

# Improved Methods for Extracting Frequent Itemsets from Interim-Support Trees

Frans Coenen

Department of Computer Science, University of Liverpool  
Chadwick Building, Peach Street, Liverpool L69 7ZF, UK  
`frans@csc.liv.ac.uk`

Paul Leng

Department of Computer Science, University of Liverpool  
Chadwick Building, Peach Street, Liverpool L69 7ZF, UK  
`phl@csc.liv.ac.uk`

Aris Pagourtzis

Department of Computer Science, National Technical University of Athens  
15780 Zografou, Athens, Greece  
`pagour@cs.ntua.gr`

Wojciech Rytter

Department of Computer Science, New Jersey Institute of Technology, US and  
Institute of Informatics, Warsaw University, Poland  
`rytter@oak.njit.edu`

Dora Souliou\*

Department of Computer Science, National Technical University of Athens  
15780 Zografou, Athens, Greece  
`dsouliou@cslab.ece.ntua.gr`

---

<sup>0</sup>Corresponding author. Tel: +30210 7721175. Fax: +30210 7721645.

## Abstract

Mining association rules in relational databases is a significant computational task with lots of applications. A fundamental ingredient of this task is the discovery of sets of attributes (*itemsets*) whose frequency in the data exceeds some threshold value. In this paper we describe two algorithms for completing the calculation of frequent sets using a tree structure for storing partial supports, called interim-support tree. The first of our algorithms (TTF) uses a novel tree pruning technique, based on the notion of (*fixed-prefix*) *potential inclusion*, which is specially designed for trees that are implemented using only two pointers per node. This allows to implement the interim-support tree in a space efficient manner. The second algorithm (PTF) explores the idea of storing the frequent itemsets in a second tree structure, called the *total support tree* (*T-tree*); the main innovation lies in the use of multiple pointers per node which provides rapid access to the nodes of the *T-tree* and makes it possible to design a new, usually faster, method for updating them.

Experimental comparison shows that these techniques result in considerable speedup for both algorithms comparing to earlier approaches that also use interim-support trees [8, 10]. Further comparison between the two new algorithms, shows that PTF is generally faster on instances with a large number of frequent itemsets, provided that they are relatively short, while TTF is more appropriate whenever there exist few or quite long frequent itemsets; in addition, TTF behaves well on instances in which the densities of the items of the database have a high variance.

**Keywords:** association rules, frequent itemsets, data mining, set-enumeration trees.

## 1 Introduction

An important data mining task initiated in [3] is the discovery of association rules over huge listings of sales data, also known as *basket data*. This task initially involves the extraction of *frequent* sets of items from a database of transactions, i.e. from a collection of sets of such items. An example of a database with transactions that are subsets of  $\{a, b, c, d, e, f, g, h\}$  is given in Table 1. The number of times that an itemset appears in transactions of the database is called its *support*. The minimum support an itemset must have in order to be considered as frequent is called the *support threshold*, a nonnegative integer denoted by  $t$ . The support of an association rule  $A \implies B$ , where  $A$  and  $B$  are sets of items, is the support of the set  $A \cup B$ . The *confidence* of rule  $A \implies B$  is equal to  $support(A \cup B) / support(A)$  and represents the fraction of transactions that contain  $B$  among transactions that contain  $A$ . A *valid* rule is one with support at least the support threshold  $t$  and with confidence at least a confidence threshold  $c$ .

**Examples of association rules.** Let  $\mathcal{D}$  be the database shown in Table 1. Let also  $t = 4$  be the support threshold and  $c = 0.5$  be the confidence threshold. Rules  $\{b\} \implies \{c\}$  and  $\{a\} \implies \{d\}$  have both adequate support, because  $support(\{b, c\}) = 7$  and  $support(\{a, d\}) = 9$ . However, the former rule is not valid since  $confidence(\{b, c\}) = support(\{b, c\}) / support(\{b\}) =$

$7/17 < 0.5$ ; on the other hand, the latter rule is valid since  $confidence(\{a, d\}) = support(\{a, d\})/support(\{a\}) = 9/15 > 0.5$ .

| <i>a b c d e f g h</i> | <i>a b c d e f g h</i> | <i>a b c d e f g h</i> |
|------------------------|------------------------|------------------------|
| 1 1 0 1 1 0 1 1        | 1 1 1 1 1 1 0 0        | 0 1 1 0 1 1 1 1        |
| 1 1 0 0 1 1 1 0        | 0 0 0 1 0 1 0 0        | 0 1 0 1 1 1 0 1        |
| 0 1 0 1 0 0 0 1        | 0 1 1 0 0 1 0 0        | 1 0 0 0 1 0 1 0        |
| 0 1 1 0 0 0 1 0        | 0 1 0 0 1 1 1 0        | 1 1 0 1 1 0 0 0        |
| 1 1 1 0 0 1 1 0        | 0 0 1 1 1 1 0 0        | 0 0 1 1 0 1 0 0        |
| 1 1 0 1 1 0 1 0        | 1 0 1 1 1 0 1 1        | 0 1 1 1 0 1 1 0        |
| 0 1 0 1 1 1 1 1        | 0 1 0 0 0 0 1 1        | 0 1 0 0 1 0 0 0        |
| 1 0 0 1 0 1 1 1        | 0 1 1 1 0 1 0 1        | 1 0 0 0 1 0 1 0        |
| 1 0 0 1 1 0 0 1        | 0 0 1 1 0 0 1 0        | 1 0 1 0 1 0 0 0        |
| 1 0 0 1 0 1 1 0        | 1 0 0 1 1 0 0 0        | 0 0 0 0 1 0 0 1        |
| 0 0 0 1 0 0 1 0        | 1 0 0 0 1 0 1 1        |                        |

Table 1: A database containing 32 transactions. Each transaction is described by a subset of  $\{a, b, c, d, e, f, g, h\}$ .

Association Rule Mining, in general, involves the extraction from a database of all valid rules. The major part of this task is the discovery of the frequent itemsets; once the support of all these sets has been counted, determining valid rules can be done as follows. For each frequent itemset  $X$  ( $support(X) \geq t$ ), consider all itemsets  $Y \subseteq X$  (all such subsets are necessarily frequent as well). If  $support(X)/support(Y) \geq c$  it turns out that the following rule is valid:

$$Y \implies X \setminus Y$$

It is not hard to see that the above procedure finds all valid rules.

Of course, there is no polynomial-time (w.r.t. the input size) algorithm for generating all frequent itemsets, since their number can be exponential in the size of the database. For example, consider a database with  $n$  items and  $n$  transactions; if there exist  $m$  transactions of the form  $111 \dots 1$ , then all  $2^n - 1$  possible itemsets have support at least  $m$  and are consequently frequent if  $m > t$ . Therefore, this problem has motivated a continuing search for effective heuristics.

The best-known algorithm, from which most others are derived, is Apriori [5]. Apriori performs repeated passes of the database, successively counting the support for single items, pairs, triples, etc.. At the end of each pass, itemsets that fail to reach the support threshold are eliminated, and *candidate* itemsets for the next pass are constructed as supersets of the remaining frequent sets. As no frequent set can have an infrequent subset, this heuristic ensures

that all sets that may be frequent are considered. The algorithm terminates when no further candidates can be constructed.

Apriori remains potentially very costly because of its multiple database passes and, especially, the possible large number of candidates in some passes. Attempts to reduce the scale of the problem include methods that begin by partitioning [16] or sampling [17] the data, and those that attempt to identify *maximal* frequent sets [7, 6] or *closed* frequent sets [18] from which all others can be derived. A number of researchers have made use of *set-enumeration tree* structures to organise candidates for more efficient counting. The FP-growth algorithm of Han et al. [13, 12] counts frequent sets using a structure, the *FP-tree*, in which tree nodes represent individual items and branches represent itemsets. FP-growth reduces the cost of support-counting because branches of the tree that are subsets of more than one itemset need only be counted once. In contemporaneous work, commencing with [11], we have also employed set-enumeration tree structures to exploit this property. Our approach begins by constructing a tree, the *P-tree*, [10, 9], which contains an incomplete summation of the support of sets found in the data. The *P-tree*, described in more detail below, shares the same performance advantage of the FP-tree but is a more compact structure. Results presented in [8] demonstrate that algorithms employing the *P-tree* can achieve comparable or superior speed to FP-growth, with lower memory requirements.

Unlike the FP-tree, which was developed specifically to facilitate the FP-growth algorithm, the *P-tree* is a generic structure which can be the basis of many possible algorithms for completing the summation of frequent sets. In this paper we describe and compare two algorithms for this purpose, namely:

1. The *T-Tree-First* (TTF) algorithm.
2. The *P-Tree-First* (PTF) algorithm.

Both algorithms make use of the incomplete summation contained in the *P-tree* to construct a second set-enumeration tree, the *T-tree*, which finally contains frequent itemsets together with their total support. The algorithms differ in the way they compute the total support: algorithm *T-Tree-First* iterates over the nodes of *T-tree*, and for each of them it traverses the *P-tree*; algorithm *P-Tree-First* starts by traversing the *P-tree* and for each node that it visits, it updates all relevant nodes at the current level of the *T-tree*.

Earlier algorithms that use similar tree structures are *Apriori-TFP* [8] and an anonymous algorithm presented in [10]; here we will refer to the latter as “Interim-Support”.

The contribution of this work lies in the introduction of techniques that can considerably accelerate the process of computing frequent itemsets. In particular, the main innovation in the first of our algorithms (TTF) is a tree pruning technique, based on the notion of *fixed-prefix potential inclusion*, which is specially designed for trees that are implemented using only two pointers per node. This allows to implement the interim-support tree in a space efficient

manner. The second algorithm (PTF) introduces the use of multiple pointers per node in the  $T$ -tree; this accelerates the access of the nodes of the  $T$ -tree and makes it possible to find and update appropriate  $T$ -tree nodes following a new, usually faster, strategy.

We perform experimental comparison of the two algorithms against the earlier algorithms Interim-Support and Apriori-TFP and show that in most cases the speedup is considerable. We also compare the two new algorithms to each other and discuss the merits of each. Our results show that PTF is faster than TTF if there are a lot of frequent itemsets in the database (small support threshold), provided that they are *short*, i.e., that they contain few items. On the other hand TTF gains ground as the support threshold increases and behaves even better for instances of variable item density which have been pre-sorted according to these densities; it also behaves much better than PTF in instances with long frequent itemsets.

## 2 Notation and Preliminaries

A database  $\mathcal{D}$  is represented by an  $m \times n$  binary matrix. The columns of  $D$  correspond to items (attributes), and the rows correspond to the transactions (records). The columns are indexed by consecutive letters  $a, b, \dots$  of the alphabet (see Table 1 for an example). The set of columns (items) is denoted by  $\mathcal{C}$ . An *itemset*  $I$  is a set of items  $I \subseteq \mathcal{C}$ . For an itemset  $I$  we define:

- $E(I)$  ( $E$ -value of  $I$ ) is the number of transactions that are exactly equal to  $I$ . This value is also called *exact support of  $I$* .
- $P(I)$  ( $P$ -value of  $I$ ) is the number of transactions that have  $I$  as a prefix. Also called *interim support of  $I$* .
- $T(I)$  ( $T$ -value of  $I$ ) is the number of transactions that contain  $I$ . Also called *total support* or, simply, *support of  $I$* .

Both the terms  $P$ -value and  $P$ -tree have been used in other contexts with other meanings. Here, the derivation is the notion of a *partially* counted support value. In this paper we consider the problem of finding all itemsets  $I$  with total support  $T(I) \geq t$ , for a given database  $\mathcal{D}$  and threshold  $t$ , starting with a  $P$ -tree containing  $P$ -values for all sets present as transactions in  $\mathcal{D}$ .

For an item  $x$  we define the *density of  $x$  in  $\mathcal{D}$*  to be the fraction of transactions of  $\mathcal{D}$  that contain  $x$ , that is  $T(\{x\})/m$ . We also define the *density of a database  $\mathcal{D}$*  to be the average density of the items of  $\mathcal{D}$ ; note that the density of  $\mathcal{D}$  is equal to the fraction of the total number of items appearing in the transactions of  $\mathcal{D}$  over the size of  $\mathcal{D}$  ( $= nm$ ).

We will make use of the following order relations:

- *Inclusion order*:  $I \subseteq J$ , the usual set inclusion relation,

- *Lexicographic order*:  $I \leq J$ ,  $I$  is lexicographically smaller or equal to  $J$  if seen as strings,
- *Prefix order*:  $I \sqsubseteq J$ ,  $I$  is a prefix of  $J$  if seen as strings. Note that  $I \sqsubseteq J \Leftrightarrow I \subseteq J \ \& \ I \leq J$ .

We will also use the corresponding operators without equality:  $I \subset J$ ,  $I \sqsubset J$  and  $I < J$ .

Notice that for any itemset  $I$ :

$$T(I) = \sum_{J:I \subseteq J} E(J)$$

and therefore:

$$T(I) = \sum_{J:I \subseteq J \ \& \ I \leq J} E(J) + \sum_{J:I \subseteq J \ \& \ J < I} E(J) = P(I) + \sum_{J:I \subseteq J \ \& \ J < I} E(J) \quad (1)$$

This property will play an important role in our algorithms.

### 3 The Interim-Support Tree

Both new algorithms TTF and PTF have a common first part which is a pre-processing of the database that results in the storage of the whole information into a structure called the  $P$ -tree or *interim-support tree*. The  $P$ -tree is a set-enumeration tree the nodes of which are distinct itemsets of the database as well as some common prefixes of these itemsets. For each node, the interim support ( $P$ -value) of the corresponding itemset is also stored.

The notion of interim-support trees was introduced in [10], where details of the construction of the  $P$ -tree were given, and more fully in [9]. The algorithm is summarised below.

*Algorithm P-Tree-Build*

*Input*: Database  $\mathcal{D}$ .  
*Output*:  $P$ -Tree of itemsets in  $\mathcal{D}$ .

(\* Start with  $P$ -tree of a single node representing the empty set \*)

**for** each transaction  $i$  in  $\mathcal{D}$  **do**  
   $c := P\text{-tree\_rootnode}$   
   $inserted := false$   
  **while not**  $inserted$  **do**  
    **if**  $c = i$  **then** increment  $P(c)$ ;  $inserted := true$   
    **else if**  $c \subset i$  **then** increment  $P(c)$ ;  $c := eldest\_child\_of.c$   
    **else if**  $c < i$  **then**  $c := next\_sibling\_of.c$   
    **else** create new node for  $i$ ;  $inserted := true$   
**return**  $P$ -tree;

Note that in this algorithm, for clarity, we use the notation  $i$  and  $c$  to denote an itemset which also is or will become the label of a node in the tree. The tree is constructed in a single pass of  $\mathcal{D}$ . As each transaction is examined, the tree is traversed in a top-down (preorder) manner until either a node with identical itemset is found or the traversal passes the position in the tree at which the new itemset should be located. During this traversal, the support of all ancestors (preceding subsets) of the itemset is incremented.

If the itemset is not found in the tree, a new node is added to the tree to represent it. At this point the traversal has reached a node  $c$  which is either null (ie a nonexistent child or sibling) or lexicographically follows the new itemset  $i$ . A node labelled  $i$  is inserted at the position in the tree structure occupied by  $c$ . The following three different cases apply for dealing with the previous node  $c$  and recording the interim support of  $i$ :

- $c$  is null: The new node  $i$  is given support  $P(i) = 1$ .
- $i \subset c$ :  $c$  becomes the child of  $i$ .  $P(i) = P(c) + 1$ .
- Otherwise:  $c$  becomes the next sibling of  $i$ .  $P(i) = 1$ .

Finally, if  $i$  has been added as a sibling of  $c$ , and  $i$  and  $c$  share a leading substring  $d$  that is not already in the tree, a node  $d$  is inserted at the position now occupied by  $i$ , with  $i$  and  $c$  becoming its children, and  $P(d) = P(i) + P(c)$ .

In any case, during the insertion of an itemset at most two new nodes will be created in the  $P$ -tree. On the other hand, if the database contains several identical itemsets, the  $P$ -tree can be much smaller than the original database.

The  $P$ -tree that corresponds to the database of Table 1 is shown in Figure 1. Note that figure 1 shows the logical structure of the  $P$ -tree. However, for the sake of memory efficiency the  $P$ -tree is implemented using two pointers per node: *down* and *right*. For a node  $v$ , its down pointer links  $v$  to one of its children — the lexicographically smaller. This child’s right pointer points to another child of  $v$ , and so on. For example, in the implementation of a  $P$ -tree containing itemsets ‘a’, ‘ab’, ‘ac’, and ‘abc’ node ‘a’ points down to ‘ab’ which in turn points down to ‘abc’ and right to ‘ac’.

The significance of the  $P$ -tree is that it performs a large part of the counting of support totals very efficiently in a single database pass. The size of the  $P$ -tree is linearly related to the original database, and will be smaller in cases where the data includes many duplicated itemsets. Most importantly, it involves no loss of relevant information, so the  $P$ -tree can be used as a surrogate for the original database in any chosen algorithm.

The FP-tree of Han et.al. [13, 12] was developed independently and contemporaneously with our  $P$ -tree [11, 10] and shares similar performance advantages. There are three significant differences between the two structures. Firstly, the construction of the FP-tree requires two database passes, the first of which eliminates attributes that fail to meet the required support threshold, so it no longer contains a complete representation of the information in the database.

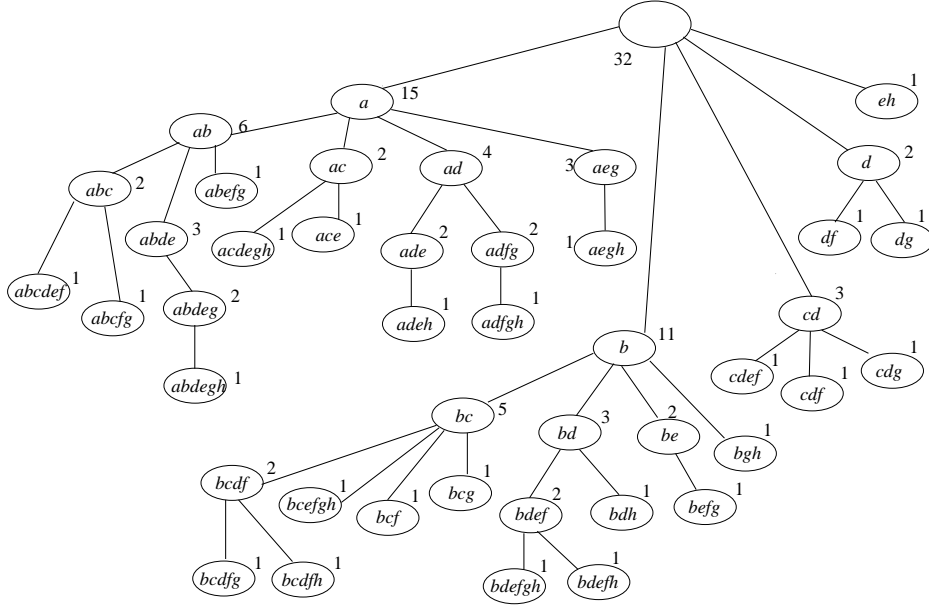


Figure 1:  $P$ -tree with interim supports for the database of Table 1.

Secondly, the nodes of the FP-tree correspond to individual items, whereas in the  $P$ -tree a sequence of items which is partially closed (i.e. which has no leading subsequence with greater support in the tree) will be stored as a single tree node. Thus, for example, two transactions  $\{a, b, c, d, e\}$  and  $\{a, b, c, x, y\}$ , which share a common prefix  $\{a, b, c\}$ , would require in all 7 nodes in the FP-tree. In the  $P$ -tree, conversely, only 3 nodes would necessarily be created: a parent for  $\{a, b, c\}$ , and child nodes for  $\{d, e\}$  and  $\{x, y\}$ . Finally, in order to implement the FP-growth algorithm, the FP-tree must store pointers at each node to link all nodes representing the same item, and also to link a node to its parent and child nodes. The nodes of the  $P$ -tree, conversely, requires only pointers to the eldest child and next sibling. Both the latter differences lead to a more compact tree structure and hence faster traversal.

More importantly, the simpler and less pointer-rich organisation of the  $P$ -tree makes it a more flexible structure than the FP-tree, which was developed specifically to implement the FP-growth algorithm. This flexibility, for example, allows us to implement an algorithm, **Apriori-TFP**, which applies an **Apriori**-like procedure to the nodes in the  $P$ -tree. In results presented in [8] both the memory requirements and construction time for the  $P$ -tree were less than for a corresponding FP-tree, and the execution time for **Apriori-TFP** was similar or less than **FP-growth** and much less than **Apriori**. In this paper we will use **Apriori-TFP** as a benchmark against which to measure the performance of the new algorithms proposed.

A further advantage of the relatively simple  $P$ -tree structure is that it facilitates scaling to deal with data that cannot be contained in main memory. In this case, the original database is segmented into partitions for each of which a separate  $P$ -tree is constructed. This process



again requires only a single pass of the database to produce a set of Partition- $P$ -trees ( $PP$ -trees). Subsequently, algorithms that require to traverse the  $P$ -tree can operate by separately traversing each of the  $PP$ -trees, accumulating support counts from each to produce the overall totals. Partitioning the FP-tree is necessarily more complex, although methods for doing this are described in [12]. In [1] we described an implementation of **Apriori-TFP** using a tree partitioning strategy. The results obtained showed that segmenting the data enabled effective scaling of the method, and demonstrated improved performance over a partitioned version of FP-growth. Similar partitioning strategies can be applied to the algorithms described in this paper, and thus, although the experiments described relate to databases that can be contained in main memory, the methods can be applied on a larger scale.

A number of other researchers have made use of the FP-tree and similar structures. The CFP-tree described in [15] stores frequent closed itemsets in a form that facilitates subsequent query processing. The main contribution of this work is a structure that can be re-used efficiently, rather than the efficiency of the construction algorithm. Reusability is also a feature of the  $P$ -tree, which, as we have mentioned, retains all relevant information from the original data. In [14] a structure is described, also (coincidentally) called a  $P$ -tree, which is quite similar to our  $P$ -tree, but (like the FP-tree) stores only one item at each node. The approach described constructs FP-trees from the  $P$ -tree rather than from the original data, producing a single overall  $FP$ -tree, for which further partitioning might become necessary if the data is too large to contain in main memory.

## 4 The $T$ -Tree-First (TTF) algorithm

The  $T$ -Tree-First (TTF) algorithm first iterates over the nodes of  $T$ -tree and for each of them it traverses the  $P$ -tree. In this section we give a detailed description of TTF.

The algorithm first scans the database and creates the  $P$ -tree, as explained in the previous section.

It then starts building the  $T$ -tree (recall that the  $T$ -tree will finally contain all frequent itemsets together with their total supports). Each level of the  $T$ -tree is implemented as a linear list, where itemsets appear in lexicographic order; nodes of such a list neither point to nor are pointed from nodes that are in the list of another level. In the beginning, the algorithm builds level 1 of the  $T$ -tree, which contains all frequent singletons; to this end it counts their support traversing the  $P$ -tree. It then builds the remaining  $T$ -tree level by level using procedure **Iteration**( $k$ ).

The algorithm is presented below. A fundamental ingredient of TTF is function **CountSupport** which is described separately.

*Algorithm T-Tree-First (TTF)*

*Input:* Database  $\mathcal{D}$ , threshold  $t$ .

*Output:* The family  $\mathcal{F}$  of frequent itemsets.

Build  $P$ -tree from database  $\mathcal{D}$ ;

(\* Build the 1-st level of  $T$ -tree \*)

**for**  $i = 1$  **to**  $n$  **do**

**if** **CountSupport**( $P$ -tree,  $\{i\}$ )  $\geq t$  **then** add  $\{i\}$  to  $\mathcal{F}_1$ ;

(\* Build the remaining levels of  $T$ -tree \*)

**for**  $k = 2$  **to**  $n$  **do**

**Iteration**( $k$ );

**if**  $\mathcal{F}_k = \emptyset$  **then** **exit**

**else**  $\mathcal{F} = \mathcal{F} \cup \mathcal{F}_k$ ;

**return**  $\mathcal{F}$ ;

Some details of procedure **Iteration**( $k$ ) need to be clarified. Its goal is to build  $\mathcal{F}_k$ , that is, the  $k$ -th level of the  $T$ -tree. The procedure uses the heuristic first described in [5]. Itemsets in  $\mathcal{F}_k$  must have all their  $(k - 1)$ -size subsets in  $\mathcal{F}_{k-1}$ . Therefore, one can start from existing itemsets in  $\mathcal{F}_{k-1}$  and try to augment them with one more item in order to create all potentially frequent itemsets. To avoid duplications the algorithm may proceed by considering for each frequent itemset  $X_{k-1}$  in  $\mathcal{F}_{k-1}$  all  $X_{k-1}$ 's supersets  $X_k = \{x\} \cup X_{k-1}$  for items  $x$  that are greater than any item of  $X_{k-1}$ .

As observed already in [5], it makes sense to consider such supersets only if  $X_{k-1}$  and the node following it, denoted  $X'_{k-1}$ , differ at the last item. The candidate superset  $X_k$  is then the union of  $X_{k-1}$  and  $X'_{k-1}$ . Then it is checked whether all the  $(k - 2)$  many remaining  $(k - 1)$ -subsets of  $X_k$  are frequent; this task is carried out by a special function called **ExistSubsets**, which we will not describe in detail here. If some of the examined subsets of  $X_k$  is not present in  $\mathcal{F}_{k-1}$ ,  $X_k$  is not added to  $\mathcal{F}_k$ .

**Procedure** **Iteration**( $k$ ) (\* Building the  $k$ -th level of  $T$ -tree \*)

**for each** itemset  $X_{k-1} \in \mathcal{F}_{k-1}$  **do**

$X'_{k-1} := next(X_{k-1})$ ;

**while**  $X'_{k-1} \neq \text{NULL}$  **do**

**if**  $X_{k-1}$  and  $X'_{k-1}$  differ only at the last item **then**

$X_k := X_{k-1} \cup X'_{k-1}$ ;

**if** **ExistSubsets**( $X_k, \mathcal{F}_{k-1}$ ) **then**

$T(X_k) := \text{CountSupport}(P\text{-tree}, X_k)$ ;

**if**  $T(X_k) \geq t$  **then** add  $X_k$  to  $\mathcal{F}_k$ ;

$X'_{k-1} := next(X'_{k-1})$ ;

**else** **exit while**;

In order to complete the description of TTF it remains to describe its most critical part, that

is, function **CountSupport**, which counts the total support of an itemset  $X$  in the  $P$ -tree in a recursive manner. An essential ingredient of **CountSupport** is the notion of *fixed-prefix potential inclusion*:

*Fixed-Prefix Potential Inclusion.*  $I \overset{\text{pot}}{\subseteq}_K J: \exists J', \text{commonprefix}(J, J') = K \ \& \ I \subseteq J'$ .

Examples: 'bdf'  $\overset{\text{pot}}{\subseteq}_{\text{'ab'}}$  'abc', 'bdf'  $\overset{\text{pot}}{\not\subseteq}_{\text{'ab'}}$  'abd'.

In words,  $I \overset{\text{pot}}{\subseteq}_K J$  means that there is an itemset greater than  $J$ , sharing with  $J$  a common prefix  $K$ , that contains  $I$ .

A second interesting inclusion relation can be defined in terms of  $\overset{\text{pot}}{\subseteq}_K$ :

*Potential Inclusion.*  $I \overset{\text{pot}}{\subseteq} J \stackrel{\text{def}}{=} I \overset{\text{pot}}{\subseteq}_J J$ , i.e.  $\exists J', J \sqsubseteq J' \ \& \ I \subseteq J'$ .

Examples: 'bdf'  $\overset{\text{pot}}{\subseteq}$  'abde', 'bdf'  $\overset{\text{pot}}{\not\subseteq}$  'abdg'.

In words,  $I \overset{\text{pot}}{\subseteq} J$  means that there is an extension of  $J$  that contains  $I$ .

The use of the above inclusion relations can significantly reduce the number of moves needed to count the support of an itemset in trees with two pointers per node. Suppose that we are looking for appearances (i.e. supersets) of an itemset  $I$  in the  $P$ -tree and we are currently visiting a node that contains itemset  $J$ :

- Nodes that are below the current node contain itemsets  $J'$  which have  $J$  as prefix. Therefore, if  $I \overset{\text{pot}}{\not\subseteq} J$  there is no point visiting the subtree rooted at the current node.
- Nodes that are to the right of the current node (siblings) contain itemsets that have  $\text{par}(J)$  (parent of  $J$ ) as prefix — and so does  $J$  — and are greater than  $J$ . If  $I \overset{\text{pot}}{\not\subseteq}_{\text{par}(J)} J$  there is no point visiting the subtrees rooted at these nodes.

These two tests result in much better tree pruning comparing to the one applied by the Interim-Support algorithm [10]. As an example, suppose that we are trying to find the support of itemset  $X=\text{'bd'}$  in a  $P$ -tree in which there is a node 'ab' with children 'abde' and 'abefg'. Then, once the tree traversal reaches node 'abde' it adds its support to  $T(X)$  and does not move to the right, that is, it avoids visiting 'abefg'. On the other hand, the Interim-Support algorithm would also examine 'abefg' (and other siblings if such existed) because it only terminates its search whenever it finds itemsets lexicographically equal or greater than  $X$ .

```

Function CountSupport(pnode, X): integer
(* Counts the total support of itemset X
in the subtree of P-tree rooted at pnode*)
T := 0;
if pnode ≠ NULL then
  J := pnode → itemset;
  if  $X \overset{\text{pot}}{\subseteq} J$  then (* makes sense to search children *)
    if  $X \subseteq J$  then T := T + P(J)
      (* inclusion is a special case of potential inclusion *)
    else T := T + CountSupport(pnode → down, X);
  if  $X \overset{\text{pot}}{\subseteq}_{\text{par}(J)} J$  then (* makes sense to search right siblings *)
    T := T + CountSupport(pnode → right, X);
return T;

```

Finally, let us explain how to check potential inclusion and fixed prefix potential inclusion. It can be shown that the following tests suffice. The proof is omitted.

- $X \overset{\text{pot}}{\subseteq} J$ : if  $X \subseteq J$  then  $X \overset{\text{pot}}{\subseteq} J$  is true. Otherwise let  $x$  be the lexicographically smaller item of  $X$  that is not item of  $J$  (such  $x$  exists). If for all items  $j$  of  $J$  are lexicographically smaller than  $x$  then  $X \overset{\text{pot}}{\subseteq} J$  is true otherwise it is false.
- $X \overset{\text{pot}}{\subseteq}_K J$ : assume  $K \subseteq J$  (otherwise the inclusion  $X \overset{\text{pot}}{\subseteq}_K J$  is obviously false). Let  $x$  be the first item of  $X \setminus K$  and  $j$  be the first item of  $J \setminus K$ . If  $x > j$  the inclusion  $X \overset{\text{pot}}{\subseteq}_K J$  holds otherwise it is false.

## 5 The *P*-Tree-First (PTF) Algorithm

The *P*-Tree-First (PTF) algorithm also begins by constructing the *P*-tree exactly as TTF, but then it follows an inverse approach in order to update the *T*-tree. In particular, during the processing of level- $k$  of the *T*-tree, each node of the *P*-tree is visited once. Let  $I$  be the itemset of a visited node; the algorithm updates all nodes of level- $k$  that are subsets of  $I$ , except for those that are also subsets of  $\text{par}(I)$  (parent of  $I$ ) — the latter have already been updated while visiting  $\text{par}(I)$ .

Level- $k$  itemsets of the *T*-tree are constructed from the itemsets of level- $(k - 1)$ , by adding single items to each of them. This is done without checking the frequency of all subsets of a candidate. This is in contrast to TTF where special care was taken in order to create as few candidates as possible; here it is more important to save time by avoiding checking the subsets. Then, the *P*-tree is traversed as described above in order to compute support for all nodes of level- $k$ . Nodes with support smaller than the threshold are removed before the generation of level- $(k + 1)$ . An illustration of this process for the database of Table 1 is shown in Figure 2.

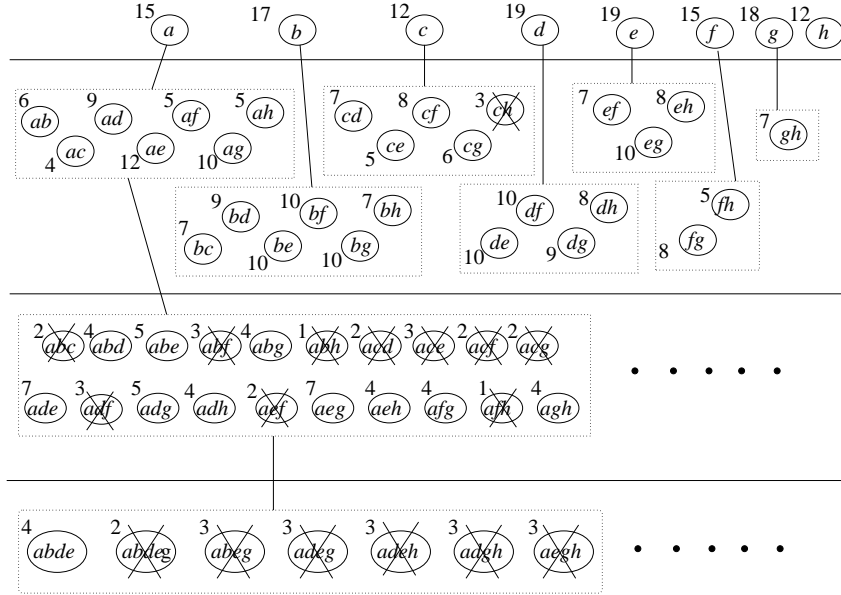


Figure 2:  $T$ -tree with total supports for the database of Table 1 .

*Algorithm P-Tree-First (PTF)*

*Input:* Database  $\mathcal{D}$ , threshold  $t$ .  
*Output:* The family  $\mathcal{F}$  of frequent itemsets.

Build  $P$ -tree from database  $\mathcal{D}$ ;  
 add  $\emptyset$  to  $\mathcal{F}_0$ ; (\* create a dummy level with one empty itemset \*)  
 (\* Build level- $k$  of the  $T$ -tree \*)  
**for**  $k = 1$  **to**  $n$  **do**  
     **Iteration**( $k$ );  
     **if**  $\mathcal{F}_k = \emptyset$  **then exit for**  
     **else**  $\mathcal{F} = \mathcal{F} \cup \mathcal{F}_k$ ;  
**return**  $\mathcal{F}$ ;

Our innovation here is the use of multiple pointers at each node of the  $T$ -tree in contrast to earlier approaches (e.g. Apriori-TFP [8]) where two pointers per node are used. In particular, each node of the  $T$ -tree contains  $n - k$  pointers, where  $n$  is the number of items and  $k$  is the level of the node; there is one pointer for each item that is lexicographically greater than the greatest item of the node. For example, in the  $T$ -tree for the database of Table 1, a node that contains itemset ‘bde’ must also contain three pointers, one for each of ‘f’, ‘g’, ‘h’. If ‘bdeg’ is found to be frequent, it will be stored in the node pointed by the ‘g’ pointer of node ‘bde’.

The use of multiple pointers provides rapid access to the nodes of the  $T$ -tree, allowing for a new strategy for  $T$ -tree update. In particular, while building level  $k$ , once a node  $I$  of the  $P$ -tree

is visited, all its  $k$ -subsets (subsets of size  $k$ ) are generated; once such a  $k$ -subset is generated, it is sought in the  $T$ -tree and, if present, its support is updated accordingly. Whenever such an itemset  $J$  has a prefix  $J'$  which is not frequent (hence neither  $J$  can be frequent) the algorithm discovers this quite early and the update process terminates. For example, if the algorithm visits a node of the  $P$ -tree with itemset 'acdfghk' and the current level of the  $T$ -tree is level-6 the algorithm should update all size-6 subsets of 'acdfghk'. Consider 'acdfgh'; the algorithm will try to find this node starting from 'a' in level-1, continuing to 'ac' in level-2, and then to 'acd', 'acdf' and 'acdfg'. If 'acd' is non-frequent, i.e. does not exist in level-3, the algorithm stops and considers the lexicographically next size-6 subset of 'acdfghk'. In fact, PTF saves even more comparisons by considering 'acdfghk' as next subset because there is no need to check any subset that contains 'acd'. Note that, in such a case, we use a 'non-frequent' itemset, called  $NF$ , which keeps the last prefix that was found to be missing from the  $T$ -tree. On the other hand, Apriori-TFP traverses a potentially large list of candidate itemsets in order to check whether any of them is a  $k$ -subset of  $I$  (note that this also happens in the original Apriori algorithm [5] where  $I$  is the current transaction scanned). This could be much slower than the above described procedure, especially if  $I$  has few  $k$ -subsets in that list. A detailed description of the update of level- $k$  of the  $T$ -tree is given below.

```

Procedure Iteration( $k$ ) (* Building  $k$ -th level of  $T$ -tree *)
for each itemset  $X_{k-1} \in \mathcal{F}_{k-1}$  do
  for each item  $x$  greater than all items of  $X_{k-1}$  do
    add  $X_k := X_{k-1} \cup \{x\}$  to  $\mathcal{F}_k$ ;
    let the  $x$ -th down pointer of  $X_{k-1}$  point to  $X_k$ ;
  (* Update total supports of nodes in  $\mathcal{F}_k$  *)
for each node  $I$  of the  $P$ -tree do
   $NF := \{\}$ ;
  for each itemset  $J \subseteq I$  with  $|J| = k$  in lex. order do
    if  $J \subseteq \text{par}(I)$  or ( $NF \subseteq J$  and  $NF \neq \{\}$ ) then
      proceed to the lex. next  $J \subseteq I$  such that
         $J$  is not subset of  $\text{par}(I)$  and does not contain  $NF$ 
    else
      repeat
        descend the  $T$ -tree following prefixes of  $J$ 
      until  $J$  is found or some  $J' \sqsubseteq J$  is missing;
      if  $J$  is found then  $T(J) := T(J) + P(I)$ 
      else  $NF := J'$ ; (*  $J'$  is missing and  $NF$  is set so that
        no itemset  $J$  containing  $J'$  will be considered
        in any subsequent inner for-loop *)
  remove from  $\mathcal{F}_k$  all nodes with support  $< t$  (threshold);

```

## 6 Complexity of the algorithms

We will next provide some bounds on the complexity of algorithms TTF and PTF. Let us first remind the reader that there can be no algorithm for this problem that runs in time polynomial w.r.t. the size of the database (equivalently, w.r.t.  $m$  and  $n$ ), since there are instances in which the number of frequent itemsets is  $2^n - 1$ . We shall therefore examine whether our algorithms' running time is polynomial w.r.t.  $m$ ,  $n$  and the number  $R$  of frequent itemsets; note that the output size is at most  $Rn$ . We will prove that this holds for TTF, but probably not for PTF.

**Theorem 1** TTF has time complexity  $O(mn^2R)$  and PTF has time complexity  $O(mn2^n)$ .

*Proof:* The dominating term in the complexity, for both algorithms, is the frequency calculation process.

We first show that the complexity of this process is  $O(mn^2R)$  for TTF. The proof is based on the fact that for each frequent itemset  $I_k$  of size  $k$ , at most  $n - k$  supersets of size  $k + 1$  may be added to the list of potentially frequent itemsets of size  $k + 1$ . This is because these itemsets are of the form  $I_{k+1} = \{x\} \cup I_k$  where  $x$  can be any item lexicographically greater than all items of  $I_k$ . Thus, the total number of potentially frequent itemsets of size  $k + 1$  is bounded by  $R_{k+1} \leq R_k(n - k) \leq R_k n$  and their overall number is thus bounded by  $Rn$ . TTF, for each of the (at most  $Rn$ ) potentially frequent itemsets, examines a part of the  $P$ -tree, which contains at most  $2m$  nodes in total. Again, comparison of the corresponding itemsets requires  $O(n)$  time and the bound follows.

On the other hand, during  $\text{Iteration}(k)$ , PTF does the following: for any of the (at most  $2m$  many) nodes of the  $P$ -tree that contains, e.g., an itemset  $I_t$  of size  $t$ , it considers all possible  $k$ -size subsets of  $I_t$ , i.e.  $\binom{t}{k} \leq \binom{n}{k}$  many itemsets. For each of these itemsets, it performs at most  $k$  moves in the  $T$ -tree in order to locate the itemset and update its frequency (if present). Summing over all levels, the frequency calculation costs at most  $2m \sum_{k=1}^n \binom{n}{k} k = mn2^n$ .  $\square$

Although the above result suggests that TTF is of lower complexity than PTF this is not always the case, as can be demonstrated by appropriate examples. In fact, the presented bounds are not directly comparable because if  $R$  is large (e.g.,  $\Theta(2^n)$ ) then the complexity of PTF is smaller, whereas if  $R \ll 2^n$  then it is larger but probably too overestimated. Experimental comparison of the two algorithms is therefore meaningful.

## 7 Experimental Comparison

We implemented four algorithms in ANSI-C: TTF, Interim-Support (IS), PTF and Apriori-TFP (ATFP). We run several experiments using a Pentium 1.6 GHz PC. We have used four types of datasets: synthetic, synthetic of variable-density, realistic datasets, and sparse datasets. The obtained results are presented below.

**Synthetic datasets.** We first experimented with datasets created by using the IBM Quest Market-Basket Synthetic Data Generator (described in [5]). We follow a standard notation according to which a dataset is described by four parameters:  $T$  represents the average transaction length (roughly equal to the database density times the number of items),  $I$  represents the average length of maximal frequent itemsets,  $N$  represents the number of items, and  $D$  represents the number of transactions in the database. We generated datasets T10.I4.N50.D10K and T10.I4.N20.D100K and run experiments with all four algorithms. The execution time of each algorithm for these two datasets and threshold varying from 1% to 5% is shown in Figure 3.

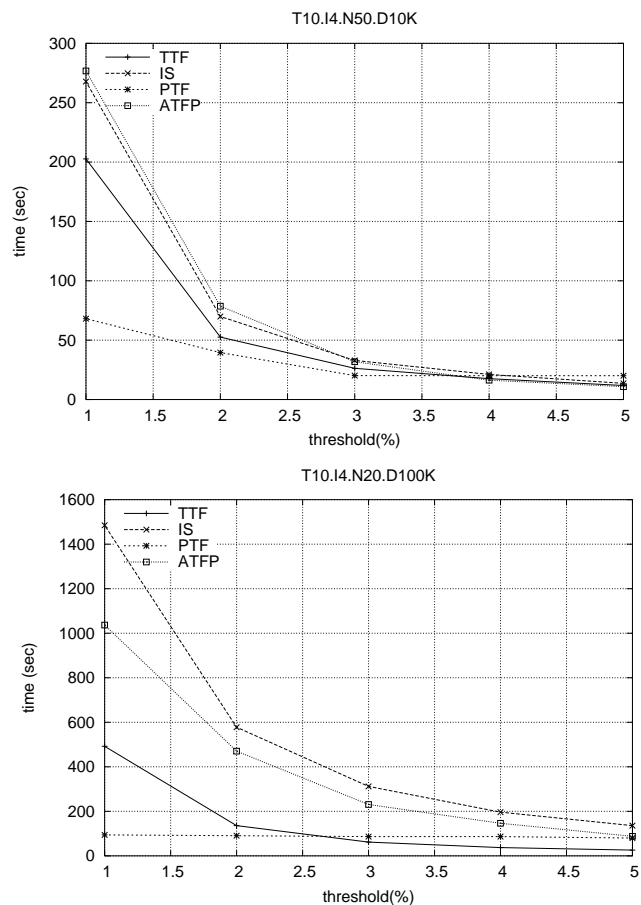


Figure 3: Results for datasets T10.I4.N50.D10K (top) and T10.I4.N20.D100K (bottom).

These results show that both algorithms TTF and PTF are faster than the earlier algorithms IS and ATFP, except for rather large thresholds. As regards TTF and IS (which also iterates over the  $T$ -Tree first), the reason for this behaviour is that IS performs fewer tests at each  $P$ -tree node that it visits; thus, whenever a contiguous part of the tree is traversed by both TTF and IS, it is IS the one that does it faster. Now, whenever the frequent itemsets are few,



they are also (most probably) of small size; a small itemset has higher chances to appear in a contiguous part of the  $P$ -tree which therefore cannot be pruned by TTF. As regards PTF and ATFP (which iterates over the  $P$ -tree first), we observe that ATFP can be faster than PTF if there are only few frequent itemsets because in such a case it can be faster to traverse the list of candidate itemsets than generating all subsets of a node.

Comparing now the two new algorithms, we observe that PTF is faster than TTF for small thresholds ( $\leq 2\%$ ). This is due to the fact that whenever the number of frequent itemsets is large, TTF performs a lot of  $P$ -tree traversals, while PTF performs only one full  $P$ -tree traversal per  $T$ -tree level. Since the size of the  $P$ -tree can be rather large (even comparable to the size of the database) its traversal is quite slow; hence, whenever TTF performs many traversals, even partial, the overall slowdown is considerable. On the other hand, PTF performs several  $T$ -tree traversals at each level but these are fast thanks to the use of multiple pointers. The two algorithms have comparable running time for thresholds above 2%. This is because for relatively sparse  $T$ -tree the  $P$ -tree traversals performed by TTF are few; in this case the economizing techniques of TTF balance, or even beat the advantages of PTF.

**Variable-Density Datasets.** To further compare TTF and PTF we implemented a probabilistic generator in order to create datasets of *variable item density* (each item has a different expected density). This generator fills the  $i$ -th item of a row with probability  $p_f - (i - 1)p_s$ , i.e., the probability decreases linearly as we move from the first to the last item of a row;  $p_f$  represents the probability of appearance of the first item and  $p_s$  is the decrement step. The expected density of the database is equal to  $p_f - \frac{(n-1)}{2}p_s = \frac{p_f - p_l}{2}$ , where  $p_l$  is the probability of appearance of the last item and  $n$  is the number of items in each row.

We have generated four variable-density datasets, one for each of the following four types (where letter ‘V’ stands for ‘variable-density’): V.T4.N20.D10K, V.T6.N20.D10K, V.T4.N20.D100K, and V.T6.N20.D100K; the corresponding first item selection probabilities and decrement steps (in parentheses) are 0.4 (0.02), 0.6 (0.03), 0.4 (0.02), and 0.6 (0.03) respectively.

We run experiments with support thresholds ranging from 0.5% to 5%. For each dataset type / threshold combination we have measured the execution time of PTF and TTF, averaging over ten experiments, one for each dataset of the type.

Results for the datasets with 10K transactions appear in Figure 4. Figure 5 shows results for the datasets with 100K transactions.

The comparison of the two algorithms is much more interesting when it comes to variable-density data sets. As before, PTF behaves better for small thresholds (roughly smaller than 2%) but TTF is faster for larger thresholds. Besides, PTF exhibits almost constant running time in most experiments. Now, whenever the  $T$ -tree is small and sparse, it happens that the few full  $P$ -tree traversals performed by PTF can take longer than the (more but not too many) partial  $P$ -tree traversals of TTF. The main reason is that potentially frequent itemsets consist mainly of lexicographically smaller items, hence the partial  $P$ -tree traversals of TTF are

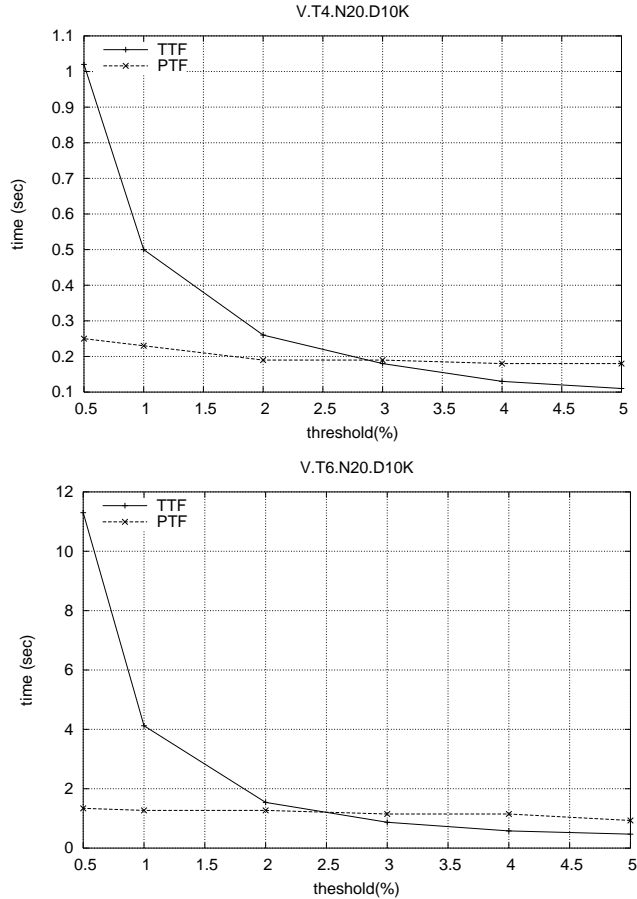


Figure 4: Results for datasets V.T4.N20.D10K (top) and V.T6.N20.D10K (bottom).

limited to a small part of the  $P$ -tree and are therefore much faster. On the other hand, TTF performs a full  $P$ -tree traversal at each level of the  $T$ -tree that contains potentially frequent itemsets, regardless of the number of these itemsets, hence it needs almost the same time as before, since it considers a similar number of levels.

Comparing the performance of the two algorithms with respect to uniformity of item densities one observes that while PTF exhibits roughly the same performance for both uniform and variable item densities, TTF is considerably faster on instances of variable item density; indeed, our results show that for variable-density datasets, TTF outperforms PTF for support thresholds above 3%, even above 2% or 1% in some cases. This is due to the fact that the performance of PTF is mainly determined by the rank of the higher level of frequent itemsets, while the performance of TTF depends heavily on the part of the  $P$ -tree that must be visited each time — which is much smaller for variable density instances, because frequent itemsets consist mainly of lexicographically smaller items.

Let us note here that for our experiments we built the variable-density datasets in such a

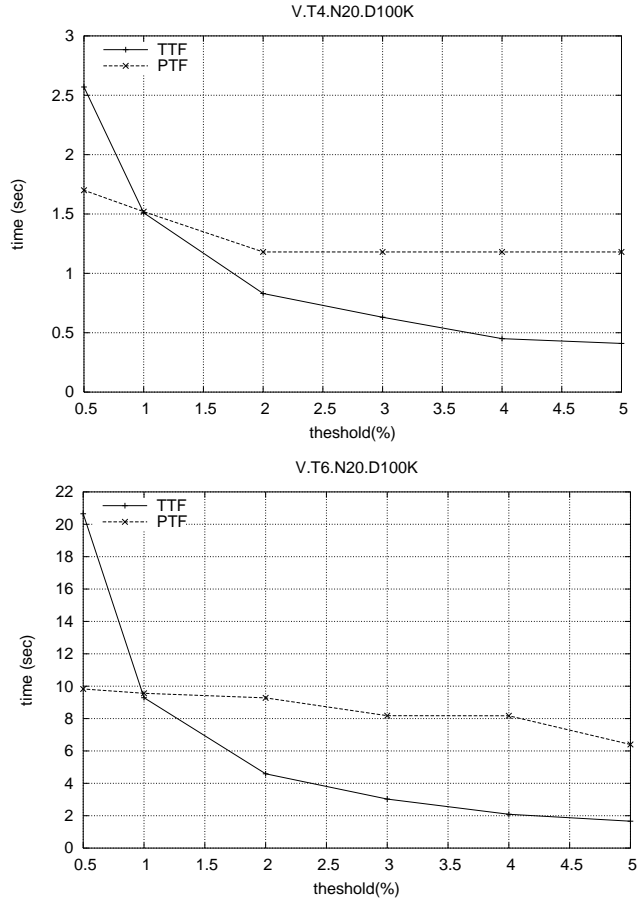


Figure 5: Results for datasets V.T4.N20.D100K (top) and V.T6.N20.D100K (bottom).

way that the lexicographically greater items are of smaller density. This property is essential for the performance of TTF, since it guarantees that most frequent itemsets consist mainly of lexicographically small items which appear in a small part of the  $P$ -tree. Therefore, to make TTF work well for real datasets, a sorting of the items in order of decreasing density should be performed in a preprocessing step.

**Realistic datasets.** In a third set of experiments we tested the behavior of both algorithms against widely used datasets, such as the ones contained in UCI Irvine Machine Learning repository. We have used two UCI datasets, namely `chess` and `mushroom` with typical suggested support threshold values (70, 75, 80, 85 % for `chess` and 20, 25, 30, 35 % for `mushroom`). Figures 6 and 7 show the time performance of the TTF algorithm. We observe that the decrease on execution time, when increasing the support threshold, is much steeper on `chess` than it is on `mushroom` dataset. This is due to the lower similarity between transactions `mushroom` compared to transactions of `chess`. This results to a larger variety of itemset frequencies for the former

dataset, while for the latter the vast majority of itemsets has a frequency of around 20%.

Unfortunately, we did not manage to obtain results for the PTF algorithm on these datasets because of memory overflow. This is mainly due to the particular structure of these datasets: both datasets (especially `chess`) contain transactions that are very similar to each other and so there are many frequent itemsets that are quite long (i.e. contain a large number of items) even in the case of large thresholds. Therefore, the  $T$ -tree becomes too large to fit into the main memory, since the whole tree must be stored and there are multiple pointers for each node; recall that, in contrast, in the case of TTF only the last level of the  $T$ -tree needs to be kept. Moreover, the existence of ‘deep’ levels in the  $T$ -tree results in huge numbers of generated subsets while traversing nodes of the  $P$ -tree which causes considerable slowdown.

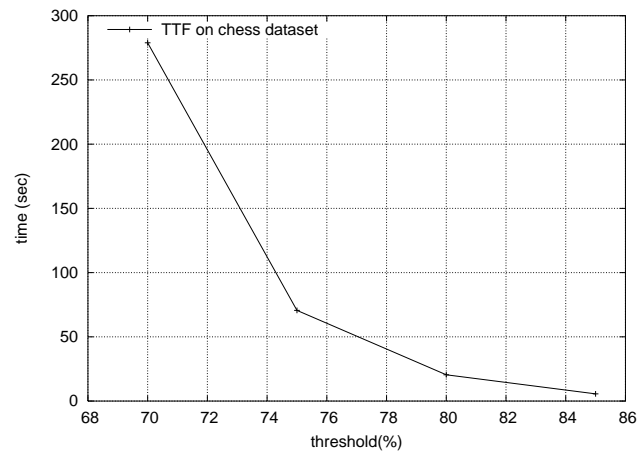


Figure 6: TTF performance on `chess` dataset




Figure 7: TTF performance on `mushroom` dataset

**Sparse datasets.** PTF however, behaves very well when we have to do with large datasets of smaller density than the one of `chess` and `mushroom` datasets. We used the IBM Quest Generator in order to generate datasets of such structure. Figure 8 shows the behavior of both algorithms; the superiority of PTF is clear when we have to do with such datasets.




Figure 8: PTF and TTF performance on synthetic datasets with 100 items

## 8 Conclusions

In this work we have developed and implemented two Apriori-style algorithms for the problem of frequent itemsets generation, called *T-Tree-First* (TTF) and *P-Tree-First* (PTF), that are based on the interim-support tree approach [10]. The two algorithms follow inverse approaches: TTF iterates over the itemsets of *T-tree*, and for each of them traverses the relevant part of the *P-tree* in order to count its total support; PTF starts by traversing the *P-tree* and for each visited node it updates all relevant nodes at the current level of the *T-tree*.

We have introduced several new techniques that result in faster algorithms comparing to earlier attempts that use similar tree structures [8, 10]. The most important of them are the *fixed-prefix potential inclusion* technique, which is used in algorithm TTF, and the use of *multiple pointers* in the *T-tree*, employed by PTF. The former allows faster support counting for *P-trees* that are built using only two pointers per node, thus being particularly memory-efficient. The latter provides fast access to the *T-tree* and makes PTF a generally efficient algorithm. We show experimentally that our new algorithms achieve considerable speedup comparing to the earlier algorithms.

The main difference between the two algorithms is that TTF performs a partial *P-tree* traversal for each potentially frequent itemset, while PTF performs only one, but full, *P-tree* traversal for each level of potentially frequent itemsets. As a result, PTF is considerably faster than TTF in instances where there are a lot of frequent itemsets, while TTF gains ground in

instances where there are fewer potentially frequent itemsets, especially if for each of them it suffices to check only a small part of the  $P$ -tree. For example, the latter case may occur whenever item densities have a high variance. However, PTF fails to perform well in the case of long frequent itemsets because the size of the  $T$ -tree becomes prohibitive; this calls for further optimisation techniques.

In conclusion, each of the two heuristics has its own merits and deserves further exploration. As a suggestion for further research, it would be interesting to investigate possible combinations of the two inverse approaches of TTF and PTF. For example, it seems reasonable to use PTF as long as the current level of the  $T$ -tree contains a lot of frequent itemsets and the level depth is small, while it may be wise to turn to TTF once the current level becomes sparse or if the level depth increases above a certain value.

## References

- [1] S. Ahmed, F. Coenen and P. Leng. Tree-based partitioning of data for association rule mining. *Knowledge and Information Systems*, 10, 2006, pp. 315-331
- [2] F. Angiulli, G. Ianni, L. Palopoli. On the complexity of inducing categorical and quantitative association rules, arXiv:cs.CC/0111009 vol. 1, Nov. 2001
- [3] R. Agrawal, T. Imielinski, and A. Swami. Mining Association Rules between Sets of Items in Large Databases. In *Proc. of ACM SIGMOD Conference on Management of Data*, Washington DC, May 1993.
- [4] R. Agrawal, T. Imielinski, and A. Swami. Database mining: a performance perspective. *IEEE Transactions on Knowledge and Data Engineering*, 5(6):914–925, Dec 1993. Special Issue on Learning and Discovery in Knowledge-Based Databases.
- [5] R. Agrawal and R. Srikant. Fast Algorithms for mining association rules. In *VLDB'94*, pp. 487–499.
- [6] R. Agrawal, C. Aggarwal and V. Prasad. Depth First Generation of Long Patterns. In *KDD 2000*, ACM, pp. 108-118.
- [7] R.J.Bayardo. Efficiently Mining Long Patterns from Databases. In *Proc. SIGMOD 1998*, pp.85-93.
- [8] F. Coenen, G. Goulbourne, and P. Leng. Computing Association Rules using Partial Totals. In L. De Raedt and A. Siebes eds, *Principles of Data Mining and Knowledge Discovery* (Proc 5th European Conference, PKDD 2001, Freiburg, Sept 2001), Lecture Notes in AI 2168, Springer-Verlag, Berlin, Heidelberg: pp. 54–66.
- [9] F. Coenen, G. Goulbourne and P. Leng. Tree Structures for Mining Association Rules. *Data Mining and Knowledge Discovery*, 8 (2004), pp. 25-51.
- [10] G. Goulbourne, F. Coenen and P. Leng. Algorithms for Computing Association Rules using a Partial-Support Tree. *Journal of Knowledge-Based Systems* 13 (2000), pp. 141–149.
- [11] G Goulbourne, F. Coenen and P Leng. Algorithms for Computing Association Rules using a Partial-Support Tree. In *Proc ES99*, pp. 132-147.

- [12] J. Han, J. Pei, Y. Yin and R. Mao. Mining Frequent Patterns without Candidate Generation: A Frequent-Pattern Tree Approach. *Data Mining and Knowledge Discovery*, 8 (2004), pp. 53-87.
- [13] J. Han, J. Pei and Y. Yin. Mining Frequent patterns without candidate generation. In *SIGMOD 2000*, pp. 1-12.
- [14] H. Huang, X. Wu and R. Relue R. Association Analysis with One Scan of Databases. In *Proc ICDM'02*, pp. 629-632.
- [15] G. Liu, H. Lu, W. Lou and J. Yu. On Computing, Storing and Querying Frequent Patterns. In *Proc KDD 2003*, pp. 607-612.
- [16] A. Savasere, E. Omiecinski and S. Navathe. An Efficient Algorithm for Mining Association Rules in Large Databases. In *VLDB 1995*, pp. 432-444.
- [17] H. Toivonen. Sampling Large Databases for Association Rules. In *VLDB 1996*, pp. 1-12.
- [18] M. J. Zaki. Generating Non-Redundant Association Rules. In *Proc. SIGKDD-2000*, pp. 34-43, 2000.