

A tight upper bound on the number of candidate patterns*

Floris Geerts, Bart Goethals, Jan Van den Bussche
University of Limburg, Belgium

Abstract

In the context of mining for frequent patterns using the standard levelwise algorithm, the following question arises: given the current level and the current set of frequent patterns, what is the maximal number of candidate patterns that can be generated on the next level? We answer this question by providing a tight upper bound, derived from a combinatorial result from the sixties by Kruskal and Katona. Our result is useful to reduce the number of database scans.

*A preliminary report on this work was presented at the 2001 IEEE International Conference on Data Mining [13].

1 Introduction

The frequent pattern mining problem [3] is by now well known. We are given a set of items \mathcal{I} and a database \mathcal{D} of subsets of \mathcal{I} called transactions. A *pattern* is some set of items; its *support* in \mathcal{D} is defined as the number of transactions in \mathcal{D} that contain the pattern; and a pattern is called *frequent* in \mathcal{D} if its support exceeds a given minimal support threshold. The goal is now to find all frequent patterns in \mathcal{D} .

The search space of this problem, all subsets of \mathcal{I} , is clearly huge. Instead of generating and counting the supports of all these patterns at once, several solutions have been proposed to perform a more directed search through all patterns. However, this directed search enforces several scans through the database, which brings up another great cost, because these databases tend to be very large, and hence they do not fit into main memory.

The standard Apriori algorithm [4] for solving this problem is based on its monotonicity property, that all subsets of a frequent pattern must be frequent. A pattern is thus considered potentially frequent, also called a *candidate* pattern, if its support is yet unknown, but all of its subsets are already known to be frequent. In every step of the algorithm, all candidate patterns are generated and their supports are then counted by performing a complete scan of the transaction database. This is repeated until no new candidate patterns can be generated. Hence, the number of scans through the database equals the maximal size of a candidate pattern. Several improvements on the Apriori algorithm try to reduce the number of scans through the database by estimating the number of candidate patterns that can still be generated.

At the heart of all these techniques lies the following purely combinatorial problem, that must be solved first before we can seriously start applying them: *given the current set of frequent patterns at a certain pass of the algorithm, what is the maximal number of candidate patterns that can be generated in the passes yet to come?*

Our contribution is to solve this problem by providing a hard and tight combinatorial upper bound. By computing our upper bound after every pass of the algorithm, we have at all times a watertight guarantee on the size of what is still to come, on which we can then base various optimization decisions, depending on the specific algorithm that is used.

In the next Section, we will discuss existing techniques to reduce the number of database scans, and point out the dangers of using existing heuristics for this purpose. Using our upper bound, these techniques can be made watertight. In Section 3, we derive our upper bound, using a combinatorial result from the sixties by Kruskal and Katona. In Section 4, we show how

to get even more out of this upper bound by applying it recursively. We will then generalize the given upper bounds such that they can be applied by a wider range of algorithms in Section 5. In Section 6, we discuss several issues concerning the implementation of the given upper bounds on top of Apriori-like algorithms. In Section 7, we give experimental results, showing the effectiveness of our result in estimating, far ahead, how much will still be generated in the future. Finally, we conclude the paper in Section 8.

2 Related Work

Nearly all frequent pattern mining algorithms developed after the proposal of the Apriori algorithm, rely on its levelwise candidate generation and pruning strategy. Most of them differ in how they generate and count candidate patterns.

One of the first optimizations was the DHP algorithm proposed by Park et al. [26]. This algorithm uses a hashing scheme to collect upper bounds on the frequencies of the candidate patterns for the following pass. Patterns of which it is already known that they will turn up infrequent can then be eliminated from further consideration. The effectiveness of this technique only showed for the first few passes. Since our upper bound can be used to eliminate passes at the end, both techniques can be combined in the same algorithm.

Other strategies, discussed next, try to reduce the number of passes. However, such a reduction of passes often causes an increase in the number of candidate patterns that need to be explored during a single pass. This trade-off between the reduction of passes and the number of candidate patterns is important since the time needed to process a transaction is dependent on the number of candidates that are covered in that transaction, which might blow up exponentially. Our upper bound can be used to predict whether or not this blowup will occur.

The Partition algorithm, proposed by Savasere et al. [27], reduces the number of database passes to two. Towards this end, the database is partitioned into parts small enough to be handled in main memory. The partitions are then considered one at a time and all frequent patterns for that partition are generated using an Apriori-like algorithm. At the end of the first pass, all these patterns are merged to generate a set of all potential frequent patterns, which can then be counted over the complete database. Although this method performs only two database passes, its performance is heavily dependent on the distribution of the data, and could generate much too many candidates.

The sampling algorithm proposed by Toivonen [29] performs at most two scans through the database by picking a random sample from the database, then finding all frequent patterns that probably hold in the whole database, and then verifying the results with the rest of the database. In the cases where the sampling method does not produce all frequent patterns, the missing patterns can be found by generating all remaining potentially frequent patterns and verifying their frequencies during a second pass through the database. The probability of such a failure can be kept small by decreasing the minimal support threshold. However, for a reasonably small probability of failure, the threshold must be drastically decreased, which can again cause a combinatorial explosion of the number of candidate patterns.

The DIC algorithm, proposed by Brin et al. [11], tries to reduce the number of passes over the database by dividing the database into intervals of a specific size. First, all candidate patterns of size 1 are generated. The frequencies of the candidate sets are then counted over the first interval of the database. Based on these frequencies, candidate patterns of size 2 are generated and are counted over the next interval together with the patterns of size 1. In general, after every interval k , candidate patterns of size $k + 1$ are generated and counted. The algorithm stops if no more candidates can be generated. Again, this technique can be combined with our technique in the same algorithm.

Another type of algorithms generate frequent patterns using a depth-first search [30, 1, 2, 17]. Generating patterns in a depth-first manner implies that the monotonicity property cannot be exploited anymore. Hence, a lot more patterns will be generated and need to be counted, compared to the breadth-first algorithms. The FPGrowth algorithm from Han et al. solves this problem by loading a compressed form of the database in main memory using the proposed FPtree. This memory-resident FPtree benefits from a very fast counting mechanism of all generated patterns.¹ Obviously, it is not always possible to load the compressed form of the database into main memory.

Other strategies try to push certain constraints into the candidate pattern generation as deeply as possible to reduce the number of candidate patterns that must be generated [14, 21, 24, 28]. Still others try to find only the set of *maximal* frequent patterns, i.e. those frequent patterns that have no superset which is also frequent [8, 30, 22]. Of course, these techniques do not give us all frequencies of all frequent patterns as required by the general pattern mining problem we consider in this paper.

The first heuristic specifically proposed to estimate the number of candidate patterns that can still be generated was used in the AprioriHybrid

¹Note that the patterns in the FPtree are represented in the so called header tables.

algorithm [6, 7]. This algorithm uses Apriori in the initial iterations and switches to AprioriTid if it expects it to run faster. This AprioriTid algorithm does not use the database at all for counting the support of candidate patterns. Rather, an encoding of the candidate patterns used in the previous iteration is employed for this purpose. The AprioriHybrid algorithm switches to AprioriTid when it expects this encoding of the candidate patterns to be small enough to fit in main memory. The size of the encoding grows with the number of candidate patterns. Therefore, it calculates the size the encoding would have in the current iteration. If this size is small enough and there were fewer candidate patterns in the current iteration than the previous iteration, the heuristic decides to switch to AprioriTid.

This heuristic (like all heuristics) is not waterproof, however. Take, for example, two disjoint datasets. The first dataset consists of all subsets of a frequent pattern of size 20. The second dataset consists of all subsets of 1 000 disjoint frequent patterns of size 5. If we merge these two datasets, we get $\binom{20}{3} + 1\,000\binom{5}{3} = 11\,140$ patterns of size 3 and $\binom{20}{4} + 1\,000\binom{5}{4} = 9\,845$ patterns of size 4. If we have enough memory to store the encoding for all these patterns, then the heuristic decides to switch to AprioriTid. This decision is premature, however, because the number of new patterns in each pass will start growing exponentially afterwards.

Also, current state-of-the-art algorithms for frequent itemset mining, such as Opportunistic Project [23] and DCI [25] use several techniques within the same algorithm and switch between these techniques using several simple, but not waterproof heuristics.

Another improvement of the Apriori algorithm, which is part of the folklore, tries to combine as many iterations as possible in the end, when only few candidate patterns can still be generated. The potential of such a combination technique was realized early on [6, 4], but the modalities under which it can be applied were never further examined. Our work does exactly that.

3 The basic upper bounds

In all that follows, L is some family of patterns of size k .

Definition 1. A *candidate pattern* for L is a pattern (of size larger than k) of which all k -subsets are in L . For a given $p > 0$, we denote the set of all size- $k + p$ candidate patterns for L by $C_{k+p}(L)$.

For any $p \geq 1$, we will provide an upper bound on $|C_{k+p}(L)|$ in terms of $|L|$. The following lemma is central to our approach. (A simple proof was given by Katona [18].)

Lemma 1. *Given n and k , there exists a unique representation*

$$n = \binom{m_k}{k} + \binom{m_{k-1}}{k-1} + \cdots + \binom{m_r}{r},$$

with $r \geq 1$, $m_k > m_{k-1} > \cdots > m_r$, and $m_i \geq i$ for $i = r, r+1, \dots, k$.

This representation is called the k -canonical representation of n and can be computed as follows: The integer m_k satisfies $\binom{m_k}{k} \leq n < \binom{m_k+1}{k}$, the integer m_{k-1} satisfies $\binom{m_{k-1}}{k-1} \leq n - \binom{m_k}{k} < \binom{m_{k-1}+1}{k-1}$, and so on, until $n - \binom{m_k}{k} - \binom{m_{k-1}}{k-1} - \cdots - \binom{m_r}{r}$ is zero.

We now establish:

Theorem 2. *If*

$$|L| = \binom{m_k}{k} + \binom{m_{k-1}}{k-1} + \cdots + \binom{m_r}{r}$$

in k -canonical representation, then

$$|C_{k+p}(L)| \leq \binom{m_k}{k+p} + \binom{m_{k-1}}{k-1+p} + \cdots + \binom{m_{s+1}}{s+p+1},$$

where s is the smallest integer such that $m_s < s+p$. If no such integer exists, we set $s = r - 1$.

Proof. Suppose, for the sake of contradiction, that

$$|C_{k+p}(L)| \geq \binom{m_k}{k+p} + \binom{m_{k-1}}{k-1+p} + \cdots + \binom{m_{s+1}}{s+p+1} + \binom{s+p}{s+p}.$$

Note that this is in $k+p$ -canonical representation. A theorem by Kruskal and Katona [12, 18, 20] says that

$$|L| \geq \binom{m_k}{k} + \binom{m_{k-1}}{k-1} + \cdots + \binom{m_{s+1}}{s+1} + \binom{s+p}{s}.$$

But this is impossible, because

$$\begin{aligned} |L| &= \binom{m_k}{k} + \binom{m_{k-1}}{k-1} + \cdots + \binom{m_{s+1}}{s+1} + \binom{m_s}{s} + \cdots + \binom{m_r}{r} \\ &\leq \binom{m_k}{k} + \binom{m_{k-1}}{k-1} + \cdots + \binom{m_{s+1}}{s+1} + \sum_{1 \leq i \leq s} \binom{i+p-1}{i} \\ &< \binom{m_k}{k} + \binom{m_{k-1}}{k-1} + \cdots + \binom{m_{s+1}}{s+1} + \sum_{0 \leq i \leq s} \binom{i+p-1}{i} \\ &= \binom{m_k}{k} + \binom{m_{k-1}}{k-1} + \cdots + \binom{m_{s+1}}{s+1} + \binom{s+p}{s}. \end{aligned}$$

The first inequality follows from the observation that $m_s \leq s + p - 1$ implies $m_i \leq i + p - 1$ for all $i = s, s - 1, \dots, r$. The last equality follows from a well-known binomial identity. \square

Notation We will refer to the upper bound provided by the above theorem as $KK_k^{k+p}(|L|)$ (for Kruskal-Katona). The subscript k , the level at which we are predicting, is important, as the only parameter is the cardinality $|L|$ of L , not L itself. The superscript $k + p$ denotes the level we are predicting.

Proposition 3 (Tightness). *The upper bound provided by Theorem 2 is tight: for any given n and k there always exists an L with $|L| = n$ such that for any given p , $|C_{k+p}(L)| = KK_k^{k+p}(|L|)$.*

Proof. Let us write a finite set of natural numbers as a string of natural numbers by writing its members in decreasing order. We can then compare two such sets by comparing their strings in lexicographic order. The resulting order on the sets is known as the *colexicographic* (or *colex*) order. An intuitive proof of the Kruskal-Katona theorem, based on this colex order, was given by Bollobás [10]. Let

$$\binom{m_k}{k} + \binom{m_{k-1}}{k-1} + \dots + \binom{m_r}{r}$$

be the k -canonical representation of n . Then, Bollobás has shown that all $k - p$ -subsets of the first n k -sets of natural numbers in colex order, are exactly the first

$$\binom{m_k}{k-p} + \binom{m_{k-1}}{k-1-p} + \dots + \binom{m_s}{r-s}$$

$k - p$ -sets of natural numbers in colex order, with s the smallest integer such that $s > p$. Using the same reasoning as above, we can conclude that all $k + p$ -supersets of the first n k -sets of natural numbers in colex order are exactly the first $KK_k^{k+p}(n)$ $k + p$ -sets of natural numbers in colex order. \square

Analogous tightness properties hold for all upper bounds we will present in this paper, but we will no longer explicitly state this.

Example 1. Let L be the set of 13 patterns of size 3:

$$\begin{aligned} & \{\{3, 2, 1\}, \{4, 2, 1\}, \{4, 3, 1\}, \{4, 3, 2\}, \\ & \{5, 2, 1\}, \{5, 3, 1\}, \{5, 3, 2\}, \{5, 4, 1\}, \{5, 4, 2\}, \{5, 4, 3\}, \\ & \{6, 2, 1\}, \{6, 3, 1\}, \{6, 3, 2\}\}. \end{aligned}$$

The 3-canonical representation of 13 is $\binom{5}{3} + \binom{3}{2}$ and hence the maximum number of candidate patterns of size 4 is $KK_3^4(13) = \binom{5}{4} + \binom{3}{3} = 6$ and the maximum number of candidate patterns of size 5 is $KK_3^5(13) = \binom{5}{5} = 1$. This is tight indeed, because

$$C_4(L) = \{\{4, 3, 2, 1\}, \{5, 3, 2, 1\}, \{5, 4, 2, 1\}, \\ \{5, 4, 3, 1\}, \{5, 4, 3, 2\}, \{6, 3, 2, 1\}\}$$

and

$$C_5(L) = \{\{5, 4, 3, 2, 1\}\}.$$

Estimating the number of levels The k -canonical representation of $|L|$ also yields an upper bound on the maximal size of a candidate pattern, denoted by $\text{maxsize}(L)$. Recall that this size equals the number of iterations the standard Apriori algorithm will perform. Indeed, since $|L| < \binom{m_k+1}{k}$, there cannot be a candidate pattern of size $m_k + 1$ or higher, so:

Proposition 4. *If $\binom{m_k}{k}$ is the first term in the k -canonical representation of $|L|$, then $\text{maxsize}(L) \leq m_k$.*

We denote this number m_k by $\mu_k(|L|)$. From the form of KK_k^{k+p} as given by Theorem 2, it is immediate that μ also tells us the last level before which KK becomes zero. Formally:

Proposition 5.

$$\mu_k(|L|) = k + \min\{p \mid KK_k^{k+p}(|L|) = 0\} - 1.$$

Estimating all levels As a result of the above, we can also bound, at any given level k , the *total* number of candidate patterns that can be generated, as follows:

Proposition 6. *The total number of candidate patterns that can be generated from a set L of k -patterns is at most*

$$KK_k^{\text{total}}(|L|) := \sum_{p \geq 1} KK_k^{k+p}(|L|).$$

4 Improved upper bounds

The upper bound KK on itself is neat and simple as it takes as parameters only two numbers: the current size k , and the number $|L|$ of current frequent

patterns. However, in reality, when we have arrived at a certain level k , we do not merely have the cardinality: we have the actual set L of current k -patterns! For example, if the frequent patterns in the current pass are all disjoint, our current upper bound will still estimate their number to a certain non-zero figure. However, by the pairwise disjointness, it is clear that no further patterns will be possible at all. In sum, because we have richer information than a mere cardinality, we should be able to get a better upper bound.

To get inspiration, let us recall that the candidate generation process of the Apriori algorithm works in two steps. In the *join* step, we join L with itself to obtain a superset of C_{k+1} . The union $p \cup q$ of two patterns $p, q \in L$ is inserted in C_{k+1} if they share their $k - 1$ smallest items:

```
insert into  $C_{k+1}$ 
select  $p[1], p[2], \dots, p[k], q[k]$ 
from  $L_k$   $p, L_k$   $q$ 
where  $p[1] = q[1], \dots, p[k - 1] = q[k - 1], p[k] < q[k]$ 
```

Next, in the *prune* step, we delete every pattern $c \in C_{k+1}$ such that some k -subset of c is not in L .

Let us now take a closer look at the join step from another point of view. Consider a family of all frequent patterns of size k that share their $k - 1$ smallest items, and let its cardinality be n . If we now remove from each of these patterns all these shared $k - 1$ smallest items, we get exactly n distinct single-item patterns. The number of pairs that can be formed from these single items, being $\binom{n}{2}$, is exactly the number of candidates the join step will generate for the family under consideration. We thus get an obvious upper bound on the total number of candidates by taking the sum of all $\binom{n_f}{2}$, for every possible family f .

This obvious upper bound on $|C_{k+1}|$, which we denote by $obvious_{k+1}(L)$, can be recursively computed in the following manner. Let I denote the set of items occurring in L . For an arbitrary item x , define the set L^x as

$$L^x = \{s - \{x\} \mid s \in L \text{ and } x = \min s\}.$$

Then

$$obvious_{k+1}(L) := \begin{cases} \binom{|L|}{2} & \text{if } k = 1; \\ \sum_{x \in I} obvious_k(L^x) & \text{if } k > 1. \end{cases}$$

This upper bound is much too crude, however, because it does not take the prune step into account, only the join step. The join step only checks two k -subsets of a potential candidate instead of all $k + 1$ k -subsets.

However, we can generalize this method such that more subsets will be considered. Indeed, instead of taking a family of all frequent patterns sharing their $k-1$ smallest items, we can take all frequent patterns sharing only their k' smallest items, for some $k' \leq k-1$. If we then remove these k' shared items from each pattern in the family, we get a new set L' of n patterns of size $k-k'$. If we now consider the set C' of candidates (of size $k-k'+1$) for L' , and add back to each of them the previously removed k' items, we obtain a pruned set of candidates of size $k+1$, where instead of just two (as in the join step), $k-k'+1$ of the k -subsets were checked in the pruning. Note that we can get the estimate $KK_{k-k'}^{k-k'+1}(|L'|)$ on the cardinality of C' from our upper bound Theorem 2.

Doing this for all possible values of k' yields an improved upper bound on $|C_{k+1}|$, which we denote by $improved_{k+1}(L)$, and which is computed by refining the recursive procedure for the obvious upper bound as follows:

$$improved_{k+1}(L) := \begin{cases} \binom{|L|}{2} & \text{if } k = 1; \\ \min\{KK_k^{k+1}(|L|), \sum_{x \in I} improved_k(L^x)\} & \text{if } k > 1. \end{cases}$$

Actually, as in the previous section, we can do this not only to estimate $|C_{k+1}|$, but also more generally to estimate $|C_{k+p}|$ for any $p \geq 1$. Henceforth we will denote our general improved upper bound by $KK_{k+p}^*(L)$. The general definition is as follows:

$$KK_{k+p}^*(L) := \begin{cases} KK_k^{k+p}(|L|) & \text{if } k = 1; \\ \min\{KK_k^{k+p}(|L|), \sum_{x \in I} KK_{k+p-1}^*(L^x)\} & \text{if } k > 1. \end{cases}$$

(For the base case, note that $KK_k^{k+p}(|L|)$, when $k = 1$, is nothing but $\binom{|L|}{p+1}$.)

By definition, KK_{k+p}^* is always smaller than KK_k^{k+p} . We now prove formally that it is still an upper bound on the number of candidate patterns of size $k+p$:

Theorem 7.

$$|C_{k+p}(L)| \leq KK_{k+p}^*(L).$$

Proof. By induction on k . The base case $k = 1$ is clear. For $k > 1$, it suffices to show that for all $p > 0$

$$C_{k+p}(L) \subseteq \bigcup_{x \in I} C_{k+p-1}(L^x) + x. \quad (1)$$

(For any set of patterns H , we denote $\{h \cup \{x\} \mid h \in H\}$ by $H + x$.)

From the above containment we can conclude

$$\begin{aligned}
|C_{k+p}(L)| &\leq \left| \bigcup_{x \in I} C_{k+p-1}(L^x) + x \right| \\
&\leq \sum_{x \in I} |C_{k+p-1}(L^x) + x| \\
&= \sum_{x \in I} |C_{k+p-1}(L^x)| \\
&\leq \sum_{x \in I} KK_{k+p-1}^*(L^x)
\end{aligned}$$

where the last inequality is by induction.

To show (1), we need to show that for every $p > 0$ and every $s \in C_{k+p}(L)$, $s - \{x\} \in C_{k+p-1}(L^x)$, where $x = \min s$. This means that every subset of $s - \{x\}$ of size $k - 1$ must be an element of L^x . Let $s - \{x\} - \{y_1, \dots, y_p\}$ be such a subset. This subset is an element of L^x iff $s - \{y_1, \dots, y_p\} \in L$ and $x = \min(s - \{y_1, \dots, y_p\})$. The first condition follows from $s \in C_{k+p}(L)$, and the second condition is trivial. Hence the theorem. \square

A natural question is why we must take the minimum in the definition of KK^* . The answer is that the two terms of which we take the minimum are incomparable. The example of an L where all patterns are pairwise disjoint, already mentioned in the beginning of this section, shows that, for example, $KK_k^{k+1}(|L|)$ can be larger than the summation $\sum_{x \in I} KK_k^*(L^x)$. But the converse is also possible: consider $L = \{\{1, 2\}, \{1, 3\}\}$. Then $KK_2^3(L) = 0$, but the summation yields 1.

Example 2. Let L consist of $\{5, 7, 8\}$ and $\{5, 8, 9\}$ plus all 19 3-subsets of $\{1, 2, 3, 4, 5\}$ and $\{3, 4, 5, 6, 7\}$. Because $21 = \binom{6}{3} + \binom{2}{2}$, we have $KK_3^4(21) = 15$, $KK_3^5(21) = 6$ and $KK_3^6(21) = 1$. On the other hand,

$$\begin{aligned}
KK_4^*(L) &= KK_3^*(L^1) + KK_3^*(L^2) + KK_3^*(L^3) + KK_3^*(L^4) \\
&\quad + KK_2^*((L^5)^6) + KK_2^*((L^5)^7) + KK_2^*((L^5)^8) + KK_2^*((L^5)^9) \\
&\quad + KK_3^*(L^6) + KK_3^*(L^7) + KK_3^*(L^8) + KK_3^*(L^9) \\
&= 4 + 1 + 4 + 1 + 0 + \dots + 0 \\
&= 10
\end{aligned}$$

and

$$\begin{aligned}
KK_5^*(L) &= KK_4^*(L^1) + KK_4^*(L^2) + KK_4^*(L^3) + KK_4^*(L^4) \\
&\quad + KK_3^*((L^5)^6) + KK_3^*((L^5)^7) + KK_3^*((L^5)^8) + KK_3^*((L^5)^9) \\
&\quad + KK_4^*(L^6) + KK_4^*(L^7) + KK_4^*(L^8) + KK_4^*(L^9) \\
&= 1 + 0 + 1 + 0 + 0 + \cdots + 0 \\
&= 2.
\end{aligned}$$

Indeed, we have 10 4-subsets of $\{1, 2, 3, 4, 5\}$ and $\{3, 4, 5, 6, 7\}$, and the two 5-sets themselves.

We can also improve the upper bound $\mu_k(|L|)$ on $\text{maxsize}(L)$. In analogy with Proposition 5, we define:

$$\mu_k^*(L) := k + \min\{p \mid KK_{k+p}^*(L) = 0\} - 1.$$

We then have:

Proposition 8.

$$\text{maxsize}(L) \leq \mu_k^*(L) \leq \mu_k(L).$$

We finally use Theorem 7 for improving the upper bound KK_k^{total} on the total number of candidate patterns. We define:

$$KK_{\text{total}}^*(L) := \sum_{p \geq 1} KK_{k+p}^*(L).$$

Then we have:

Proposition 9. *The total number of candidate patterns that can be generated from a set L of k -patterns is bounded by $KK_{\text{total}}^*(L)$. Moreover,*

$$KK_{\text{total}}^*(L) \leq KK_k^{\text{total}}(L).$$

5 Generalized upper bounds

The upper bounds presented in the previous sections work well for algorithms that generate and test candidate patterns of one specific size at a time. However, a lot of algorithms generate and test patterns of different sizes within the same pass of the algorithm [11, 8, 29]. Hence, these algorithms know in advance that several patterns of size larger than k are frequent or not. Since our upper bound is solely based on the patterns of a certain length k , it does not use information about patterns of length larger than k .

Nevertheless, these larger sets could give crucial information. More specifically, suppose we have generated all frequent patterns of size k , and we also already know in advance that a certain set of size larger than k is not frequent. Our upper bound on the total number of candidate patterns that can still be generated, would disregard this information. We will therefore generalize our upper bound such that it will also incorporate this additional information.

5.1 Generalized KK -bound

From now on, L is some family of sets of patterns $L_k, L_{k+1}, \dots, L_{k+q}$ which are known to be frequent, such that L_{k+p} contains patterns of size $k+p$, and all $k+p-1$ -subsets of all patterns in L_{k+p} are in L_{k+p-1} . We denote by $|L|$ the sequence of numbers $|L_k|, |L_{k+1}|, \dots, |L_{k+q}|$.

Similarly, let I be a family of sets of patterns $I_k, I_{k+1}, \dots, I_{k+q}$ which are known to be infrequent, such that I_{k+p} contains patterns of size $k+p$ and all $k+p-1$ -subsets of all patterns in I_{k+p} are in L_{k+p-1} . We denote by $|I|$ the sequence of numbers $|I_k|, |I_{k+1}|, \dots, |I_{k+q}|$. Note that for each $p \geq 0$, L_{k+p} and I_{k+p} are disjoint.

Before we present the general upper bounds, we also generalize our notion of a candidate pattern.

Definition 2. A *candidate pattern* for (L, I) of size $k+p$ is a pattern which is not in L_{k+p} or I_{k+p} , all of its k -subsets are in L_k , and none of its subsets of size larger than k is included in $I_k \cup I_{k+1} \cup \dots \cup I_{k+q}$. For a given p , we denote the set of all $k+p$ -size candidate patterns for (L, I) by $C_{k+p}(L, I)$.

We note:

Lemma 10.

$$C_{k+p}(L, I) = \begin{cases} C_{k+1}(L_k) \setminus (L_{k+1} \cup I_{k+1}) & \text{if } p = 1; \\ C_{k+p}(C_{k+p-1}(L, I) \cup L_{k+p-1}) \setminus (L_{k+p} \cup I_{k+p}) & \text{if } p > 1. \end{cases}$$

Proof. The case $p = 1$ is clear. For $p > 1$, we show the inclusion in both directions.

\supseteq For every set in $C_{k+p}(C_{k+p-1}(L, I) \cup L_{k+p-1})$, we know that all of its k -subsets are always contained in a $k+p-1$ subset, and these are in $C_{k+p-1}(L, I) \cup L_{k+p-1}$. By definition, we know that for every set in $C_{k+p-1}(L, I)$, all of its k -subsets are in L_k . Also, for every set in L_{k+p-1} , all of its k -subsets are in L_k . By definition, for every set in $C_{k+p-1}(L, I)$, all of its $k+p-i$ -subsets are not in I_{k+p-i} . Also, for

every set in L_{k+p-1} , all of its $k+p-i$ -subsets are in L_{k+p-i} and hence they are not in I_{k+p-i} since they are disjoint. By definition, none of the patterns in $L_{k+p} \cup I_{k+p}$ are in $C_{k+p}(L, I)$.

\subseteq It suffices to show that for every set in $C_{k+p}(L, I)$, every $k+p-1$ -subset s is in $C_{k+p-1}(L, I) \cup L_{k+p-1}$. Obviously, this is true, since if it is not already in L_{k+p-1} , still all k -subsets of s must be in L_k , s can not be in I_{k+p-1} and none of its subsets can be in any $I_{k+p-\ell}$ with $\ell > 1$.

□

Hence, we define

$$gKK_k^{k+p}(|L|, |I|) := \begin{cases} KK_k^{k+1}(|L_k|) - |L_{k+1}| - |I_{k+1}| & \text{if } p = 1; \\ KK_{k+p-1}^{k+p}(gKK_k^{k+p-1}(|L|, |I|) + |L_{k+p-1}|) - |L_{k+p}| - |I_{k+p}| & \text{if } p > 1, \end{cases}$$

and obtain:

Theorem 11.

$$|C_{k+p}(L, I)| \leq gKK_k^{k+p}(|L|, |I|) \leq KK_k^{k+p}(|L_k|) - |L_{k+p}| - |I_{k+p}|.$$

Proof. The first inequality is clear by Lemma 10. The second inequality is by induction on p . The base case $p = 1$ is by definition. For $p > 1$, we have:

$$\begin{aligned} gKK_k^{k+p}(|L|, |I|) &= KK_{k+p-1}^{k+p}(gKK_k^{k+p-1}(|L|, |I|) + |L_{k+p-1}|) \\ &\quad - |L_{k+p}| - |I_{k+p}| \\ &\leq KK_{k+p-1}^{k+p}(KK_k^{k+p-1}(|L_k|) - |I_{k+p-1}|) - |L_{k+p}| - |I_{k+p}| \\ &\leq KK_{k+p-1}^{k+p}(KK_k^{k+p-1}(|L_k|)) - |L_{k+p}| - |I_{k+p}| \\ &= KK_k^{k+p}(|L_k|) - |L_{k+p}| - |I_{k+p}| \end{aligned}$$

where the first inequality is by induction and because of the monotonicity of KK , the second inequality also because of the monotonicity of KK and the last equality follows from

$$KK_k^{k+p}(|L_k|) = KK_{k+p-1}^{k+p}(KK_k^{k+p-1}(|L_k|)).$$

□

Again, we can also generalize the upper bound on the maximal size of a candidate pattern, denoted by $\text{maxsize}(L, I)$, and the upper bound on the total number of candidate patterns, both also incorporating (L, I) :

$$g\mu(|L|, |I|) := k + \min\{p \mid gKK_k^{k+p}(|L|, |I|) = 0\} - 1$$

$$gKK_k^{\text{total}}(|L|, |I|) := \sum_{p \geq 1} gKK_k^{k+p}(|L|, |I|).$$

We obtain:

Proposition 12.

$$\text{maxsize}(L, I) \leq g\mu(|L|, |I|) \leq \mu(|L|).$$

Proposition 13. *The total number of candidate patterns that can be generated from (L, I) is bounded by $gKK_k^{\text{total}}(|L|, |I|)$. Moreover,*

$$gKK_k^{\text{total}}(|L|, |I|) \leq KK_k^{\text{total}}(|L_k|).$$

Example 3. Suppose L_3 consists of all subsets of size 3 of the set $\{1, 2, 3, 4, 5, 6\}$. Now assume we already know that I_4 contains patterns $\{1, 2, 3, 4\}$ and $\{3, 4, 5, 6\}$. The KK upper bound presented in the previous section would estimate the number of candidate patterns of sizes 4, 5, and 6 to be at most $\binom{6}{4} = 15$, $\binom{6}{5} = 6$, and $\binom{6}{6} = 1$ respectively. Nevertheless, using the additional information, gKK can already reduce these numbers to 13, 3, and 0. Also, μ would predict the maximal size of a candidate pattern to be 6, while $g\mu$ can already predict this number to be at most 5. Similarly, KK_{total} would predict the total number of candidate patterns that can still be generated to be at most 22, while gKK_{total} can already deduce this number to be at most 16.

5.2 Generalized KK^* -bound

Using the generalized basic upper bound, we can now also generalize our improved upper bound KK^* . For an arbitrary item x , define the family of sets L^x as $L_k^x, L_{k+1}^x, \dots, L_{k+q}^x$, and I^x as $I_k^x, I_{k+1}^x, \dots, I_{k+q}^x$. We define:

$$gKK_{k+p}^*(L, I) := \begin{cases} gKK_k^{k+p}(|L|, |I|) & \text{if } k = 1; \\ \min\{gKK_k^{k+p}(|L|, |I|), \sum_{x \in I} gKK_{k+p-1}^*(L^x, I^x)\} & \text{if } k > 1. \end{cases}$$

We then have:

Theorem 14.

$$|C_{k+p}(L, I)| \leq gKK_{k+p}^*(L, I) \leq KK_{k+p}^*(L_k) - |L_{k+p}| - |I_{k+p}|.$$

Proof. The proof of the first inequality is similar to the proof of Theorem 7, instead that we now need to show that for all $p > 0$,

$$C_{k+p}(L, I) \subseteq \bigcup_{x \in I} C_{k+p-1}(L^x, I^x) + x.$$

Therefore, we need to show for every $s \in C_{k+p}(L, I)$, $s - \{x\} \in C_{k+p-1}(L^x, I^x)$, where $x = \min s$. First, this means that every subset of $s - \{x\}$ of size $k - 1$ must be in L_k^x . Let $s - \{x\} - \{y_1, \dots, y_p\}$ be such a subset. This subset is an element of L_k^x if and only if $s - \{y_1, \dots, y_p\} \in L_k$ and $x = \min(s - \{y_1, \dots, y_p\})$. The first condition follows from $s \in C_{k+p}(L, I)$, and the second condition is trivial. Second, we need to show that $s - \{x\}$ is not in L_{k+p}^x . Since $s \in C_{k+p}(L, I)$, s is not in L_{k+p} and hence $s - \{x\}$ cannot be in L_{k+p}^x . Finally, we need to show that none of the subsets of $s - \{x\}$ of size greater than $k - 1$ are in $I_{k+1}^x, \dots, I_{k+p-1}^x$. Let $s - \{x\} - \{y_1, \dots, y_m\}$ be such a subset. Since $s \in C_{k+p}(L, I)$, $s - \{y_1, \dots, y_m\}$ is not in I_{k+p-m} , and hence $s - \{x\} - \{y_1, \dots, y_m\}$ cannot be in I_{k+p-m}^x .

We prove the second inequality by induction on k . The base case $k = 1$ is clear. For all $k > 0$, we have

$$\begin{aligned} & gKK_{k+p}^*(L, I) \\ &= \min\{gKK_k^{k+p}(|L|, |I|), \sum_{x \in I} gKK_{k+p-1}^*(L^x, I^x)\} \\ &\leq \min\{KK_k^{k+p}(|L_k|) - |L_{k+p}| - |I_{k+p}|, \sum_{x \in I} KK_{k+p-1}^*(L_k^x) - |L_{k+p}^x| - |I_{k+p}^x|\} \\ &= \min\{KK_k^{k+p}(|L|), \sum_{x \in I} KK_{k+p-1}^*(L^x)\} - |L_{k+p}| - |I_{k+p}| \\ &= KK_{k+p}^*(L_k) - |L_{k+p}| - |I_{k+p}| \end{aligned}$$

where the left hand side of the minimum in the inequality is by Theorem 11 and the right hand side is by induction. \square

Again, we get an upper bound on $\text{maxsize}(L, I)$:

$$g\mu^*(L, I) := k + \min\{p \mid gKK_{k+p}^*(L, I) = 0\} - 1,$$

and on the total number of candidate patterns that can still be generated:

$$gKK_{\text{total}}^*(L, I) := \sum_{p \geq 1} gKK_{k+p}^*(L, I).$$

We then have the following analogous propositions to 8 and 9:

Proposition 15.

$$\text{maxsize}(L, I) \leq g\mu^*(L, I) \leq \mu^*(L).$$

Proposition 16. *The total number of candidate patterns that can be generated from (L, I) is bounded by $gKK_{\text{total}}^*(L, I)$. Moreover,*

$$gKK_{\text{total}}^*(L, I) \leq KK_{\text{total}}^*(L_k).$$

Example 4. Consider the same set of patterns as in the previous example. I.e., L_3 consists of all subsets of size 3 of the set $\{1, 2, 3, 4, 5, 6\}$ and $\{1, 2, 3, 4\}$ and $\{3, 4, 5, 6\}$ are included in I_4 . The KK^* upper bound presented in the previous section would also estimate the number of candidate patterns of sizes 4, 5, and 6 to be at most $\binom{6}{4} = 15$, $\binom{6}{5} = 6$, and $\binom{6}{6} = 1$ respectively. Nevertheless, using the additional information, gKK^* can perfectly predict these numbers to be 13, 2, and 0. Again, μ^* would predict the maximal size of a candidate pattern to be 6, while $g\mu^*$ can already predict this number to be at most 5. Similarly, KK_{total}^* would predict the total number of candidate patterns that can still be generated to be at most 22, while gKK_{total}^* can already deduce this number to be at most 15.

6 Efficient Implementation

For simplicity reasons, we will restrict ourselves to the explanation of how the improved upper bounds can be implemented. The proposed implementation can be easily extended to support the computation of the general upper bounds.

To evaluate our upper bounds we implemented an optimized version of the Apriori algorithm using a trie data structure to store all generated patterns, similar to the one described by Brin et al. [11]. This trie structure makes it cheap and straightforward to implement the computation of all upper bounds. Indeed, a top-level subtrie (rooted at some singleton pattern $\{x\}$) represents exactly the set L^x we defined in Section 4. Every top-level subtrie of this subtrie (rooted at some two-element pattern $\{x, y\}$) then represents $(L^x)^y$, and so on. Hence, we can compute the recursive bounds while traversing the trie, after the frequencies of all candidate patterns are counted, and we have to traverse the trie once more to remove all candidate patterns that turned out to be infrequent. This can be done as follows.

Remember, at that point, we have the current set of frequent patterns of size k stored in the trie. For every node at depth d smaller than k , we compute the $k - d$ -canonical representation of the number of descendants this node has

at depth k , which can be used to compute μ_{k-d} (cf. Proposition 4), KK_{k-d}^ℓ for any $\ell \leq \mu_{k-d}$ (cf. Theorem 2) and hence also KK_{k-d}^{total} (cf. Proposition 6). For every node at depth $k-1$, its KK^* and μ^* values are equal to its KK and μ values respectively. Then compute for every $p > 0$, the sum of the $KK_{k-d+p-1}^*$ values of all its children, and let KK_{k-d+p}^* be the smallest of this sum and KK_{k-d}^{k-d+p} until this minimum becomes zero, which also gives us the value of μ^* . Finally, we can compute KK_{total}^* for this node. If this is done for every node, traversed in a depth-first manner, then finally the root node will contain the upper bounds on the number of candidate patterns that can still be generated, and on the maximum size of any such pattern. The soundness and completeness of this method follows directly from the theorems and propositions of the previous sections.

We should also point out that, since the numbers involved can become exponentially large (in the number of items), an implementation should take care to use arbitrary-length integers such as provided by standard mathematical packages. Since the length of an integer is only logarithmic in its value, the lengths of the numbers involved will remain polynomially bounded.

7 Experimental Evaluation

All experiments were performed on a 400MHz Sun Ultra Sparc with 512 MB main memory, running Sun Solaris 8. The algorithm was implemented in C++ and uses the GNU MP library for arbitrary-length integers [15].

Data sets We have experimented using three real data sets, of which two are publicly available, and one synthetic data set generated by the program provided by the Quest research group at IBM Almaden [5]. The mushroom data set contains characteristics of various species of mushrooms, and was originally obtained from the UCI repository of machine learning databases [9]. The BMS-WebView-1 data set contains several months worth of clickstream data from an e-commerce web site, and is made publicly available by Blue Martini Software [19]. The basket data set contains transactions from a Belgian retail store, but can unfortunately not be made publicly available. Table 1 shows the number of items and the number of transactions in each data set. The table additionally shows the minimal support threshold we used in our experiments for each data set, together with the resulting number of iterations and the time (in seconds) which the Apriori algorithm needed to find all frequent patterns.

The results from the experiment with the real data sets were not immediately as good as the results from the synthetic data set. The reason for this,

Data set	#Items	#Transactions	MinSup	#It's	Time
T40I10D100K	1 000	100 000	700	18	1 700s
mushroom	120	8 124	813	16	663s
BMS-Webview-1	498	59 602	36	15	86s
basket	13 103	41 373	5	11	43s

Table 1: Database Characteristics

however, turned out to be the bad ordering of the items, as explained next.

Reordering From the form of L^x , it can be seen that the order of the items can affect the recursive upper bounds. By computing the upper bound only for a subset of all frequent patterns (namely L^x), we win by incorporating the structure of the current collection of frequent patterns, but we also lose some information. Indeed, whenever we recursively restrict ourselves to a subtrie L^x , then for every candidate pattern s with $x = \min s$, we lose the information about exactly one subpattern in L , namely $s - x$.

We therefore would like to make it likely that many of these excluded patterns are frequent. A good heuristic, which has already been used for several other optimizations in frequent pattern mining [8, 11, 2], is to force the most frequent items to appear in the most candidate patterns, by reordering the single item patterns in increasing order of frequency.

After reordering the items in the real life data set, using this heuristic, the results became very analogous with the results using the synthetic datasets.

Efficiency The cost for the computation of the upper bounds is negligible compared to the cost of the complete algorithm. Indeed, the time T needed to calculate the upper bounds is largely dictated by the number n of currently known frequent sets. We have shown experimentally that T scales linearly with n . Moreover, the constant factor in our implementation is very small (around 0.00001). We ran several experiments using the different data sets and varying minimal support thresholds. After every pass of the algorithm, we registered the number of known frequent sets and the time spent to compute all upper bounds, resulting in 145 different data points. Figure 1 shows these results.

Upper bounds

- Figure 2 shows, after each level k , the computed upper bound KK and improved upper bound KK^* for the number of candidate patterns of size $k + 1$, as well as the actual number $|C_{k+1}|$ it turned out to be.





Figure 1: Time needed to compute upper bounds is linear in the number of nodes.

We omitted the upper bound for $k + 1 = 2$, since the upper bound on the number of candidate patterns of size 2 is simply $\binom{|L|}{2}$, with $|L|$ the number of frequent items.


- Figure 3 shows the upper bounds on the total number of candidate patterns that could still be generated, compared to the actual number of candidate patterns, $|C_{\text{total}}|$, that were effectively generated. Again, we omitted the upper bound for $k = 1$, since this number is simply $2^{|L|} - |L| - 1$, with $|L|$ the number of frequent items.
- Figure 4 shows the computed upper bounds μ and μ^* on the maximal size of a candidate pattern. Also here we omitted the result for $k = 1$, since this number is exactly the number of frequent items.

The results are pleasantly surprising:

- Note that the improvement of KK^* over KK , and of μ^* over μ , anticipated by our theoretical discussion, is indeed dramatic.
- Comparing the computed upper bounds with the actual numbers, we observe the high accuracy of the estimations given by KK^* . Indeed, the estimations of KK_{k+1}^* match almost exactly the actual number of candidate patterns that has been generated at level $k + 1$. Also note that the number of candidate patterns in T40I10D100K is decreasing




(a) basket




(b) BMS-Webview-1

Figure 2: Actual and estimated number of candidate patterns.




(c) T40I10D100K




(d) mushroom

Figure 2: Actual and estimated number of candidate patterns.




(a) basket




(b) BMS-Webview-1

Figure 3: Actual and estimated total number of future candidate patterns.




(c) T40I10D100K




(d) mushroom

Figure 3: Actual and estimated total number of future candidate patterns.




(a) basket




(b) BMS-Webview-1

Figure 4: Estimated size of the largest possible candidate pattern.



(c) T40I10D100K



(d) mushroom


Figure 4: Estimated size of the largest possible candidate pattern.

in the first four iterations and then increases again. This perfectly illustrates that the heuristic used for AprioriHybrid, as explained in the related work section, would not work on this data set. Indeed, any algorithm that exploits the fact that the current number of candidate patterns is small enough and there were fewer candidate patterns in the current iteration than in the previous iteration, would falsely interpret these observations, since the number of candidate patterns in the next iterations increases again. The presented upper bounds perfectly predict this increase.


- The upper bounds on the total number of candidate patterns are still very large when estimated in the first few passes, which is not surprising because at these initial stages, there is not much information yet. For the mushroom and the artificial data sets, the upper bound is almost exact when the frequent patterns of size 3 are known. For the basket data set, this result is obtained when the frequent patterns of size 4 are known and size 6 for the BMS-Webview-1 data set.
- We also performed experiments for varying minimal support thresholds. The results obtained from these experiments were entirely similar to those presented above.

Combining iterations As discussed in the Introduction, the proposed upper bound can be used to protect several improvements of the Apriori algorithm from generating too many candidate patterns. One such improvement tries to combine as many iterations as possible in the end, when only few candidate patterns can still be generated. We have implemented this technique within our implementation of the Apriori algorithm.

We performed several experiments on each data set and limited the number of candidate patterns that is allowed to be generated. If the upper bound on the total number of candidate patterns is below this limit, the algorithm generates and counts all possible candidate patterns within the next iteration. Figure 5 shows the results. The x -axis shows the total number of iterations in which the algorithm completed, and the y -axis shows the total time the algorithm needed to complete. As can be seen, for all datasets, the algorithm can already combine all remaining iterations into one very early in the algorithm. For example, the BMS-Webview-1 dataset, which normally performs 15 iteration, could be reduced to six iterations to give an optimal performance. If the algorithm already generated all remaining candidate patterns in the fifth iteration, the number of candidate patterns that turned out to be infrequent was too large, such that the gain of reducing iterations has




(a) basket




(b) BMS-Webview-1

Figure 5: Combining iterations.



(c) T40I10D100K



(d) mushroom

Figure 5: Combining iterations.

been consumed by the time needed to count all these candidate patterns. Nevertheless, it is still more effective than not combining any passes at all. If we allowed the generation of all candidate patterns to occur in even earlier iterations, although the upper bound predicted a to large number of candidate patterns, this number became indeed to large keep in main memory.

8 Conclusion

Motivated by several heuristics to reduce the number of database scans in the context of frequent pattern mining, we provide a hard and tight combinatorial upper bound on the number of candidate patterns and on the size of the largest possible candidate pattern, given a set of frequent patterns. Our findings are not restricted to a single algorithm, but can be applied to any frequent pattern mining algorithm which is based on the levelwise generation of candidate patterns. Using the standard Apriori algorithm, on which most frequent pattern mining algorithms are based, our experiments showed that these upper bounds can be used to considerably reduce the number of database scans without taking the risk of getting a combinatorial explosion of the number of candidate patterns.

9 Acknowledgement

We wish to thank Blue Martini Software for contributing the KDD Cup 2000 data, the machine learning repository librarians Catherine Blake and Chris Mertz for providing access to the mushroom data, and Tom Brijs for providing the Belgian retail market basket data.

References

- [1] R.C. Agarwal, C.C. Aggarwal, and V.V.V. Prasad. Depth first generation of long patterns. In R. Ramakrishnan, S. Stolfo, R.J. Bayardo, Jr., and I. Parsa, editors, *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 108–118. ACM Press, 2000.
- [2] R.C. Agarwal, C.C. Aggarwal, and V.V.V. Prasad. A tree projection algorithm for generation of frequent itemsets. *Journal of Parallel and Distributed Computing*, 61(3):350–371, March 2001.

- [3] R. Agrawal, T. Imielinski, and A.N. Swami. Mining association rules between sets of items in large databases. In P. Buneman and S. Jajodia, editors, *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, volume 22(2) of *SIGMOD Record*, pages 207–216. ACM Press, 1993.
- [4] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A.I. Verkamo. Fast discovery of association rules. In U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 307–328. MIT Press, 1996.
- [5] R. Agrawal and R. Srikant. *Quest Synthetic Data Generator*. IBM Almaden Research Center, San Jose, California, <http://www.almaden.ibm.com/cs/quest/syndata.html>.
- [6] R. Agrawal and R. Srikant. Fast algorithms for mining association rules. In J.B. Bocca, M. Jarke, and C. Zaniolo, editors, *Proceedings 20th International Conference on Very Large Data Bases*, pages 487–499. Morgan Kaufmann, 1994.
- [7] R. Agrawal and R. Srikant. Fast algorithms for mining association rules. IBM Research Report RJ9839, IBM Almaden Research Center, San Jose, California, June 1994.
- [8] R.J. Bayardo, Jr. Efficiently mining long patterns from databases. In Haas and Tiwary [16], pages 85–93.
- [9] C.L. Blake and C.J. Merz. *UCI Repository of machine learning databases*. University of California, Irvine, Dept. of Information and Computer Sciences, <http://www.ics.uci.edu/~mlearn/MLRepository.html>, 1998.
- [10] B. Bollobás. *Combinatorics*. Cambridge University Press, 1986.
- [11] S. Brin, R. Motwani, J.D. Ullman, and S. Tsur. Dynamic itemset counting and implication rules for market basket data. In *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, volume 26(2) of *SIGMOD Record*, pages 255–264. ACM Press, 1997.
- [12] P. Frankl. A new short proof for the Kruskal–Katona theorem. *Discrete Mathematics*, 48:327–329, 1984.
- [13] F. Geerts, B. Goethals, and J. Van den Bussche. A tight upper bound on the number of candidate patterns. In N. Cercone, T.Y. Lin, and

- X. Wu, editors, *Proceedings of the 2001 IEEE International Conference on Data Mining*, pages 155–162. IEEE Computer Society, 2001.
- [14] B. Goethals and J. Van den Bussche. On supporting interactive association rule mining. In Y. Kambayashi, M.K. Mohania, and A.M. Tjoa, editors, *Proceedings of the Second International Conference on Data Warehousing and Knowledge Discovery*, volume 1874 of *Lecture Notes in Computer Science*, pages 307–316. Springer, 2000.
- [15] T. Granlund and K. Ryde. *GNUmp, Library for arithmetic on arbitrary precision numbers*. <http://www.gnu.org/directory/gnump.html>.
- [16] L.M. Haas and A. Tiwary, editors. *Proceedings of the 1998 ACM SIGMOD International Conference on Management of Data*, volume 27(2) of *SIGMOD Record*. ACM Press, 1998.
- [17] 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*, 2003. To appear.
- [18] G.O.H. Katona. A theorem of finite sets. In *Theory Of Graphs*, pages 187–207. Akadémia Kiadó, 1968.
- [19] R. Kohavi, C. Brodley, B. Frasca, L. Mason, and Z. Zheng. KDD-Cup 2000 organizers’ report: Peeling the onion. *SIGKDD Explorations*, 2(2):86–98, 2000. <http://www.ecn.purdue.edu/KDDCUP>.
- [20] J.B. Kruskal. The number of simplices in a complex. In *Mathematical Optimization Techniques*, pages 251–278. Univ. of California Press, 1963.
- [21] L.V.S. Lakshmanan, R.T. Ng, J. Han, and A. Pang. Optimization of constrained frequent set queries with 2-variable constraints. In A. Delis, C. Faloutsos, and S. Ghandeharizadeh, editors, *Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data*, volume 28(2) of *SIGMOD Record*, pages 157–168. ACM Press, 1999.
- [22] D. Lin and Z.M. Kedem. Pincer-search: A new algorithm for discovering the maximum frequent set. In H.-J. Schek, F. Saltor, I. Ramos, and G. Alonso, editors, *EDBT*, volume 1377 of *Lecture Notes in Computer Science*, pages 105–119. Springer, 1998.
- [23] J. Liu, Y. Pan, K. Wang, and J. Han. Mining frequent item sets by opportunistic projection. In D. Hand, D. Keim, and R.T. Ng, editors, *Proceedings of the Eight ACM SIGKDD International Conference*

- on Knowledge Discovery and Data Mining*, pages 229–238. ACM Press, 2002.
- [24] R.T. Ng, L.V.S. Lakshmanan, J. Han, and A. Pang. Exploratory mining and pruning optimizations of constrained association rules. In Haas and Tiwary [16], pages 13–24.
 - [25] S. Orlando, P. Palmerini, R. Perego, and F. Silvestri. Adaptive and resource-aware mining of frequent sets. In V. Kumar, S. Tsumoto, P.S. Yu, and N.Zhong, editors, *Proceedings of the 2002 IEEE International Conference on Data Mining*. IEEE Computer Society, 2002. To appear.
 - [26] J.S. Park, M.-S. Chen, and P.S. Yu. An effective hash based algorithm for mining association rules. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data*, volume 24(2) of *SIGMOD Record*, pages 175–186. ACM Press, 1995.
 - [27] A. Savasere, E. Omiecinski, and S. Navathe. An efficient algorithm for mining association rules in large databases. In U. Dayal, P.M.D. Gray, and S. Nishio, editors, *Proceedings 21th International Conference on Very Large Data Bases*, pages 432–444. Morgan Kaufmann, 1995.
 - [28] R. Srikant, Q. Vu, and R. Agrawal. Mining association rules with item constraints. In D. Heckerman, H. Mannila, and D. Pregibon, editors, *Proceedings of the Third International Conference on Knowledge Discovery and Data Mining*, pages 66–73. AAAI Press, 1997.
 - [29] H. Toivonen. Sampling large databases for association rules. In T.M. Vijayaraman, A.P. Buchmann, C. Mohan, and N.L. Sarda, editors, *Proceedings 22nd International Conference on Very Large Data Bases*, pages 134–145. Morgan Kaufmann, 1996.
 - [30] M.J. Zaki. Scalable algorithms for association mining. *IEEE Transactions on Knowledge and Data Engineering*, 12(3):372–390, May/June 2000.