

# Finding Good Itemsets by Packing Data

Nikolaj Tatti

HIIT, Department of Information and Computer Science  
Helsinki University of Technology  
ntatti@cc.hut.fi

Jilles Vreeken

Department of Computer Science  
Universiteit Utrecht  
jillesv@cs.uu.nl

## Abstract

*The problem of selecting small groups of itemsets that represent the data well has recently gained a lot of attention. We approach the problem by searching for the itemsets that compress the data efficiently. As a compression technique we use decision trees combined with a refined version of MDL. More formally, assuming that the items are ordered, we create a decision tree for each item that may only depend on the previous items. Our approach allows us to find complex interactions between the attributes, not just co-occurrences of 1s. Further, we present a link between the itemsets and the decision trees and use this link to export the itemsets from the decision trees. In this paper we present two algorithms. The first one is a simple greedy approach that builds a family of itemsets directly from data. The second one, given a collection of candidate itemsets, selects a small subset of these itemsets. Our experiments show that these approaches result in compact and high quality descriptions of the data.*

## 1 Introduction

One of the major topics in data mining research is the discovery of interesting patterns in data. From the introduction of frequent itemset mining and association rules [2], the pattern explosion was acknowledged: at high frequency thresholds only common knowledge is revealed, while at low thresholds prohibitively many patterns are returned.

Part of this problem can be solved by reducing these collections either lossless or lossy, however even then the resulting collections are often so large that they cannot be analyzed by hand or even machine. Recently, it was therefore argued [14] that while the efficiency of the search process has received ample attention, there still exists a strong need

for pattern mining approaches that deliver compact, yet high quality, collections of patterns (see Section 6 for a more detailed discussion). Our goal is to identify the family of itemsets that form the best description of the data. Recent proposals to this end all consider just part of the data, by either only considering co-occurrences [30] or being lossy in nature [5, 7, 20]. In this paper, we present two methods that do describe all interactions in the data. Although different in approach, both methods return small families of itemsets, which are selected to provide high-quality lossless descriptions of the data in terms of local patterns. Importantly, our parameterless methods regard the data symmetrically. That is, we consider not just the 1s in the data, but also the 0s. Therefore, we are able to find patterns that describe *all* interactions between items in the data, not just co-occurrences.

As a measure of quality for the collection of itemsets we employ the practical variant of Kolmogorov Complexity [23], the Minimum Description Length (MDL) principle [13]. This principle implies that we should do induction through compression. It states that the best model is the model that provides the best compression of the data: it is the model that captures best the regularities of the data, with as little redundancy as possible.

The main idea of our approach is to use decision trees to determine the shortest possible encoding of an attribute, by using the values of already transmitted attributes. For example, let us assume two binary attributes  $A$  and  $B$ . Now say that for 90% of the time when the attribute  $A$  has a value of 1, the attribute  $B$  has a value of 0. If this situation occurs frequently, we recognize this dependency, and include the item  $A$  in the tree deciding how to encode  $B$ .

Using such trees allows us to find complex interactions between the items while at the same time MDL provides us with a parameter-free framework for removing fake interactions that are due to the noise in the data. The main outcome of our methods is not the decision trees, but the group of itemsets that form their paths: these are the impor-

tant patterns in the data since they capture the dependencies between the attributes implied by the decision trees.

The two algorithms we introduce to this end are orthogonal in approach. Our first method builds the encoding decision trees directly from the data; it greedily introduces splits until no split can help to compress the data further. Just as naturally as we can extract itemsets from these trees, we can consider the trees that can be built from a collection of itemsets. That link is exploited by our second method, which tries to select the best itemsets from a larger collection.

Experimental evaluation shows that both methods return small collections of itemsets that provide high quality data descriptions. These sets allow for very short encoding of the data, which inherently shows that the most important patterns in the data are captured. As the number of itemsets are small, we can easily expose the resulting itemsets to further analysis, either by hand or by machine.

The rest of this paper is as follows. After the covering preliminaries in Section 2, we discuss how to use decision trees to optimally encode the data succinct in Section 3. Next, in Section 4, we explain the connection between decision trees and itemsets. Section 5 introduces our method with which good itemsets can be selected by weighing these through our decision tree encoding. Related work is discussed in Section 6, after which we present the experiments on our methods in Section 7. We round up with discussion and conclusions in Sections 8 and 9.

## 2 Preliminaries and Notation

In this section we introduce preliminaries and notations used in subsequent sections.

A *binary dataset*  $D$  is a collection of  $|D|$  *transactions*, binary vectors of length  $K$ . The  $i$ th element of a random transaction is represented by an *attribute*  $a_i$ , a Bernoulli random variable. We denote the collection of all the attributes by  $A = \{a_1, \dots, a_K\}$ . An *itemset*  $X = \{x_1, \dots, x_L\} \subseteq A$  is a subset of attributes. We will often use the dense notation  $X = x_1 \dots x_L$ .

Given an itemset  $X$  and a binary vector  $v$  of length  $L$ , we use the notation  $p(X = v)$  to express the probability of  $p(x_1 = v_1, \dots, x_L = v_L)$ . If  $v$  contains only 1s, then we will use the notation  $p(X = 1)$ , if  $v$  contains only 0s, then we will use the notation  $p(X = 0)$ .

Given a binary dataset  $D$  we define  $q_D$  to be an *empirical distribution*,

$$q_D(A = v) = |\{t \in D \mid t = v\}|/|D|.$$

We define the frequency of an itemset  $X$  to be  $fr(X) = q_D(X = 1)$ .

In the paper we use the common convention  $0 \log 0 = 0$ . All logarithms in the paper are of base 2.

In the subsequent sections we will need some knowledge of graphs. All the graphs in the paper are directed. Given a graph  $G$  we denote by  $V(G)$  the set of vertices and by  $E(G)$  the edges of  $G$ . A directed graph is said to be *acyclic* (DAG) if there is no cycle in the graph. A directed graph is said to be *directed spanning tree* if each node (except one special node) has exactly one outgoing edge. The special node has no outgoing edge and is called *sink*.

## 3 Packing Binary Data with Decision Trees

In this section we present our model for packing the data and a greedy algorithm for searching good models.

### 3.1 The Definition of the Model

Our goal in this section is to define a model that is used to transmit a binary dataset  $D$  from a transmitter to a receiver. We do this by transmitting one transaction at the time, the order of which does not matter. Within a single transaction we transmit the items one at the time.

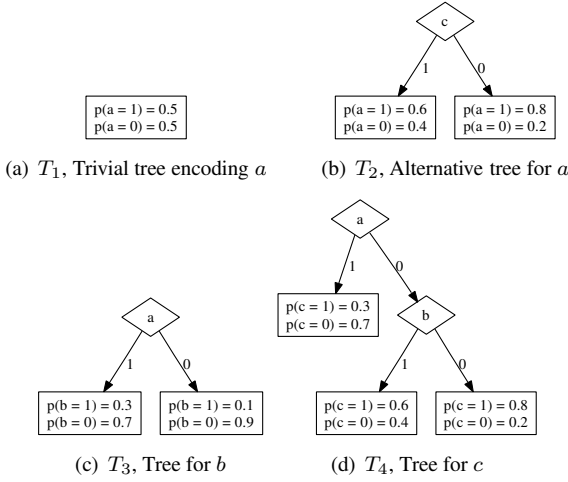
Assume that we are transmitting an attribute  $a_t$ . As the attribute may have two values, we need to have two codes to indicate its value. We define the table in which these two codes are stored to be a *coding table*. Obviously, the codes need to be optimal, that is, as short as possible. From information theory [10], we have the optimal Shannon codes of length  $-\log(p(x))$ . Here, the optimal code lengths are thus  $-\log q_D(a_t = 1)$  and  $-\log q_D(a_t = 0)$ . We need to transmit the attribute  $|D|$  times. The cost of these transmissions is

$$-|D| \sum_{v=\{0,1\}} q_D(a_t = v) \log q_D(a_t = v).$$

This is the simplest case of encoding  $a_t$ . Note that we are not interested in the actual codes, but only in their lengths: they allow us to determine the complexity of a model.

A more complex and more interesting approach to encode  $a_t$  succinct is to have several coding tables from which the transmitter chooses one for transmission. Choosing the coding table is done via a decision tree that branches on the values of other attributes in the same transaction. That is, we have a decision tree used for encoding  $a_t$  in which each leaf node is associated with a different coding table of  $a_t$ . The leaf is selected by testing the values of other attributes within the same transaction.

**Example 1.** Assume that we have three attributes,  $a$ ,  $b$ , and  $c$  and consider the trees given in Figure 1. In Figure 1(a) we have the simplest tree, a simple coding table with no dependencies at all. A more complex tree is given in Figure 1(b) where the transmitter chooses from two coding table for a based on the value of  $c$ . Similarly in, Figure 1(d) we have



**Figure 1. Toy decision trees.**

three different coding tables for  $c$ . The choice of the coding table in this case is based on the values of  $a$  and  $b$ .

Let us introduce some notation. Let  $T$  be a tree encoding  $a_t$ . We use the notation  $t(T) = a_t$ . We set  $src(T)$  to be the set of all items used in  $T$  for choosing the coding table.

**Example 2.** For the tree  $T_3$  in Figure 1(c) we have  $t(T_3) = b$  and  $src(T_3) = \{a\}$  and for  $T_4$  in Figure 1(d) we have  $t(T_4) = c$  and  $src(T_4) = \{a, b\}$ .

To define the cost of transmitting  $a_t$  we first define  $lvs(T)$  to be the set of all leaves in  $T$ . Let  $L \in lvs(T)$  be a leaf and  $q_D(L)$  be the probability of  $L$  being chosen. Further,  $q_D(a_t = v | L)$  is the probability of  $a_t = v$  given that  $L$  is chosen. We now know that the optimal cost, denoted by  $c_D(T)$ , is

$$-|D| \sum_{L \in lvs(T)} \sum_{v \in \{0,1\}} q_D(a_t = v, L) \log q_D(a_t = v | L).$$

**Example 3.** The number of bits needed by  $T_1$  in Figure 1(a) to transmit  $a$  in a random transaction is

$$-0.5 \log 0.5 - 0.5 \log 0.5 = 1.$$

Similarly, if we assume that  $q_D(a = 1) = q_D(a = 0) = 0.5$ , the number of bits needed by  $T_3$  to transmit  $c$  in a random transaction is

$$0.5 (-0.3 \log 0.3 - 0.7 \log 0.7) + 0.5 (-0.1 \log 0.1 - 0.9 \log 0.9) = 0.62.$$

In order for the receiver to decode the attribute  $a_t$  he must know what coding table was used. Thus, he must be able to use the same decision tree that the transmitter used for encoding  $a_t$ . To ensure this, the transmitter must know

$src(T)$  when decoding  $a_t$ . So, the attributes must have an order in which they are sent and the decision trees may only use the attributes that have already been transmitted.

The aforementioned requirement is easily characterized by the following construction. Let  $G$  be a directed graph with  $K$  nodes, each node corresponding to an attribute. The graph  $G$  contains all the edges of form  $(a_t, a_s)$  where  $a_s \in src(T)$ , where  $T$  is the tree encoding  $a_t$ . We call  $G$  the *dependency graph*. It is easy to see that there exists an order of the attributes if and only if  $G$  is an acyclic graph (DAG). If  $G$  constructed from a set of trees  $\mathcal{T} = \{T_1, \dots, T_K\}$  is indeed DAG we call the set  $\mathcal{T}$  a *decision tree model*.

**Example 4.** Consider a graph given in Figure 2(a) constructed from the trees  $T_2, T_3$ , and  $T_4$  (Figure 1). We cannot use this combination of trees for encoding since there is a cycle in the graph. On the other hand if we use trees  $T_1, T_3$ , and  $T_4$ , then the resulting graph (given in Figure 2(b)) is acyclic and thus these trees can be used for the transmission.



(a) Dependency graph with cycles (b) Dependency acyclic graph

**Figure 2. Dependency graphs constructed from the trees given in Figure 1.**

### 3.2 Encoding the Data

In order for the receiver to be able to decode the attributes, he must know both the coding tables and the trees. Hence, we need to transmit both of these. First, we cover how the coding tables, the leafs of the decision trees, are transmitted.

To transmit the coding tables we use the concept of Refined MDL [13]. Refined MDL is an improved version of the more traditional two-part MDL (sometimes referred to as the crude MDL). The basic idea of the refined variant is that instead of transmitting the coding tables, the transmitter and the receiver use so called universal codes. Universal codes are the cornerstone of Refined MDL. As these are codes that can be derived without any further shared information, this allows for a good weighing of the actual complexity of the data and model, with virtually no overhead. While the practicality of applying such codes depends on the type of the model, our decision trees are particularly well-suited.

These universal codes provide a cost called the *complexity* of the model. This cost can be calculated as follows: let  $L$  be a leaf in the decision tree (i.e. coding table), and  $M$  be the number of transactions for which  $L$  is used. Then the complexity of this leaf, denoted by  $c_{\text{MDL}}(L)$ , is

$$c_{\text{MDL}}(L) = \log \sum_{k=0}^M \binom{M}{k} \left(\frac{k}{M}\right)^k \left(\frac{M-k}{M}\right)^{M-k}.$$

In general, there is no known closed formula for the complexity of the model. Hence estimates are usually employed [29]. However, for our tree models we can apply an existing linear-time algorithm that solves the complexity for multinomial models [21]. We should also point out that the Refined MDL is asymptotically equivalent to Bayes Information Criteria (BIC) if the number of transactions goes to infinity and the number of free parameters stays fixed. However, for moderate numbers of transactions there may be significant differences [13].

Now that the coding tables can be transmitted, we need to know how to transmit the actual tree  $T$ . To encode the tree we simply transmit the nodes of the tree in a sequence. We use one bit to indicate whether the node is a leaf, or an intermediate node  $N \in \text{intr}(T)$ . For an intermediate node we additionally use  $\log K$  bits, where  $K$  is the number of attributes in  $D$ , to indicate the item that is used for the split.

The combined cost of a tree  $T$ , denoted by  $c(T)$ , is

$$c(T) = \sum_{N \in \text{intr}(T)} (1 + \log K) + c_D(T) + \sum_{L \in \text{lvs}(T)} (1 + c_{\text{MDL}}(L)),$$

that is, the cost  $c(T)$  is the number of bits needed to transmit the tree *and* the attribute  $a_t$  in each transaction of  $D$ .

**Example 5.** Assume that we have a dataset with 100 transactions and 3 items. Assume also that  $q_D(a=0) = q_D(a=1) = 0.5$ . We know that the complexity of the leaves in this case is  $c_{\text{MDL}}(L) = 3.25$ . The cost of the tree  $T_3$  (Figure 1(c)) is

$$\begin{aligned} c(T_3) &= 1 + \log 3 \\ &\quad + 1 + 3.25 + 50(-0.3 \log 0.3 - 0.7 \log 0.7) \\ &\quad + 1 + 3.25 + 50(-0.1 \log 0.1 - 0.9 \log 0.9) \\ &= 69.8. \end{aligned}$$

Given a decision tree model  $\mathcal{T} = \{T_1, \dots, T_K\}$  we define the cost  $c(\mathcal{T}) = \sum_i c(T_i)$ . The cost  $c(\mathcal{T})$  is the number of bits needed to transmit the trees, one for each attribute, and the complete dataset  $D$ .

We should point out that for data with many items, the term  $\log K$  grows and hence the threshold increases for selecting an attribute into any decision tree. This is an interesting behavior, as due to the finite number of transactions,

for datasets with many items there is an increased probability that two items will correlate, even though they are independent according to the generative distribution.

### 3.3 Greedy Algorithm

Our goal is to find the decision tree model with the lowest complexity cost. However, since many problems related to the decision trees are **NP**-complete [26] we will resort to a greedy heuristic to approximate the decision tree model  $\mathcal{T}$  with the lowest  $c(\mathcal{T})$ . It is based on the ID3 algorithm.

To fully introduce the algorithm we need some notation: By  $\text{TRIVIALTREE}(a_t)$  we mean the simplest tree packing  $a_t$  without any other attributes (see Figure 1(a)). Given a tree  $T$ , a leaf  $L \in \text{lvs}(T)$ , and an item  $c$  not occurring in the path from  $L$  to the root of  $T$ , we define  $\text{SPLITTREE}(T, L, c)$  to be a new tree where  $L$  is replaced by a non-leaf node testing the value of  $c$  and having two leaves as the branches.

The algorithm  $\text{GREEDYPACK}$  starts with a tree model consisting only of trivial trees. The algorithm finds the tree which saves the most bits by splitting. To ensure that the decision tree model is valid,  $\text{GREEDYPACK}$  builds a dependency graph  $G$  describing the dependencies of the trees and makes sure that  $G$  is acyclic. The algorithm terminates when no further split can be made that saves any bits.

---

**Algorithm 1**  $\text{GREEDYPACK}$  algorithm constructs a decision tree model  $\mathcal{T} = \{T_1, \dots, T_K\}$  from a binary data  $D$ .

---

```

1:  $V \leftarrow \{v_1, \dots, v_K\}, E \leftarrow \emptyset.$ 
2:  $G \leftarrow (V, E).$ 
3:  $T_i \leftarrow \text{TRIVIALTREE}(a_i),$  for  $i = 1, \dots, K.$ 
4: while there are changes do
5:   for  $i = 1, \dots, K$  do
6:      $O_i \leftarrow T_i.$ 
7:     for  $L \in \text{lvs}(T_i), j = 1, \dots, K$  do
8:       if  $E \cup (v_i, v_j)$  is acyclic and  $a_j \notin \text{path}(L)$  then
9:          $U \leftarrow \text{SPLITTREE}(T_i, L, a_j).$ 
10:        if  $c(U) < c(O_i)$  then
11:           $O_i \leftarrow U, s_i \leftarrow j.$ 
12:        end if
13:      end if
14:    end for
15:  end for
16:   $k \leftarrow \arg \min_i \{c(O_i) - c(T_i)\}.$ 
17:  if  $c(O_k) < c(T_k)$  then
18:     $T_k \leftarrow O_k.$ 
19:     $E \leftarrow E \cup (v_k, v_{s_k}).$ 
20:  end if
21: end while
22: return  $\{T_1, \dots, T_K\}.$ 

```

---

## 4 Itemsets and Decision Trees

So far we have discussed how to transmit binary data by using decision trees. In this section we present how to select the itemsets representing the dependencies implied by the decision trees. We will use this link in Section 5. A similar link between itemsets and decision trees is explored in [27] although our setup and goals are different.

Given a leaf  $L$ , the dependency of the item  $a_t$  is captured in the coding table of  $L$ . Hence we are interested in finding itemsets that carry the same information. That is, itemsets from which we can compute the coding table. To derive the codes for the leaf  $L$  it is sufficient to compute the probability

$$q_D(a_t = 1 \mid L) = q_D(a_t = 1, L) / q_D(L). \quad (1)$$

Our goal is to express the probabilities on the right side of the equation using itemsets. In order to do that let  $P$  be the path from  $L$  to its root. Let  $pos(L)$  be the items along the path  $P$  which are tested positive. Similarly, let  $neg(L)$  be the attributes which are tested negative. Using the inclusion-exclusion principle we see that

$$\begin{aligned} q_D(L) &= q_D(pos(L) = 1, neg(L) = 0) \\ &= \sum_{V \subseteq neg(L)} (-1)^{|V|} fr(pos(L) \cup V). \end{aligned} \quad (2)$$

We compute  $q_D(a_t = 1, L)$  in a similar fashion. Let us define  $sets(L)$  for a given leaf  $L$  to be

$$\begin{aligned} sets(L) &= \{V \cup pos(L) \mid V \subseteq neg(L)\} \\ &\cup \{V \cup pos(L) \cup \{a_t\} \mid V \subseteq neg(L)\}. \end{aligned}$$

Combining Eqs. 1–2 we see that the collection  $sets(L)$  satisfies our goal.

**Proposition 6.** *The coding table associated with the leaf  $L$  can be computed from the frequencies of  $sets(L)$ .*

**Example 7.** *Let  $L_1$ ,  $L_2$ , and  $L_3$  be the leaves (from left to right) of  $T_4$  in Figure 1(d). Then the corresponding families of itemsets are  $sets(L_1) = \{a, ac\}$ ,  $sets(L_2) = \{b, ab, bc, abc\}$ , and  $sets(L_3) = \{\emptyset, a, b, ab, c, ac, bc, abc\}$ .*

We can easily see that the family  $sets(L)$  is essentially the smallest family of itemsets from which the coding table can be derived uniquely.

**Proposition 8.** *Let  $\mathcal{G} \neq sets(L)$  be a family of itemsets. Then there are two data sets, say  $D_1$  and  $D_2$ , for which  $q_{D_1}(a_t = 1 \mid L) \neq q_{D_2}(a_t = 1 \mid L)$  but  $fr(\mathcal{G}; D_1) = fr(\mathcal{G}; D_2)$ .*

Given a tree  $T$  we define  $sets(T)$  to be  $sets(T) = \bigcup_{L \in lvs(T)} sets(L)$ . We also define  $sets(\mathcal{T}) = \bigcup_i sets(T_i)$  where  $\mathcal{T} = \{T_1, \dots, T_K\}$  is a decision tree model.

## 5 Choosing Good Itemsets

The connection between itemsets and decision trees made in the previous section allows us to consider an orthogonal approach to identify good itemsets. Informally, our goal is to construct decision trees from a family of itemsets  $\mathcal{F}$ , selecting the subset from  $\mathcal{F}$  that provides the best compression of the data. More formally, our new approach is as follows: given a downward closed family of itemsets  $\mathcal{F}$ , we build a decision tree model  $\mathcal{T} = \{T_1, \dots, T_K\}$  providing a good compression of the data, with  $sets(\mathcal{T}) \subseteq \mathcal{F}$ .

Before we can describe our main algorithm, we need to introduce some further notation. Firstly, given two trees  $T_p$  and  $T_n$  not using attribute  $c$ , we define  $JOINTREE(c, T_p, T_n)$  to be the join tree with  $c$  as the root node,  $T_p$  as the positive branch of  $c$ , and  $T_n$  as the negative branch of  $c$ . Secondly, to define our search algorithm we need to find the best tree

$$\begin{aligned} bt(a_t; S, \mathcal{F}) &= \arg \min_T \{c(T) \mid t(T) = a_t, \\ &\quad src(T) \subseteq S, sets(T) \subseteq \mathcal{F}\}, \end{aligned}$$

that is,  $bt(a_t; S, \mathcal{F})$ , returns the best tree for  $a_t$  for which the related sets are in  $\mathcal{F}$  and only splits on attributes in  $S$ .

To compute the optimal tree  $bt(a_t; S, \mathcal{F})$ , we use the exhaustive method (presented originally in [27]) given in Algorithm 2. The algorithm is straightforward: it tests each valid item as the root and recurses itself on both branches.

---

**Algorithm 2** GENERATE algorithm for calculating  $bt(a_t; S, \mathcal{F})$ , that is, the best tree  $T$  for  $a_t$  using only  $S$  as source and having  $sets(T) \subseteq \mathcal{F}$ .

---

```

1:  $B \leftarrow S \cap (\bigcup \mathcal{F})$ .
2:  $\mathcal{C} \leftarrow \text{TRIVIALTREE}(a_t)$ .
3: for  $b \in B$  do
4:    $\mathcal{G} \leftarrow \{X - b \mid b \in X \in \mathcal{F}\}$ .
5:    $(D_p, D_n) \leftarrow \text{SPLIT}(D, b)$ .
6:    $T_p \leftarrow \text{GENERATE}(a_t, \mathcal{G}, S, D_p)$ .
7:    $T_n \leftarrow \text{GENERATE}(a_t, \mathcal{G}, S, D_n)$ .
8:    $\mathcal{C} \leftarrow \mathcal{C} \cup \text{JOINTREE}(b, T_p, T_n)$ .
9: end for
10: return  $\arg \min_T \{c(T) \mid T \in \mathcal{C}\}$ .
```

---

We can now describe the actual algorithm for constructing decision tree models with a low cost. Our method automatically discovers the order in which the attributes can be transmitted most succinct. For this, it needs to find sets of attributes  $S_i$  for each attribute  $a_i$  such that these should be encoded before  $a_i$ . The collection  $\mathcal{S} = \{S_1, \dots, S_K\}$  should define an acyclic graph and the actual trees are  $bt(a_i; S_i, \mathcal{F})$ . We use  $c(\mathcal{S})$  as a shorthand for the total complexity  $\sum_i c(bt(a_i; S_i, \mathcal{F}))$  of the best model built from  $\mathcal{S}$ .

We construct the set  $\mathcal{S}$  iteratively. At the beginning of the algorithm we have  $S_i = \emptyset$  and we increase the sets  $S_i$  one attribute at a time. We allow ourselves to mark the attributes. The idea is that once the attribute  $a_i$  is marked, then we are not allowed to augment  $S_i$  any longer. At the beginning none of the nodes are marked.

To describe a single step in the algorithm we consider a graph  $H = (v_0, \dots, v_K)$ , where  $v_1, \dots, v_K$  represent the attributes and  $v_0$  is a special auxiliary node. We start by adding edges  $(v_i, v_0)$  having the weight  $c(bt(a_i; S_i, \mathcal{F}))$ , thus the cost of the best tree possible from  $\mathcal{F}$  using only the attributes in  $S_i$ . Then, for each unmarked node  $v_i$  we find out what other extra attribute will help most to encode it succinct. To do this, we add the edge  $(v_i, v_j)$  for each  $v_j$  with the weight  $c(bt(a_i; S_i \cup \{a_j\}, \mathcal{F}))$ . Now, let  $U$  be the minimum directed spanning tree of  $H$  having  $v_0$  as the sink. Consider an unmarked node  $v_i$  such that  $(v_i, v_0) \in E(U)$ . That node is now the best choice to be fixed, as it helps to encode the data best. We therefore mark attribute  $a_i$  and add  $a_i$  to each  $S_j$  for each ancestor  $v_j$  of  $v_i$  in  $U$ . This process is repeated until all attributes are marked. The details of the algorithm are given in Algorithm 3.

---

**Algorithm 3** The algorithm SETPACK constructs a decision tree model  $\mathcal{T}$  given a family of itemsets  $\mathcal{F}$  such that  $sets(\mathcal{T}) \subseteq \mathcal{F}$ . Returns a DAG, a family  $S = (S_1, \dots, S_K)$  of sets of attributes. The trees are  $T_i = bt(a_i, S_i, \mathcal{F})$ .

---

```

1:  $\mathcal{S} = (S_1, \dots, S_K) \leftarrow (\emptyset, \dots, \emptyset)$ .
2:  $r = (r_1, \dots, r_K) \leftarrow (\mathbf{false}, \dots, \mathbf{false})$ .
3:  $V \leftarrow \{v_0, \dots, v_K\}$ .
4: while there exists  $r_i = \mathbf{false}$  do
5:    $E \leftarrow \emptyset$ .
6:   for  $i = 1, \dots, K$  do
7:      $E \leftarrow E \cup (v_i, v_0)$ .
8:      $w(v_i, v_0) \leftarrow c(bt(a_i; S_i, \mathcal{F}))$ .
9:     if  $r_i = \mathbf{false}$  then
10:      for  $j = 1, \dots, K$  do
11:         $T \leftarrow bt(a_i; S_i \cup \{a_j\}, \mathcal{F})$ .
12:        if  $c(T) \leq w(v_i, v_0)$  then
13:           $E \leftarrow E \cup (v_i, v_j)$ ,  $w(v_i, v_j) \leftarrow c(T)$ .
14:        end if
15:      end for
16:    end if
17:  end for
18:   $U \leftarrow dmst(V, E)$  {Directed Min. Spanning Tree.}
19:  for  $(v_i, v_0) \in E(U)$  and  $r_i = \mathbf{false}$  do
20:     $r_i \leftarrow \mathbf{true}$ .
21:    for  $v_j$  is a parent of  $v_i$  in  $U$  do
22:       $S_j \leftarrow S_j + a_i$ .
23:    end for
24:  end for
25: end while
26: return  $S$ .
```

---

The marking of the attributes guarantees that there can be no cycles in  $\mathcal{S}$ . In fact, the marking order also tells us a valid order for transmitting the attributes. Further, as at least one attribute is marked at each step, this guarantees that the algorithm terminates in  $K$  steps.

Let  $\mathcal{S}$  be the collection of sources. The following proposition tells us that the augmentation performed by SETPACK does not compromise the optimality of collections next to  $\mathcal{S}$ .

**Proposition 9.** *Assume the collection of sources  $\mathcal{S} = \{S_1, \dots, S_K\}$ . Let  $\mathcal{O} = \{O_1, \dots, O_K\}$  be the collection of sources such that  $S_i \subseteq O_i$  and  $|O_i| \leq |S_i| + 1$ . Let  $\mathcal{S}'$  be the collection that Algorithm 3 produces from  $\mathcal{S}$  in a single step. Then there is a collection  $\mathcal{S}^*$  such that  $S'_i \subseteq S_i^*$  and that  $c(\mathcal{S}^*) \leq c(\mathcal{O})$ .*

*Proof.* Let  $G$  be the graph constructed by Algorithm 3 for the collection  $\mathcal{S}$ . Construct the following graph  $W$ : For each  $O_i$  such that  $O_i = S_i$  add the edge  $(v_i, v_0)$ . For each  $O_i \neq S_i$  add the edge  $(v_i, v_j)$ , where  $\{a_j\} = O_i - S_i$ . But  $W$  is a directed spanning tree of  $G$ . Let  $U$  be the directed minimum spanning tree returned by the algorithm. Let  $S_i^* = S'_i$  if  $(v_i, v_0) \in E(U)$  and  $S_i^* = S'_i \cup \{a_j\}$  if  $(v_i, v_j) \in E(U)$ . Note that  $\mathcal{S}^*$  defines a valid model and because  $U$  is optimal we must have  $c(\mathcal{S}^*) \leq c(\mathcal{O})$ .  $\square$

**Corollary 10.** *Assume that  $\mathcal{F}$  is a family of itemsets having 2 items, at maximum. The algorithm SETPACK returns the optimal tree model.*

Let us consider the complexity of the algorithms. The algorithm SETPACK runs in a polynomial time. By using dynamic programming we can show that GENERATE runs in  $O(|F|^2)$  time. We also tested a faster variant of the algorithm in which the exhaustive search in GENERATE is replaced by the greedy approach similar to the ID3 algorithm. We call this variant SETPACKGREEDY.

## 6 Related Work

Finding interesting itemsets is a major research theme in data mining. To this end, many measures have been suggested over time. A classic measure for ranking itemsets is frequency, for which there exist efficient search algorithms [2, 15]. Other measures involve comparing how much an itemset deviates from the independence assumption [1, 3, 4, 11]. In yet other approaches more flexible models are used, such as, Bayes networks [17, 18], Maximum Entropy estimates [24, 31]. Related are also low-entropy sets: itemsets for which the entropy of the data is low [16].

Many of these approaches suffer from the fact that they require a user-defined threshold and further that at low thresholds extremely many itemsets are returned, many of which convey the same information. To address the latter

problem we can use closed [28] or non-derivable [6] itemsets that provide a concise representation of the original itemsets. However, these methods deteriorate even under small amounts of noise.

Alternative to these approaches of describing the pattern set, there are methods that instead pick groups of itemsets that describe the data well. As such, we are not the first to embrace the compression approach to data mining [12]. Recently, Siebes et al. [30] introduced the MDL-based KRIMP algorithm to battle the frequent itemset explosion at low support thresholds. It returns small subsets of itemsets that together capture the distribution of the data well. These *code tables* have been successfully applied in classification [22], measuring the dissimilarity of data [33], and data generation [34]. While these applications shows the practicality of the approach, KRIMP can only describe the patterns between the items that are present in the dataset. On the other hand, we consider the 0s and the 1s in the data symmetrically and hence we are able to provide more detailed descriptions of the data; including patterns between the presence and absence of items.

More different from our methods are the lossy data description approaches. These strive to describe just part of the data, and as such may overlook important interactions. Summarization [7] is a compression approach that identifies a group of itemsets such that each transaction is summarized by one set with as little loss of information as possible. Yet different are *pattern teams* [20], which are groups of most-informative length- $k$  itemsets [19], selected through an external interestingness measure. As this approach is computationally intensive, the number of team members is typically  $< 10$ . Bringmann et al. [5] proposed a similar selection method that can consider larger pattern sets. However, it also requires the user to choose a quality measure to which the pattern set has to be optimized, unlike our parameter-free and lossless method.

Alternatively we can view the approach in this paper as building a global model for data and then selecting the itemsets that describe the model. This approach then allows us to use MDL as a model selection technique. In a related work [32] the authors build decomposable models in order to select a small family of itemsets that model the data well.

The decision trees returned by our methods, and particularly the DAG that they form, have a passing resemblance to Bayes networks [9]. However, as both the model construction and complexity weighing differ strongly, so do the outcomes. To be more precise, in our case the distributions  $p(x, par(x))$  are modeled and weighted via decision trees whereas in the Bayes network setup any distribution is weighted equally. Furthermore, we use the correspondence between the itemsets and the decision trees to output local patterns, as opposed to Bayes networks which are traditionally used as global models.

## 7 Experiments

This section contains the results of the empirical evaluation of our methods using toy and real datasets.

### 7.1 Datasets

For the experimental validation of the two packing strategies we use a group of datasets with strongly differing statistics. From the LUCS/KDD repository [8] we took a number of often used databases to allow for comparison to other methods. To test our methods on real data we used the Mammals presence database and the Helsinki CS-courses dataset. The latter contains the enrollment records of students taking courses at the Department of Computer Science of the University of Helsinki. The *mammals* dataset consists of the absence/presence of European mammals [25] in geographical areas of 50x50 kilometers.<sup>1</sup> The details of these datasets are provided in Table 1.

**Table 1. Statistics of the datasets used in the experiments.**

<i>Dataset</i>	$ D $	$K$	% of 1's
<i>anneal</i>	898	71	20.1
<i>breast</i>	699	16	62.4
<i>courses</i>	3506	98	4.6
<i>mammals</i>	2183	40	46.9
<i>mushroom</i>	8124	119	19.3
<i>nursery</i>	12960	32	28.1
<i>pageblocks</i>	5473	44	25.0
<i>tic-tac-toe</i>	958	29	34.5

### 7.2 Experiments with Toy Datasets

To evaluate whether our method correctly identifies (in)dependencies, we start our experimentation using two artificial datasets of 2000 transactions and 10 items. For both databases, the data is generated per transaction, and the presence of the first item is based on a fair coin toss. For the first database, the other items are similarly generated. However, for the second database, the presence of an item is 90% dependent on the previous item. As such, both datasets have item densities of about 50%.

If we apply GREEDYPACK, our greedy decision tree building method, to these datasets we see that it is unable to compress the independent database at all. Opposing, the dependently generated dataset can be compressed into only

<sup>1</sup>The full version of the dataset is available for research purposes upon request, <http://www.european-mammals.org>.

**Table 2. Compression, number of trees and numbers of extracted itemsets for the greedy algorithm.**

Dataset	GREEDYPACK					KRIMP			
	$c(\mathcal{T}_b)$ (bits)	$c(\mathcal{T})$ (bits)	$\frac{c(\mathcal{T})}{c(\mathcal{T}_b)}$ (%)	# trees	# sets	min-sup	# sets	# bits	ratio (%)
<i>anneal</i>	23104	12342	<b>53.4</b>	71	1203	1	102	22154	34.6
<i>breast</i>	8099	2998	<b>37.0</b>	16	17	1	30	4613	16.9
<i>courses</i>	76326	61685	<b>80.8</b>	98	1230	2	148	71019	79.3
<i>mammals</i>	78044	50068	<b>64.2</b>	40	845	200	254	90192	42.3
<i>mushroom</i>	442062	115347	<b>26.1</b>	119	999	1	424	231877	20.9
<i>nursery</i>	337477	180803	<b>53.6</b>	32	3409	1	260	258898	45.5
<i>pageblocks</i>	15280	7611	<b>49.8</b>	44	219	1	53	10911	5.0
<i>tic-tac-toe</i>	25123	14137	<b>56.3</b>	29	619	1	162	28812	62.3

50% of the original number of bits. Inspection of the resulting itemsets show that the resulting model correctly describes the dependencies in detail: The resulting 19 itemsets are  $\{a_1, \dots, a_{10}, a_1a_2, \dots, a_9a_{10}\}$ .

### 7.3 The Greedy Method

Recall that our goal is to find high quality descriptions of the data. Following the MDL principle, the quality of the found descriptions can objectively be measured by the compression of the data. We present the compressed sizes for GREEDYPACK in Table 2. The encoding costs  $c(\mathcal{T})$  include the size of the encoded data and the decision trees. The initial costs, as denoted by  $c(\mathcal{T}_b)$ , are those of encoding the data using naïve single-node TRIVIALTREES. Each of these experiments required 1–10 seconds runtime, with an exception of 60s for *mushroom*.

From Table 2, we see that all models returned by GREEDYPACK strongly reduce the number of bits required to describe the data; this implicitly shows that good models are returned. The quality can be gauged by taking the compression ratios into account. In general, our greedy method reduces the number of bits to only half of what the independent model requires. As two specific examples of the found dependencies, in the *courses* dataset the course *Data Mining* was packed using *Machine Learning*, *Software Engineering*, *Information Retrieval Methods* and *Data Warehouses*. Likewise, *AI* and *Machine Learning* were used to pack the *Robotics* course.

Like discussed above, our approach and the KRIMP [30] algorithm have stark differences in what part of the data is considered. However, as both methods use compression, and result good itemsets, it is insightful to compare the algorithms. For the latter we here allow it to compress as well as possible, and thus, consider candidates up to as low min-sup thresholds as feasible.

Let us compare between the outcomes of either method. For KRIMP these are itemsets, for ours it is the combina-

tion of the decision trees and the related itemsets. We see that KRIMP typically returns fewer itemsets than GREEDYPACK. However, our method returns itemsets that describe interactions between both present *and* absent items.

Next, we observed that especially the initial KRIMP compression requires many more bits than ours, and as such KRIMP attains better compression ratios. However, if we disregard the ratios and look at the raw number of bits the two methods require, we see that KRIMP generally requires twice as many bits to describe *only* the 1’s in the data than GREEDYPACK does to represent *all* of the data.

### 7.4 Validation through Classification

To further assess the quality of our models we use a simple classification scheme [22]. First, we split the training database into separate class-databases. We pack each of these. Next, the class labels of the unseen transactions were assigned according to the model that compressed it best.

We ran these experiments for three databases, viz. *mushroom*, *breast* and *anneal*. A random 90% of the data was used to train the models, leaving 10% to test the accuracy on. The accuracy scores we noted, resp. 100%, 98.0% and 93.4%, are fully comparable to (and for the second, even better than) the classifiers considered in [22].

### 7.5 Choosing Good Itemsets

In this subsection we evaluate SETPACK, our itemset selection algorithm. Recall that this algorithm selects itemsets such that they allow for building succinct encoding decision trees. The difference with GREEDYPACK is that in this setup the resulting itemsets should be a subset of a given candidate family. Here, we consider frequent itemsets as candidates. We set the support threshold such that the experiments with SETPACK were finished within  $\frac{1}{2}$ –2 hours, with an exception of 23 hours for considering the



**Table 3. Compressed sizes and number of extracted itemsets for the itemset selection algorithms.**

Dataset	Candidate Itemsets		SETPACK			SETPACKGREEDY			KRIMP	
	min-sup	# sets	$c(\mathcal{T})$	$\frac{c(\mathcal{T})}{c(\mathcal{T}_b)}(\%)$	# sets	$c(\mathcal{T})$	$\frac{c(\mathcal{T})}{c(\mathcal{T}_b)}(\%)$	# sets	# bits	# sets
<i>anneal</i>	175	8837	20777	<b>89.9</b>	103	20781	<b>89.9</b>	69	31196	53
<i>breast</i>	1	9920	5175	<b>63.7</b>	42	5172	<b>63.9</b>	49	4613	30
<i>courses</i>	55	5030	64835	<b>84.9</b>	268	64937	<b>85.1</b>	262	73287	93
<i>mammals</i>	700	7169	65091	<b>83.4</b>	427	65622	<b>84.1</b>	382	124737	125
<i>mushroom</i>	1000	123277	313428	<b>70.9</b>	636	262942	<b>59.5</b>	1225	474240	140
<i>nursery</i>	50	25777	314081	<b>93.0</b>	276	314295	<b>93.1</b>	218	265064	225
<i>pageblocks</i>	1	63599	11961	<b>78.3</b>	92	11967	<b>78.3</b>	95	10911	53
<i>tic-tac-toe</i>	7	34019	23118	<b>92.0</b>	620	23616	<b>94.0</b>	277	28957	159

large candidate family for *mushroom*. For comparison we use the same candidates for KRIMP. We also compare to SETPACKGREEDY, which required 1–12 minutes, 7 minutes typically, with an exception of 2½ hours for *mushroom*.

Comparing the results of this experiment (Table 3) with the results of GREEDYPACK in the previous experiment, we see that the selection process is more strict: now even fewer itemsets are regarded as interesting enough. Large candidate collections are strongly reduced in number: up to three orders of magnitude. On the other hand, the compression ratios are still very good. The reason that GREEDYPACK produces smaller compression ratios is because it is allowed to consider any itemset.

Further, the fact alone that even with this very strict selection the compression ratios are generally well below 90% show that these few sets are indeed of high importance to describing the major interactions in the data.

If we compare the number of selected sets to KRIMP, we see that our method returns in the same order as many itemsets. These descriptions require far less bits than those found by KRIMP. As such, ours are a better approximation of the Kolmogorov complexity of the data.

Between SETPACK and SETPACKGREEDY the outcomes are very much alike; this goes for both the obtained compression as well as the number of returned itemsets. However, the greedy search of SETPACKGREEDY allows for much shorter running times.

## 8 Discussion

The experimentation on our methods validates the quality of the returned models. The models correctly detect dependencies in the data while ignoring independencies. Only a small number of itemsets is returned, which are shown to provide strong compression of the data. By the MDL principle we then know these describes all important regularities in the data distribution in detail efficiently and without

redundancy. This claim is further supported by the high classification accuracies our models achieve.

The GREEDYPACK algorithm generally uses more itemsets and obtains better packing ratios than SETPACK. While GREEDYPACK is allowed to use any itemset, SETPACK may only use frequent itemsets. This suggests that we may be able to achieve better ratios if we use different candidates, for example, low-entropy sets [16].

The running times of the experiments reported in this work range from seconds to hours and depend mainly on the number of attributes and rows of the datasets. The exhaustive version SETPACK may be slow on very large candidate sets, however, the greedy version SETPACKGREEDY can even handle such families well. Considering that our current implementation is rather naïve and the fact that both methods are easily parallelized, both GREEDYPACK and SETPACKGREEDY are suited for the analysis of large databases.

The main outcomes of our models are the itemsets that identify the encoding paths. However, the decision trees from which these sets are extracted can also be regarded as interesting as these provide an easily interpretable view on the major interactions in the data. Further, just considering the attributes used in such a tree as an itemset also allows for simple inspection of the main associations.

In this work we employ the MDL criterion to identify the optimal model. Alternatively, one could consider using either BIC or AIC, both of which can easily be applied to judge between our decision tree-based models.

## 9 Conclusions

In this paper we presented two methods that find compact sets of high quality itemsets. Both methods employ compression to select the group of patterns that describe *all* interactions in the data best. That is, the data is considered symmetric and thus both the 0s and 1s are taken into account in these descriptions. Experimentation with our methods

showed that high quality models are returned. Their compact size, typically tens to thousands of itemsets, allow for easy further analysis of the found interactions.

## References

- [1] C. C. Aggarwal and P. S. Yu. A new framework for itemset generation. In *Proceedings of the ACM SIGACT-SIGMOD-SIGART symposium on Principles of Database Systems (PODS)*, 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, 1996.
- [3] S. Brin, R. Motwani, and C. Silverstein. Beyond market baskets: Generalizing association rules to correlations. In *ACM SIGMOD International Conference on Management of Data*, pages 265–276. ACM Press, 1997.
- [4] S. Brin, R. Motwani, J. D. Ullman, and S. Tsur. Dynamic itemset counting and implication rules for market basket data. In *ACM SIGMOD International Conference on Management of Data*, pages 255–264, 1997.
- [5] B. Bringmann and A. Zimmermann. The chosen few: On identifying valuable patterns. In *IEEE International Conference on Data Mining (ICDM)*, pages 63–72, 2007.
- [6] T. Calders and B. Goethals. Mining all non-derivable frequent itemsets. In *Proceedings of the 6th European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, pages 74–85, 2002.
- [7] V. Chandola and V. Kumar. Summarization - compressing data into an informative representation. In *Proceedings of the IEEE Conference on Data Mining*, pages 98–105, 2005.
- [8] F. Coenen. The LUCS-KDD discretised/normalised ARM and CARM data library. 2003.
- [9] G. F. Cooper and E. Herskovits. A Bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9:309–347, 1992.
- [10] T. Cover and J. Thomas. *Elements of Information Theory*, 2nd ed. John Wiley and Sons, 2006.
- [11] W. DuMouchel and D. Pregibon. Empirical bayes screening for multi-item associations. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 67–76, 2001.
- [12] C. Faloutsos and V. Megalooikonomou. On data mining, compression and kolmogorov complexity. In *Data Mining and Knowledge Discovery*, volume 15, pages 3–20. Springer, 2007.
- [13] P. D. Grünwald. *The Minimum Description Length Principle*. MIT Press, 2007.
- [14] J. Han, H. Cheng, D. Xin, and X. Yan. Frequent pattern mining: Current status and future directions. In *Data Mining and Knowledge Discovery*, volume 15. Springer, 2007.
- [15] J. Han and J. Pei. Mining frequent patterns by pattern-growth: methodology and implications. *SIGKDD Explorations Newsletter*, 2(2):14–20, 2000.
- [16] H. Heikinheimo, E. Hinkkanen, H. Mannila, T. Mielikäinen, and J. K. Seppänen. Finding low-entropy sets and trees from binary data. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 350–359, 2007.
- [17] S. Jaroszewicz and T. Scheffer. Fast discovery of unexpected patterns in data, relative to a bayesian network. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 118–127, 2005.
- [18] S. Jaroszewicz and D. A. Simovici. Interestingness of frequent itemsets using bayesian networks as background knowledge. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 178–186, 2004.
- [19] A. J. Knobbe and E. K. Y. Ho. Maximally informative k-itemsets and their efficient discovery. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 237–244, 2006.
- [20] A. J. Knobbe and E. K. Y. Ho. Pattern teams. In *Proceedings of the 10th European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, pages 577–584, 2006.
- [21] P. Kontkanen and P. Myllymäki. A linear-time algorithm for computing the multinomial stochastic complexity. *Information Processing Letters*, 103(6):227–233, 2007.
- [22] M. van Leeuwen, J. Vreeken, and A. Siebes. Compression picks the item sets that matter. In *Proceedings of the 10th European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, pages 585–592, 2006.
- [23] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and its Applications*. Springer-Verlag, 1993.
- [24] R. Meo. Theory of dependence values. *ACM Trans. Database Syst.*, 25(3):380–406, 2000.
- [25] A. J. Mitchell-Jones, G. Amori, W. Bogdanowicz, B. Krystufek, P. J. H. Reijnders, F. Spitzenberger, M. Stubbe, J. B. M. Thissen, V. Vohralik, and J. Zima. *The Atlas of European Mammals*. Academic Press, 1999.
- [26] K. V. S. Murthy. *On growing better decision trees from data*. PhD thesis, Johns Hopkins Univ., Baltimore, 1996.
- [27] S. Nijssen and É. Fromont. Mining optimal decision trees from itemset lattices. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 530–539, 2007.
- [28] N. Pasquier, Y. Bastide, R. Taouil, and L. Lakhal. Discovering frequent closed itemsets for association rules. *Lecture Notes in Computer Science*, 1540:398–416, 1999.
- [29] J. Rissanen. Fisher information and stochastic complexity. *IEEE Transactions on Information Theory*, 42(1):40–47, 1996.
- [30] A. Siebes, J. Vreeken, and M. van Leeuwen. Item sets that compress. In *Proceedings of the SIAM Conference on Data Mining*, pages 393–404, 2006.
- [31] N. Tatti. Maximum entropy based significance of itemsets. *Knowledge and Information Systems (KAIS)*, 2008. Accepted for publication.
- [32] N. Tatti and H. Heikinheimo. Decomposable families of itemsets. In *Proceedings of the 12th European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2008.
- [33] J. Vreeken, M. van Leeuwen, and A. Siebes. Characterising the difference. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 765–774, 2007.
- [34] J. Vreeken, M. van Leeuwen, and A. Siebes. Preserving privacy through data generation. In *Proceedings of the IEEE Conference on Data Mining*, pages 685–690, 2007.