

Efficiently Mining Frequent Closed Partial Orders (Extended Abstract)

Jian Pei¹ Jian Liu² Haixun Wang³ Ke Wang¹ Philip S. Yu³ Jianyong Wang⁴

¹ Simon Fraser Univ., Canada, {jpei, wangk}@cs.sfu.ca ² State Univ. of New York at Buffalo, USA, liu8@cse.buffalo.edu

³ IBM T.J. Watson Research Center, USA, {haixun, psyu}@us.ibm.com ⁴ Tsinghua Univ., China, jianyong@mail.tsinghua.edu.cn

1 Introduction

Mining ordering information from sequence data is an important data mining task. Sequential pattern mining [1] can be regarded as mining frequent segments of total orders from sequence data. However, sequential patterns are often insufficient to concisely capture the general ordering information.

Example 1 (Motivation) Suppose MapleBank in Canada wants to investigate whether there is some orders which customers often follow to open their accounts. A database DB in Table 1 about four customers' sequences of opening accounts in MapleBank is analyzed.

Given a support threshold min_sup , a sequential pattern is a sequence s which appears as subsequences of at least min_sup sequences. For example, let $min_sup = 3$. The following four sequences are sequential patterns since they are subsequences of three sequences, 1, 2 and 4, in DB .

CHK → MMK → MORT → RESP;
CHK → MMK → MORT → BROK;
CHK → RRSP → MORT → RESP;
CHK → RRSP → MORT → BROK

The sequential patterns capture the frequent account opening patterns shared by customers. However, the four sequential patterns cannot completely capture the ordering shared by customers 1, 2 and 4. It is easy to see that a partial order R as shown in Figure 1 is shared by the three account opening sequences. The partial order R summarizes the four sequential patterns – the four sequential patterns are paths in partial order R . It also provides more information about the ordering than the sequential patterns. ■

The knowledge about ordering, especially the frequent partial orders in string databases, has many applications. For example, ordering information is often important in analysis of biological experiment data. To discover patterns in gene expression matrices, one promising approach [2] is to look for order-preserving submatrices (OPSMs). That is, in an n by m gene expression matrix for n genes and m experiments, each element $v_{i,j}$ gives the expression level of a gene g_i in an experiment e_j . A submatrix is order-

Account codes and explanation

Account code	Account type
CHK	Checking account
MMK	Money market
RRSP	Retirement Savings Plan
MORT	Mortgage
RESP	Registered Education Savings Plan
BROK	Brokerage

Customer Records

Cid	Sequence of account opening
1	CHK → MMK → RRSP → MORT → RESP → BROK
2	CHK → RRSP → MMK → MORT → RESP → BROK
3	MMK → CHK → BROK → RESP → RRSP
4	CHK → MMK → RRSP → MORT → BROK → RESP

Table 1. A database DB of sequences of account opening.

preserving if the expression levels of all genes in the submatrix induce the same (linear or partial) ordering of the experiments. Such a pattern may arise if the experiments in the order-preserving submatrix represent distinct stages in the progress of a disease or in a cellular process, and the expression levels of all genes in the submatrix vary across the stages in the same way.

The problem of mining partial orders in sequence data has been studied before from angles different from our study. Two major categories of previous work exist. The first batch of studies only look at some specific kinds of frequent partial orders, such as serial and parallel orders [6]. Another category of research is on mining global partial orders [5, 4]. However, very different from the problem studied here, [5] tried to find one or a (small) set of partial orders that fit the *whole* data set as well as possible, which was modeled as an optimization problem. An implicit assumption is that the whole data set somehow follows a global trend. For some applications, such as the DNA microarray data analysis and network packet routing, there is usually no non-trivial order that can be expected globally. This paper addresses such situations. That is, we want to find the

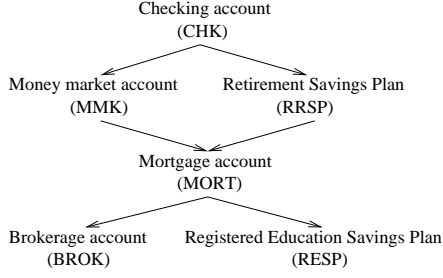


Figure 1. A frequent partial order R in Example 1.

complete set of partial orders that are frequent in a database, but not necessarily dominate the database.

Simultaneously to this study, Casas-Garriga [3] also proposed a similar idea to use closed partial orders to summarize sequential data. However, [3] focuses on the concept and does not provide an efficient algorithm for the mining. The method in [3] first mines the complete set of frequent closed sequential patterns and then derives closed partial orders from closed sequential orders. As shown in our experimental results, the method may not be efficient for large data sets. Therefore, mining frequent partial orders from sequence data remains a challenging problem.

2 Problem Definition

A *partial order* is a binary relation that is reflexive, antisymmetric, and transitive. A *total order* (or called linear order) is a partial order R such that for any two items x and y , if $x \neq y$ then either $R(x, y)$ or $R(y, x)$ holds.

A partial order R can be expressed in its *transitive closure* as a directed acyclic graph (DAG for short): the items are the vertices in the graph and $x \rightarrow y$ is an edge if and only if $(x, y) \in R$ and $x \neq y$. We also write an edge (x, y) as xy or $x \rightarrow y$.

A transitive closure may contain redundant edges. Generally, an edge $x \rightarrow y$ is *redundant* if there is a path from x to y that does not contain the edge. For a partial order R , the *transitive reduction* of R can be drawn in a *Hasse diagram*: for $(x, y) \in R$ and $x \neq y$, x is positioned higher than y ; edge $x \rightarrow y$ is drawn if and only if the edge is not redundant. The transitive closure may have much more edges than the transitive reduction.

Let V be a set of items, which serves as the domain of our string database. A *string* defines a global order on a subset of V . In this paper, we focus on strings instead of general sequences, and assume that each item appears in a string at most once, but every item not necessarily appears in a string.

A string can be written as $s = x_1 \cdots x_l$, where $x_1, \dots, x_l \in V$. l is called the *length* of string s , i.e., $len(s) = l$.

For strings $s = x_1 \cdots x_l$ and $s' = y_1 \cdots y_m$, s is said a *super-string* of s' and s' a *sub-string* of s if (1) $m \leq l$ and (2) there exist integers $1 \leq i_1 < \cdots < i_m \leq l$ such that $x_{i_j} = y_j$ ($1 \leq j \leq m$). We also say s contains s' . For a string database SDB , the *support* of a string s , denoted by $sup(s)$, is the number of strings in SDB that are super-strings of s .

The total order defined by string s can also be written in its *transitive closure* $\mathcal{C}(s) = \{(x_i, x_j) | 1 \leq i < j \leq l\}$. Please note that, in a transitive closure, we omit the trivial pairs (x_i, x_i) .

The *order containment relation* is defined as, for two partial orders R_1 and R_2 , if $R_1 \subset R_2$, then R_1 is called *weaker* than R_2 and R_2 is *stronger* than R_1 . By intuition, a partially ordered set (or poset for short) satisfying R_2 will also satisfy R_1 .

A *string database* SDB is a multiset of strings. For a partial order R , a string s is said to *support* R if $R \subseteq \mathcal{C}(s)$. The *support of R in SDB* , denoted by $sup(R)$, is the number of strings in SDB that support R . Given a minimum support threshold min_sup , a partial order R is called *frequent* if $sup(R) \geq min_sup$.

In a string database, for a minimum support threshold, there can be many frequent partial orders. To avoid the triviality, instead of reporting all frequent partial orders, we can only mine the representative ones.

Example 2 (Frequent closed partial orders) Consider string database DB in Table 1 again. There exists no another partial order R' such that R' is stronger than R in Figure 1 and is also shared by strings 1, 2 and 4. In other words, R is the strongest one among all frequent partial orders shared by strings 1, 2 and 4. Thus, the partial order R is not redundant and can be used as a representative of the frequent partial orders shared by strings 1, 2 and 4. Technically, R is a frequent closed partial order. ■

A partial order R is *closed* in a string database SDB if there exists no partial order $R' \supset R$ such that $sup(R) = sup(R')$. A partial order R is a *frequent closed partial order* if it is both frequent and closed.

Problem Definition. The problem of *mining frequent closed partial orders from strings* is to find the complete set of frequent closed partial orders in a given string database SDB with respect to a minimum support threshold min_sup . ■

3 Algorithm Frecpo

Frecpo (for Frequent closed partial orders) searches a set enumeration tree of transitive reductions of partial orders in a depth-first manner. In principle, a partial order can be uniquely represented as the set of edges in its transitive reduction. Moreover, all edges in a set can be sorted in the

dictionary order and thus can be written as a list. In fact, any global order on the edges works. For the sake of convenience, we choose dictionary order as an example here. Therefore, we can enumerate all partial orders in the dictionary order. A set enumeration tree of partial orders can be formed: for orders R_1 and R_2 , R_1 is an ancestor of R_2 and R_2 is a descendant of R_1 in the tree if and only if the list of edges in R_1 is a prefix of the list of edges in R_2 .

For example, consider a set of items $\{a, b, c\}$. The transitive reductions of all possible partial orders on the three items can be enumerated in a set enumeration tree shown in Figure 2.

By a depth-first search of the set enumeration tree of transitive reductions of partial orders, *Frecpo* will not miss any frequent partial order. The depth-first search instead of the breadth-first search is used because many previous studies strongly suggest that a depth-first search with appropriate pseudo-projection techniques often achieves a better performance than a breadth-first search when mining large databases.

To be efficient and scalable, *Frecpo* prunes the futile branches and narrows the search space sharply. Basically, three types of techniques are used.

Pruning infrequent items, edges and partial orders. If a partial order R in the set enumeration tree is infrequent, then the partial orders in the subtree rooted at R , which is stronger than R , cannot be frequent. The subtree can be pruned. Hence, *Frecpo* often does not have to search the complete set enumeration tree. Instead, only the upper part of the tree which contains all the frequent partial orders is searched. Moreover, only frequent closed partial orders will be output.

Pruning forbidden edges. Not every edge can appear in the transitive reduction of a partial order. For example, if both edges ab and bc are in a partial order R , then edge ac cannot appear in the transitive reduction of R . In this case, edge ac is called a forbidden edge. Removing the forbidden edges can also narrow the search space.

Extracting transitive reductions of frequent partial orders directly. In *Frecpo*, we develop an efficient method to identify frequent closed partial orders and also extract their transitive reductions from various subsets of strings. Thus, *Frecpo* does not need to compute the transitive reductions.

Algorithm *Frecpo* is shown in Figure 3. We explain the critical details using an example.

Example 3 (*Frecpo*) Let us consider mining frequent closed partial orders from the string database *SDB* in Table 2 with respect to minimum support threshold $min_sup = 2$.

By scanning the database only once, *Frecpo* computes the supports of the items. Infrequent items are pruned, such as f in our running example ($sup(f) = 1$).

To prune the infrequent edges and the forbidden edges,

Input: a string database *SDB* and a minimum support threshold min_sup ;

Output: the complete set of frequent CPOs;

Method:

- 1: scan database once, find frequent items;
- 2: scan database again, find global feasible edges;
// Note: if the total number of items in *SDB* // is not large, the first two scans can be combined.
- 3: let R be the set of global feasible edges with support $|SDB|$;
- 4: if $R \neq \emptyset$ then output R as a frequent CPO;
- 5: let $L = e_1, \dots, e_n$ be the list of global feasible edges with support less than $|SDB|$;
- 6: for each edge e_i in L do
- 7: if $R \cup \{e_i\}$ does not contain any redundant edge and there exists no FCPO R' found before such that $R' \supset (R \cup \{e_i\})$ and $sup(R') = sup(e_i)$ then
- 8: form $R \cup \{e_i\}$ -projected database $SDB|_{R \cup \{e_i\}}$;
- 9: recursively mine $SDB|_{R \cup \{e_i\}}$

Figure 3. The *Frecpo* algorithm.

Sid	String	Transitive closure $\mathcal{C}(s)$
1	abcdef	ab, ac, ad, ae, af, bc, bd, be bf, cd, ce, cf, de, df, ef
2	acbde	ac, ab, ad, ae, cb, cd, ce, bd, be, de
3	dabce	da, db, dc, de, ab, ac, ae, bc, be, ce
4	dcabe	dc, da, db, de, ca, cb, ce, ab, ae, be

Table 2. String database *SDB* as the running example.

Frecpo scans the database again and fills in a matrix $\{cnt[x, y]\}$, where x and y are both frequent items, and $cnt[x, y]$ registers both $sup(xy)$ and the list of items that appear between x and y in all strings having been scanned so far that contain xy . The list is called the anchor list. The matrix is called the *detection matrix* as shown below.

	a	b	c	d	e
a		4, \emptyset	3, \emptyset	2, $\{b, c\}$	4, $\{b\}$
b	0, \emptyset		2, \emptyset	2, \emptyset	4, \emptyset
c	1, \emptyset	2, \emptyset		2, \emptyset	4, \emptyset
d	2, \emptyset	2, $\{a\}$	2, \emptyset		4, \emptyset
e	0, \emptyset	0, \emptyset	0, \emptyset	0, \emptyset	

From the detection matrix, we can immediately prune the infrequent edges (those with support less than 2, such as ca). An edge is a forbidden edge if its anchor list is not empty. Forbidden edges ad , ae and db can be pruned as well.

In this example, *SDB* contains 6 different items. There are $6 \times 5 = 30$ possible different edges. 20 different edges appear in the database. Only 11 edges survive from the pruning.

Only the edges ab , ac , bc , bd , be , cb , cd , ce , da , dc and




Figure 2. The enumeration tree of the transitive reductions of possible partial orders on a, b, c .

de can be used to construct the transitive reduction of a frequent closed partial orders. They are called the *global feasible edges*. Among them, ab, be, ce and de have support 4, i.e., they appear in every string in SDB . The four edges form a frequent closed partial order. In other words, the set of global feasible edges that appear in every string forms a frequent CPO. Interestingly, the set is in fact the transitive reduction, since any redundant edge in the set is identified as a forbidden edge by the detection matrix.

Once a frequent closed partial order R is found, *Frecpo* expands R to its children. Only frequent, non-forbidden edges in the R -projected database should be used to expand R to its children in the enumeration tree.

Frequent closed partial order R_1 is the order shared by all strings. Thus, any other frequent closed partial order will be stronger than R_1 .

The other frequent closed partial orders in transitive reduction can be partitioned into the following subsets according to the dictionary order of the remaining global feasible edges (i.e., ac, bc, bd, cb, cd, da , and dc): (1) the ones having edge ac in their transitive reduction; (2) the ones having edge bc but no ac in their transitive reduction; ...; and (7) the one having dc but no other edges in its transitive reduction (if it is a frequent closed partial order). These subsets can be mined one by one in a depth-first search manner.

We first consider the subset of frequent closed partial orders having edge ac in their transitive reductions. They also contain R_1 . The strings in SDB that are super-strings of ac , namely strings 1, 2 and 3, are collected as the $(R_1 \cup \{ac\})$ -projected database.

We prune the local infrequent items, infrequent edges and forbidden edges by scanning the $(R_1 \cup \{ac\})$ -projected database once and filling in the local detection matrix. The feasible edges in this projected database are bc, bd and cd . Since each feasible edge has support 2, which is less than the number of strings in the projected database, we extract $R_2 = R_1 \cup \{ac\}$ as the transitive reduction of a frequent closed partial order. Any frequent partial order having ac must be stronger than R_2 .

Since we have three local feasible edges in the $(R_1 \cup \{ac\})$ -projected database, the remaining frequent closed partial orders having ac in their transitive reduction can be further partitioned into three sub-subsets: the ones having ac and bc , the ones having ac and bd but no bc , and the ones having ac and cd but no bc nor bd .

R_2 has an edge ab , and any frequent partial order having ac is a superset of R_2 . Clearly, edges ab, ac and bc cannot stay together in a transitive reduction, since ac is redundant in such a case. Thus, we immediately determine that the first sub-subset is empty without checking the database.

The remaining frequent closed partial orders can be found by recursive depth-first search. Limited by space, we omit the details here. ■

In implementation, we use the pseudo-projection technique [7]. Instead of producing a physical copy of strings for every projected database, we use hyperlinks (implemented as pointers) to link the strings in the projected database together. At any time, there is only one copy of the database. Scanning and deriving projected database are efficient with the help of hyperlinks.

Frecpo has three distinct advantages. First, the mining in transitive reduction avoids substantial space and I/O overhead. Second, it directly extracts frequent closed partial orders in transitive reduction. Last, it aggressively and progressively prunes futile branches in recursive depth-first search.

We tested *Frecpo* and some related methods (including the one in [3]) extensively on both real and synthetic data sets. Our experimental results show that the frequent closed partial order mining is efficient – we can find application meaningful patterns from real data sets. Furthermore, *Frecpo* is efficient and scalable in mining large data sets. The experimental results are reported in the full version of the paper.

References

- [1] R. Agrawal and R. Srikant. Mining sequential patterns. In *ICDE'95*.
- [2] A. Ben-Dor et al. Discovering local structure in gene expression data: the order-preserving submatrix problem. In *RECOMB'02*.
- [3] G. Casas-Garriga. Summarizing sequential data with closed partial orders. In *Proc. 2005 SIAM Int. Conf. Data Mining*.
- [4] A. Gionis et al. Fragments of order. In *KDD'03*.
- [5] H. Mannila and C. Meek. Global partial orders from sequential data. In *KDD'00*.
- [6] H. Mannila et al. Discovery of frequent episodes in event sequences. *Data Mining and Knowledge Discovery*, 1:259–289, 1997.
- [7] J. Pei et al. PrefixSpan: Mining sequential patterns efficiently by prefix-projected pattern growth. In *ICDE'01*.