

Mining Time-constrained Sequential Patterns with Constraint Programming

John O.R. Aoga · Tias Guns · Pierre Schaus

Received: date / Accepted: date

Abstract Constraint Programming (CP) has proven to be an effective platform for constraint based sequence mining. Previous work has focused on standard frequent sequence mining, as well as frequent sequence mining with a maximum 'gap' between two matching events in a sequence. The main challenge in the latter is that this constraint can not be imposed independently of the omnipresent frequency constraint. Indeed, the gap constraint changes whether a subsequence is included in a sequence, and hence its frequency. In this work, we go beyond that and investigate the integration of timed events and constraining the minimum/maximum gap as well as minimum/maximum span. The latter constrains the allowed time between the first and last matching event of a pattern. We show how the three are interrelated, and what the required changes to the frequency constraint are. Key in our approach is the concept of an extension window defined by gap/span and we develop techniques to avoid scanning the sequences needlessly, as well as using a backtracking-aware data structure. Experiments demonstrate that the proposed approach outperforms both specialized and CP-based approaches in almost all cases and that the advantage increases as the minimum frequency threshold decreases. This paper is an extension of the original manuscript presented at CPAIOR'17 [5].

Keywords Data mining · Sequential pattern mining · Constraint programming · Global constraint · Gap constraint · Span constraint · Time constraint

J.O.R Aoga (✉)

Institute of Information and Communication Technologies, Electronics and Applied Mathematics (ICTEAM), UCLouvain, Belgium

Ecole Doctorale Science de l'Ingénieur (ED-SDI), Université d'Abomey-Calavi (UAC), Bénin

Orcid Id: 0000-0002-7213-146X

E-mail: john.aoga@{uclouvain.be, gmail.com}

T. Guns

Vrije Universiteit Brussel (VUB), Brussels, Belgium

Katholieke Universiteit Leuven, Belgium

E-mail: tias.guns@{vub.be, cs.kuleuven.be}

P. Schaus

Institute of Information and Communication Technologies, Electronics and Applied Mathematics (ICTEAM), UCLouvain, Belgium

E-mail: pierre.schaus@uclouvain.be

1 Introduction

Sequential pattern mining (SPM) is an important research domain within data mining and widely used in applications such as web log mining, disease diagnoses mining, event sequence mining, etc [1]. The problem of SPM is to find frequent sequence patterns (also called sequential patterns) in a database of sequences, i.e. an ordered list of events which together occur in the data more than a given number of times. This task is a great challenge since the search space is extremely large; $O(m^n)$ solutions are available for patterns with length at most n and for sequences with an average number of m events.

In practice, finding all sequential patterns is typically not enough, as often an overwhelming number of patterns is returned. Hence, there is a need to guide the search towards patterns of interest to the practitioner. This calls for techniques which can incorporate preferences or restrictions on the length and content of the patterns (constraints). In many applications, the time elapsed between events is also important to take into account.

Assume for instance a database containing sequences of web pages visited by users on a given web-site. One could be interested in access patterns within a single browsing session, for example with no more than 20 minutes time between two pages. Also in biological sequence mining the position and distance of the symbols in the sequence matter. A constraint on the maximum time between any two consecutive symbols in the pattern is called a *gap* constraint, while a constraint on the time from the first to the last event is called a *span* constraint. In this work, we assume all sequences have explicit timestamps and the goal is to support gap and span constraints as well as constraints on frequency and syntax of patterns.

Related work The problem of sequential pattern mining, first introduced by Agrawal et al. [2], is widely studied [2, 13, 28, 35, 38, 39, 40] with many applications as well [15, 16]. These works can be categorized into 1) apriori-based (horizontal/vertical formatting) [6, 35, 40] and 2) projection-based [28] methods. In general, users only need a small subset of the found patterns. Hence, a number of works have focused on the addition of user-defined constraints such as inclusion/exclusion items, pattern length (minimum/maximum), super-pattern, aggregate function (sum, average, maximum, minimum, standard deviation, ...), regular expression and span/gap. They are discussed in more detail in [29].

GSP [35] was the first approach including gap and span constraints. This method is not very efficient since it requires to generate all candidate patterns and to scan the dataset several times. Some approaches added the constraints in a post-processing step [27]. In the *cSPADE* algorithm [40], the constraints are directly integrated into the sequential pattern search process. It efficiently takes into account constraints such as length and width restrictions on the pattern, item constraints, minimum and maximum gaps between events, as well as a maximum span. Unlike *cSPADE*, *GenPrefixSpan* [3] is an extension of the depth-first *PrefixSpan* [28] algorithm to allow gap constraints. Time constraints on the sequences (instead of events) have also been investigated in [10]. Special classes of SPM problem or constraints was also tackled: the closed/maximal SPM [11, 21, 22, 37, 38] the multi-dimensional SPM [31], the episodes events [23], etc. However, all the above-mentioned approaches lack flexibility at the algorithmic level, since adding a new constraint often involves changing the whole algorithm and may hinder scalability. For instance, methods that

can efficiently take regular expression constraints into account together with time constraints are rare in specialized methods. An exception is *PG* [29], which starts from the observation that many constraints are prefix-monotone. This property is weaker than standard (anti)monotonicity when used in pruning, but still valid. *PG*'s pruning principles are specific to prefix-monotonicity however, which does not allow it to fully exploit regular expression constraints for example.

As an alternative, the use of Constraint Programming (CP) has been investigated [4, 9, 12, 18, 19, 20, 24, 25]. Kemmar et al. [19] have subsequently shown that this approach can be made more scalable by grouping all low-level constraints involving the frequency computation into one global constraint. Moreover, they investigated the top-k sequential pattern mining problem [18]. More recently, we have shown that combining this approach with algorithmic techniques from both the CP community and the data mining community can result in a global constraint that outperforms generic as well as specialized methods [4].

While the above CP methods can handle constraints on the pattern syntax, gap and span constraints are only supported by the much less efficient approach of [25]. The reason is that the timing information is hidden in the global *frequency* constraint. Hence, Kemmar et al. [20] extended their global constraint for the gap constraint specifically.

Contributions In this work, we wish to improve on [20] and [4] by modifying the global frequency constraint to capture the most common time-related constraints: explicitly timed events, minimum/maximum gap, and minimum/maximum span. To maintain scalability, we must ensure that we do not needlessly scan the sequences in the database during the search. Our contributions can be summarized as follows: 1) we adapt the backtracking-aware data structure introduced in [4] to store all possible occurrences of the pattern in a sequence, including the first matching symbol to support *span* constraints; 2) we avoid scanning a sequence for a symbol beyond the (precomputed) last occurrence of that symbol in the sequence; 3) we introduce the concept of *extension window* of an embedding and avoid to scan overlapping windows multiple times; 4) we avoid scanning for the start of an extension window, which is specific to the minimum gap constraint, by precomputing these in advance; and finally 5) we experimentally show that using this global constraint we outperform other sequence mining algorithms in all but a few cases. Furthermore, we show that in a CP framework this global constraint can be combined with a number of other independent constraints: item inclusion/exclusion constraints, pattern length constraints or string constraints [14] such as regular expression [30] and grammar [17, 32] constraints.

Limitations The type of patterns found by our approach are the commonly used sequences of single events. We do not consider other related pattern types such as sequences of sets of events (sequences of itemsets), multivariate temporal patterns [7] or episodes [36]. These require changes to the pattern representation and the embedding (constraint). One could build on the principles and data structures investigated in this paper for those purposes, but we do not investigate that here.

2 Preliminaries

In this section, we revisit the preliminary concepts for both Sequential Pattern Mining (SPM) and Constraint Programming (CP). Most of these concepts can be found for SPM in [1, 41] and in [34] for CP.

2.1 Sequential Pattern Mining Background

Assume \mathcal{L} is an alphabet, that is, a list of possible symbols. Table 1a represents an example sequence database (SDB) with timestamps. The database is a set of tuples (sid, s) where sid is the sequence identifier and $s = \langle (s_1, t_1)(s_2, t_2) \dots (s_n, t_n) \rangle$ is a sequence; an ordered list of symbols/events (s_k) occurred at time t_k , where $t_1 \leq t_2 \leq \dots \leq t_n$. We use s_i^s , respectively s_i^t , to represent just the list of symbols, respectively timestamps, of sequence i . In the rest of the paper, we assume a sequence database has timestamps, and when the exact timing is not important we will write $\langle ABC \rangle$ to mean $\langle (A, 1)(B, 2)(C, 3) \rangle$.

Example 1 Assume $s = \langle (A, 2)(B, 4)(D, 6)(D, 8)(B, 10)(E, 12)(C, 14) \rangle$ is a sequence, $s^s = \{A, B, D, D, B, E, C\}$, $s^t = \{2, 4, 6, 8, 10, 12, 14\}$ and its length $size(s) = 7$.

A sequence can be a subsequence of another sequence. For example $\alpha = \langle ADC \rangle$ is a subsequence of s . More formally, the subsequence relation is:

Definition 1 Subsequence relation (\preceq). Sequence $\alpha = \langle \alpha_1 \alpha_2 \dots \alpha_k \rangle$ is a subsequence of $s = \langle (s_1, t_1)(s_2, t_2) \dots (s_n, t_n) \rangle$ denoted by $\alpha \preceq s$ iff (i) $k \leq n$ and (ii) there exists a list of integers (e_1, \dots, e_k) , an embedding, with $1 \leq e_1 \leq \dots \leq e_k \leq n$ such that $s^s[e_i] = \alpha_i$.

Example 2 Sequence $\alpha = \langle ADC \rangle$ is a subsequence of s with embedding $(1, 3, 7)$. Another valid embedding would be $(1, 4, 7)$. Note that for this standard subsequence relation, timing is not important.

The *coverage* of a pattern in a sequence database is the set of sequences in SDB for which our pattern is a subsequence: $Cover_{SDB}(\alpha) = \{(sid, s) \in SDB \mid \alpha \preceq s\}$. We denote by *frequency* of a sequence the size of its coverage ($|Cover_{SDB}(\alpha)|$) and by *support* the relative frequency ($Support_{SDB}(\alpha) = |Cover_{SDB}(\alpha)|/size(SDB)$).

Example 3 Sequence $\alpha = \langle ADC \rangle$ is a subsequence of sequences 1, 2 and 3 in Table 1a, hence $Cover_{SDB}(\langle ADC \rangle) = \{sid_1, sid_2, sid_3\}$, its frequency is 3 and $Support_{SDB}(\langle ADC \rangle) = 75\%$.

The problem of SPM, first introduced by Agrawal et al. [2], is as follows:

Definition 2 SPM problem. Find all subsequences (α) in sequence database (SDB) such that $Support_{SDB}(\alpha) \geq \theta$ where θ is the given support threshold. Each such subsequence α is called a frequent sequence pattern or simply sequential pattern.

In the remaining of the paper, we will use *sequential pattern* to mean *frequent sequence pattern* and *sequence pattern* if the frequency of the pattern does not matter.

There exist multiple algorithms for the SPM problem. The *PrefixSpan* algorithm [28] is among the most famous ones and relies on the idea of the prefix-projected database. Our approach will build on this concept.

a) Sequence database (SDB)		b) <i>nextPosGap</i>						c) <i>lastPosMap</i>					
sid	sequence	1	2	3	4	5	6	7	A	B	C	D	E
<i>sid</i> ₁	$\langle(A, 2)(B, 5)(D, 6)(C, 10)(B, 11)\rangle$	2	4	4	6	6			1	5	4	3	0
<i>sid</i> ₂	$\langle(B, 1)(A, 2)(A, 9)(D, 12)(C, 15)(A, 18)(B, 24)\rangle$	3	3	4	5	6	7	8	6	7	5	4	0
<i>sid</i> ₃	$\langle(A, 2)(B, 4)(D, 6)(D, 8)(B, 10)(E, 12)(C, 14)\rangle$	3	4	5	6	7	8	8	1	5	7	4	6
<i>sid</i> ₄	$\langle(A, 1)(C, 2)(C, 3)(B, 4)\rangle$	4	5	5	5				1	4	3	0	0

Table 1: **a)** A sequence database SDB, **b)** a structure for the next position of minimum gap time N (precomputed) and **c)** the last position map.

Definition 3 Prefix, prefix-projected database. Let $\alpha = \langle\alpha_1 \dots \alpha_k\rangle$ be a pattern. If a sequence $\beta = \langle\beta_1 \dots \beta_n\rangle$ is a super-sequence of α : $\alpha \preceq \beta$, then the *prefix* of β w.r.t. α is the ‘*smallest prefix*’ of β that is still a super-sequence of α : $\langle\beta_1 \dots \beta_j\rangle$ s.t. $\alpha \preceq \langle\beta_1 \dots \beta_j\rangle$ and $\nexists j' < j : \alpha \preceq \langle\beta_1 \dots \beta_{j'}\rangle$. The sequence $\langle\beta_{j+1} \dots \beta_n\rangle = \text{suffix}_\alpha(\beta)$ is called the suffix and it represents the prefix-projection obtained by *projecting* the prefix away. A prefix-projected database of a pattern α , denoted by $SDB|_\alpha$, is the set of prefix-projections of all sequences in SDB that are super-sequences of α : $SDB|_\alpha = \{(sid_i, \text{suffix}_\alpha(sid_i)) \mid \alpha \preceq SDB[sid_i]\}$.

Example 4 Consider our running example in Table 1a, where we omit timing information. Assume $\alpha = \langle A \rangle$, then $SDB|_\alpha = \{(sid_1, \langle BDCB \rangle), (sid_2, \langle ADCAB \rangle), (sid_3, \langle BDDBEC \rangle), (sid_4, \langle CCB \rangle)\}$ (details in Table 2a).

The *prefix-projected frequency* of the symbol $a \in I$ ($\text{freqs}(a, SDB|_\alpha)$) is the number of sequences in $SDB|_\alpha$ where this symbol appears in the suffix: $\text{freqs}(a, SDB|_\alpha) = |\{(sid, s) \in SDB|_\alpha \mid a \in \text{suffix}_\alpha(s)\}|$. We use $\text{freqs}(a)$ instead of $\text{freqs}(a, SDB|_\alpha)$ when no ambiguity is possible about the database. Thus, the prefix-projected frequencies of $SDB|_{\langle A \rangle}$ are: $\text{freqs}(A) = 1$, $\text{freqs}(B) = 4$, $\text{freqs}(C) = 4$, $\text{freqs}(D) = 3$, $\text{freqs}(E) = 1$.

Performing a depth-first search, the *PrefixSpan* algorithm starts with an empty prefix and extends a pattern with one symbol at each step. Then, it computes the projected database and prefix-projected frequencies and extends it again (using only the prefix-projected frequent items). This process continues until all sequential patterns are found. This method is efficient because it avoids extending patterns with infrequent symbols. Also, instead of storing all suffixes explicitly, it maintains just one pointer to the suffix *starting position* for each sequence. This is called the *pseudo-projected* database $pSDB|_\alpha = \{(sid_i, j+1) \in SDB\}$ such that $\langle\beta_{j+1} \dots \beta_n\rangle = \text{suffix}_\alpha(sid_i)$.

Example 5 Extending prefix $\langle A \rangle$ with $\langle D \rangle$ over $SDB|_{\langle A \rangle}$ gives $SDB|_{\langle AD \rangle} = \{(sid_1, \langle CB \rangle), (sid_2, \langle CAB \rangle), (sid_3, \langle DBEC \rangle)\}$ and can be represented as the pseudo-projected database: $pSDB|_{\langle A \rangle} = \{(sid_1, 4), (sid_2, 5), (sid_3, 4)\}$.

We now look at the subsequence relation under a $\text{gap}^{[M, N]}$ constraints, with M the minimum and N the maximum gap between two subsequent events, and under a $\text{span}^{[W, Y]}$ constraints with W the minimum and Y the maximum span between the first and last event. This requires changing the subsequence definition in Definition 1.

Definition 4 Subsequence relation under *gap* ($\preceq^{\text{gap}^{[M, N]}}$). A sequence $\alpha = \langle\alpha_1 \alpha_2 \dots \alpha_k\rangle$ is a subsequence of $s = \langle(s_1, t_1)(s_2, t_2) \dots (s_n, t_n)\rangle$ under $\text{gap}^{[M, N]}$ constraint ($\alpha \preceq^{\text{gap}^{[M, N]}} s$) iff (i) $k \leq n$; (ii) there exists a list of integers (e_1, \dots, e_k) , an

embedding, with $1 \leq e_1 \dots \leq e_k \leq n$ such that $s^s[e_i] = \alpha_i$; and (iii) the time between two consecutive events $t_{e_{i-1}}$ and t_{e_i} must be between M and N for all $i \in [2, k]$, $M \leq t_{e_i} - t_{e_{i-1}} \leq N$. An embedding (e_1, \dots, e_k) for $\alpha \preceq^{gap^{[M, N]}}$ s is called a *gap^[M, N]-embedding*.

We can similarly define $\preceq^{span^{[W, Y]}}$ where condition (iii) becomes: the time between the first event t_{e_1} and the last t_{e_k} must be between W and Y i.e. $W \leq t_{e_k} - t_{e_1} \leq Y$. We can similarly define the *gap + span* subsequence relation which contains both conditions.

Example 6 Given sid_1 with $s_1^t = \{2, 5, 6, 10, 11\}$ and sid_2 with $s_2^t = \{1, 2, 9, 12, 15, 18, 24\}$ in Table 1a. Then, $\langle ADB \rangle \preceq^{gap^{[3, 7]}}$ $SDB[sid_1]$ with embedding $(e_1, e_2, e_3) = (1, 3, 5)$ because $\{3 \leq (s_1^t[e_2] - s_1^t[e_1]) = 4 \leq 7 \text{ and } 3 \leq (s_1^t[e_3] - s_1^t[e_2]) = 5 \leq 7\}$. Note the difference between the positions e_i of the embedding and its time $s_1^t[e_i]$. As another example, $\langle ADB \rangle \not\preceq^{gap^{[3, 7]}}$ $SDB[sid_2]$ because $3 \leq (s_2^t[e_3] - s_2^t[e_2]) = 12 \not\leq 7$. Similarly, this embedding and hence the sequence respects a *span^[8, 10]* constraint in sid_1 : $8 \leq s_1^t[e_3] - s_1^t[e_1] = 9 \leq 10$.

The definition of $Cover_{SDB}(\alpha)$, $Support_{SDB}(\alpha)$ and the SPM problem can be easily adapted to use the gap/span subsequence relation instead of the original relation \preceq .

Example 7 Assume $\alpha = \langle ADC \rangle$, $\theta = 3$ and $gap^{[3, 7]}$, $Cover_{SDB}^{gap^{[3, 7]}}(\alpha) = \{sid_1, sid_2, sid_3\}$ and hence $Support_{SDB}^{gap^{[3, 7]}}(\alpha) = 3$. Thus, α is a gap-constrained sequential pattern for the given threshold.

The search space to find the sequential patterns can become intractably large. Hence, to reduce this space several algorithms rely on the *anti-monotonicity property*.

Property 1 (Anti-monotonicity). Assume C is a constraint. If a sequence s satisfies an anti-monotone constraint C then all its subsequences also satisfy C . Equivalently, if any sequence s violates an anti-monotone constraint C all its super-sequences also violate C .

The frequency with minimum *gap* (M) and maximum *span* (Y) are anti-monotone constraints but the frequency with maximum *gap* (N) constraint violates this property.

Example 8 Assuming our running example, $\langle ADC \rangle$ is frequent under $gap^{[3, 7]}$ with $\theta = 3$ but $\langle AC \rangle$ is not frequent ($Support_{SDB}^{gap^{[3, 7]}}(\langle AC \rangle) = 2 < 3$).

Fortunately, the maximum gap constraint is *prefix anti-monotone* i.e. every prefix of p satisfies the maximum *gap* constraint if p satisfies it [29]. We use this property to filter embeddings which helps us discover if a prefix is infrequent. The minimum *span* (W) does not satisfy this property, hence we use it only to verify embeddings of fully assigned patterns.

2.2 CP-based model for SPM problem

A constraint satisfaction problem [34] is defined as a triplet (V, D, C) where V is a set of decision variables with their domains D (possible values of V). C is a set of

constraints, each constraint is defined over V and restricts the possible values that these variables can take. Solving the problem of SPM using Constraint Programming (CP) consists of defining the model (V, D, C) .

We present *CP model* of sequential pattern mining introduced in [25] and the *GapSeq* [20] and *PPIC* (*Prefix Projection Incremental Counting*) [4] global constraints.

Definition 5 Variables and Domains for SPM [25]. Let l be the length of the longest sequence in SDB ($l = \max(\{size(s) \mid s \in SDB\})$); $P = [P_1, P_2, \dots, P_l]$ is an array of variables, representing a pattern, where each P_i represents the i^{th} symbol in the pattern. The domain D_i of P_i is the set of symbols \mathcal{L} plus the empty symbol ϵ : $D_i(P_i) = \{\epsilon\} \cup \mathcal{L}$.

Example 9 For instance for the dataset in Table 1a, $l = 7$, $P = [P_1, \dots, P_7]$ and for all $i \in [1, l]$, $D_i = \{\epsilon, A, B, C, D, E\}$. $\langle A, D, C \rangle$ corresponds to $P = [A, D, C, \epsilon, \epsilon, \epsilon, \epsilon]$.

Definition 6 Filtering rules. Assume $\forall i \leq l$, $p = \langle p_1, \dots, p_i \rangle$ is the assigned values of variables $\{P_1, \dots, P_i\}$, a CP model over P represents the SPM problem given a threshold θ , $gap^{[M,N]}$ and $span^{[W,Y]}$ iff the following conditions are satisfied by every valid assignment to P :

1. $P_1 \neq \epsilon$ (to avoid an empty pattern);
2. $\forall i \in \{2, \dots, l-1\} : P_i = \epsilon \Rightarrow P_{i+1} = \epsilon$ (to allow pattern with *length* $< l$);
3. Frequency constraint: $Support_{SDB}(p) \geq \theta$;
4. Frequency under $gap^{[M,N]}$ constraint: $Support_{SDB}^{gap^{[M,N]}}(p) \geq \theta$;
5. Frequency under $span^{[W,Y]}$ constraint: $Support_{SDB}^{span^{[W,Y]}}(p) \geq \theta$.

PPIC [4] global constraint. $PPIC(P, SDB, \theta)$ is a global constraint for the SPM problem without gap/span, built on the prefix-projection principle, which enforces filtering rules 1,2,3 in a single propagator. It improved on the state-of-the-art with four elements: (a) a backtracking-aware data structure inspired by *trailing-based CP technique*, (b) efficient support counting by precomputing the last positions of each symbol, (c) not scanning sequences whose prefix can not contain the symbol (precomputed) and (d) removing the infrequent symbols of the projected database only from the next domain D_{i+1} .

GapSeq [20] global constraint. $GapSeq(P, SDB, \theta, M, N)$ is a global constraint for SPM problem under $gap^{[M,N]}$ which enforces filtering rules 1,2,3,4 in a single propagator with the limitation that gap constraints are expressed in terms of position distances i.e. the gap are measured according to the number of events hence time does not matter.

Our global constraint. $PPICt(P, SDB, \theta, M, N, W, Y)$ is a global constraint for SPM problem under $gap^{[M,N]}$ and $span^{[M,N]}$ over time-stamped databases which enforces filtering rules 1 to 5 in a single propagator. Hence, gap/span constraints are expressed in terms of time, but it can be used for position distances as well.

3 Trailing-based data structure for the embedding database

In this section, we introduce the notions of *embedding database* and *extension windows* which reconsider the concept of projected database to incorporate time constraints. Then, we present a trailing-based data structure to store this database.

a) Without time constraints: Projected Database (since time does not matter we omit it)				
sid	$pSDB _{\langle A \rangle}$	$SDB _{\langle A \rangle}$'s of $pSDB _{\langle A \rangle}$	$pSDB _{\langle AD \rangle}$	$SDB _{\langle AD \rangle}$'s of $pSDB _{\langle AD \rangle}$
sid_1	1	$\langle BDCB \rangle$	3	$\langle CB \rangle$
sid_2	2	$\langle ADCAB \rangle$	4	$\langle CAB \rangle$
sid_3	1	$\langle BDDBEC \rangle$	2	$\langle DDBEC \rangle$
sid_4	1	$\langle CCB \rangle$		
b) Considering time constraints: Embedding database and extension windows				
sid	$embSDB _{\langle A \rangle}^{[3,7]}$	$ew_e^{gap^{[3,7]}}(s)$'s of $embSDB _{\langle A \rangle}^{[3,7]}$	$embSDB _{\langle AD \rangle}^{[3,7]}$	$ew_e^{gap^{[3,7]}}(s)$'s of $embSDB _{\langle AD \rangle}^{[3,7]}$
sid_1	(1)	$\langle (B, 5)(D, 6) \rangle$	(1,3)	$\langle (C, 10)(B, 11) \rangle$
sid_2	(2),(3),(6)	$\langle (A, 9), (D, 12)(C, 15), (B, 24) \rangle$	(3,4)	$\langle (C, 15)(A, 18) \rangle$
sid_3	(1)	$\langle (D, 6)(D, 8) \rangle$	(1,3),(1,4)	$\langle (B, 10)(E, 12), (E, 12)(C, 14) \rangle$
sid_4	(1)	$\langle (B, 4) \rangle$		

Table 2: Embedding database and extension windows for patterns $\langle A \rangle$ and $\langle AD \rangle$: **a)** without time constraints **b)** with time constraints ($gap^{[3,7]}$). Note that embeddings are positions in s , not timings and these positions start from 0.

3.1 Embedding database and extension windows

In fact, when having a gap constraint and using prefix-projection (see Definition 3), the assumption that a pattern can be extended with the symbol appearing after the *smallest* matching prefix does not hold anymore. That is, given a sequence, if the first embedding of the prefix cannot be extended because the gap is too small or large, there could exist another embedding that can be extended.

Example 10 Assume the pattern $\alpha = \langle ADC \rangle$ and $gap^{[3,7]}$. There are two embeddings of α in sid_3 : (1, 3, 7) and (1, 4, 7). The first embedding is not a $gap^{[3,7]}$ -embedding since $3 \leq (t_7 - t_3 = 8) \not\leq 7$ while the second one is.

Hence, it is not sufficient to store just the (suffix of the) smallest embedding as is done in the pseudo-projected database. Instead, we can store all embeddings. One can draw the parallel of this notion with the notion the pseudo-projected database $pSDB|_{\alpha}$ but instead of only storing the first embedding, we store all available embeddings:

Definition 7 Embedding database ($embSDB|_{\alpha}$). Assume a sequence $s = \langle (s_1, t_1) (s_2, t_2) \dots (s_n, t_n) \rangle$ and a subsequence $\alpha = \langle \alpha_1 \alpha_2 \dots \alpha_k \rangle$ with $k \leq n$. The set of all embeddings of α in s is $emb_{\alpha}(s) = \{(e_1, \dots, e_k) | 1 \leq e_1 \leq e_k \leq n \text{ such that } s^s[e_i] = \alpha_i\}$. The embedding database of α is now defined as $embSDB|_{\alpha} = \{(sid, emb_{\alpha}(s)) | (sid, s) \in SDB\}$.

The embedding database under $gap^{[M,N]}$ of a sequence s is the set of transactions identifiers together with all embeddings of s that satisfy the $gap^{[M,N]}$ constraint; denoted as $embSDB|_{\alpha}^{[M,N]}$. Similarly for the embedding database under $span^{[W,Y]}$ and the combination of gap and span.

GapSeq [20] stores for each embedding the position after the last embedding, called *right pattern extensions*, but this is not sufficient to support a *span* constraint. Our method will store the start and stop position of each embedding, which is sufficient for span, gap and the combination of the two.

Given a span and/or gap constraint, an embedding can only be extended with events whose timing satisfies the span/gap constraints. We call this subsequence of events the extension window of an embedding:

Definition 8 Extension Window (ew). Assume a given sequence $s = \langle (s_1, t_1) (s_2, t_2) \dots (s_n, t_n) \rangle$, a subsequence $\alpha = \langle \alpha_1 \alpha_2 \dots \alpha_k \rangle$ and a $gap^{[M, N]}$ constraint. Let $e = (e_1, e_2, \dots, e_k)$ be any valid $gap^{[M, N]}$ -embedding of α in s . The *extension window* of this embedding, denoted $ew_e^{gap^{[M, N]}}(s)$, is the subsequence $\langle (s_u, t_u) (s_{u+1}, t_{u+1}) \dots (s_{v-1}, t_{v-1}) (s_v, t_v) \rangle$ such that $(t_{e_k} + M \leq t_u) \wedge (t_v \leq t_{e_k} + N) \wedge (\nexists t'_u \in s^t, t_{e_k} + M \leq t'_u < t_u) \wedge (\nexists t'_v \in s^t, t_v < t'_v \leq t_{e_k} + N)$. The start and the end position of this extension window are respectively u and v .

Example 11 Assume $gap^{[3, 7]}$ and $\alpha = \langle A \rangle$. For sid_3 in Table 1a, there is one $gap^{[3, 7]}$ -embedding: (1) with extension window $\langle (D, 6) (D, 8) \rangle$; hence, if $\langle A \rangle$ is extended with any symbol other than D it will no longer be covered. For $\alpha = \langle A, D \rangle$ there are now two possible embeddings: (1, 3) and (1, 4). Their extension windows are: $\langle (B, 10) (E, 12) \rangle, \langle (E, 12) (C, 14) \rangle$. Table 2b shows the embeddings and extension windows for these two patterns for all sequences in the SDB of Table 1a. A comparison can be done with the same versions without time restrictions presented in Table 2a.

3.2 Trailing-based data structures

Trailing as a mechanism to restore the state. CP-Solvers implementing the depth-first search backtracking algorithms need an efficient state restoration system [34]. This system is based on the *trail* and *time-stamping* mechanism¹. *Trailing* consists of recording the changes in a node to be able to restore it later on backtrack. The main advantage of trailing is that it makes it possible to focus on the design of the filtering without worrying about the state restoration. In each node, the solver propagates each of the constraints until a fix-point is reached; hence, a state can change several times within one search node. The trail keeps a time-stamp associated with a memory location to avoid storing a same memory on the trail more than once per search node. CP-Solvers typically expose some “reversible” objects externally using this mechanism. The *reversible integer*, denoted by *rint*, is an example of “reversible” objects for the primitive type *int*. Our trailing-based data structure also uses this mechanism.

Trailing the embedding database. We introduce a trailing-based data structure to efficiently store and restore the embedding database. The key idea is to store the embedding database in ‘backtracking aware’ vectors.

This idea was introduced in *PPIC* [4] allowing to drastically speeding up the search for sequential patterns without time constraints. See an illustration of this data structure based on the projected database examples of the Table 2a in the Fig. 1a. Two reversible integers store respectively the start position in the vector (ϕ) and the number of entries (φ) in the embedding database. When branching, data is appended at the $\phi + \varphi$ position and ϕ and φ are updated. When backtracking, only the start position and number of elements need to be restored/trailed; the vector can stay unchanged in memory, with the parts after $\phi + \varphi$ overwritten later. In [4], only the sequence ids (*sids* vector) and the start position of the suffixes (*embs* vector) must be stored. This is not sufficient to handle time constraints information.

¹ Except some CP-Solvers such as Gecode, Oz/Mozart and Figaro.

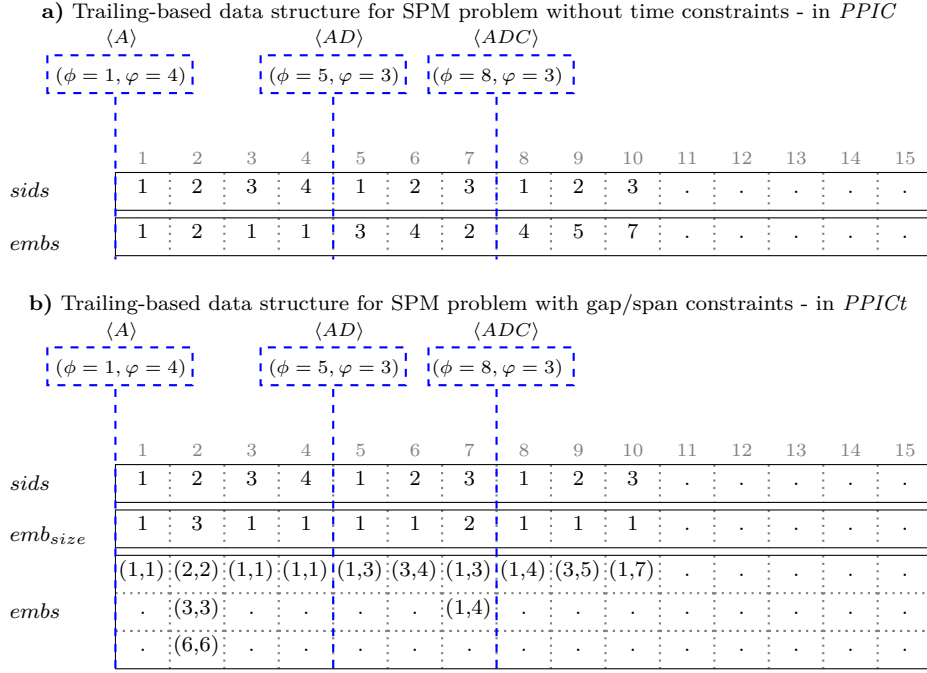


Fig. 1: Trailing-based data structure to store and restore the embeddings database: **a)** without time constraints **b)** with $gap^{[3,7]}$ constraints

Trailing-based embedding database. We use reversible vectors to store the start and the end positions of all the possible time-constrained embeddings for every sequence. The vectors: *sids*, *emb_{size}* and *embs* respectively represent the sequence ids, the number of embeddings and the start/end positions of the embeddings. These start and end positions are sufficient to verify the span and gap constraints during pattern extension.

Example 12 Figure 1b depicts an example of this data structure: for pattern $\langle A \rangle$, the data of $embSDB|_{\langle A \rangle}^{[3,7]}$ is stored between indices $\phi = 1$ and $\phi + \varphi - 1 = 4$. This pattern $\langle A \rangle$ is then further extended by the symbol *D* and $embSDB|_{\langle AD \rangle}^{[3,7]}$ is stacked next to it, between indices $\phi = 5$ and $\phi + \varphi - 1 = 7$. This pattern $\langle AD \rangle$ is then further extended by the symbol *C* and $embSDB|_{\langle ADC \rangle}^{[3,7]}$ is stacked between indices $\phi = 8$ and $\phi + \varphi - 1 = 10$. The $gap^{[3,7]}$ -embedding of $\langle ADC \rangle$ in *sid₂* is (3, 4, 5) but we only store the start (3) and end (5) as (3, 5) we can compute the valid extension window based on gap and span constraint.

4 *PPICt* global constraint under time constraints

This section presents *PPICt* (Prefix Projection Incremental Counting with time restrictions), our filtering algorithm for finding sequential patterns under gap and span

Listing 1: PPICt($P, SDB, \theta, M, N, W, Y$)

```

1 // pre: variables  $\langle P_1, \dots, P_i \rangle_{i \in [1, l]}$  are bound,  $SDB$  is given
2 //  $P_i$  is the new assigned variable since previous call (let's  $P_i = a$ ).
3 if ( $a == \epsilon$ )
4   foreach ( $j \in \{i+1, \dots, l\}$ )
5      $P_j.assign(\epsilon)$  // Filtering rule.2
6   Remove all embeddings that do not satisfy minimum span  $W$  and fail
   should the pattern no longer be frequent
7 else
8    $freqs = ProjectAndGetFreqs(i, SDB, a, M, N, W, Y)$ 
9   foreach ( $v \in D(P_{i+1})$ ) if ( $v \neq 0$  and  $freqs[v] < \theta$ )  $P_{i+1}.remove(v)$ 

```

constraints. This algorithm support sequences with timestamps (as presented in Table 1a). Constraints such as regular expression, item inclusion/exclusion, pattern length and all other constraints that do not depend on the embeddings can be added to the model. Constraints that can change what a valid embedding is would need to be added to the propagator, as we do for the time-based gap and span constraints.

4.1 PPICt filtering algorithm and improvements

The PPICt($P, SDB, \theta, M, N, W, Y$) global constraint is given in Listing 1. The filtering procedure is triggered whenever a pattern is extended by a new symbol (P_{i+1}). If that symbol is ϵ then by Rule 2 all $P_j, j > i$ should also be ϵ . The pattern is hence fully assigned and so we can filter with the minimum span constraint which is not prefix-monotone. If the pattern is still frequent then this is a new solution.

If P_{i+1} was assigned a *non- ϵ* value, then the procedure ProjectAndGetFreqs counts for each symbol what the size of the projected database would be if the new pattern is extended with that symbol. The computed result, denoted *freqs* in the pseudo-code, is used to prune the domain of the next pattern variable P_{i+1} by removing infrequent symbols; this is valid because the constraints filtered in ProjectAndGetFreqs are *prefix anti-monotone*.

The main difference with PPIC [4] is that all embeddings must be stored instead of just the prefix (see Section 3.1). Embeddings can only be extended by symbols appearing in its *extension windows*. The projected frequency counting should only count symbols appearing in an extension window. Indeed, a symbol not appearing in an extension window of any embedding of the sequence would not be a valid support for extending the current pattern as it would not satisfy the time constraints.

This leads to the following key ingredients of the ProjectAndGetFreqs function: 1) as presented in Section 3.2, we adapt the backtracking-aware data structure introduced in [4] to store all possible occurrences of the pattern in a sequence, including the starting symbol to support *span* constraints; 2) we avoid scanning a sequence for a symbol beyond the (precomputed) last occurrence of that symbol in the sequence; 3) we introduce the concept of *extension window* of an embedding and avoid to scan overlapping windows multiple times; 4) we avoid searching for the position of the start of an extension window, which depends on the minimum gap time, by precomputing these position in advance. These ingredients are detailed next.

4.1.1 Ingredient 1. Avoid scanning all sequences

We reuse the *lastPosMap* precomputed structure of *PPIC* to avoid scanning a sequence if the last position of that sequence is before the start of the extension window. For a symbol a , the *lastPosMap*[a] is the last position of this symbol in the sequence: $\text{lastPosMap}[a] = \max\{p \leq \text{size}(s) : s[p] = a\}$.

Example 13 Assuming the *lastPosMap* precomputed structure provided in Table 1c and the symbol A , *lastPosMap*[A] is $\{1, 6, 1, 1\}$. Hence, when searching for A , we must stop at the first position for the sequences sid_1 , sid_3 and sid_4 but for the sequence sid_2 we stop at position 6.

However, we cannot use the same structure for support counting (which also need to search symbols over sequences) as *PPIC* did, since this assumes that all symbols up to the end of the sequence must be counted, while we should only count symbols in the extension windows. This can have a big impact if the sequences contain many duplicates symbols as shown in [4]. In our case, this problem is minimized since extension windows are often smaller.

4.1.2 Ingredient 2. Avoid scanning more than once the events occurring in overlapping extension windows

The extension windows of a sequence can possibly overlap. For instance in Table 2b with $\alpha = \langle AD \rangle$, in sid_3 , $\langle (E, 12) \rangle$ is present in both extension windows. Then when computing the *freqs* vector, some positions could be revisited several times. This source of inefficiency can be avoided by keeping track of the current largest position visited so far in any extension window. This position is denoted *pos* in Listings 2 and 3. When the next extension window for the current sequence is considered by *updateSupport* in Listing 3, all symbols before *pos* have already been counted, so only positions after *pos* should be visited and afterwards *pos* is updated.

4.1.3 Ingredient 3. Avoid scanning the sequences when given a minimum gap

Given the current pattern $\alpha = \langle \alpha_1 \alpha_2 \dots \alpha_k \rangle$, a sequence s and a valid $\text{gap}^{[M, N]} \text{span}^{[W, Y]}$ -embedding $e = (e_1, e_2, \dots, e_k)$, all the symbols in the extension window $\text{ew}_e^{\text{gap}^{[M, N]} \text{span}^{[W, Y]}}(s)$ must be visited for updating the frequency counters. While it is easy to compute the start *time* of the extension window using the minimum gap M : $t_{e_k} + M$; finding the first position u in the sequence such that $t_u \geq t_{e_k} + M$ requires scanning the sequence starting from e_k . To avoid this, we propose to precompute, for the given minimum gap, the position of the beginning of the extension window from any possible position. This can be done with one linear scan over each sequence when the propagator is initialized; and the precomputed positions are stored in a structure called *nextPosGap*.

Definition 9 Building the nextPosGap structure. Assume $s = \langle (s_1, t_1)(s_2, t_2) \dots (s_n, t_n) \rangle$ is a sequence. Given $k \in [1, n]$ a position in s and M a minimum gap, the *nextPosGap*[s][k] is the position of the *smallest time* satisfying the minimum gap: $\text{nextPosGap}[s][k] = i$ such that $(i > k) \wedge (t_i \geq t_k + M) \wedge (\nexists i' < i : t_{i'} \geq t_k + M)$.

Listing 2: ProjectAndGetFreqs(i, SDB, a, M, N, W, Y)

```

1 // Internal state:  $\phi, \varphi, sids, emb_{size}, emb_s$ 
2  $\phi' = \phi + \varphi; \varphi' = 0; freqs[b] = 0 \forall b \in \mathcal{L}$ 
3 if  $i == 1$ : // first assigned symbol, scan for symbol
4   for  $sid = 1$  to  $size(SDB)$ : // for every sequence in SDB
5      $seq = SDB[sid]; nEmb = 0; pos = 0; visitedI[b] = false \forall b \in \mathcal{L}$ 
6     for  $j = 0$  to  $lastPosMap[sid][a]$ : // find each symbol  $a$ 
7       if  $seq[j] == a$ : // new match
8          $emb_s[\phi' + \varphi'][nEmb] = (j, j); nEmb = nEmb + 1$ 
9          $pos = updateSupport(j, j, sid, pos, visitedI)$ 
10        if  $(pos \geq size(seq))$  break // window ends with sequence
11      if  $(nEmb > 0)$  // store sequence meta-data
12         $sids[\phi' + \varphi'] = sid; emb_{size}[\phi' + \varphi'] = nEmb; \varphi' = \varphi' + 1$ 
13    else: // non-empty prefix
14      for  $c = \phi$  to  $\phi + \varphi - 1$ : // for all sequence in projected database
15         $sid = sids[c]; seq = SDB[sid]; nEmb = 0; pos = 0$ 
16         $visitedI[b] = false \forall b \in \mathcal{L}$ 
17        for  $k = 1$  to  $emb_{size}[c]$ : // for each prefix embedding
18           $(b, e) = emb_s[c][k]$  // begin and end position of embedding
19           $maxT = \min(seq^t[sid][b] + Y, seq^t[sid][e] + N)$  // max time window
20           $j = nextPosGap[sid][e]$  // precomputed position of minT
21          while  $(j < lastPosMap[sid][a]$  and  $seq^t[sid][j] \leq maxT)$  :
22            if  $seq[j] == a$ : // new embedding
23               $emb_s[\phi' + \varphi'][nEmb] = (b, j); nEmb = nEmb + 1$ 
24               $pos = updateSupport(b, j, sid, pos, visitedI)$ 
25              if  $(pos \geq size(seq))$  break // window ends with sequence
26             $j = j + 1$ 
27          if  $(nEmb > 0)$  // store sequence meta-data
28             $sids[\phi' + \varphi'] = sid; emb_{size}[\phi' + \varphi'] = nEmb; \varphi' = \varphi' + 1$ 
29         $\phi = \phi'; \varphi = \varphi'$ 
30    return  $freqs$ 

```

Listing 3: updateSupport($b, e, sid, pos, visitedI$)

```

1  $s = SDB[sid]; k = \max(nextPosGap[sid][e], pos)$ 
2  $maxT = \min(s^t[sid][e] + N, s^t[sid][b] + Y)$ 
3 while  $(k < size(s)$  and  $s^t[sid][k] \leq maxT)$  :
4   if  $(!visitedI[s[k]])$   $freqs[s[k]] = freqs[s[k]] + 1; visitedI[s[k]] = true$ 
5    $k = k + 1$ 
6 return  $k$ 

```

Example 14 Assume $s = \langle (A, 2)(B, 5), (D, 6), (C, 10)(B, 11) \rangle$, $k = 2$ and $M = 3$. $nextPosGap[s][k] = 4$ because $t_4 = 10$ is the smallest time such that $t_4 \geq 5 + 3 = 8$. Table 1b shows the $nextPosGap$ of SDB (the values $nextPosGap[s][k] > size(s) + 1$ means the minimum gap is not available for that position).

4.1.4 Putting it all together

The core of the algorithm is in the ProjectAndGetFreqs procedure (presented in Listing 2) that gathers all the ingredients. We distinguish two cases. Assuming $i \in [1, l]$, if $i == 1$ it means that the pattern was previously empty and is now composed of one unique symbol. If $(i > 1)$ the pattern is composed of at least two symbols which means that the gap/span must be considered.

In the first case ($i == 1$), all sequences of SDB are considered. For every sequence, all the positions having the symbol a are stored as an embedding. As the embedding is a singleton, there is not need to consider the gap/span constraints at this point.

The call to `updateSupport` (Listing 3) will update the *freqs* vector by visiting each symbol present in the extension window of the current embedding (position of symbol a). Variable *pos* is used to avoid incrementing the frequency of a symbol twice in the same sequence.

In the second case ($i > 1$), the main loop at line 14 iterates over the previous (*parent*) projected database stored between ϕ and $\phi + \varphi - 1$ and builds the new one starting at index $\phi + \varphi$. For each embedding of a sequence, s (line 17), the maximum of the time window is computed. We search all positions only in extension window ensuring to have time greater than minimum gap time and lower than the maximum gap and span times computed based on the first and the last element of the embedding (line 19). The `updateSupport` is also called to update the *freqs* vector for every extended embedding created.

Finally, lines 12 and 28 update *sids* and *emb_size* in order to ensure the consistency of the data structure. Then, line 29 updates the reversible integers ϕ and φ to reflect the newly computed projected database.

4.2 Additional constraints

The advantage of CP based sequence mining is its capacity to accept additional constraints. The global constraint approach *PPICt* is less flexible than the decomposition approach of [25] as it does not expose the embeddings. Nevertheless many useful syntax constraints [14] can be added on the sequence pattern variables: $P = [P_1, P_2 \dots, P_l]$. Note that these constraints are directly² posted as independent constraints in the solver constraint store and run together with the mining process. We can indeed post any such constraint, for example:

Pattern length constraints. One can impose a minimum and a maximum over the length of the pattern. These constraints are easy to handle considering all patterns are terminated by the empty symbol (ϵ). Hence, the minimum pattern length (*Lmin*) is defined as $\forall i \in [1, Lmin]$ and $Lmin < l$, $P_i \neq \epsilon$. The Maximum pattern length (*Lmax*) is obtained by limiting the length of P to *Lmax*: $P = [P_1, P_2 \dots, P_{Lmax}]$ with $Lmax < l$.

Symbol inclusion/exclusion. The number of occurrences of symbols in the sequence pattern can be modeled with *Among* [8] and global cardinality (*GCC*) [33] constraints largely available in CP-Solvers. Considering $v \geq 0$ a given number of occurrences of symbol $i \in \mathcal{L}$, the *Among*(P, i, v) constraint will only allow a number v times i ($v = 0$ prohibits i). To handle several symbols inclusion/exclusion, one can use several *Among* constraints or *GCC*(P, \mathcal{V}) constraint with \mathcal{V} a collection of pairs (v, i) . *GCC* also offers the possibility of defining range of values of occurrences.

Regular expression and grammar. The *Regular* global constraint [30] can be used to enforce that P satisfies a given regular expression. Most of CP-based approaches [4, 20, 19] hence easily support regular expression constraints. Considering e a user-supplied regular expression and \mathcal{A} its deterministic finite automaton,

² It is neither post-processing nor hard-coded

the $Regular(P, \mathcal{A})$ constraint only allows patterns matching \mathcal{A} . Since a regular expression is a formal language, the grammar constraint can also be used. However, grammar constraint could also use to define context-free languages [17] or general languages [32].

4.3 Time and space complexity

Let us denote by $m = size(SDB)$ the number of sequences, l the length of the longest sequence, $L = size(\mathcal{L})$ the size of item alphabet. In the worst case, the time complexity of our propagator is in $O(m \times l^2 + L)$ and the space complexity is in $O(m \times l^2)$.

Intuition (Space complexity) Our data structure needs $O(l)$ memory entries to store the embeddings for one sequence of length one. The projected database for such a sequence thus requires $O(m \times l)$. During search, the search will branch over each sequence variable in turn which corresponds to extending the pattern by one symbol at a time. Each time a symbol is added, an extra layer of embeddings is stored in our trailing-based data structure (Fig. 1). Given that each pattern has a length of at most l , the space complexity of our data structure is hence $O(l \times m \times l) = O(m \times l^2)$. \square

Intuition (Time complexity) $PPICt$ needs $O(l + (m \times l^2) + L) = O(m \times l^2 + L)$ time to be complete since the lines 4-5 of the Listing 1 cost $O(l)$, the ProjectAndGetFreqs method $O(m \times l^2)$ and line 9 costs $O(L)$. The complexity of the loops 6 and 21 including the updateSupport in Listing 2 is $O(l)$ since we avoid scanning overlapping symbols. Hence, the complexity of the ProjectAndGetFreqs method is $O(m \times l + m \times l \times l) = O(m \times l^2)$. \square

5 Experiments

This section reports the experiments we made to evaluate our approach in comparison with other CP-based and specialized methods. More specially, we answer the following questions: *Q1*. What is the performance of the state-of-the-art for sequential pattern mining without time constraints? *Q2*. What is the difference in performance of $PPICt$ for sequential patterns mining with time restrictions? *Q3*. What is the effect of a standalone constraint in the mining process? *Q4*. What is the impact of the computation of the additional embeddings in $PPICt$?

Before answering these questions, we present in Table 3 the features of the seven real-life datasets that we use, as well as the experimental protocol and the alternative sequential pattern miners used for the comparisons. Note that the data and framework are available online and open source³.

Experimental protocol

$PPICt$ is implemented in the Scala language in the CP-Solver OcaR [26]. All experiments are run in the JVM with maximum memory set to 8GB. All the experiments

³ <http://sites.uclouvain.be/cp4dm/spm/ppict/>

SDB	$size(\mathcal{L})$	$size(SDB)$	$allsymbols(SDB)$	$max\ size(s)$	$avg.\ size(s)$	density
BIBLE	13905	36369	787066	100	21.64	0.0016
FIFA	2990	20450	741092	100	36.24	0.0121
Kosarak	21144	69999	558373	796	7.98	0.0004
LEVIATHAN	9025	5834	197251	100	33.81	0.0037
MSNBC	17	31790	423776	100	13.33	0.7841
PubMed	19931	17237	509440	198	29.56	0.0015
protein	25	103120	49729890	600	482.25	19.2901

Table 3: Seven real-life datasets features. Respectively: dataset name, number of distinct symbols, number of sequences, total number of symbols in the dataset, maximum sequence length, average sequence length, and density computed by $\frac{allsymbols(SDB)}{size(\mathcal{L}) \times size(SDB)}$ (Kosarak is the sparsest dataset and Protein is the densest).

Methods	Frequency	Gap	Span	Regular	Among	Length	other constraints ¹
<i>PPICt</i>	x	x	x	x	x	x	x
<i>GapSeq</i>	x	x*			x	x	x
<i>cSPADE</i>	x	x	x**			x	

Table 4: Sequential pattern miners with supported constraints. ¹ is any other user-supplied constraint that does not depend on the embeddings (not implemented but could be). **GapSeq* does not consider time but position of events, ***cSPADE* does not support minimum span constraint.

are conducted using a 2.7GHz Intel Core i5 64 bit processor and 8GB of RAM with Linux 3.19.0-32-generic from Mint 17.3. We set the execution time limit to 3600 seconds (1 hour). We also restrict the output of all software to only the mining statistics and do not print the patterns found. The minimum support θ is denoted by *Minsup*.

Alternative sequential patterns miners

We make comparisons with *GapSeq*⁴ [20], a CP approach that outperforms other CP-based methods supporting gap constraints; *cSPADE*⁵ [40] a highly scalable specialized sequence miner that supports gap and span constraints. Its search is not based on pattern extension as *GapSeq* and *PPICt* are, but on repeated (temporal) joins of embeddings. We also provide a comparison to *PPIC*⁶ [4] without gap constraints, *PPIC* has shown to outperform both specialized and generic miners for standard frequent sequence mining. Table 4 shows the supported constraints for these miners.

⁴ <https://sites.google.com/site/cp4spm/>

⁵ <http://www.cs.rpi.edu/~zaki/www-new/pmwiki.php/Software>

⁶ <http://sites.uclouvain.be/cp4dm/spm/>




Fig. 2: CPU times for PPIC (without time constraints) with several minsup (missing points indicate a timeout) [4]

5.1 Performances results

5.1.1 Q1: GapSeq vs PPIC vs cSPADE for SPM without time restriction

As shown in [4] and illustrated in Fig. 2, *PPIC* clearly outperforms both CP-based and specialized approaches for many datasets (with several different features) except for the sparsest dataset Kosarak-70k where it is competitive with *cSPADE*. The protein dataset is large and dense with many patterns even at high support. For such a challenging dataset, *PPIC* is at least one hundred times faster.

5.1.2 Q2: Time performance for PPICt under gap and span constraints

We first compare *PPICt* with *GapSeq* and *PPIC* for gap constraints. Then, we combine gap and span constraints.

Figure 3 shows the CPU time for the sequence mining task under minimum and maximum gap for several θ (*Minsup*) values over six datasets. *PPICt* dominates both CP-based and specialized methods. Except for the Kosarak dataset, *PPICt* is often faster, and increasingly so for low frequency thresholds. Upon inspecting the output of the Kosarak dataset we see that several frequent patterns have the same size and cover the same set of sequences. The temporal join approach used by *cSPADE* is very fast in this case. This was also the case for *PPIC* in the non-time constrained case.

We also combine the *gap* and *span* constraints, which is not supported by *GapSeq*. The results are presented in Fig. 4. Our approach outperforms *cSPADE* by a wide margin in this case. These results show that *PPICt* is still efficient when combining time constraints.

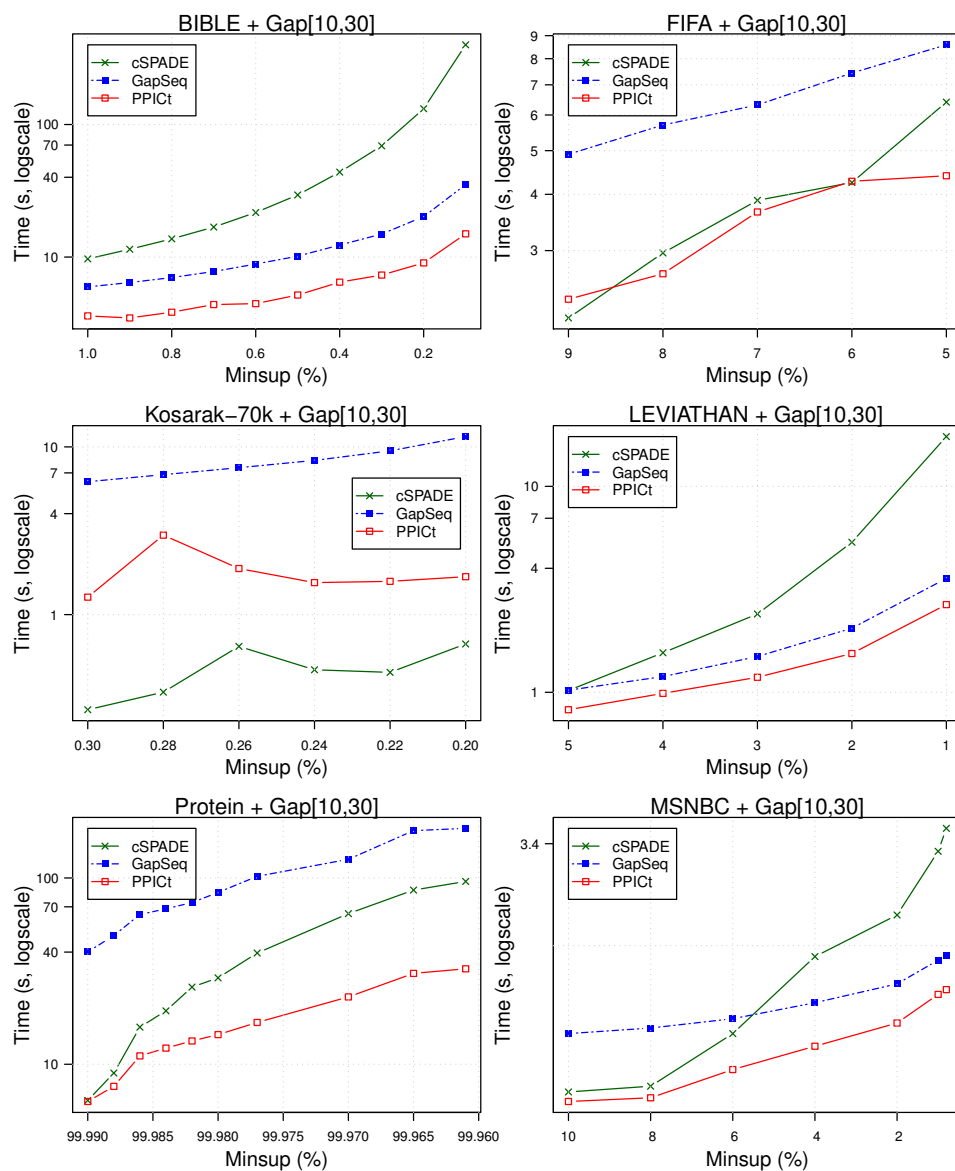


Fig. 3: CPU times when considering minimum and maximum gap constraints for several minsup (missing points indicate a timeout)

5.1.3 Q3: Effect of maximum gap constraint

We now look at the sensitivity of the methods to the threshold of the maximum gap constraint. We fix the frequency threshold to a low value that makes mining without further constraints challenging and increase the maximum gap constraint from 1 to

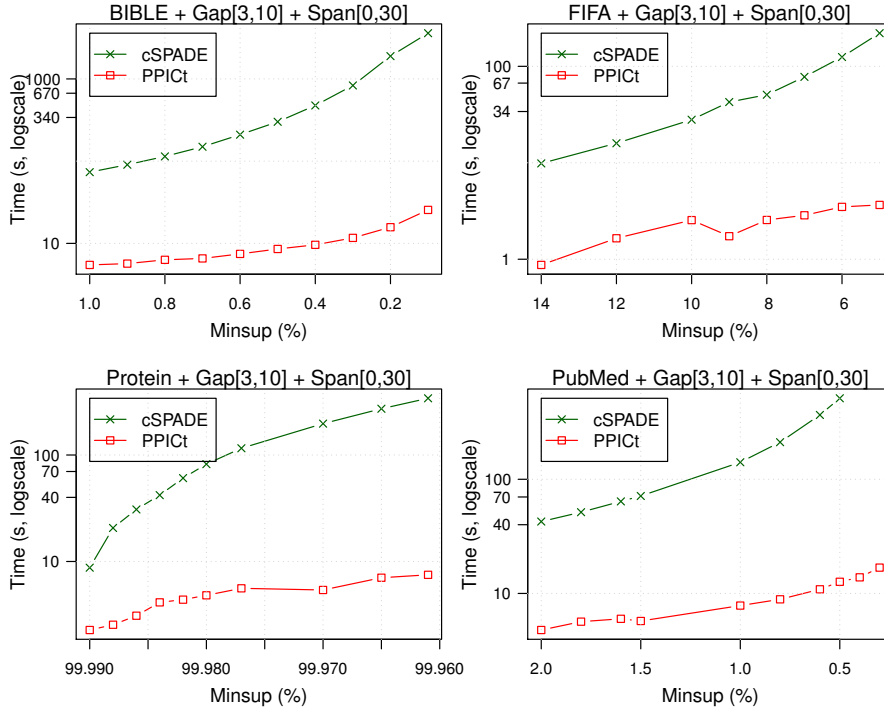


Fig. 4: CPU times when considering both *gap* and *span* constraints for several minsup (missing points indicate a timeout).

9. As can be seen in Fig. 5, the runtime of *cSPADE* increases much more quickly with increasing maximum gap. For *GapSeq* it depends on the dataset, but *PPICt*'s performance is more stable and increases more moderately compared to the other methods.

5.1.4 Q4: Experiments over databases without time restrictions

To answer *Q4*., we use *PPICt* to find only the sequential patterns without any time considerations. That is *PPICt* where minimum gap/span is 0 and maximum gap/span is the infinity, denoted by *PPICt*[0,Inf]. Hence, we compare *PPIC* with *PPICt*[0,Inf]. The results are reported in Fig. 6. We can notice that *PPIC* is always faster. This is possible since such *PPIC* improvements could not be used under time restrictions. Moreover, to preserve the structure of datasets the reduction of datasets by preprocessing is forbidden.

5.2 Handling additional Constraints

To demonstrate the ability to accommodate additional constraints we experiment the combination of *PPICt* with some other syntax constraints. The result is shown

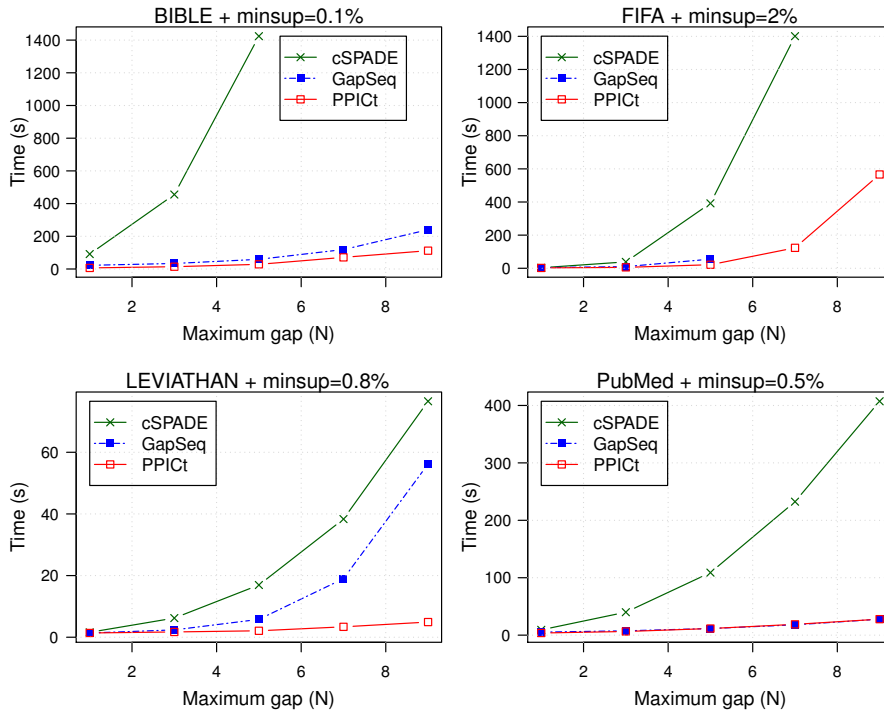


Fig. 5: CPU times for several maximum gap with fixed minsup over Bible, Fifa, Leviathan and PubMed datasets (missing points indicate a timeout).

	Gap		+Pattern Length		+Among		+Regular	
	nSols	time(s)	nSols	time(s)	nSols	time(s)	nSols	time(s)
BIBLE*	32307	46.181	1542	45.622	171	43.390	8	0.191
PubMed**	13086	22.632	1304	21.600	235	19.889	3	0.091

* $\theta = 0.1\% \wedge \text{Gap}[10, 30] \wedge (Lmin = Lmax = 5) \wedge$ the number of A equal to 1 $\wedge E$ is forbidden $\wedge \text{Regular}(A+(B\{2,\} | C* | D+) B*C*D*)$ where ($A = 11829, B = 2, C = 8212, D = 6556, E = 5590$)

** $\theta = 0.3\% \wedge \text{Gap}[10, 30] \wedge (Lmin = Lmax = 4) \wedge$ the number of A and B equal to 1 $\wedge \text{Regular}(B+A*C*A*)$ where ($A = 3335, B = 12155, C = 16599$)

Table 5: Combination of pattern length, item inclusion/exclusion, regular expression constraints with gap constraint.

in Table 5; the constraint parameters were artificially constructed in an interactive setting. We can observe that the addition of the constraints reduces the number of solutions and the computation time. A *generate-and-filter* approach using a specialized algorithm would not be able to benefit from stronger filtering.




Fig. 6: Comparing $PPICt$ without time restriction ($PPICt[0,Inf]$) with $PPIC$

6 Conclusion

We introduced $PPICt$, a global constraint to solve sequential pattern mining problem under time constraints. It integrates *gap* and *span* constraints for databases with or without timestamps. Our approach often outperforms *cSPADE*, the state-of-the-art specialized method and always outperforms *GapSeq*, the state-of-the-art CP based approach allowing to handle time constraints. This was made possible thanks to the backtracking-aware data structure to store embeddings of pattern based on trailing techniques. Also, algorithmic ingredients help to improve further: the precomputed next position of minimum gap, the avoidance of scanning all dataset and the avoidance of the overlapping between extension windows when computing the frequencies of symbols. Moreover, we report experimental results over several real-life datasets which demonstrate that our proposal is mostly competitive with or outperforms both specialized and CP-based methods. Additional constraints such as regular expression, item inclusion/exclusion, pattern length constraints are also available to increase the flexibility of users and practitioners.

Acknowledgements The research is supported by the FRIA-FNRS (Fonds pour la Formation à la Recherche dans l'Industrie et dans l'Agriculture, Belgium) and FWO (Research Foundation – Flanders).

References

1. Aggarwal, C.C., Han, J.: Frequent pattern mining. Springer (2014)
2. Agrawal, R., Srikant, R.: Mining sequential patterns. In: Data Engineering, 1995. Proceedings of the Eleventh International Conference on, pp. 3–14. IEEE (1995)
3. Antunes, C., Oliveira, A.L.: Generalization of pattern-growth methods for sequential pattern mining with gap constraints. In: P. Perner, A. Rosenfeld (eds.) Machine Learning and Data Mining in Pattern Recognition: 3rd International Conference, MLDM 2003 Leipzig, Germany, July 5–7, 2003 Proceedings, pp. 239–251. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
4. Aoga, J.O.R., Guns, T., Schaus, P.: An efficient algorithm for mining frequent sequence with constraint programming. In: P. Frasconi, N. Landwehr, G. Manco, J. Vreeken (eds.) Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19–23, 2016, Proceedings, Part II, pp. 315–330. Springer International Publishing, Cham (2016)
5. Aoga, J.O.R., Guns, T., Schaus, P.: Mining time-constrained sequential patterns with constraint programming. In: D. Salvagnin, M. Lombardi (eds.) Integration of AI and OR Techniques in Constraint Programming - 13th International Conference, CPAIOR 2017, Padova, Italy, June 5 - 8, 2017, Proceedings, Lecture Notes in Computer Science. Springer (2017)
6. Ayres, J., Flannick, J., Gehrke, J., Yiu, T.: Sequential pattern mining using a bitmap representation. In: Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, July 23–26, 2002, Edmonton, Alberta, Canada, pp. 429–435. ACM (2002)
7. Batal, I., Fradkin, D., Harrison, J., Moerchen, F., Hauskrecht, M.: Mining recent temporal patterns for event detection in multivariate time series data. In: Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 280–288. ACM (2012)
8. Beldiceanu, N., Contejean, E.: Introducing global constraints in chip. *Mathematical and computer Modelling* **20**(12), 97–123 (1994)
9. Coquery, E., Jabbour, S., Saïs, L., Salhi, Y.: A sat-based approach for discovering frequent, closed and maximal patterns in a sequence. In: L.D. Raedt, C. Bessière, D. Dubois, P. Doherty, P. Frasconi, F. Heintz, P.J.F. Lucas (eds.) ECAI 2012 - 20th European Conference on Artificial Intelligence. Montpellier, France, August 27–31, 2012, *Frontiers in Artificial Intelligence and Applications*, vol. 242, pp. 258–263. IOS Press (2012)
10. Desai, N.A.K., Ganatra, A.: Efficient constraint-based sequential pattern mining (spm) algorithm to understand customers buying behaviour from time stamp-based sequence dataset. *Cogent Engineering* **2**(1), 1072,292 (2015)
11. Fournier-Viger, P., Wu, C.W., Tseng, V.S.: Mining maximal sequential patterns without candidate maintenance. In: *Advanced Data Mining and Applications*, pp. 169–180. Springer (2013)
12. Guns, T., Nijssen, S., De Raedt, L.: k-pattern set mining under constraints. *Knowledge and Data Engineering, IEEE Transactions on* **25**(2), 402–418 (2013)
13. Han, J., Pei, J., Yin, Y., Mao, R.: Mining frequent patterns without candidate generation: A frequent-pattern tree approach. *Data mining and knowledge discovery* **8**(1), 53–87 (2004)
14. He, J., Flener, P., Pearson, J., Zhang, W.M.: Solving string constraints: The case for constraint programming. In: *International Conference on Principles and Practice of Constraint Programming*, pp. 381–397. Springer (2013)
15. Henriques, R., Antunes, C., Madeira, S.C.: Methods for the efficient discovery of large item-indexable sequential patterns. In: A. Appice, M. Ceci, C. Loglisci, G. Manco, E. Masciari, Z.W. Ras (eds.) *New Frontiers in Mining Complex Patterns: Second International Workshop, NFMCP 2013, Held in Conjunction with ECML-PKDD 2013, Prague, Czech Republic, September 27, 2013, Revised Selected Papers*, pp. 100–116. Springer International Publishing, Cham (2014)
16. Henriques, R., Madeira, S.C.: Bicspam: flexible biclustering using sequential patterns. *BMC Bioinformatics* **15**(1), 130 (2014)
17. Kadioglu, S., Sellmann, M.: Grammar constraints. *Constraints* **15**(1), 117–144 (2010)
18. Kemmar, A., Lebbah, Y., Loudni, S., Boizumault, P., Charnois, T.: Prefix-projection global constraint and top-k approach for sequential pattern mining. *Constraints* **22**(2), 265–306 (2017)

19. Kemmar, A., Loudni, S., Lebbah, Y., Boizumault, P., Charnois, T.: Prefix-projection global constraint for sequential pattern mining. In: G. Pesant (ed.) *Principles and Practice of Constraint Programming: 21st International Conference, CP 2015, Cork, Ireland, August 31 – September 4, 2015, Proceedings*, pp. 226–243. Springer International Publishing, Cham (2015)
20. Kemmar, A., Loudni, S., Lebbah, Y., Boizumault, P., Charnois, T.: A global constraint for mining sequential patterns with GAP constraint. In: C. Quimper (ed.) *Integration of AI and OR Techniques in Constraint Programming - 13th International Conference, CPAIOR 2016, Banff, AB, Canada, May 29 - June 1, 2016, Proceedings, Lecture Notes in Computer Science*, vol. 9676, pp. 198–215. Springer (2016)
21. Li, C., Wang, J.: Efficiently mining closed subsequences with gap constraints. In: *Proceedings of the SIAM International Conference on Data Mining, SDM 2008, April 24–26, 2008, Atlanta, Georgia, USA*, pp. 313–322. SIAM (2008)
22. Lu, S., Li, C.: Aprioriadjust: An efficient algorithm for discovering the maximum sequential patterns. In: *Proc. Intern. Workshop Knowl. Grid and Grid Intell (2004)*
23. Mannila, H., Toivonen, H., Verkamo, A.I.: Discovery of frequent episodes in event sequences. *Data mining and knowledge discovery* **1**(3), 259–289 (1997)
24. Metivier, J., Boizumault, P., Crémilleux, B., Khiari, M., Loudni, S.: A constraint-based language for declarative pattern discovery. In: *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*, pp. 1112–1119. IEEE (2011)
25. Nègrevergne, B., Guns, T.: Constraint-based sequence mining using constraint programming. In: L. Michel (ed.) *Integration of AI and OR Techniques in Constraint Programming - 12th International Conference, CPAIOR 2015, Barcelona, Spain, May 18–22, 2015, Proceedings, Lecture Notes in Computer Science*, vol. 9075, pp. 288–305. Springer (2015)
26. Oscala Team: Oscala: Scala in OR (2012). Available from <https://bitbucket.org/oscarlib/oscar>
27. Parthasarathy, S., Zaki, M.J., Ogiwara, M., Dwarkadas, S.: Incremental and interactive sequence mining. In: *Proceedings of the 8th international conference on Information and knowledge management*, pp. 251–258. ACM (1999)
28. Pei, J., Han, J., Mortazavi-Asl, B., Pinto, H., Chen, Q., Dayal, U., Hsu, M.C.: Prefixspan: Mining sequential patterns efficiently by prefix-projected pattern growth. In: *proceedings of the 17th international conference on data engineering*, pp. 215–224. IEEE (2001)
29. Pei, J., Han, J., Wang, W.: Constraint-based sequential pattern mining: the pattern-growth methods. *Journal of Intelligent Information Systems* **28**(2), 133–160 (2007)
30. Pesant, G.: A regular language membership constraint for finite sequences of variables. In: *International conference on principles and practice of constraint programming*, pp. 482–495. Springer (2004)
31. Pinto, H., Han, J., Pei, J., Wang, K., Chen, Q., Dayal, U.: Multi-dimensional sequential pattern mining. In: *Proceedings of the tenth international conference on Information and knowledge management*, pp. 81–88. ACM (2001)
32. Quimper, C.G., Walsh, T.: Global grammar constraints. In: *International Conference on Principles and Practice of Constraint Programming*, pp. 751–755. Springer (2006)
33. Régim, J.C.: Generalized arc consistency for global cardinality constraint. In: *Proceedings of the thirteenth national conference on Artificial intelligence-Volume 1*, pp. 209–215. AAAI Press (1996)
34. Rossi, F., Van Beek, P., Walsh, T.: *Handbook of CP*. Elsevier (2006)
35. Srikant, R., Agrawal, R.: Mining sequential patterns: Generalizations and performance improvements. Springer (1996)
36. Tatti, N., Cule, B.: Mining closed episodes with simultaneous events. In: *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '11*, pp. 1172–1180. ACM, New York, NY, USA (2011).
37. Wang, J., Han, J., Li, C.: Frequent closed sequence mining without candidate maintenance. *IEEE Transactions on Knowledge and Data Engineering* **19**(8), 1042–1056 (2007)
38. Yan, X., Han, J., Afshar, R.: Clospan: Mining: Closed sequential patterns in large datasets. In: *Proceedings of the 2003 SIAM International Conference on Data Mining*, pp. 166–177. SIAM (2003)
39. Zaki, M.J.: Efficient enumeration of frequent sequences. In: *Proceedings of the seventh international conference on Information and knowledge management*, pp. 68–75. ACM (1998)
40. Zaki, M.J.: Sequence mining in categorical domains: incorporating constraints. In: *Proceedings of the ninth international conference on Information and knowledge management*, pp. 422–429. ACM (2000)

-
41. Zhao, Q., Bhowmick, S.S.: Sequential pattern mining: A survey. ITechnical Report CAIS Nayang Technological University Singapore pp. 1–26 (2003)