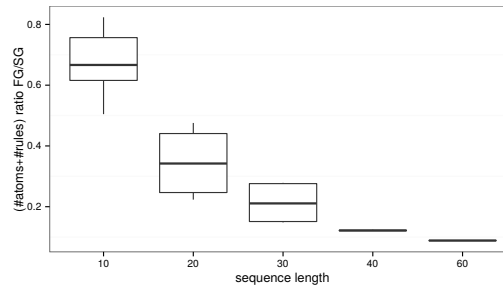
- all proposed improvements do improve runtime for high frequency thresholds on these small synthetic databases. For $f_{min} = 20\%$, the curve of every proposed improvement is below the curve of the basic encoding. For high thresholds, prefix-projection and itemsets improvements are significantly better. Nonetheless, the lower the threshold, the lower the difference between computation times. This shows that, except for prefix-projection, the improvements are not so efficient for hard mining tasks.
- the prefix-projection improvement is the fastest and reduces significantly the computation time (by 2 to 3 orders of magnitude).
- the skip-gaps strategy is more efficient than the fill-gaps strategy for these small datasets. The skip-gaps strategy requires less time to extract the same set of patterns than the fill-gaps strategy, for the same encoding improvements.

We will see below that this last result does not accurately predict which strategy should be preferred for mining real datasets. Before, we analyse the memory requirements of the different encodings.

We first note that the memory consumption is not related to the frequency threshold. This is a specificity of declarative pattern mining. Thus, Fig. 6 compares the embedding strategies only for a unique frequency threshold $f_{min} = 20\%$. The curves show the number of grounded atoms and rules. As it represents a tight approximation of the memory requirement, we will refer to memory in the sequel.

Unsurprisingly, the richer the encoding is, the more memory is required. But the differences are not really significant, except for the prefix-projection programs (proj) which requires the highest number of atoms. We can see that using frequent itemsets (itms) is efficient to reduce the memory requirement. This means that the grounding

Fig. 7 Dispersion of ratios of memory consumption obtained for the skip-gaps strategy to those obtained for the fill-gaps strategy. Boxplots were computed for problem instances with threshold at 20% and for all lengths and all encodings.



step was able to exploit the additional rules to avoid the creation of useless atoms and rules. Such a kind of rules is really interesting because, as the algorithmic complexity of the mining task is not high, the efficiency of the ASP program is related to his grounding size.

In addition, from this last point of view, we can note that the fill-gaps strategy requires several order less memory than the skip-gaps strategy. The longer the sequences, the larger the difference. This result is illustrated by Fig. 7. For each problem instance, the ratio of memory usage is computed by dividing the memory required by encoding with skip-gaps strategy with the memory required by the similar encoding with the fill-gaps strategy. Fig. 7 illustrates with boxplots the dispersion of these ratios for different sequence lengths. Fig. 7 clearly shows that the longer the sequences are, the more efficient the fill-gaps strategy is for memory consumption.

To end this overall comparison, it is interesting to come back to runtime. The overall results of Fig. 5 show that the skip-gaps strategy seems better, but considering that the fill-gaps strategy requires less memory, it is interesting to analyse the evolution of computation time with respect to database size.

Fig. 8 illustrates the ratio of runtimes in both strategies when the database size increases. The support threshold, f_{min} , is fixed to 10% and the sequence mean length to 20. We used the prefix-projection encoding for this experiment. Similarly to the previous figure, the ratios were individually computed for each pair of results (fill-gaps/skip-gaps) and the figure shows statistics about these ratio.

Fig. 8 shows clearly that when the database size increases, the fill-gaps strategy becomes more efficient than the skip-gaps strategy.

From these experiments, we can conclude that combining prefix-projection with the fill-gaps strategy gives the best encoding. Thus, in the next subsection, we will compare this encoding with CPSM.

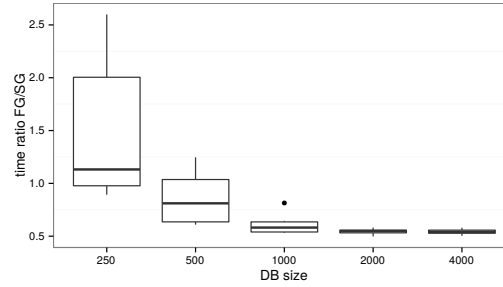


Fig. 8 Dispersion of ratios of runtime of the skip-gaps strategy to the runtime of the fill-gaps strategy.

Table 2 Dataset characteristics: alphabet size, number of sequences and items, max and mean length of sequences, dataset density.

Dataset	$ \mathcal{S} $	$ D $	$ D $	$\max T $	$\text{avg} T $	density
Unix user	265	484	10935	1256	22.59	0.085
JMLR	3847	788	75646	231	96.00	0.025
iPRG	21	7573	98163	13	12.96	0.617
FIFA	20450	2990	741092	100	36.239	0.012

7.2 Real dataset analysis

In these experiments, we analyse the proposed encodings on processing real datasets. We use the same real datasets as selected in [Negrevergne and Guns, 2015] to have a representative panel of application domains:

- JMLR: a natural language processing dataset; each transaction is a paper abstract from the Journal of Machine Learning Research,
- UNIX: each transaction is a series of shell commands executed by a user during one session,
- iPRG: each transaction is a sequence of peptides that is known to cleave in presence of a Trypsin enzyme,
- FIFA: each transaction is a sequence of webpages visited by a user during a single session.

The dataset characteristics are sum up in Table 2. Some of them are similar to those of simulated datasets.

7.2.1 Comparison of frequent pattern mining with CPSM

Fig. 9 compares the runtimes of ASP-based sequence mining (using the ASP system *clingo*) and CPSM (based on the CP solver *gecode*). We ran the two versions of CPSM. CPSM makes use of global constraints to compute embeddings. This version is known to be very efficient, but it cannot cope with embedding constraints, while CPSM-emb does but is less efficient. We do not compare our approach with dedicated algorithms, which are known to be more efficient than declarative mining

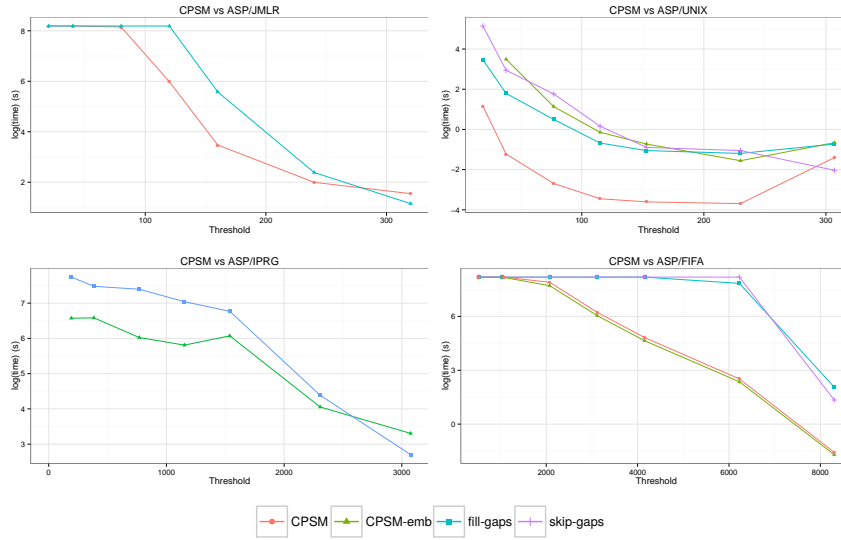


Fig. 9 Runtime for mining frequent patterns with four approaches: CPSM, CPSM-emb, ASP with fill-gaps, ASP with skip-gaps.

approaches (see [Negrevergne and Guns, 2015] for such comparisons). The timeout was set to 1 hour.

The results show that the runtimes obtained with *clingo* are comparable to CPSM-emb. It is lower for IPGR, very similar for UNIX and larger for JMLR. These results are consistent to those presented in [Gebser et al., 2016] for synthetic datasets. When sequences become large, the efficiency of our encoding decreases somewhat. The mean length for JMLR is 96 while it is only 12.96 for iPRG. For CPSM with global constraints, the runtime-efficiency is several order of magnitude faster. To be fair, it should be noted that ASP approach ran with four parallel threads while CPSM-emb ran with no multi-threading since it does not support it. It should also be noted that CPSM requires a lot of memory, similarly to ASP-based solving.

7.2.2 Comparison of constrained frequent pattern mining with CPSM

In this section, we detail the performance on constrained pattern mining tasks. We compare our approach with CPSM-emb, which enables *max-gap* and *max-span* constraints. In this experiments we took the same setting as the experiments of [Negrevergne and Guns, 2015]: we add first a constraint *max-gap*=2 and then we combine it with a second constraint *max-span*=10. For each setting, we compute the frequent patterns with our ASP encoding and with CPSM for the four datasets.

Fig. 10 shows the runtime and the number of patterns for each experiment. This figure illustrates results for completed searches. A first general remark is that adding

constraints to ASP encodings reduces computation times. Surprisingly for CPSM, for some thresholds the computation with some constraints requires more time than without constraints. This is the case for example for the iPRG dataset: CPSM could not solve the mining problem within the timeout period for thresholds 769 and 384. Surprisingly, it could complete the task for lower thresholds whereas the task should be more difficult. ASP required also more time to solve the same problem instances, but it could complete them. Again, we can note that the mean sequence length impacts the performance of ASP encodings. CPSM has lower runtime on JMLR than ASP while it is the opposite on iPRG.

The curves related to the number of patterns demonstrate that the number of extracted pattern decreases when the number of constraints increases. Since we present only the results of completed solving, CPSM and ASP yield the same set of patterns.

7.2.3 Analysis of condensed pattern extraction

Fig. 11 illustrates the results for condensed pattern mining. This approach cannot be compared to CPSM since it does not propose means for encoding such kind of patterns.

This experiment compares the resource requirements in time and memory for mining closed/maximal and backward-closed/maximal patterns. For each of these mining task, we compared the skip-gaps and fill-gaps strategies. The main encoding is still based on prefix-projection. Three real datasets have been processed (JMLR, UNIX and IPRG). The FIFA dataset was not processed due to its heavy memory requirement for some of these tasks.

We can first note that the difference between the number of extracted patterns is low. As expected, all encodings that complete a given mining task extract the same number of patterns. This result supports the correctness of our approach. From the memory point of view, we see that the encodings extracting condensed patterns requires several order of magnitude more memory, especially for (backward-)closed patterns. It is also interesting to note that the memory requirement for the fill-gaps strategy is not linked to the threshold, contrary to the skip-gaps strategy. Again, the fill-gaps strategy seems to be more convenient for small thresholds. We can note that there is a big difference between datasets concerning runtime. For instance, frequent patterns are faster to extract for JMLR and UNIX, but maximal patterns are faster to compute on IPRG. The density of this last dataset makes maximal pattern extraction easier. Uniformly, we can conclude that fill-gaps is faster than skip-gaps. The complexity of the encoding related to insertable items with skip-gaps makes the problem difficult to solve. Opposed to the experiments presented in [Guyet et al., 2016], we did not use any solving heuristic. For maximal patterns, a huge improvement of runtime was observed when using the *subset-minimal* heuristic⁷.

⁷ The use of subset-minimal heuristic keeps solving the maximal patterns problem complete.

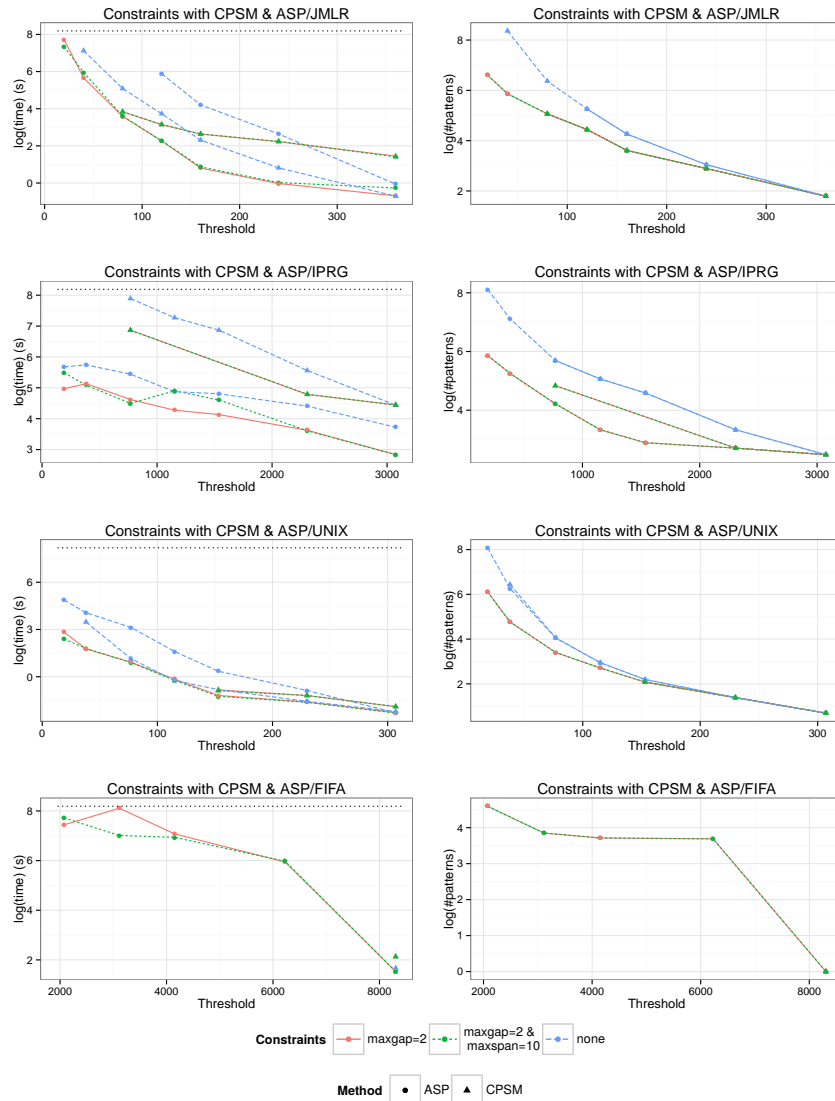


Fig. 10 Results for constrained sequence mining tasks with ASP vs CPSM. On the left: runtime; on the right: number of patterns. The two figures are in log scales. From top to bottom, JMLR, IPRG, UNIX and FIFA. For each plot, the curves illustrate the results for different type of constraints (see legend). The horizontal dashed line figure out the 1 hour timeout.

8 Conclusion and perspectives

This article has presented a declarative approach of sequential pattern mining based on answer set programming. We have illustrated how to encode a broad range of

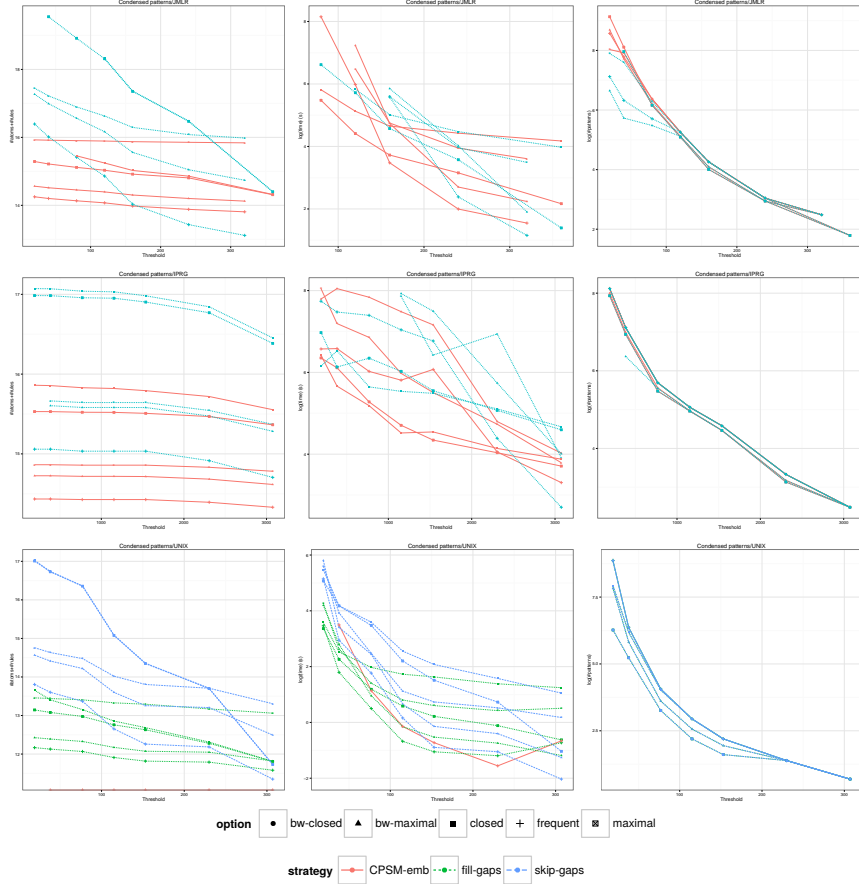


Fig. 11 From left to right, problem size, runtime and number of extracted patterns with respect to the frequency threshold. Runtimes are shown only if the solving was complete, contrary to pattern numbers which show the number of extracted patterns within the timeout period. From top to bottom, JMLR, IPRG and UNIX. For each plot, the curves illustrate the results for different types of condensed patterns (see legend) and for the two embedding strategies (*fill-gaps* in red-plain line, *skip-gaps* in blue-dashed line).

mining tasks (including condensed representations and constrained patterns) in pure ASP. Thus, we shown the first advantage of declarative pattern mining: for most well-specified tasks, the development effort is significantly lower than for procedural approaches. The integration of new constraints within our framework requires only few lines of code. This was made possible thanks to the flexibility of both ASP language and solvers.

Nonetheless, another objective of this paper was to give the intuition to the reader that while encoding a straightforward solution to a problem can be easy in ASP, writing efficient programs may be complex. Developing competitive encodings requires

a good understanding of the solving process. To this end, we have presented several possible improvements of basic sequential pattern mining and two alternatives for encoding the main complex task, i.e. computing embeddings. These encodings have been extensively evaluated on synthetic and real datasets to draw conclusions about the overall efficiency of this approach (especially compared to the constraint programming approach CPSM) and about which are the best encodings among the proposed ones and in which context.

The first conclusion of these experiments is that our ASP approach has comparable computing performances with CPSM-emb as long as the length of the sequence remains reasonable. This can be explained considering that solving the embedding problem is a difficult task for pure declarative encodings while CPSM relies on dedicated propagators. The propagators of CPSM solve the embedding problem using additional procedural code. It turns that, for solving the embedding problem in ASP, encoding using a *fill-gaps* strategy appears to be better than using the *skip-gaps* strategy on real datasets, thanks to lower memory requirements.

The second conclusion is that adding constraints on patterns reduces runtime, but increases memory consumption. For real datasets, the more constraints are added, the more memory is required. This is due, to encoding the constraints, but also to encoding the information that may be required to compute constraints. For example, encodings using the *maxspan* constraint require more complex embeddings (*occS/4* atoms) than encodings without this constraint.

To fully benefit from the flexibility of our approach to proceed large datasets, we need to improve the efficiency of the computation of embeddings. Our objective is now to mimic the approach of CPSM consisting in using propagators within the solver to solve the part of the problems for which procedural approaches are efficient. The new *clingo 5* series will integrate “ASP modulo theory” solving processes. This new facilities will enable to combine ASP and propagators in an efficient way.

Acknowledgements We would like to thanks Roland Kaminski and Max Ostrowski for their significant inputs and comments about ASP encodings; and Benjamin Negrevergne and Tias Guns for their suggestions about the experimental part. We also thank the anonymous reviewers for their valuable comments and constructive suggestions.

Appendix

Listing 10 illustrates how the encoding of the *skip-gaps* strategy can be transformed to mine sequential patterns that are sequences of itemsets.

```

1 item(I) :- seq(_, _, I). %set of items
2
3 %sequential pattern generation
4 patpos(1).
5 0 { patpos(X+1) } 1 :- patpos(X), X<maxlen.
6 patlen(L) :- patpos(L), not patpos(L+1).
7
```

```

8 %generate an itemset for each position
9 l { pat(X,I): item(I) } :- patpos(X).
10
11 %pattern occurrences
12 occS(T,l,P) :- seq(T,P,I):pat(l,I); seq(T,P,_).
13 occS(T,Pos+1,P) :- occS(T,Pos,Q), Q<P, seq(T,P,J),
14                     pat(Pos+1,J), seq(T,P,I):pat(Pos+1,I).
15
16 support(T) :- occS(T, L, _), patlen(L).
17 :- { support(T) } th-1.

```

Listing 10 Mining sequences with patterns as sequences of itemsets

The first difference with the encoding of Listing 2 concerns the generation of patterns. The upper bound constraint of the choice rule in Line 9 has been removed, enabling the possible generation of every non-empty subset of \mathcal{I} .

The second difference is that the new ASP rules verify the inclusion of all items in itemsets. Line 14, $\text{seq}(T, P, I) : \text{pat}(l, I)$ indicates that for each atom $\text{pat}(l, I)$ there should exist an atom $\text{seq}(T, P, I)$ to satisfy the rule body. A similar expression is used Line 15.

References

- [Agrawal et al., 1993] Agrawal, R., Imielinski, T., and Swami, A. (1993). Mining association rules between sets of items in large databases. In *Proceedings of the ACM SIGMOD Conference on Management of Data*, pages 207–216.
- [Agrawal and Srikant, 1995] Agrawal, R. and Srikant, R. (1995). Mining sequential patterns. In *Proceedings of the International Conference on Data Engineering*, pages 3–14.
- [Biere et al., 2009] Biere, A., Heule, M., van Maaren, H., and Walsh, T. (2009). *Handbook of satisfiability. Frontiers in Artificial Intelligence and Applications*, volume 185. IOS Press.
- [Bonchi et al., 2006] Bonchi, F., Giannotti, F., Lucchese, C., Orlando, S., Perego, R., and Trasarti, R. (2006). Conquest: a constraint-based querying system for exploratory pattern discovery. In *Proceedings of the International Conference on Data Engineering*, pages 159–159.
- [Boulicaut and Jeudy, 2005] Boulicaut, J.-F. and Jeudy, B. (2005). Constraint-based data mining. In Maimon, O. and Rokach, L., editors, *Data Mining and Knowledge Discovery Handbook*, pages 399–416. Springer US.
- [Brewka et al., 2015] Brewka, G., Delgrande, J. P., Romero, J., and Schaub, T. (2015). asprin: Customizing answer set preferences without a headache. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*, pages 1467–1474.
- [Bruynooghe et al., 2015] Bruynooghe, M., Blockeel, H., Bogaerts, B., De Cat, B., De Pooter, S., Jansen, J., Labarre, A., Ramon, J., Denecker, M., and Verwer, S. (2015). Predicate logic as a modeling language: Modeling and solving some machine learning and data mining problems with IDP3. *Theory and Practice of Logic Programming*, 15(06):783–817.
- [Coletta and Negrevergne, 2016] Coletta, R. and Negrevergne, B. (2016). A SAT model to mine flexible sequences in transactional datasets. *arXiv preprint arXiv:1604.00300*.
- [Coquery et al., 2012] Coquery, E., Jabbour, S., Saïs, L., and Salhi, Y. (2012). A SAT-Based approach for discovering frequent, closed and maximal patterns in a sequence. In *Proceedings of European Conference on Artificial Intelligence (ECAI)*, pages 258–263.
- [Dao et al., 2015] Dao, T., Duong, K., and Vrain, C. (2015). Constrained minimum sum of squares clustering by constraint programming. In *Proceedings of Principles and Practice of Constraint Programming*, pages 557–573.

- [De Raedt, 2015] De Raedt, L. (2015). Languages for learning and mining. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*, pages 4107–4111.
- [Garofalakis et al., 1999] Garofalakis, M., Rastogi, R., and Shim, K. (1999). SPIRIT: Sequential pattern mining with regular expression constraints. In *Proceedings of the International Conference on Very Large Data Bases*, pages 223–234.
- [Gebser et al., 2016] Gebser, M., Guyet, T., Quiniou, R., Romero, J., and Schaub, T. (2016). Knowledge-based sequence mining with ASP. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 1497–1504.
- [Gebser et al., 2011] Gebser, M., Kaminski, R., Kaufmann, B., Ostrowski, M., Schaub, T., and Schneider, M. (2011). Potassco: The Potsdam answer set solving collection. *AI Communications*, 24(2):107–124.
- [Gebser et al., 2014] Gebser, M., Kaminski, R., Kaufmann, B., and Schaub, T. (2014). *Clingo = ASP + control*: Preliminary report. In *Technical Communications of the Thirtieth International Conference on Logic Programming*.
- [Gelfond and Lifschitz, 1991] Gelfond, M. and Lifschitz, V. (1991). Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9:365–385.
- [Guns et al., 2015] Guns, T., Dries, A., Nijssen, S., Tack, G., and De Raedt, L. (2015). MiningZinc: A declarative framework for constraint-based mining. *Artificial Intelligence*, page In press.
- [Guns et al., 2011] Guns, T., Nijssen, S., and De Raedt, L. (2011). Itemset mining: A constraint programming perspective. *Artificial Intelligence*, 175(12-13):1951–1983.
- [Gupta and Han, 2013] Gupta, M. and Han, J. (2013). *Data Mining: Concepts, Methodologies, Tools, and Applications*, chapter Applications of Pattern Discovery Using Sequential Data Mining, pages 947–970. IGI-Global.
- [Guyet et al., 2014] Guyet, T., Moinard, Y., and Quiniou, R. (2014). Using answer set programming for pattern mining. In *Proceedings of conference “Intelligence Artificielle Fondamentale” (IAF)*.
- [Guyet et al., 2016] Guyet, T., Moinard, Y., Quiniou, R., and Schaub, T. (2016). Fouille de motifs séquentiels avec ASP. In *Proceedings of conference “Extraction et la Gestion des Connaissances” (EGC)*, pages 39–50.
- [Imielinski and Mannila, 1996] Imielinski, T. and Mannila, H. (1996). A database perspective on knowledge discovery. *Communications of the ACM*, 39(11):58–64.
- [Janhunen and Niemelä, 2016] Janhunen, T. and Niemelä, I. (2016). The answer set programming paradigm. *AI Magazine*, 37:13–24.
- [Järvisalo, 2011] Järvisalo, M. (2011). Itemset mining as a challenge application for answer set enumeration. In *Proceedings of the conference on Logic Programming and Nonmonotonic Reasoning*, pages 304–310.
- [Lallouet et al., 2013] Lallouet, A., Moinard, Y., Nicolas, P., and Stéphan, I. (2013). Programmation logique. In Marquis, P., Papini, O., and Prade, H., editors, *Panorama de l’intelligence artificielle : ses bases méthodologiques, ses développements*, volume 2. Cépaduès.
- [Lefèvre and Nicolas, 2009] Lefèvre, C. and Nicolas, P. (2009). The first version of a new ASP solver: ASPeRiX. In *Proceedings of the conference on Logic Programming and Nonmonotonic Reasoning*, pages 522–527.
- [Leone et al., 2006] Leone, N., Pfeifer, G., Faber, W., Eiter, T., Gottlob, G., Perri, S., and Scarcello, F. (2006). The DLV system for knowledge representation and reasoning. *ACM Trans. Comput. Logic*, 7(3):499–562.
- [Lhote, 2010] Lhote, L. (2010). Number of frequent patterns in random databases. In Skiadas, C. H., editor, *Advances in Data Analysis*, Statistics for Industry and Technology, pages 33–45.
- [Lifschitz, 2008] Lifschitz, V. (2008). What is answer set programming? In *Proceedings of the Conference on Artificial Intelligence (AAAI)*, pages 1594–1597.
- [Low-Kam et al., 2013] Low-Kam, C., Raïssi, C., Kaytoue, M., and Pei, J. (2013). Mining statistically significant sequential patterns. In *Proceedings of the IEEE International Conference on Data Mining*, pages 488–497.
- [Métivier et al., 2013] Métivier, J.-P., Loudni, S., and Charnois, T. (2013). A constraint programming approach for mining sequential patterns in a sequence database. In *Proceedings of the*

- Workshops of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD).*
- [Mooney and Roddick, 2013] Mooney, C. H. and Roddick, J. F. (2013). Sequential pattern mining – approaches and algorithms. *ACM Comput. Surv.*, 45(2):1–39.
- [Muggleton and De Raedt, 1994] Muggleton, S. and De Raedt, L. (1994). Inductive logic programming: Theory and methods. *The Journal of Logic Programming*, 19:629–679.
- [Negrevertne et al., 2013] Negrevertne, B., Dries, A., Guns, T., and Nijssen, S. (2013). Dominance programming for itemset mining. In *Proceedings of the International Conference on Data Mining*, pages 557–566.
- [Negrevertne and Guns, 2015] Negrevertne, B. and Guns, T. (2015). Constraint-based sequence mining using constraint programming. In *Proceedings of International Conference on Integration of AI and OR Techniques in Constraint Programming, CPAIOR*, pages 288–305.
- [Nethercote et al., 2007] Nethercote, N., Stuckey, P. J., Becket, R., Brand, S., Duck, G. J., and Tack, G. (2007). MiniZinc: Towards a standard CP modelling language. In *Proceedings of the conference on Principles and Practice of Constraint Programming*, pages 529–543.
- [Pei et al., 2004] Pei, J., Han, J., Mortazavi-Asl, B., Wang, J., Pinto, H., Chen, Q., Dayal, U., and Hsu, M.-C. (2004). Mining sequential patterns by pattern-growth: the PrefixSpan approach. *IEEE Transactions on Knowledge and Data Engineering*, 16(11):1424–1440.
- [Pei et al., 2007] Pei, J., Han, J., and Wang, W. (2007). Constraint-based sequential pattern mining: The pattern-growth methods. *Journal of Intelligent Information Systems*, 28(2):133–160.
- [Perer and Wang, 2014] Perer, A. and Wang, F. (2014). Frequence: interactive mining and visualization of temporal frequent event sequences. In *Proceedings of the international conference on Intelligent User Interfaces*, pages 153–162.
- [Rossi et al., 2006] Rossi, F., Van Beek, P., and Walsh, T. (2006). *Handbook of constraint programming*. Elsevier.
- [Shen et al., 2014] Shen, W., Wang, J., and Han, J. (2014). Sequential pattern mining. In Aggarwal, C. C. and Han, J., editors, *Frequent Pattern Mining*, pages 261–282. Springer International Publishing.
- [Simons et al., 2002] Simons, P., Niemelä, I., and Soinen, T. (2002). Extending and implementing the stable model semantics. *Artificial Intelligence*, 138(1-2):181–234.
- [Srikant and Agrawal, 1996] Srikant, R. and Agrawal, R. (1996). Mining sequential patterns: Generalizations and performance improvements. In *Proceedings of the 5th International Conference on Extending Database Technology*, pages 3–17.
- [Syrjänen and Niemelä, 2001] Syrjänen, T. and Niemelä, I. (2001). The smodels system. In *Proceedings of the conference on Logic Programming and Nonmonotonic Reasoning*, pages 434–438.
- [Ugarte et al., 2015] Ugarte, W., Boizumault, P., Crémilleux, B., Lepailleur, A., Loudni, S., Plantevit, M., Raïssi, C., and Soulet, A. (2015). Skypattern mining: From pattern condensed representations to dynamic constraint satisfaction problems. *Artificial Intelligence*, page In press.
- [Uno, 2004] Uno, T. (2004). http://research.nii.ac.jp/uno/code/lcm_seq.html.
- [Vautier et al., 2007] Vautier, A., Cordier, M., and Quiniou, R. (2007). Towards data mining without information on knowledge structure. In *Proceedings of the Conference on Principles and Practice of Knowledge Discovery in Databases*, pages 300–311.
- [Wang and Han, 2004] Wang, J. and Han, J. (2004). BIDE: Efficient mining of frequent closed sequences. In *Proceedings of the International Conference on Data Engineering*, pages 79–90.
- [Yan et al., 2003] Yan, X., Han, J., and Afshar, R. (2003). CloSpan: Mining closed sequential patterns in large datasets. In *Proceedings of the SIAM Conference on Data Mining*, pages 166–177.
- [Zaki, 2001] Zaki, M. J. (2001). SPADE: An efficient algorithm for mining frequent sequences. *Journal of Machine Learning*, 42(1/2):31–60.
- [Zhang et al., 2015] Zhang, L., Luo, P., Tang, L., Chen, E., Liu, Q., Wang, M., and Xiong, H. (2015). Occupancy-based frequent pattern mining. *ACM Transactions on Knowledge Discovery from Data*, 10(2):1–33.