

# Probably the Best Itemsets

Nikolaj Tatti

ADReM, University of Antwerp, Antwerpen, Belgium  
nikolaj.tatti@ua.ac.be

## ABSTRACT

One of the main current challenges in itemset mining is to discover a small set of high-quality itemsets. In this paper we propose a new and general approach for measuring the quality of itemsets. The method is solidly founded in Bayesian statistics and decreases monotonically, allowing for efficient discovery of all interesting itemsets. The measure is defined by connecting statistical models and collections of itemsets. This allows us to score individual itemsets with the probability of them occurring in random models built on the data.

As a concrete example of this framework we use exponential models. This class of models possesses many desirable properties. Most importantly, Occam's razor in Bayesian model selection provides a defence for the pattern explosion. As general exponential models are infeasible in practice, we use decomposable models; a large sub-class for which the measure is solvable. For the actual computation of the score we sample models from the posterior distribution using an MCMC approach.

Experimentation on our method demonstrates the measure works in practice and results in interpretable and insightful itemsets for both synthetic and real-world data.

## Categories and Subject Descriptors

H.2.8 [Database management]: Database applications—Data mining; G.3 [Probability and Statistics]: Markov processes

## General Terms

Algorithms, Theory

## Keywords

Itemset mining, exponential models, decomposable models, junction trees, Bayesian model selection, MCMC

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

KDD'10, July 25–28, 2010, Washington, DC, USA.

Copyright 2010 ACM 978-1-4503-0055-1/10/07 ...\$10.00.

## 1. INTRODUCTION

Discovering frequent itemsets is one of the most active fields in data mining. As a measure of quality, frequency possesses a lot of positive properties: it is easy to interpret and as it decreases monotonically there exist efficient algorithms for discovering large collections of frequent itemsets [2]. However, frequency also has serious drawbacks. A frequent itemset may be uninteresting if its elevated frequency is caused by frequent singletons. On the other hand, some non-frequent itemsets could be interesting. Another drawback is the problem of pattern explosion when mining with a low threshold.

Many different quality measures have been suggested to overcome the mentioned problems (see Section 5 for a more detailed discussion). Usually these measures compare the observed frequency to some expected value derived, for example, from the independence model. Using such measures we may obtain better results. However, these approaches still suffer from pattern explosion. To point out the problem, assume that two items, say  $a$  and  $b$  are correlated, and hence are considered significant. Then any itemset containing  $a$  and  $b$  will be also considered significant.

*EXAMPLE 1.* Assume a dataset with  $K$  items  $a_1, \dots, a_K$  such that  $a_1$  and  $a_2$  always yield identical value and the rest of the items are independently distributed. Assume that we apply some statistical method to evaluate the significance of itemsets by using the independence model as the ground truth. If an itemset  $X$  contains  $a_1 a_2$ , its frequency will be higher than then the estimate of the independence model. Hence, given enough data, the  $P$ -value of the statistical test will go to 0, and we will conclude that the itemset  $X$  is interesting. Consequently we will find  $2^{K-2}$  interesting itemsets.

In this work we approach the problem of defining quality measure from a novel point of view. We construct a connection between itemsets and statistical models and use this connection to define a new quality measure for itemsets. To motivate this approach further, let us consider the following example.

*EXAMPLE 2.* Consider a binary dataset with 5 items, say  $a_1, \dots, a_5$ , generated from the independence model. We argue that if we know that the data comes from the independence model, then the only interesting itemsets are the singletons. The reasoning behind this claim is that the frequencies of singletons correspond exactly to the column margins, the parameters of the independence model. Once we know the singleton frequencies, there is nothing left in the data that would be statistically interesting.

Let us consider a more complicated example. Assume that data is generated from a Chow-Liu tree model [4], say

$$p(A) = p(a_1)p(a_2 | a_1)p(a_3 | a_1)p(a_4 | a_2)p(a_5 | a_4).$$

Again, if we know that data is generated from this model, then we argue that the interesting itemsets are  $a_1$ ,  $a_2$ ,  $a_3$ ,  $a_4$ ,  $a_5$ ,  $a_1a_2$ ,  $a_1a_3$ ,  $a_2a_4$ , and  $a_4a_5$ . The reasoning here is the same as with the independence model. If we know the frequencies of these itemsets we can derive the parameters of the distribution. For example,  $p(a_2 = 1 | a_1 = 1) = fr(a_1a_2) / fr(a_1)$ .

Let us now demonstrate that this approach will produce much smaller and more meaningful output than the method given in Example 1.

**EXAMPLE 3.** Consider the data given in Example 1. To fully describe the data we only need to know the frequencies of the singletons and the fact that  $a_1$  and  $a_2$  are identical. This information can be expressed by outputting the frequencies of singleton itemsets and the frequency of itemset  $a_1a_2$ . This will give us  $K + 1$  interesting patterns in total.

Our approach is to extend the idea pitched in the preceding example to a general itemset mining framework. In the example we knew which model generated the data, in practice, we typically do not. To solve this we will use the Bayesian approach, and instead of considering just one specific model, we will consider a large collection of models, namely exponential models. A virtue of these models is that we can naturally connect each model to certain itemsets. A model  $M$  has a posterior probability  $P(M | D)$ , that is, how probable is the model given the data. The score of a single itemset then is just the probability of it being a parameter of a random model given the data. This setup fits perfectly the given example. If we have strong evidence that data is coming from the independence model, say  $M$ , then the posterior probability  $P(M | D)$  will be close to 1, and the posterior probability of any other model will be close to 0. Since the independence model is connected to the singletons, the score for singletons will be 1 and the score for any other itemset will be close to 0.

Interestingly, using statistical models for defining significant itemsets provides an approach to the problem of the pattern set explosion (see Section 3.2 for more technical details). Bayesian model selection has an in-built Occam's razor, favoring simple models over complex ones. Our connection between models and itemsets is such that simple models will correspond to small collections of itemsets. In result, only a small collection of itemsets will be considered interesting, unless the data provides sufficient amount of evidence.

Our contribution in the paper is two-fold. First, we introduce a general framework of using statistical models for scoring itemsets in Section 2. Secondly, we provide an example of this framework in Section 3 by using exponential models and provide solid theoretical evidence that our choices are well-founded. We provide the sampling algorithm in Section 4. We discuss related work in Section 5 and present our experiments in Section 6. Finally, we conclude our work with Section 7. The proofs are given in Appendix [19]. The implementation is provided for research purposes<sup>1</sup>.

<sup>1</sup><http://adrem.ua.ac.be/implementations>

## 2. SIGNIFICANCE OF ITEMSETS BY STATISTICAL MODELS

As we discussed in the introduction, our goal is to define a quality measure for itemsets using statistical models. In this section we provide a general framework for such a score. We will define the actual models in the next section.

We begin with some preliminary definitions and notations. In our setup a *binary dataset* is a collection of  $N$  transactions, binary vectors of length  $K$ . We assume that these vectors are independently generated from some unknown distribution. Such a dataset can be easily represented by a binary matrix of size  $N \times K$ . By an attribute  $a_i$  we mean a Bernoulli random variable corresponding to the  $i$ th column of the data. We denote the set of *attributes* by  $A = \{a_1, \dots, a_K\}$ .

An itemset  $X$  is simply a subset of  $A$ . Given an itemset  $X = \{a_{i_1}, \dots, a_{i_L}\}$  and a transaction  $t$  we denote by  $t_X = (t_{i_1}, \dots, t_{i_L})$  the projection of  $t$  into  $X$ . We say that  $t$  covers  $X$  if all elements in  $t_X$  are equal to 1.

We say that a collection of itemsets  $\mathcal{F}$  is *downward closed* if for each member  $X \in \mathcal{F}$  any sub-itemset is also included. This property plays a crucial point in mining frequent patterns since it allows effective candidate pruning in level-wise approach and branch pruning in a DFS approach.

Our next step is to define the correspondence between statistical models and families of itemsets. Assume that we have a set  $\mathcal{M}$  of statistical models for the data. We will discuss in later sections what specific models we are interested in, but for the moment, we will keep our discussion on a high level. Each model  $M \in \mathcal{M}$  has a posterior probability  $P(M | D)$ , that is, how probable the model  $M$  is given data  $D$ . To link the models to families of itemsets we assume that we have a function *fam* that identifies a model  $M$  with a *downward closed* family of itemsets. As we will see later on, there is one particular natural choice for such a function.

Now that we have our models that are connected to certain families of itemsets, we are ready to define a score for individual itemsets. The score for an itemset  $X$  is the posterior probability of  $X$  being a member in a family or itemsets,

$$sc(X) = \sum_{\substack{X \in fam(M), \\ M \in \mathcal{M}}} P(M | D). \quad (1)$$

The motivation for such score is as follows. If we are sure that some particular model  $M$  is the correct model for  $D$ , then the posterior probability for that model will be close to 1 and the posterior probabilities for other models will be close to 0. Consequently, the score for an itemset  $X$  will be close to 1 if  $X \in fam(M)$ , and 0 otherwise.

Naturally, the pivotal choice of this score lies in the mapping *fam*. Such a mapping needs to be statistically well-founded, and especially the size of an itemset family should be reflected in the complexity of the corresponding model. We will see in the following section that a particular choice for the model family and mapping *fam* has these properties, and leads to certain important qualities.

**PROPOSITION 4.** *The score decreases monotonically, that is,  $X \subseteq Y$  implies  $sc(X) \geq sc(Y)$ .*

**PROOF.** We are allowing *fam* only to map on downward closed families. Hence the inequality

$$\sum_{\substack{X \in fam(M), \\ M \in \mathcal{M}}} P(M | D) \geq \sum_{\substack{Y \in fam(M), \\ M \in \mathcal{M}}} P(M | D)$$

holds, and consequently we have  $sc(X) \geq sc(Y)$ . This completes the proof.  $\square$

### 3. EXPONENTIAL MODELS

In this section we will make our framework more concrete by providing a specific set  $\mathcal{M}$  of statistical models and the function  $fam$  identifying the models with families of itemsets. We will first give the definition of the models and the mapping. After this, we point out the main properties of our model and justify our choices. However, as, it turns out that computing the score for these models is infeasible, so instead, we solve this problem by considering decomposable models.

#### 3.1 Definition of the models

Models of exponential form have been studied exhaustively in statistics, and have been shown to have good theoretical and practical properties. In our case, using exponential models provide a natural way of describing the dependencies between the variables. In fact, the exponential model class contains many natural models such as, the independence model, the Chow-Liu tree model, and the discrete Gaussian model. Finally, such models have been used successfully for predicting itemset frequencies [14] and ranking itemsets [18].

In order to define our model, let  $\mathcal{F}$  be a downward closed family of itemsets containing all singletons. For an itemset  $X \in \mathcal{F}$  we define an indicator function  $S_X(t)$  mapping a transaction  $t$  into binary value. If the transaction  $t$  covers  $X$ , then  $S_X(t) = 1$ , and 0 otherwise. We define an exponential model  $M$  associated with  $\mathcal{F}$  to be the collection of distributions having the exponential form

$$P(A = t \mid M, r) = \exp\left(\sum_{X \in \mathcal{F}} r_X S_X(t)\right),$$

where  $r_X$  is a parameter, a real value, for an itemset  $X$ . Model  $M$  also contains all the distributions that can be obtained as a limit of the distribution having the exponential form. This technicality is needed to handle distributions with zero probabilities. Since the indicator function  $S_\emptyset(t)$  is equal to 1 for any  $t$ , the parameter  $r_\emptyset$  acts like a normalization constant. The rest of the parameters form a parameter vector  $r$  of length  $|\mathcal{F}| - 1$ . Naturally, we set  $\mathcal{F} = fam(M)$ .

**EXAMPLE 5.** Assume that  $\mathcal{F}$  consists only of singleton itemsets. Then the corresponding model has the form

$$\exp\left(r_\emptyset + \sum_{i=1}^K r_{a_i} S_{a_i}(t)\right) = \exp(r_\emptyset) \prod_{i=1}^K \exp(r_{a_i} S_{a_i}(t)). \quad (2)$$

Since  $S_{a_i}(t)$  depends only on  $t_i$ , the model is actually the independence model. The other extreme is when  $\mathcal{F}$  consists of all itemsets. Then we can show that the corresponding model contains all possible distributions, that is, the model is in fact the parameter-free model.

As an intermediate example, the tree model in Example 2 is also an exponential model with a corresponding family  $\{a_1, \dots, a_5\} \cup \{a_1a_2, a_2a_3, a_2a_4, a_4a_5\}$ .

The intuition behind the model is that when an itemset  $X$  is an element of  $\mathcal{F}$ , then the dependencies between the items in  $X$  are considered important in the corresponding model.

For example, if  $\mathcal{F}$  consists only of singletons, then there are no important correlations, hence the corresponding model should be the independence model. On the other hand, in the tree model given in Example 2 the important correlations are the parent-child item pairs, namely,  $a_1a_2, a_2a_3, a_2a_4, a_4a_5$ . These are exactly, the itemsets (along with the singletons) that correspond to the model.

Our choice for the models is particularly good since the complexity of models reflects the size of itemset family. Since Bayesian approach has an in-built tendency to punish complex families (we will see this in Section 3.2), we are punishing large families of itemsets. If the data states that simple models are sufficient, then the probability of complex models will be low, and consequently the score for large itemsets will also be low. In other words, we casted the problem of pattern set explosion into a model overfitting problem and used Occam's razor to punish the complex models!

#### 3.2 Computing the Model

Now that we have defined our model  $M$ , our next step is to compute the posterior probability  $P(M \mid D)$ . That is, the probability of  $M$  given the data set  $D$ . We select the model prior  $P(M)$  to be uniform. Recall that in Eq. 2 a model  $M$  has a set of parameters, that is, to pinpoint a single distribution in  $M$  we need a set of parameters  $r$ . Following the Bayesian approach to compute  $P(M \mid D)$  we need to marginalize out the nuisance parameters  $r$ ,

$$P(M \mid D) = \int_r P(M, r \mid D) \propto \int_r P(r \mid M) \prod_{t \in D} P(t \mid M, r).$$

In the general case, this integral is too complex to solve analytically so we employ the popular BIC estimate [15],

$$P(M \mid D) \approx C \times P(D \mid M, r^*) \times \exp\left(-\log |D| \frac{|\mathcal{F}| - 1}{2}\right), \quad (3)$$

where  $C$  is a constant and  $r^*$  is the maximum likelihood estimate of the model parameters. This estimate is correct when  $|D|$  approaches infinity [15]. So instead of computing a complex integral our challenge is to discover the maximum likelihood estimate  $r^*$  and compute the likelihood of the data. Unfortunately, using such model is an **NP**-hard problem (see, for example, [17]). We will remedy this problem in Section 3.4 by considering a large subclass of exponential models for which the maximum likelihood can be easily computed.

#### 3.3 Justifications for Exponential Model

In this section we will provide strong theoretical justification for our choices and show that our score fulfills the goals we set in the introduction.

We saw in Example 2 that if the data comes from the independence model, then we only need the frequencies of the singleton itemsets to completely explain the underlying model. The next theorem shows that this holds in general case.

**THEOREM 6.** Assume that data  $D$  is generated from a distribution  $p$  that comes from an exponential model  $M$ . Let  $\mathcal{F} = fam(M)$  be the family of itemsets. We can derive the maximum likelihood estimate from the frequencies of  $\mathcal{F}$ . Moreover, as the number of transactions goes to infinity, we can derive the true distribution  $p$  from the frequencies of  $\mathcal{F}$ .

The preceding theorem showed that  $fam(M)$  is sufficient family of itemsets in order to derive the correct true distribution. The next theorem shows that we favor small families: if the data can be explained with a simpler model, that is, using less itemsets, then the simpler model will be chosen and, consequently, redundant itemsets will have a low score.

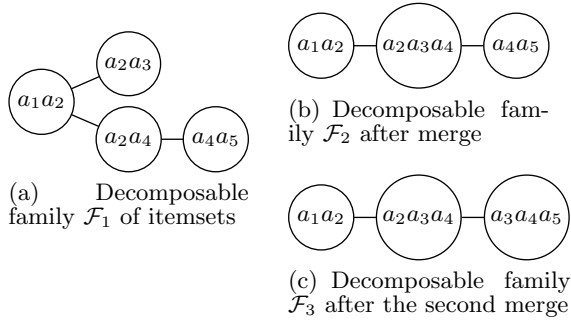
**THEOREM 7.** *Assume that data  $D$  is generated from a distribution  $p$  that comes from a model  $M$ . Assume also that if any other model, say  $M'$ , contains this distribution, then  $|fam(M')| > |fam(M)|$ . Then the following holds: as the number of data points in  $D$  goes into infinity,  $sc(X) = 1$  if  $X \in fam(M)$ , otherwise  $sc(X) = 0$ .*

### 3.4 Decomposable Models

We saw in Section 3.2 that in practice we cannot compute the score for general exponential models. In this section we study a subclass of exponential models, for which we can easily compute the needed score. Roughly speaking, a decomposable model is an exponential model where the corresponding maximal itemsets can be arranged to a specific tree, called junction tree. By considering only decomposable models we obviously will lose some models, for example, the discrete Gaussian model, that is, a model corresponding to all itemsets of size 1 and 2 is not decomposable. On the other hand, many interesting and practically relevant models are decomposable, for example Chow-Liu trees. Finally, these models are closely related to Bayesian networks and Markov Random Fields (see [5] for more details).

To define a decomposable model, let  $\mathcal{F}$  be a downward closed family of itemsets. We write  $\mathcal{G} = \max(\mathcal{F})$  to be the set of maximal itemsets from  $\mathcal{F}$ . Assume that we can build a tree  $\mathcal{T}$  using itemsets from  $\mathcal{G}$  as nodes with the following property: If  $X, Y \in \mathcal{G}$  have a common item, say  $a$ , then  $X$  and  $Y$  are connected in  $\mathcal{T}$  (by a unique path) and every itemset along that path contains  $a$ . If this property holds for  $\mathcal{T}$ , then  $\mathcal{T}$  is called *junction tree* and  $\mathcal{F}$  is *decomposable*. We will use  $E(\mathcal{T})$  to denote the edges of the tree.

Not all families have junction trees and some families may have multiple junction trees.



**Figure 1:** Figure 1(a) shows that the itemset family given in Example 2 is decomposable. Figure 1(b) shows the junction tree for the family after  $\text{MERGE}(\{a_2\}, a_3, a_4)$  and Figure 1(c) shows the junction tree after  $\text{MERGE}(\{a_4\}, a_3, a_5)$ .

**EXAMPLE 8.** *Let  $\mathcal{F}_1$  be the family of itemsets connected to the Chow-Liu model given in Example 2. The maximal itemsets  $\mathcal{F}_1$  are  $\{a_1a_2, a_2a_3, a_2a_4, a_4a_5\}$ . Figure 1(a) shows*

*a junction tree for this family, making the family decomposable. On the other hand, family  $\{a_1a_2, a_1a_3, a_2a_3\}$  is not decomposable since there is no junction tree for this family.*

The most important property of decomposable families is that we can compute the maximum likelihood efficiently. We first define the entropy of an itemset  $X$ , denoted by  $H(X)$ , as

$$H(X) = - \sum_{t \in \{0,1\}^{|X|}} q_D(X=t) \log q_D(X=t),$$

where  $q_D$  is the empirical distribution of the data.

**THEOREM 9.** *Let  $\mathcal{F}$  be a decomposable family and let  $\mathcal{T}$  be its junction tree. The maximum log-likelihood is equal to*

$$-\frac{\log P(D | M, r^*)}{|D|} = \sum_{X \in \max(\mathcal{F})} H(X) - \sum_{\substack{(X,Y) \\ \in E(\mathcal{T})}} H(X \cap Y).$$

**EXAMPLE 10.** *Assume that our model space  $\mathcal{M}$  consists only of two models, namely the tree model  $M_1$  given in Example 2 and the independence model, which we denote  $M_2$ . Assume also that we have a dataset with 9 transactions,*

$$\begin{aligned} t_1 &= (1, 0, 0, 0, 0), t_2 = (1, 1, 0, 0, 1), t_3 = (0, 0, 0, 0, 0), \\ t_4 &= (1, 1, 1, 1, 1), t_5 = (1, 1, 0, 1, 1), t_6 = (0, 1, 1, 0, 0), \\ t_7 &= (0, 0, 1, 0, 0), t_8 = (0, 0, 1, 0, 1), t_9 = (0, 1, 1, 1, 1). \end{aligned}$$

*To compute the probabilities  $P(M_1 | D)$  and  $P(M_2 | D)$ , we need to know the entropies of certain itemsets*

$$\begin{aligned} H(a_1) &= H(a_2) = H(a_3) = H(a_5) = 0.68, H(a_4) = 0.64, \\ H(a_1a_2) &= 1.31, H(a_2a_3) = 1.37, \\ H(a_2a_4) &= H(a_4a_5) = 1.06. \end{aligned}$$

*The log-likelihood of the independence model is equal to*

$$\log P(D | M_2, r^*) = -9(4 \times 0.68 + 0.64).$$

*We use the junction tree given in Figure 1(a) and Theorem 9 to compute the log-likelihood of  $M_1$ ,*

$$\log P(D | M_1, r^*) = -9(1.31 + 1.37 + 2 \times 1.06 - 2 \times 0.68 - 0.64).$$

*Note that  $|fam(M_2)| = 6$  and  $|fam(M_1)| = 10$ . Thus Eq. 3 implies that*

$$\begin{aligned} P(M_1 | D) &\propto P(D | M_1, r^*) \exp(-9/2 \log 9) = 6.27 \times 10^{-14}, \\ P(M_2 | D) &\propto P(D | M_2, r^*) \exp(-5/2 \log 9) = 2.43 \times 10^{-14}. \end{aligned}$$

*We get the final probabilities by noticing that  $P(M_1 | D) + P(M_2 | D) = 1$  so that we have  $P(M_1 | D) = 0.72$  and  $P(M_2 | D) = 0.28$ . Consequently, the scores for itemsets are equal to  $sc(a_1a_2) = sc(a_2a_3) = sc(a_2a_4) = sc(a_4a_5) = 0.72$ ,  $sc(a_i) = 1$ , for  $i = 1, \dots, 5$ , and  $sc(X) = 0$  otherwise.*

## 4. SAMPLING MODELS

Now that we have means for computing the posterior probability of a single decomposable model, our next step is to compute the score of an itemset namely, the sum in Eq. 1. The problem is that this sum has an exponential number of terms, and hence we cannot solve by enumerating all possible families. We approach this problem from a different point of view. Instead of computing the score for each itemset individually, we will divide our mining method into two steps:

1. Sample random decomposable models from the posterior distribution  $P(M | D)$ .
2. Estimate the true score of an itemset by computing the number of sampled families of itemsets in which the itemset occurs.

## 4.1 Moving from One Model to Another

In order to sample we will use a MCMC approach by modifying the current decomposable family by two possible operations, namely

- **MERGE**: Select two maximal itemsets, say  $X$  and  $Y$ . Let  $S = X \cap Y$ . Since  $X$  and  $Y$  are maximal,  $X - S \neq \emptyset$  and  $Y - S \neq \emptyset$ . Select  $x \in X - S$  and  $y \in Y - S$ . Add a new itemset  $S \cup \{x, y\}$  into the family  $\mathcal{F}$  along with all possible sub-itemsets. We will use notation  $\text{MERGE}(S, x, y)$  to denote this operation.
- **SPLIT**: Select an itemset  $X \in \max(\mathcal{F})$ . Select two items  $x, y \in X$ . Delete  $X$  and all sub-itemsets containing  $x$  and  $y$  simultaneously. We will denote this operation by  $\text{SPLIT}(X, x, y)$ .

Naturally, not all splits and merges are legal, since some operations may result in a family that is not decomposable, or even downward closed.

**EXAMPLE 11.** *The family  $\mathcal{F}_2$  given in Figure 1(b) is obtained from the family  $\mathcal{F}_1$  given in Figure 1(a) by performing  $\text{MERGE}(\{a_2\}, a_3, a_4)$ . Moreover,  $\mathcal{F}_3$  (Figure 1(c)) is obtained from  $\mathcal{F}_2$  by performing  $\text{MERGE}(\{a_4\}, a_3, a_5)$ . Conversely, we can go back by performing  $\text{SPLIT}(a_3a_4a_5, a_3, a_5)$  first and  $\text{SPLIT}(a_2a_3a_4, a_3, a_4)$  second.*

The next theorem tells us which splits are legal.

**THEOREM 12.** *Let  $\mathcal{F}$  be decomposable family and let  $X \in \max(\mathcal{F})$  and let  $x, y \in X$ . Then the resulting family after a split operation  $\text{SPLIT}(X, x, y)$  is decomposable if and only, there are no other maximal itemsets in  $\mathcal{F}$  containing  $x$  and  $y$  simultaneously.*

**EXAMPLE 13.** *All possible split combinations are legal in families  $\mathcal{F}_1$  and  $\mathcal{F}_2$  given in Figure 1(a) and Figure 1(b). However, for  $\mathcal{F}_3$  given in Figure 1(c)  $\text{SPLIT}(a_2a_3a_4, a_3, a_4)$  is illegal since  $a_3a_4a_5$  contains  $a_3$  and  $a_4$ . Similarly, the operation  $\text{SPLIT}(a_3a_4a_5, a_3, a_4)$  is illegal.*

In order to identify legal merges, we will need some additional structures. Let  $\mathcal{F}$  be a downward closed family and let  $\mathcal{G} = \max(\mathcal{F})$  be its maximal itemsets. Let  $S$  be an itemset. We construct a *reduced family*, denoted by  $rf(\mathcal{F}; S)$  with the following procedure. Let us first define

$$\mathcal{X} = \{X - S \mid X \in \mathcal{G}, S \subsetneq X\}.$$

To obtain the reduced family  $rf(\mathcal{F}; S)$  from  $\mathcal{X}$ , assume there are two itemsets  $X, Y \in \mathcal{X}$  such that  $X \cap Y \neq \emptyset$ . We remove these two sets from  $\mathcal{X}$  and replace them with  $X \cup Y$ . This is continued until no such replacements are possible. We ignore any reduced family that contains 0 or 1 itemsets. The reason for this will be seen in Theorem 15, which implies that such families will not induce any legal merges.

**EXAMPLE 14.** *The non-trivial reduced families of the family given in Figure 1(a) are  $rf(\mathcal{F}_1; a_2) = \{a_1, a_3, a_4\}$  and  $rf(\mathcal{F}_1; a_4) = \{a_2, a_5\}$ . Similarly, the reduced families for the family given in Figure 1(b) are  $rf(\mathcal{F}_2; a_2) = \{a_1, a_3a_4\}$ , and  $rf(\mathcal{F}_2; a_4) = \{a_2a_3, a_5\}$ . Finally, the reduced families for the family given in Figure 1(c) are  $rf(\mathcal{F}_3; a_2) = \{a_1, a_3a_4\}$  and  $rf(\mathcal{F}_3; a_3a_4) = \{a_2, a_5\}$ .*

The next theorem tells us when  $\text{MERGE}(S, x, y)$  is legal.

**THEOREM 15.** *Let  $\mathcal{F}$  be decomposable family. A merge operation is legal, that is,  $\mathcal{F}$  is still decomposable after adding  $Z = S \cup \{x, y\}$  if and only if there are sets  $V, W \in rf(\mathcal{F}; S)$ ,  $V \neq W$ , such that  $x \in V$  and  $y \in W$ .*

**EXAMPLE 16.** *Family  $\mathcal{F}_2$  in Figure 1(b) is obtained from the family  $\mathcal{F}_1$  in Figure 1(a) by  $\text{MERGE}(a_2, a_3a_4)$ . This is legal operation since  $rf(\mathcal{F}_1; a_2) = \{a_1, a_3, a_4\}$ . Similarly, merge transforming  $\mathcal{F}_2$  to  $\mathcal{F}_3$  is legal since  $rf(\mathcal{F}_2; a_4) = \{a_2a_3, a_5\}$ . However, this merge would not be legal in  $\mathcal{F}_1$  since we do not have  $a_3$  in  $rf(\mathcal{F}_1; a_4)$ .*

## 4.2 MCMC Sampling Algorithm

Sampling requires a proposal distribution  $Q(M' | M)$ . Let  $M$  be a current model. We denote the number of legal operations, either a split or a merge, by  $d(M)$ . Let  $M'$  be a model obtained by sampling uniformly one of the legal operations and applying it to  $M$ . The probability of reaching  $M'$  from  $M$  with a single step is  $Q(M' | M) = 1/d(M)$ . Similarly, the probability of reaching  $M$  from  $M'$  with a single step is  $Q(M | M') = 1/d(M')$ . Consequently, if we sample  $u$  uniformly from the interval  $[0, 1]$  and accept the step moving from  $M$  into  $M'$  if and only if  $u$  is smaller than

$$\frac{P(M' | D)Q(M | M')}{P(M | D)Q(M' | M)} = \frac{P(M' | D)d(M)}{P(M | D)d(M')}, \quad (4)$$

then the limit distribution of the MCMC will be the posterior distribution  $P(M | D)$  provided that the MCMC chain is ergodic. The next theorem shows that this is the case.

**THEOREM 17.** *Any decomposable model  $M$  can be reached from any other model  $M'$  by a sequence of legal operations.*

Our first step is to compute the ratio of the models given in Eq 4. To do that we will use the BIC estimate given in Eq. 3 and Theorem 9. Let us first define a function

$$\begin{aligned} \text{gain}(X, x, y) = \\ |D|(H(X) - H(X - x) - H(X - y) + H(X - \{x, y\})), \end{aligned}$$

where  $X$  is an itemset and  $x, y \in X$  are items.

**THEOREM 18.** *Let  $M$  be a decomposable model and let  $M' = \text{SPLIT}(X, x, y; M)$  be a model obtained by a legal split. Let  $A$  be the BIC estimate of  $P(M | D)$  and let  $B$  be the BIC estimate of  $P(M' | D)$ . Then*

$$B/A = \exp(\text{gain}(X, x, y) - \log |D|2^{|X|-3}).$$

*Similarly, if  $M' = \text{MERGE}(S, x, y; M)$ , then*

$$B/A = \exp(-\text{gain}(S \cup \{x, y\}, x, y) + \log |D|2^{|S|-1}).$$

To compute the gain we need the entropies for 4 itemsets. Let  $X$  be an itemset. To compute  $H(X)$  we first order the transactions in  $D$  such that the values corresponding to  $X$  are in lexicographical order. This is done with a radix sort  $\text{SORT}(D, X)$  given in Algorithm 1. This sort is done in  $O(|D||X|)$  time. After the data is sorted we can easily compute the entropy with a single data scan: Set  $e = 0$  and  $p = 0$ . If the values of  $X$  of the current transaction is equal to the previous transaction we increase  $p$  by  $1/|D|$ , otherwise we add  $-p \log p$  to  $e$  and set  $p$  to  $1/|D|$ . Once the scan is finished,  $e$  will be equal to  $H(X)$ . The pseudo code for computing the entropy is given in Algorithm 2.

---

**Algorithm 1:**  $\text{SORT}(D, X)$ . Routine for sorting the transactions. Used by ENTROPY as a pre-step for computing the entropy.

---

```

1 if  $X = \emptyset$  or  $D = \emptyset$  then return  $D$ ;
2  $a_i \leftarrow$  first item in  $X$ ;
3  $D_0 \leftarrow \{t \in D \mid t_i = 0\}$ ;  $D_1 \leftarrow \{t \in D \mid t_i = 1\}$ ;
4  $D_0 \leftarrow \text{SORT}(D_0, X - a_i)$ ;  $D_1 \leftarrow \text{SORT}(D_1, X - a_i)$ ;
5 return  $D_0$  concatenated with  $D_1$ .
```

---



---

**Algorithm 2:** ENTROPY( $D, X$ ). Computes the entropy of  $X$  from the dataset  $D$ .

---

```

1 SORT( $D, X$ );
2  $e \leftarrow 0$ ;  $p \leftarrow 0$ ;
3  $u \leftarrow$  first transaction in  $D$ ;
4 foreach  $t \in D$  do
5   if  $u_X \neq t_X$  then
6      $e \leftarrow e - p \log p$ ;
7      $u \leftarrow t$ ;
8      $p \leftarrow 1/|D|$ ;
9   else
10     $p \leftarrow p + 1/|D|$ ;
11  $e \leftarrow e - p \log p$ ;
12 return  $e$ ;
```

---

Our final step is to compute  $d(M)$  and actually sample the operations. To do that we first write  $sd(M)$  for the number of possible SPLIT operations and let  $sd(M, X)$  be the number of possible SPLIT operations using itemset  $X$ . Similarly, we write  $md(M, S)$  for the number of legal merges using  $S$  and also  $md(M)$  for the amount of legal merges in total.

Given a maximal itemset  $X$  we build an occurrence table, which we denote by  $st(X)$ , of size  $|X| \times |X|$ . For  $x, y \in X$ , the entry of the table  $st(X, x, y)$  is the number of maximal itemsets containing  $x$  and  $y$ . If  $st(X, x, y) = 1$ , then Theorem 12 states that  $\text{SPLIT}(X, x, y)$  is legal. Consequently, to sample a split operation we first select a maximal itemset weighted by  $sd(M, X)/sd(M)$ . Once  $X$  is selected we select uniformly one legal pair  $(x, y)$ .

To sample legal merges, recall that  $\text{MERGE}(S, x, y)$  involves selecting two maximal itemsets  $X$  and  $Y$  such that  $S = X \cap Y$ ,  $x \in X - S$ , and  $y \in Y - S$ . Instead of selecting these itemsets, we will directly sample an itemset  $S$  and then select two items  $x$  and  $y$ . This sampling will work only if two legal merges  $\text{MERGE}(S_1, x_1, y_1)$  and  $\text{MERGE}(S_2, x_2, y_2)$  result in two different outcomes whenever  $S_1 \neq S_2$ .

**THEOREM 19.** Let  $S_1$  and  $S_2$  be two different itemsets and let  $x_1, y_1 \notin S_1$ , and  $x_2, y_2 \notin S_2$  be items. Assume that  $\text{MERGE}(S_i, x_i, y_i)$  is a legal merge for  $i = 1, 2$ . Define  $Z_i = S_i \cup \{x_i, y_i\}$  for  $i = 1, 2$ . Then  $Z_1 \neq Z_2$ .

The construction of a reduced family states that, if  $V, W \in rf(\mathcal{F}; S)$ ,  $V \neq W$ , then  $V \cap W = \emptyset$ . It follows from Theorem 15 that

$$md(M, S) = \sum_{\substack{V, W \in rf(\mathcal{F}; S) \\ V \neq W}} |V||W|.$$

To sample a merge we first sample an itemset  $S$  weighted by  $md(M, S)/md(M)$ . Once  $S$  is selected, we sample two different itemsets  $V, W \in rf(\mathcal{F}; S)$  (weighted by  $|V|$  and  $|W|$ ). Finally, we sample  $x \in V$  and  $y \in W$ .

Sampling  $S$  for a merge operation is feasible only if the number of reduced families for which the merge degree is larger than zero is small.

**THEOREM 20.** Let  $K$  be the number of items. There are at most  $K$  maximal itemsets. There are at most  $K - 1$  itemsets for which the degree  $md(M, \cdot) > 0$ .

Pseudo-code for a sampling step is given in Algorithm 3.

---

**Algorithm 3:** MCMC step for sampling decomposable models.

---

```

1  $u \leftarrow$  random integer between 1 and  $d(M)$ ;
2 if  $u \leq sd(M)$  then
3   Sample  $X$  from  $\max(\mathcal{F})$  weighted by  $sd(M, X)$ ;
4   Sample  $x, y \in X$  such that  $X$  is the only maximal
   itemset containing both  $x$  and  $y$ ;
5    $M' \leftarrow \text{SPLIT}(X, x, y; M)$ ;
6    $g \leftarrow \text{gain}(X, x, y) - \log |D|2^{|X|-3}$ ;
7 else
8   Sample  $S$  weighted by  $md(M, S)$ ;
9   Sample  $V \in rf(\mathcal{F}; S)$  weighted by  $|V|$ ;
10  Sample  $W \in rf(\mathcal{F}; S)$ ,  $V \neq W$ , weighted by  $|W|$ ;
11  Sample  $x \in V$  and  $y \in W$ ;
12   $M' \leftarrow \text{MERGE}(S, x, y; M)$ ;
13   $g \leftarrow -\text{gain}(S \cup \{x, y\}, x, y) + \log |D|2^{|S|-1}$ ;
14  $z \leftarrow$  random real number from  $[0, 1]$ ;
15 if  $z \leq \exp(g) \frac{d(M)}{d(M')}$  then return  $M'$ ;
16 else return  $M$ ;
```

---

### 4.3 Speeding Up the Sampling

We have demonstrated what structures we need to compute so that we can sample legal operations. After a sample, we can reconstruct these structures from scratch. In this section we show how to optimize the sampling by constructing the structures incrementally using Algorithms 4–7.

First of all, we store only maximal itemsets of  $\mathcal{F}$ . Theorem 20 states that there can be only  $K$  such sets, hence split and merge operations can be done efficiently.

During a split or a merge, we need to update what split operations are legal after the split. We do this by updating an occurrence table  $st(X)$ . An update takes  $O(|X|^2)$  time. The next theorem shows which maximal itemsets we need to update for legal split operations after a merge.

**THEOREM 21.** *Let  $\mathcal{F}$  be a downward closed family of itemsets and let  $\mathcal{G}$  be the family after performing  $\text{MERGE}(S, x, y)$ . Let  $Y$  be a maximal itemset in  $\max(\mathcal{F}) \cap \max(\mathcal{G})$ . Then legal split operations using  $Y$  remain unchanged during the merge unless  $Y$  is the unique itemset among maximal itemsets in  $\mathcal{F}$  containing either  $S \cup \{x\}$  or  $S \cup \{y\}$ .*

The following theorem tells us how reduced families should be updated after a merge operation. To ease the notation, let us denote by  $\text{link}(\mathcal{F}, S, x)$  the unique itemset (if such exists) in  $\text{rf}(\mathcal{F}; S)$  containing  $x$ .

**THEOREM 22.** *Let  $\mathcal{F}$  be a downward closed family of itemsets and let  $\mathcal{G}$  be the family after performing  $\text{MERGE}(S, x, y)$ . Then the reduced families are updated as follows:*

1. *Itemsets  $\text{link}(\mathcal{F}, S, x)$  and  $\text{link}(\mathcal{F}, S, y)$  in  $\text{rf}(\mathcal{F}; S)$  are merged into one itemset in  $\text{rf}(\mathcal{G}; S)$ .*
2. *Itemset  $\{x\}$  is added into  $\text{rf}(\mathcal{G}; S \cup \{y\})$ . Itemset  $\{y\}$  is added into  $\text{rf}(\mathcal{G}; S \cup \{x\})$ .*
3. *Let  $T \subsetneq S$  and let  $z \in S - T$ . The itemset containing  $z$  in  $\text{rf}(\mathcal{F}; T \cup \{y\})$  is augmented with item  $x$ . Similarly, itemset containing  $z$  in  $\text{rf}(\mathcal{F}; T \cup \{x\})$  is augmented with item  $y$ .*
4. *Otherwise,  $\text{rf}(\mathcal{F}; T) = \text{rf}(\mathcal{G}; T)$  or  $\text{md}(\mathcal{F}; T) = 0$  and  $\text{md}(\mathcal{G}; T) = 0$ .*

Theorems 21 and 22 only covered the updates during merges. Since  $\text{SPLIT}(S \cup \{x, y\}, x, y)$  and  $\text{MERGE}(S, x, y)$  are opposite operations we can derive the needed updates for splits from the preceding theorems.

**COROLLARY 23 (OF THEOREM 21).** *Let  $\mathcal{F}$  be a downward closed family of itemsets and let  $\mathcal{G}$  be the family after performing  $\text{SPLIT}(X, x, y)$ . Let  $Y$  be a maximal itemset in  $\max(\mathcal{F}) \cap \max(\mathcal{G})$ . Then legal split operations using  $Y$  remain unchanged during the merge unless  $Y$  is the unique itemset among maximal itemsets in  $\mathcal{G}$  containing either  $X - \{x\}$  or  $X - \{y\}$ .*

**COROLLARY 24 (OF THEOREM 22).** *Let  $\mathcal{F}$  be a downward closed family of itemsets and let  $\mathcal{G}$  be the family after performing  $\text{SPLIT}(X, x, y)$ . Let  $S = X - \{x, y\}$ . Then the reduced families are updated as follows:*

1. *Itemset containing  $\{x, y\}$  in  $\text{rf}(\mathcal{F}; S)$  is split into two parts,  $\text{link}(\mathcal{G}, S, x)$  and  $\text{link}(\mathcal{G}, S, y)$ .*
2. *Itemset  $\{x\}$  is removed from  $\text{rf}(\mathcal{G}; S \cup \{y\})$ . Itemset  $\{y\}$  is removed from  $\text{rf}(\mathcal{G}; S \cup \{x\})$ .*
3. *Let  $T \subsetneq S$  and let  $z \in S - T$ . Item  $x$  is removed from the itemset containing  $z$  in  $\text{rf}(\mathcal{F}; T \cup \{y\})$ . Similarly, item  $y$  is removed from the itemset containing  $z$  in  $\text{rf}(\mathcal{F}; T \cup \{x\})$ .*
4. *Otherwise,  $\text{rf}(\mathcal{F}; T) = \text{rf}(\mathcal{G}; T)$  or  $\text{md}(\mathcal{F}; T) = 0$  and  $\text{md}(\mathcal{G}; T) = 0$ .*

We keep in memory only those families that have positive merge degree. Theorem 20 tells us that there are only  $K - 1$  such families. By studying the code in the update algorithm we see that, except in two cases, the update of a family is either a insertion/deletion of an element into an itemset or a merge of two itemsets. The first complex case is given on Line 2 in  $\text{MERGESIDE}$  which corresponds to Case

2 in Theorem 22. The problem is that this family may have contained only itemset before the merge, hence we did not store it. Consequently, we need to recreate the missing itemset, and this is done in  $O(\sum_{X \in \mathcal{F}} X)$  time. The second case occurs on Line 2 in  $\text{SPLITSIDE}$ . This corresponds to the case where we need to break the itemset  $W \in \text{rf}(S)$  containing  $x$  and  $y$  apart during a split (Case 1 in Corollary 24). This is done by constructing the new sets from scratch. The construction needs  $O(|W|K(|W| + M))$  time, where  $M$  is the size of largest itemset in  $\mathcal{F}$ .

---

**Algorithm 4:**  $\text{SPLITUPDATE}(X, x, y)$ . Routine for updating the structures during  $\text{SPLIT}(X, x, y)$ .

---

- 1 Update  $\mathcal{F}$ ;
  - 2 Remove  $\text{link}(S, x)$  from  $\text{rf}(S)$ ;
  - 3  $S \leftarrow X - \{x, y\}$ ;
  - 4  $\text{SPLITSIDE}(X, x, y, S)$ ;  $\text{SPLITSIDE}(X, y, x, S)$ ;
- 

---

**Algorithm 5:** Subroutine  $\text{SPLITSIDE}(X, a, b, S)$  used by  $\text{SPLITUPDATE}$ .

---

- 1  $Z \leftarrow S \cup \{a\}$ ;
  - 2 **while** changes **do**
  - 3    $Z \leftarrow Z \cup \{X \in \max(\mathcal{F}) ; S \subsetneq Z \cap X\}$ ;
  - 4 Add  $Z$  into  $\text{rf}(S)$ ;
  - 5 **if**  $\text{rf}(S \cup \{a\})$  exists **then**
  - 6   Remove  $\{b\}$  from  $\text{rf}(S \cup \{a\})$ ;
  - 7 **for**  $T \subsetneq S$ ,  $\text{rf}(T \cup \{a\})$  exists **do**
  - 8    $z \leftarrow$  (any) item in  $S - T$ ;
  - 9   Remove  $b$  from  $\text{link}(S, z)$ ;
  - 10 **if** there is unique  $Z \in \max(\mathcal{F})$  s.t.  $S \cup \{a\} \subsetneq Z$  **then**
  - 11   Update  $\text{st}(Z)$ ;
- 

---

**Algorithm 6:**  $\text{MERGEUPDATE}(S, x, y)$ . Routine for updating the structures during  $\text{MERGE}(S, x, y)$ .

---

- 1 Merge  $\text{link}(S, x)$  and  $\text{link}(S, y)$  in  $\text{rf}(S)$ ;
  - 2  $A \leftarrow$  itemset in  $\max(\mathcal{F})$  such that  $S \cup \{x\} \subseteq A$ ;
  - 3  $B \leftarrow$  itemset in  $\max(\mathcal{F})$  such that  $S \cup \{y\} \subseteq B$ ;
  - 4 Build  $\text{st}(S \cup \{x, y\})$  from  $\text{st}(A)$  and  $\text{st}(B)$ ;
  - 5  $\text{MERGESIDE}(S, x, y)$ ;  $\text{MERGESIDE}(S, y, x)$ ;
  - 6 Update  $\mathcal{F}$ ;
- 

---

**Algorithm 7:** Subroutine  $\text{MERGESIDE}(S, a, b)$  used by  $\text{MERGEUPDATE}$ .

---

- 1  $U \leftarrow S \cup \{a\}$ ;
  - 2 **if**  $U \notin \max(\mathcal{F})$  **and**  $\text{rf}(U)$  does not exist **then**
  - 3    $\text{rf}(U) \leftarrow \left\{ \bigcup_{U \subsetneq X \in \max(\mathcal{F})} X \right\}$ ;
  - 4 Add  $b$  into  $\text{rf}(U)$ ;
  - 5 **for**  $T \subsetneq S$ ,  $\text{rf}(T \cup \{a\})$  exists **do**
  - 6    $z \leftarrow$  (any) item in  $S - T$ ;
  - 7   Augment  $\text{link}(S, z)$  with  $b$ ;
  - 8 **if** there is unique  $Z \in \max(\mathcal{F})$  s.t.  $U \subsetneq Z$  **then**
  - 9   Update  $\text{st}(Z)$ ;
-

## 5. RELATED WORK

Many quality measures have been suggested for itemsets. A major part of these measures are based on how much the itemset deviates from some null hypothesis. For example, itemset measures that use the independence model as background knowledge have been suggested in [1, 3]. More flexible models have been proposed, such as, comparing itemsets against graphical models [11] and local Maximum Entropy models [12, 18]. In addition, mining itemsets with low entropy has been suggested in [9].

Our main theoretical advantage over these approaches is that we look at the itemsets as a whole collection. For example, consider that we discover that item  $a$  and  $b$  deviate greatly from the null hypothesis. Then any itemset containing both  $a$  and  $b$  will also be deemed interesting. The reason for this is that these methods are not adopting to the discovered fact that  $a$  and  $b$  are correlated, but instead they continue to use the same null hypothesis. We, on the other hand, avoid this problem by considering models: if itemset  $ab$  is found interesting that information is added into the statistical model. If this new model then explains bigger itemsets containing  $a$  and  $b$ , then we have no reason to add these itemsets, into the model, and hence such itemsets will not be considered interesting.

The idea of mining a pattern set as a whole in order to reduce the number of patterns is not new. For example, pattern reduction techniques based on minimum description length principle has been suggested [16, 21, 10]. Discovering decomposable models have been studied in [20]. In addition, a framework that incrementally adopts to the patterns approved by the user has been suggested in [8]. Our main advantage is that these methods require already discovered itemset collection as an input, which can be substantially large for low thresholds. We, on the other hand, skip this step and define the significance for itemsets such that we can mine the patterns directly.

## 6. EXPERIMENTS

In this section we present our empirical evaluation of the measure. We first describe the datasets and the setup for the experiments, then present the results with synthetic datasets and finally the results with real-world datasets.

### 6.1 Setup for the Experiments

We used  $2 \times 3$  synthetic datasets and 3 real-world datasets.

The first three synthetic datasets, called *Ind*, contained 15 independent items and 100,  $10^3$ , and  $10^4$  transactions, respectively. We set the frequency for the individual items to be 0.1. The next three synthetic datasets, called *Path*, also contained 15 items. In these datasets, an item  $a_i$  were generated from the previous one with  $P(a_i = 1 | a_{i-1} = 1) = P(a_i = 0 | a_{i-1} = 0) = 0.75$ . The probability of the first item was set to 0.5. We set the number of transactions for these datasets to 100,  $10^3$ , and  $10^4$ , respectively.

Our first real-world dataset *Paleo*<sup>2</sup> contains information of species fossils found in specific paleontological sites in Europe [7]. The dataset *Courses* contains the enrollment records of students taking courses at the Department of Computer Science of the University of Helsinki. Finally, our last dataset is *Dna* is DNA copy number amplification data collection of human neoplasms [13]. We used 100 first

items from this data and removed empty transactions. The basic characteristics of the datasets are given in Table 1.

For each data we sampled the models from the posterior distribution using techniques described in Section 3.4. We used singleton model as a starting point and did 5000 restarts. The number of required MCMC steps is hard to predict since the structure of the state space of decomposable models is complex. Further it also depends on the actual data. Hence, we settle for heuristic: for each restart we perform  $100K \log K$  MCMC steps, where  $K$  is the number of items. Doing so we obtained  $N = 5000$  random models for each dataset. The execution times for sampling are given in Table 1. Let  $\{\mathcal{F}_1, \dots, \mathcal{F}_N\}$  be the discovered models. We estimated the itemset score  $sc(X) \approx |\{\mathcal{F}_i | X \in \mathcal{F}_i\}|/N$  and mined interesting itemsets using a simple depth-first approach.

Name	$ D $	K	# of steps	time
<i>Ind</i>	100– $10^4$	15	4 063	4m–2h
<i>Path</i>	100– $10^4$	15	4 063	7m–3.5h
<i>Dna</i>	1160	100	46 052	6h
<i>Paleo</i>	501	139	68 590	5h
<i>Courses</i>	3506	90	40 499	8.5h

**Table 1: Basic characteristics of the datasets. The fourth column contains the number of sample steps and the last column is the execution time.**

### 6.2 Synthetic datasets

Our main purpose for the experiments with synthetic datasets is to demonstrate how the score behaves as a function of number of data points. To this end, we plotted the number of significant itemsets, that is itemsets whose score was higher than the threshold  $\sigma$ , as a function of the threshold  $\sigma$ . The results are shown in Figures 2(a) and 2(b).

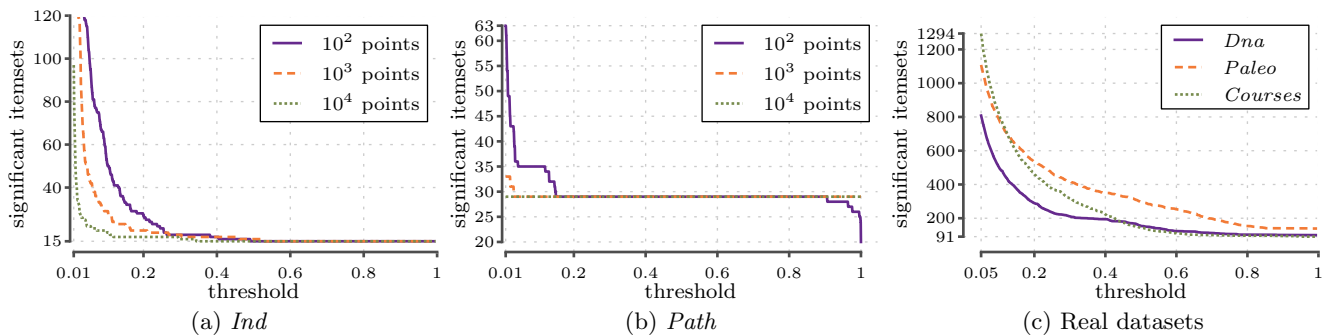
Ideally, for *Ind*, the dataset with independent variables we should have only 15 significant itemsets, that is, the singletons, for any  $\sigma > 0$ . Similarly, for *Path* we should have  $15 + 14 = 29$  itemsets, the singletons and the pairs of form  $a_i a_{i+1}$ . We can see from Figures 2(a) and 2(b) that as we increase the number of transactions in data, the number of significant itemsets approaches these ideal cases, as predicted by Theorem 7. The convergence to the ideal case is faster in *Path* than in *Ind*. The reason for this can be explained by the curse of dimensionality. In *Ind* we have  $15 \times 14/2 = 105$  combinations of pairs of items. There is a high probability that some of these item pairs appear to be correlated. On the other hand, for *Path*, let us assume that we have the correct model. That is, the singletons and the 14 pairs  $a_i a_{i+1}$ . The only valid itemsets of size 3 that we can add to this model are of the form  $a_i a_{i+1} a_{i+2}$ . There are only 13 of such sets, hence the probability of finding such itemset important is much lower. Interestingly, in *Path* we actually benefit from the fact that we are using decomposable models instead of general exponential models.

### 6.3 Use cases with real-world datasets

Our first experiment with real-world data is to study the number of significant itemsets as a function of the threshold  $\sigma$ . Figure 2(c) shows the number of significant itemsets for all three datasets. We see that the number of significant itemsets increases faster than for the synthetic datasets as

<sup>2</sup>NOW public release 030717 available from [7].





**Figure 2: Number of significant itemsets as a function of the threshold. The smallest threshold used for *Ind* and *Path* is 0.01. The smallest threshold used for real datasets is 0.05.**

the threshold decreases. The main reason for this difference, is that with real-world datasets we have more items and less transactions. This is seen especially in the *Paleo* dataset for which the number of significant itemsets increases steeply between the interval 0.4–1.0 when compared to *Dna* and *Courses*.

Our next experiment is to compare the score against baselines, namely, the frequency  $fr(X)$  and entropy  $H(X)$ . These comparisons are given in Figures 3(a) and 3(b). In addition, we computed the correlation coefficients (given in Table 2). From results, we see that  $sc(X)$  has a positive correlation with frequency and a negative correlation with entropy. The correlation with entropy is expected, since low-entropy implies that the empirical distribution of an itemset is different than the uniform distribution. Hence, using the frequency of such an itemset should improve the model and consequently the itemset is considered interesting.

Data	$fr(X)$	$H(X)$	index diff.
<i>Dna</i>	0.16	-0.27	-0.28
<i>Paleo</i>	0.45	-0.16	-0.17
<i>Courses</i>	0.18	-0.35	-0.02

**Table 2: Correlations of  $sc(X)$  and baselines.**

Both datasets *Paleo* and *Dna* have a natural order between the items, for example, for *Paleo* dataset it is the era when particular species was extant. Our next experiment is to test whether this order is represented in the discovered patterns. In our experiments, the items in these datasets were ordered using this natural order. Let  $X = a_i a_j$  be an itemset of size 2. In Figure 3(c), we plotted  $j - i$  as a function of  $sc(X)$ . We used only itemsets of size 2, since larger itemsets have inherently smaller scores and larger difference between the indices. In addition, we compute the correlation coefficients (given in Table 2). From the results we see that for *Paleo* and *Dna* the more significant itemsets tend to have a smaller difference between their indices. On the other hand, we do not discover any significant correlation in *Courses*.

Finally, we report some of the discovered patterns from *Courses*. The 4 most significant itemsets of size 2 are (*Computer Architectures, Performance Analysis*) with a score of 0.95, (*Design & Analysis of Algorithms, Principles of Functional Programming*) scoring 0.94, (*Database Systems II, In-*

*formation Storage*) scoring 0.94, (*Three concepts: probability, Machine Learning*) scoring 0.92.

## 7. CONCLUSIONS

In this paper we introduced a novel and general approach for ranking itemsets. The idea behind the approach is to connect statistical models and collections of itemsets. This connection enables us to define the score of an itemset as the probability of itemset occurring in a random model. Doing so, we transformed the problem of mining patterns into a more classical problem of modeling.

As a concrete example of the framework, we used exponential models. These models have many important theoretical and practical properties. The connection with itemsets is natural and the Occam’s razor inherent to these models can be used against the pattern explosion problem. Our experiments support the theoretical results and demonstrate that the measure works in practice.

## 8. ACKNOWLEDGMENTS

Nikolaj Tatti is funded by FWO postdoctoral mandate.

## 9. REFERENCES

- [1] C. C. Aggarwal and P. S. Yu. A new framework for itemset generation. In *Proceedings of the 17th ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems (PODS 1998)*, pages 18–24. ACM Press, 1998.
- [2] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A. I. Verkamo. Fast discovery of association rules. In *Advances in Knowledge Discovery and Data Mining*, pages 307–328. AAAI/MIT Press, 1996.
- [3] S. Brin, R. Motwani, and C. Silverstein. Beyond market baskets: Generalizing association rules to correlations. In *SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data*, pages 265–276. ACM Press, May 1997.
- [4] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Inf. Theory*, 14(3):462–467, 1968.
- [5] R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Statistics for Engineering and Information Science. Springer-Verlag, 1999.

- [6] I. Csiszár. I-divergence geometry of probability distributions and minimization problems. *The Annals of Probability*, 3(1):146–158, Feb. 1975.
- [7] M. Fortelius, A. Gionis, J. Jernvall, and H. Mannila. Spectral ordering and biochronology of european fossil mammals. *Paleobiology*, 32(2):206–214, March 2006.
- [8] S. Hanhijärvi, M. Ojala, N. Vuokko, K. Puolamäki, N. Tatti, and H. Mannila. Tell me something I don't know: randomization strategies for iterative data mining. In *Knowledge Discovery and Data Mining (KDD 2009)*, pages 379–388, 2009.
- [9] H. Heikinheimo, J. K. Seppänen, E. Hinkkanen, H. Mannila, and T. Mielikäinen. Finding low-entropy sets and trees from binary data. In *Knowledge discovery and data mining (KDD 2007)*, pages 350–359, 2007.
- [10] H. Heikinheimo, J. Vreeken, A. Siebes, and H. Mannila. Low-entropy set selection. In *Proceedings of the SIAM International Conference on Data Mining (SDM 2009)*, pages 569–580, 2009.
- [11] S. Jaroszewicz and D. A. Simovici. Interestingness of frequent itemsets using bayesian networks as background knowledge. In *Knowledge discovery and data mining (KDD 2004)*, pages 178–186, 2004.
- [12] R. Meo. Theory of dependence values. *ACM Trans. Database Syst.*, 25(3):380–406, 2000.
- [13] S. Myllykangas, J. Himberg, T. Böhlting, B. Nagy, J. Hollmén, and S. Knuutila. Dna copy number amplification profiling of human neoplasms. *Oncogene*, 25(55):7324–7332, Nov. 2006.
- [14] D. Pavlov, H. Mannila, and P. Smyth. Beyond independence: Probabilistic models for query approximation on binary transaction data. *IEEE TKDE*, 15(6):1409–1421, 2003.
- [15] G. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978.
- [16] A. Siebes, J. Vreeken, and M. van Leeuwen. Item sets that compress. In *Proceedings of the SIAM International Conference on Data Mining (SDM 2006)*, 2006.
- [17] N. Tatti. Computational complexity of queries based on itemsets. *IPL*, pages 183–187, June 2006.
- [18] N. Tatti. Maximum entropy based significance of itemsets. *KAIS*, 17(1):57–77, 2008.
- [19] N. Tatti. Probably the best itemsets. Technical Report 2010/06, University of Antwerp, 2010. <http://adrem.ua.ac.be/publications>.
- [20] N. Tatti and H. Heikinheimo. Decomposable families of itemsets. In *Knowledge Discovery in Databases: PKDD 2008*, pages 472–487, 2008.
- [21] N. Tatti and J. Vreeken. Finding good itemsets by packing data. In *Proceedings of the 8th IEEE International Conference on Data Mining (ICDM 2008)*, pages 588–597, 2008.
- [22] A. Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of the American Mathematical Society*, 52(3):426–482, 1943.

## APPENDIX

### A. PROOFS FOR SECTION 3

PROOF OF THEOREM 6. We will prove the theorem using the well-known connection between maximum entropy principle and exponential models. Let  $q_D$  be the empirical distribution of the data. Let  $\mathcal{P}$  be the collection of distributions,

$$\mathcal{P} = \{p \mid p(X=1) = fr(X) = q_D(X=1), X \in \mathcal{F}\},$$

that is, a distribution in  $\mathcal{P}$  has the same itemset frequencies for  $\mathcal{F}$  as the data. Assume for the time being that all entries in  $q_D$  are positive. Then, a well-known theorem [6] states that the unique distribution, say  $p^*$ , maximizing the entropy among  $\mathcal{P}$  belongs to the exponential model  $M$  such that  $fam(M) = \mathcal{F}$ . Let the corresponding parameters for the model be  $r^*$ .

We need to show that  $p^*$  is actually a maximum likelihood distribution. To see this, let  $q \in M$  be another distribution from the model and let  $r$  be its parameters. A straightforward calculus reveals that

$$\begin{aligned} \log \frac{p^*(D)}{q(D)} &= \sum_{t \in D} \log \frac{p^*(t)}{q(t)} = |D| \sum_{t \in \Omega} q_D(t) \log \frac{p^*(t)}{q(t)} \\ &= |D| \sum_{t \in \Omega} \sum_{X \in \mathcal{F}} q_D(X=1)(r_X^* - r_X) \\ &= |D| \sum_{t \in \Omega} \sum_{X \in \mathcal{F}} p^*(X=1)(r_X^* - r_X) \\ &= |D| \sum_{t \in \Omega} p^*(t) \log \frac{p^*(t)}{q(t)} = |D| KL(p^* \| q) \geq 0, \end{aligned}$$

where  $KL(p \| q)$  is the Kullback-Leibler divergence between  $p^*$  and  $q$ . The inequality shows that  $r^*$  has the highest likelihood. If  $q_D$  has zero probabilities, then we can consider a sequence  $\epsilon_n \rightarrow 0$  as  $n \rightarrow \infty$ . Define  $q_n = (1 - \epsilon_n)q_D + \epsilon_n 2^{-K}$ . Let  $p_n^*$  be the maximal entropy distribution computed from  $q_n$ . Let  $p^*$  be the maximal entropy distribution computed from  $q_n$  and let  $r_n = (1 - \epsilon_n)r^* + \epsilon_n 2^{-K}$ . Since  $r_n$  has the same frequencies than  $q_n$  for each  $X \in \mathcal{F}$ , it follows that  $H(p_n^*) \geq H(r_n)$ . Since  $r_n$  converges to  $p^*$ , we have that  $\liminf H(p_n^*) \geq H(p^*)$ . Any converging subsequence of  $p_n^*$  will converge to a distribution in  $\mathcal{P}$ , and hence having the entropy  $H(p^*)$ . Since maximal entropy distribution is unique, the sequence  $p_n^*$  converges to the unique value,  $p^*$ . Following the same calculus as above we can show that  $p^*$  is the maximal likelihood distribution

$$\begin{aligned} \sum_{t \in \Omega} q_D(t) \log \frac{p^*(t)}{q(t)} &= \lim_{n \rightarrow \infty} \sum_{t \in \Omega} q_n(t) \log \frac{p_n^*(t)}{q(t)} \\ &= \lim_{n \rightarrow \infty} \sum_{t \in \Omega} \sum_{X \in \mathcal{F}} q_n(X=1)(r_X^* - r_X) \\ &= \lim_{n \rightarrow \infty} \sum_{t \in \Omega} \sum_{X \in \mathcal{F}} p_n^*(X=1)(r_X^* - r_X) \\ &= \lim_{n \rightarrow \infty} \sum_{t \in \Omega} p_n^*(t) \log \frac{p_n^*(t)}{q(t)} \\ &= \sum_{t \in \Omega} p^*(t) \log \frac{p^*(t)}{q(t)} = KL(p^* \| q) \geq 0. \end{aligned}$$

Finally, as the number of transactions goes into infinity  $r^*$  converges to the true parameters of the model and  $p^*$  converges into true underlying distribution  $p$ .  $\square$

To prove Theorem 7 we will need the following lemma whose prove can be found for example in [22].

LEMMA 25. Assume that  $D$  with  $N$  data points comes from a model  $M$  with parameters  $r$ . Let  $r^N$  be the maximal likelihood estimate. Then

$$\log P(D | M, r^N) - \log P(D | M, r)$$

converges uniformly to  $\chi^2$  distribution with  $d \leq |F| - 1$  number degrees of freedom.

PROOF OF THEOREM 7. Let  $M' \neq M$  be a model and  $\mathcal{F}' = \text{fam}(M')$ . We need to show that

$$\frac{P(M' | D)}{P(M | D)} \rightarrow 0$$

as the number of data points goes to infinity.

We will use the BIC estimate given in Eq. 3 to prove the theorem. Let  $N = |D|$ . Let us define  $p^N$  to be the maximal likelihood distribution from  $M$  and  $q^N$  to be the maximal likelihood distribution in  $M'$ . Let us write

$$A_N = \log q^N(D) - \log p^N(D)$$

and

$$B_N = \log N(|\mathcal{F}'| - |\mathcal{F}'|)/2.$$

Then asymptotically

$$\log P(M' | D) - \log P(M | D) = A_N + B_N.$$

Let  $A$  be the limit of  $A_N$  and  $B$  be the limit of  $B_N$ .

Assume that  $p \notin M'$ . Let  $q \in M'$  be the maximum likelihood distribution as  $N \rightarrow \infty$ . Since the exponential models are closed sets  $q$  cannot be  $p$ . A straightforward calculation reveals that  $N^{-1}A_N \rightarrow -KL(p||q) < 0$ . In addition,  $N^{-1}B_N \rightarrow 0$ . Thus  $N^{-1}(A_N + B_N)$  approaches a negative value, and so  $A + B = -\infty$ .

Assume that  $p \in M'$ , then the conditions imply that  $|\mathcal{F}'| > |\mathcal{F}|$  so that  $B = -\infty$ . Lemma 25 implies that  $A$  is a difference of two  $\chi^2$  distributions. More importantly, since convergence in Lemma 25 is uniform we can compute the limit inside the probability so that the probability

$$P(A_N + B_N > \sigma) \rightarrow P(A + B > \sigma) = P(A > \infty) = 0$$

for any  $\sigma \in \mathbb{R}$ . Hence  $A_N + B_N$  approaches  $-\infty$ . This proves the theorem.  $\square$

PROOF OF THEOREM 9. Let  $p(A) = P(A | M, r)$  be the distribution from a model  $M$  with parameters  $r$ . To ease the notation, let us define  $\mathcal{G} = \max(\mathcal{F})$  and

$$\mathcal{S} = \{X \cap Y | (X, Y) \in E(\mathcal{T})\}.$$

A classic result states that we can decompose  $p$  to factors,

$$p(A = t) = \frac{\prod_{X \in \mathcal{G}} p(X = t_X)}{\prod_{S \in \mathcal{S}} p(S = t_S)},$$

where  $t_X$  is a projection of a binary vector  $t$  to the variables of  $X$ . Let  $q(A)$  be the empirical distribution. Let  $p$  be the maximum-likelihood distribution. The proof of Theorem 6 states that  $p(X = 1) = q(X = 1)$  for any itemset  $X \in \mathcal{F}$ . Using the exclusion-inclusion rules and the fact that  $\mathcal{F}$  is downward closed we can show that  $p(X = t) = q_D(X = t)$

for every  $X \in \mathcal{F}$  and every possible binary vector  $t$  of length  $|X|$ . Using the equality

$$\begin{aligned} \sum_{t \in D} \log p(X = t_X) &= |D| \sum_{t_X} q(X = t_X) \log p(X = t_X) \\ &= |D| \sum_{t_X} q(X = t_X) \log q(X = t_X) \\ &= -|D|H(X) \end{aligned}$$

we see that the log-likelihood is equal to

$$\begin{aligned} \log \prod_{t \in D} p(A = t) &= \sum_{t \in D} \log \frac{\prod_{X \in \mathcal{G}} p(X = t_X)}{\prod_{S \in \mathcal{S}} p(S = t_S)} \\ &= -|D| \sum_{X \in \mathcal{G}} H(X) + |D| \sum_{S \in \mathcal{S}} H(S). \end{aligned}$$

This proves the theorem.  $\square$

## B. PROOFS FOR SECTION 4

We will need the following technical lemma for several subsequent proofs.

LEMMA 26. Let  $X_1, \dots, X_N \in \max(\mathcal{F})$  with  $N \geq 3$ . Define  $S_i = X_i \cap X_{i+1}$  and  $S_N = X_N \cap X_1$ . If each  $S_i$  has a unique item (when compared to other  $S_i$ ), then  $\mathcal{F}$  is not decomposable.

PROOF OF THEOREM 12. Let us assume that  $X$  is the only itemset containing  $x$  and  $y$  simultaneously. Let  $\mathcal{T}$  be a junction tree. The adjacent itemsets of  $X$  either miss  $x$  or  $y$  or both. Thus we can remove  $X$  and replace it  $X - x$  and  $X - y$ , connected to each other, the adjacent itemsets of  $X$  can be connected to either  $X - x$  or to  $X - y$ . If  $X - x$  is not maximal, that is there is an itemset  $Z \supset X - x$ , then we can remove  $X - x$  and connect all the edges going to  $X - x$  to  $Z$ . We can repeat this for  $X - y$  as well. The resulting tree is a junction tree.

To prove the other direction, assume that there is an another itemset  $Y$  containing  $x$  and  $y$ , simultaneously. There must be an item  $z \in X$  such that  $z \notin Y$ . Then itemsets  $Y$ ,  $X - x$ , and  $X - y$  satisfy the requirements of Lemma 26, hence the new family is not decomposable.  $\square$

PROOF OF THEOREM 15. Let  $\mathcal{T}$  be a junction tree. Assume that  $V$  and  $W$  exist. By construction of  $rf(\mathcal{F}; S)$  it follows that there must be two itemsets  $X, Y \in \mathcal{F}$  such that  $S = X \cap Y$ ,  $x \in X$ , and  $y \in Y$ . If  $X$  and  $Y$  are adjacent in  $\mathcal{T}$ , we can add  $Z$  between  $X$  and  $Y$  (possibly removing  $X$  and  $Y$  if they are no longer maximal), and  $\mathcal{T}$  still remains a junction tree.

If  $X$  and  $Y$  are not adjacent, then there is a path  $P$  connecting them. There must be itemsets  $P_i$  and  $P_{i+1}$  along that path such that  $P_i \cap P_{i+1} = S$ , otherwise  $X$  and  $Y$  would have been merged during the construction of reduced family, which is a contradiction. If we remove edge  $(P_i, P_{i+1})$  and add edge  $(X, Y)$  we will obtain an alternative junction tree for  $\mathcal{F}$ . It is straightforward to see that this tree is truly a junction tree. But in this tree  $X$  and  $Y$  are adjacent, so we can add  $Z$  between them.

To prove the other direction, assume that adding  $Z$  is a legal merge. By definition it immediately implies that there are sets  $V, W \in rf(\mathcal{F}; S)$  such that  $x \in V$  and  $y \in W$ . We need to show that  $V \neq W$ . Let  $\mathcal{H}$  be the family resulted from the merge. Let  $\mathcal{U}$  be a junction tree for  $\mathcal{H}$ . Let  $X \in \max(\mathcal{F})$

be an itemset such that  $S \cup \{x\} \subseteq X$  and, similarly, let  $Y \in \max(\mathcal{F})$  such that  $S \cup \{y\} \subseteq Y$ . We may assume that  $X \neq S \cup \{x\}$  and  $Y \neq S \cup \{y\}$ , otherwise the proof is trivial. Hence,  $X, Y \in \max(\mathcal{H})$ . Let  $P$  be a path in  $\mathcal{U}$  from  $X$  to  $Y$ . Let  $R$  be a path in  $\mathcal{U}$  from  $Z$  to  $X$  and let  $P_e$  be the first entry in path  $R$  that also occurs in  $P$ . We must have  $Z \cap X \subseteq P_e$ . The path from  $Z$  to  $Y$  must also contain  $P_e$ , and so  $Z \cap Y \subseteq P_e$ . This implies that  $Z = (Z \cap X) \cup (Z \cap Y) \subseteq P_e$ . Since  $Z$  is a maximal set, we must have  $Z = P_e$ .

Itemset  $Z$  is the only maximal itemset in  $\mathcal{H}$  that contains  $x$  and  $y$  simultaneously. Otherwise, Theorem 12 states that  $\text{SPLIT}(Z, x, y)$  is illegal in  $\mathcal{H}$  so  $\text{MERGE}(S, x, y)$  is illegal in  $\mathcal{F}$ .

We can remove  $Z = P_e$  from  $\mathcal{U}$ , attach  $P_{e-1}$  and  $P_{e+1}$  to each other, and connect the rest adjacent nodes either to  $P_{e-1}$  or to  $P_{e+1}$  depending whether these nodes have  $x$  or  $y$  as a member. The outcome, say  $\mathcal{R}$ , is a junction tree for  $\mathcal{F}$ .

If  $V = W$ , then there must be a sequence

$$X = (T_1, \dots, T_N) = Y \in \mathcal{F}$$

such that  $S \subsetneq T_i \cap T_{i+1}$ . Now consider a path  $O$  in  $\mathcal{R}$  visiting each  $T_i$  in turn. Since  $\mathcal{R}$  is a junction tree we must have  $S \subsetneq O_i \cap O_{i+1}$ . Since  $X = T_1$  and  $Y = T_N$  path  $O$  must use the edge  $(P_{e-1}, P_{e+1})$ . But we must have  $P_{e-1} \cap P_{e+1} = S$ , otherwise  $\mathcal{U}$  would not have been a junction tree. This contradiction implies that  $V \neq W$ .  $\square$

**PROOF OF THEOREM 17.** We will prove the theorem by showing that all models can be brought by legal splits to the independence model. Since a split  $\text{SPLIT}(X, x, y)$  and a merge  $\text{MERGE}(X - \{x, y\}, x, y)$ , are the opposite operations we can reach model  $M$  from  $M'$  by first reaching the independence model with splits and then reaching  $M$  from the independence model by merges.

Let  $\mathcal{T}$  be a junction tree for  $\mathcal{F}' = \text{fam}(M')$  and let  $X \in \max(\mathcal{F}')$  be a maximal itemset such that  $X$  is the leaf clique in a junction tree  $\mathcal{T}$  with  $|X| \geq 2$ . If no such  $X$  exist, then  $M'$  is in fact the independence model. Let  $Y$  be the maximal itemset into which  $X$  is connected in  $\mathcal{T}$ .

There must be an item  $x \in X$  such that  $x$  is not contained in any other maximal itemset of  $\mathcal{F}'$ . To prove this, assume otherwise. Then for each item  $z$  there is a maximal itemset  $Z$  that contains  $z$ . The path from  $X$  to  $X$  must go through  $Y$  so the running intersection property implies that  $x \in Y$ . But this implies that  $X \subseteq Y$  which contradicts the maximality of  $X$ . Let  $y \in X$  such that  $x \neq y$ . Theorem 12 now states that  $\text{SPLIT}(X, x, y)$  is a legal operation. We can now iteratively apply this step until we reach the independence model. This proves the theorem.  $\square$

**PROOF OF THEOREM 18.** For notational convenience, let us write

$$a(M) = \sum_{\substack{(X, Y) \\ \in \mathcal{E}(\mathcal{T})}} H(X \cap Y) \text{ and } b(M) = \sum_{X \in \max(\mathcal{F})} H(X).$$

Theorem 9 and Eq. 3 implies that

$$\log A = |D|(a(M) - b(M)) - \frac{\log |D|(|\mathcal{F}| - 1)}{2}$$

and

$$\log B = |D|(a(M') - b(M')) - \frac{\log |D|(|\mathcal{F}'| - 1)}{2}.$$

Let us write  $d = |D|(a(M') - b(M') - a(M) + b(M))$ . We will first show that  $d = \text{gain}(X, x, y)$ . Let us denote  $S =$

$X - \{x, y\}$ ,  $U = X - x$ , and  $V = X - y$ . Consider the possibility that  $U$  is not maximal in  $\mathcal{F}'$ , in such case there must be a maximal itemset  $Q \in \mathcal{F}$  such that  $U \subset Q$  and that  $Q$  is adjacent to  $X$  with a separator  $U$ . Similarly, if  $V$  is not maximal in  $\mathcal{F}'$ , then there is a maximal itemset  $P \in \mathcal{F}$  such that  $V \subset P$  and that  $P$  is adjacent to  $X$  with a separator  $V$ . By studying the first part of the proof of Theorem 12 we see that there are four possible cases depending whether  $U$  and/or  $V$  is maximal set in  $\mathcal{F}'$ . The split operations are

$$\textcircled{X} \longrightarrow \textcircled{U} \text{---} S \text{---} \textcircled{V} \quad (\text{Case 1})$$

$$\textcircled{X} \text{---} V \text{---} \textcircled{P} \longrightarrow \textcircled{U} \text{---} S \text{---} \textcircled{P} \quad (\text{Case 2})$$

$$\textcircled{Q} \text{---} U \text{---} \textcircled{X} \longrightarrow \textcircled{Q} \text{---} S \text{---} \textcircled{V} \quad (\text{Case 3})$$

$$\textcircled{Q} \text{---} U \text{---} \textcircled{X} \text{---} V \text{---} \textcircled{P} \longrightarrow \textcircled{Q} \text{---} S \text{---} \textcircled{P} \quad (\text{Case 4})$$

In all cases, the difference  $d$  is equal to  $H(X) + H(S) - H(V) - H(U) = \text{gain}(X, x, y)$ . For example, in Case 2, the term  $H(X)$  is in  $b(M)$  but not in  $b(M')$ , similarly the term  $H(U)$  is in  $b(M')$  but not in  $b(M)$ , consequently  $b(M) - b(M') = H(X) - H(U)$ . Similarly,  $a(M) - a(M') = H(V) - H(S)$ . Thus  $d = \text{gain}(X, x, y)$  for Case 2, and similar observations show that  $d = \text{gain}(X, x, y)$  holds also for other cases.

To complete the proof we need to show that  $|\mathcal{F}| - |\mathcal{F}'| = 2^{|X|-2}$ . During a split, itemsets containing  $x$  and  $y$  are removed, but there are exactly  $2^{|X|-2}$  such itemsets. We can now combine these results

$$\begin{aligned} \log B - \log A &= d - \frac{\log |D|}{2} (|\mathcal{F}| - |\mathcal{F}'|) \\ &= \text{gain}(X, x, y) - \log |D| 2^{|X|-3}, \end{aligned}$$

which proves the theorem.  $\square$

**PROOF OF THEOREM 19.** Let  $X_i$  and  $Y_i$  be the maximal itemsets such that  $X_i \cap Y_i = S_i$ ,  $x_i \in X_i$  and  $y_i \in Y_i$ . These sets must exist since adding  $Z_i$  is a legal merge.

Assume that  $Z_1 = Z_2$ . This implies that either all items  $x_1, y_1, x_2$ , and  $y_2$  are unique or two of them are the same. Assume the latter case and assume that  $x_1 = x_2$ . Then we must have  $y_2 \in S_1 - S_2$  and  $y_1 \in S_2 - S_1$ . We have  $x_1 \in X_1 \cap X_2$ ,  $y_1 \in X_2 \cap Y_1$ , and  $y_2 \in Y_1 \cap X_1$ . Hence the conditions in Lemma 26 hold and  $\mathcal{F}$  is not decomposable.

Assume now that all four items are unique. This implies that  $x_1, y_1 \in S_2 - S_1$  and  $y_2 \in S_1 - S_2$ . Again we have  $x_1 \in X_1 \cap X_2$ ,  $y_1 \in X_2 \cap Y_1$ , and  $y_2 \in Y_1 \cap X_1$  so Lemma 26 implies that  $\mathcal{F}$  is not decomposable.  $\square$

**PROOF OF THEOREM 20.** Let  $S$  be an itemset. We will show that  $md(M, S) > 0$  only if  $S$  is a separator, that is,  $S = X \cap Y$  for two adjacent itemsets in a junction tree. The theorem will follow from this, since a junction tree contains at most  $K$  nodes and hence at most  $K - 1$  edges.

To prove the result, assume that  $md(M, S) > 0$ . This implies that there are two sets, say  $U, V \in \text{rf}(\mathcal{F}, S)$ . Let  $x \in U$  and  $y \in V$ . This implies that there are (at least) two sets  $X$  and  $Y$  such that  $S \cup \{x\} \subset X$  and  $S \cup \{y\} \subset Y$ . Let  $P = (P_1, \dots, P_L)$  be a path from  $X$  to  $Y$  in the junction tree. The running intersection property implies that

$S \subseteq P_i \cap P_{i+1}$  If there is no  $i$  such that  $S = P_i \cap P_{i+1}$ , then  $X$  and  $Y$  should have been joined together during the construction of  $rf(\mathcal{F}, S)$ . In other words, there must be a set in  $rf(\mathcal{F}, S)$  containing  $x$  and  $y$ . Hence, there are adjacent itemsets  $P_i$  and  $P_{i+1}$  such that  $S = P_i \cap P_{i+1}$ . This proves the theorem.  $\square$

PROOF OF THEOREM 21. Adding  $X = S \cup \{x, y\}$  into  $\mathcal{F}$  can only make splits illegal. Assume that  $\text{SPLIT}(Y, a, b)$  becomes illegal after adding  $X$ . Theorem 12 implies that  $a, b \in X$ . We must have  $\{a, b\} \not\subseteq S$  since otherwise the split would not be legal in the original family. Assume that  $a = x$  and let  $W \in \max(\mathcal{F})$  is the maximal itemset containing  $\{x\} \cup S$  (such itemset exists because of the definition of merge). Note that,  $b \neq y$ , so  $b \in S$  and  $\{a, b\} \in W$ . If there is another maximal itemset, say  $Z \in \max(\mathcal{F})$ , such that  $\{x, b\} = \{a, b\} \subset Z$ , then the split  $\text{SPLIT}(Y, x, b)$  is not legal in the original family. This proves the theorem.  $\square$

PROOF OF THEOREM 22. Let  $Z = S \cup \{x, y\}$ . This itemset is the only maximal itemset in  $\mathcal{G}$  that contains  $x$  and  $y$  simultaneously. The definition of reduced family now implies immediately Claims 2 and 3.

To prove Claim 1 let  $V, W \in rf(\mathcal{F}; S)$  such that  $x \in V$  and  $y \in W$ . These must be separate sets because of Theorem 19. The set  $Z$  now connects these sets into one in  $rf(\mathcal{G}; S)$ .

To see Claim 4 let  $T$  be an itemset. If  $x, y \in T$ , then  $rf(\mathcal{G}; T)$  has only one set and  $rf(\text{ifam}\mathcal{G}; T)$  has zero sets. If  $x \in T$  and  $y \notin T$ , then  $T - S \neq \emptyset$ , otherwise  $T$  is covered by Claim 2 or 3. In such case, the construction of  $rf(\mathcal{G}; T)$  does not use  $Z$ , and hence remains unchanged. Assume that  $x, y \notin T$ . Let  $X, Y \in \max(\mathcal{F})$  be the itemsets such that  $S \cup \{x\} \subseteq X$  and  $S \cup \{y\} \subseteq Y$ . If  $T - S \neq \emptyset$ , then  $Z$  is not used. Assume now that  $S \subsetneq T$ . In this case the items of  $X - T$  and  $Y - T$  are already joined into one itemset, hence  $rf(\mathcal{G}; T)$  remains unchanged.  $\square$