

# Clustering Boolean Tensors

Saskia Metzler · Pauli Miettinen

the date of receipt and acceptance should be inserted later

**Abstract** Graphs – such as friendship networks – that evolve over time are an example of data that are naturally represented as binary tensors. Similarly to analysing the adjacency matrix of a graph using a matrix factorization, we can analyse the tensor by factorizing it. Unfortunately, tensor factorizations are computationally hard problems, and in particular, are often significantly harder than their matrix counterparts. In case of Boolean tensor factorizations – where the input tensor and all the factors are required to be binary and we use Boolean algebra – much of that hardness comes from the possibility of overlapping components. Yet, in many applications we are perfectly happy to partition at least one of the modes. For instance, in the aforementioned time-evolving friendship networks, groups of friends might be overlapping, but the time points at which the network was captured are always distinct. In this paper we investigate what consequences this partitioning has on the computational complexity of the Boolean tensor factorizations and present a new algorithm for the resulting clustering problem. This algorithm can alternatively be seen as a particularly regularized clustering algorithm that can handle extremely high-dimensional observations. We analyse our algorithm with the goal of maximizing the similarity and argue that this is more meaningful than minimizing the dissimilarity. As a by-product we obtain a PTAS and an efficient 0.828-approximation algorithm for rank-1 binary factorizations. Our algorithm for Boolean tensor clustering achieves high scalability, high similarity, and good generalization to unseen data with both synthetic and real-world data sets.

## 1 Introduction

Tensors, multi-way generalizations of matrices, are becoming more common in data mining. Binary 3-way tensors, for example, can be used to record ternary relations (e.g. source IP–target IP–target port in network traffic analysis), a set of different relations between the same entities (e.g. subject–relation–object data

---

S. Metzler · P. Miettinen  
Max-Planck-Institut für Informatik, Saarbrücken, Germany  
E-mail: {smetzler,pmiettinen}@mpi-inf.mpg.de

for information extraction), or different instances of the same relation between the same entities (e.g. the adjacency matrices of a dynamic graph over time, such as who sent email to whom in which month). Given such data, a typical goal in data mining is to find structure and regularities from it in order to create an informative summary that describes the data well. To that end, tensor decomposition methods are often employed, and if the data is binary, it is natural to restrict also the factors to be binary. This gives rise to the various Boolean tensor decompositions that have recently been studied (e.g. Miettinen 2011; Erdős and Miettinen 2013b; Cerf et al 2013).

Finding good Boolean tensor decompositions is, however, computationally hard. Broadly speaking, this hardness stems from the complexity caused by the overlapping factors. But when the existing Boolean tensor decomposition algorithms are applied to real-world data sets, they often return factors that are (almost) non-overlapping in at least one mode (see, e.g. Erdős and Miettinen 2013a,b). This means, the respective factor matrix has exactly one non-zero per row. Typically the mode where these non-overlapping factors appear is unsurprising given the data: in the above examples, it would be target ports in the network traffic analysis data, relations in the information extraction data, and time (month) in the email data – in all of these cases, this mode has somewhat different behavior compared to the other two.

These observations lead to the question: what if we restrict one mode to non-overlapping factors? This removes one of the main sources of the computational complexity – the overlapping factors – from one mode, so we would expect the problem to admit better approximability results. We would also expect the resulting algorithms to be simpler. Would the computationally simpler problem transfer to better results? In this paper we study all these questions and, perhaps surprisingly, give an affirmative yes answer to all of them.

If we accept the argument that clustering one mode of the input data is preferable (and we emphasize that this is highly application-dependent), the next question is why do we not simply cluster the slices of the tensor using any off-the-shelf clustering algorithm. Our answer to that is two-fold: On one hand, having binary centroids to express binary data often improves the interpretability (and sparsity) of the results, and hence we want to develop specific algorithms for the Boolean case. On the other hand, just clustering the slices will not reveal the underlying structure *within* the slices: extracting that structure facilitates the interpretation of the clustering, and will also alleviate the curse of dimensionality, as we shall see in the experiments. With Boolean tensor clustering, we seek for a concise representation of the major patterns in the data. Given the example of subject–relation–object data, this would amount to finding clusters of relations and describing them by the most important subjects and objects for each cluster.

The general form of Boolean tensor clustering (Section 3.1) is in fact very similar to standard binary clustering problems, and therefore we concentrate on the more restricted problem of Boolean CP clustering (Section 3.2). While we approach these problems from the point of view of tensor decompositions, they can be equivalently seen from the clustering point of view (Section 3.3): The task then is to find a clustering while the cluster centroids are constrained to comply with the tensor structure. And while we concentrate on 3-way tensors, our approach can be extended to higher-order Boolean tensors (Section 3.4), although it involves certain application-specific design decisions.

Normally, the quality of a clustering or tensor decomposition is measured using the distance of the original data to its representation. In Section 3.5 we argue, however, that for many data mining problems (especially with binary data), measuring the quality using the similarity instead of distance leads to a more meaningful analysis. Motivated by this, the goal of our algorithms is to maximize the similarity (Section 4.2) instead of minimizing the dissimilarity. As a by-product of the analysis of our algorithms, we also develop and analyse algorithms for maximum-similarity binary rank-1 decompositions in Section 4.3. We show that the problem admits a polynomial-time approximation scheme (PTAS) and present a scalable algorithm that achieves a  $2(\sqrt{2} - 1) \approx 0.828$  approximation ratio. The experimental evaluation, in Section 5, shows that our algorithm achieves comparable representations of the input tensor when compared to methods that do (Boolean or normal) tensor CP decompositions – even when our algorithm is solving a more restricted problem – while being generally faster and admitting good generalization to unseen data.

## 2 Preliminaries

We provide here the basic notation and concepts used throughout this paper. For more in-depth treatise of general tensor operations and decompositions, we refer the reader to Kolda and Bader (2009), and to Miettinen (2011) for specifics of Boolean tensors.

Throughout this paper, we indicate vectors as bold lower-case letters ( $\mathbf{v}$ ), matrices as bold upper-case letters ( $\mathbf{M}$ ), and tensors as bold upper-case calligraphic letters ( $\mathcal{T}$ ). Element  $(i, j, k)$  of a 3-way tensor  $\mathcal{X}$  is denoted as  $x_{ijk}$ . A colon in a subscript denotes taking that mode entirely; for example,  $\mathcal{X}_{::k}$  is the  $k$ th *frontal slice* of  $\mathcal{X}$  ( $\mathbf{X}_k$  in short).

Fibers of a tensor are the higher order analogues of matrix rows and columns. They are obtained by fixing every index but one; for example,  $\mathbf{x}_{:jk}$  is the  $j$ th column of  $k$ th frontal slice of  $\mathcal{X}$ . An  $n$ -way tensor can be *unfolded* (or *matricized*) into a matrix by arranging its fibers (i.e. its columns, rows, or tubes in case of a 3-way tensor) as columns of a matrix. For a *mode- $n$  matricization*, the mode- $n$  fibers are used as the columns and the result is denoted by  $\mathbf{X}_{(n)}$ .

The outer product of vectors is denoted by  $\boxtimes$ . For vectors  $\mathbf{a}$ ,  $\mathbf{b}$ , and  $\mathbf{c}$  of length  $n$ ,  $m$ , and  $l$ ,  $\mathcal{X} = \mathbf{a} \boxtimes \mathbf{b} \boxtimes \mathbf{c}$  is an  $n$ -by- $m$ -by- $l$  tensor with  $x_{ijk} = a_i b_j c_k$ .

For a tensor  $\mathcal{X}$ ,  $|\mathcal{X}|$  denotes its number of non-zero elements. The Frobenius norm of a 3-way tensor  $\mathcal{X}$  is  $\|\mathcal{X}\| = \sqrt{\sum_{i,j,k} x_{ijk}^2}$ . If  $\mathcal{X}$  is binary,  $|\mathcal{X}| = \|\mathcal{X}\|^2$ . The *similarity* between two  $n$ -by- $m$ -by- $l$  binary tensors  $\mathcal{X}$  and  $\mathcal{Y}$  is defined as

$$\text{sim}(\mathcal{X}, \mathcal{Y}) = nml - |\mathcal{X} - \mathcal{Y}|. \quad (1)$$

The *Boolean tensor sum* of binary tensors  $\mathcal{X}$  and  $\mathcal{Y}$  is defined as  $(\mathcal{X} \vee \mathcal{Y})_{ijk} = x_{ijk} \vee y_{ijk}$ . For binary matrices  $\mathbf{X}$  and  $\mathbf{Y}$  where  $\mathbf{X}$  has  $r$  columns and  $\mathbf{Y}$  has  $r$  rows their *Boolean matrix product*,  $\mathbf{X} \circ \mathbf{Y}$ , is defined as  $(\mathbf{X} \circ \mathbf{Y})_{ij} = \bigvee_{k=1}^r x_{ik} y_{kj}$ . The *Boolean matrix rank* of a binary matrix  $\mathbf{A}$  is the least  $r$  such that there exists a pair of binary matrices  $(\mathbf{X}, \mathbf{Y})$  of inner dimension  $r$  with  $\mathbf{A} = \mathbf{X} \circ \mathbf{Y}$ .

A binary matrix  $\mathbf{X}$  is a *cluster assignment matrix* if each row of  $\mathbf{X}$  has exactly one non-zero element. In that case the Boolean matrix product corresponds to the

regular matrix product,

$$\mathbf{X} \circ \mathbf{Y} = \mathbf{X}\mathbf{Y}. \quad (2)$$

We can now define the Boolean tensor rank and two decompositions: Boolean CP and Tucker3.

**Definition 1** The *Boolean tensor rank* of a 3-way binary tensor  $\mathcal{X}$ ,  $\text{rank}_B(\mathcal{X})$ , is the least integer  $r$  such that there exist  $r$  triplets of binary vectors  $(\mathbf{a}_i, \mathbf{b}_i, \mathbf{c}_i)$  with  $\mathcal{X} = \bigvee_{i=1}^r \mathbf{a}_i \boxtimes \mathbf{b}_i \boxtimes \mathbf{c}_i$ .

The low-rank Boolean tensor CP decomposition is defined as follows.

**Problem 1 (Boolean CP)** Given an  $n$ -by- $m$ -by- $l$  binary tensor  $\mathcal{X}$  and an integer  $r$ , find binary matrices  $\mathbf{A}$  ( $n$ -by- $r$ ),  $\mathbf{B}$  ( $m$ -by- $r$ ), and  $\mathbf{C}$  ( $l$ -by- $r$ ) such that they minimize  $|\mathcal{X} - \bigvee_{i=1}^r \mathbf{a}_i \boxtimes \mathbf{b}_i \boxtimes \mathbf{c}_i|$ .

The standard (non-Boolean) tensor rank and CP decomposition are defined analogously (Kolda and Bader 2009). Both finding the least error Boolean CP decomposition and deciding the Boolean tensor rank are NP-hard (Miettinen 2011). Following Kolda and Bader (2009), we use  $[[\mathbf{A}, \mathbf{B}, \mathbf{C}]]$  to denote the normal 3-way CP decomposition and  $[[\mathbf{A}, \mathbf{B}, \mathbf{C}]]_B$  for the Boolean CP decomposition.

Let  $\mathbf{X}$  be  $n_1$ -by- $m_1$  and  $\mathbf{Y}$  be  $n_2$ -by- $m_2$  matrix. Their *Kronecker (matrix) product*,  $\mathbf{X} \otimes \mathbf{Y}$ , is the  $n_1 n_2$ -by- $m_1 m_2$  matrix defined by

$$\mathbf{X} \otimes \mathbf{Y} = \begin{pmatrix} x_{11}\mathbf{Y} & x_{12}\mathbf{Y} & \cdots & x_{1m_1}\mathbf{Y} \\ x_{21}\mathbf{Y} & x_{22}\mathbf{Y} & \cdots & x_{2m_1}\mathbf{Y} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n_1 1}\mathbf{Y} & x_{n_1 2}\mathbf{Y} & \cdots & x_{n_1 m_1}\mathbf{Y} \end{pmatrix}.$$

The *Khatri–Rao (matrix) product* of  $\mathbf{X}$  and  $\mathbf{Y}$  is defined as “column-wise Kronecker”. That is,  $\mathbf{X}$  and  $\mathbf{Y}$  must have the same number of columns ( $m_1 = m_2 = m$ ), and their Khatri–Rao product  $\mathbf{X} \odot \mathbf{Y}$  is the  $n_1 n_2$ -by- $m$  matrix defined as

$$\mathbf{X} \odot \mathbf{Y} = (\mathbf{x}_1 \otimes \mathbf{y}_1, \mathbf{x}_2 \otimes \mathbf{y}_2, \dots, \mathbf{x}_m \otimes \mathbf{y}_m). \quad (3)$$

Notice that if  $\mathbf{X}$  and  $\mathbf{Y}$  are binary, so are  $\mathbf{X} \otimes \mathbf{Y}$  and  $\mathbf{X} \odot \mathbf{Y}$ .

We can write the Boolean CP decomposition equivalently as matrices in three different ways using unfolding and matrix products:

$$\begin{aligned} \mathbf{X}_{(1)} &\approx \mathbf{A} \circ (\mathbf{C} \odot \mathbf{B})^T, \\ \mathbf{X}_{(2)} &\approx \mathbf{B} \circ (\mathbf{C} \odot \mathbf{A})^T, \\ \mathbf{X}_{(3)} &\approx \mathbf{C} \circ (\mathbf{B} \odot \mathbf{A})^T. \end{aligned} \quad (4)$$

The Boolean Tucker3 decomposition can be seen as a generalization of the Boolean CP decomposition:

**Problem 2 (Boolean Tucker3)** Given an  $n$ -by- $m$ -by- $l$  binary tensor  $\mathcal{X}$  and three integers  $r_1$ ,  $r_2$ , and  $r_3$ , find a binary  $r_1$ -by- $r_2$ -by- $r_3$  core tensor  $\mathcal{G}$  and binary factor matrices  $\mathbf{A}$  ( $n$ -by- $r_1$ ),  $\mathbf{B}$  ( $m$ -by- $r_2$ ), and  $\mathbf{C}$  ( $l$ -by- $r_3$ ) such that they minimize

$$\left| \mathcal{X} - \bigvee_{\alpha=1}^{r_1} \bigvee_{\beta=1}^{r_2} \bigvee_{\gamma=1}^{r_3} g_{\alpha\beta\gamma} \mathbf{a}_\alpha \boxtimes \mathbf{b}_\beta \boxtimes \mathbf{c}_\gamma \right|. \quad (5)$$

We use  $[[\mathcal{G}; \mathbf{A}, \mathbf{B}, \mathbf{C}]]_B$  as the shorthand notation for the Boolean Tucker3 decomposition.

### 3 Problem Definitions

We consider the variation of *tensor clustering* where the idea is to cluster one mode of a tensor and potentially reduce the dimensionality of the other modes. Another common approach is to do the co-clustering equivalent, that is, to cluster each mode of the tensor simultaneously. The former is the approach taken, for example, by Huang et al (2008), while the latter appears for instance in Jegelka et al (2009). Unlike either of these methods, however, we concentrate on binary data endowed with the Boolean algebra.

We will present our problems in the format that leaves the exact measure of goodness of the approximation open. In Section 3.5, we will argue that the similarity – instead of the more common dissimilarity – should be used as the measure of goodness. We will give our definitions for 3-way tensors first, provide alternative definitions as clustering tasks in Section 3.3, and discuss extensions to higher-order tensors in Section 3.4.

#### 3.1 General Problem

Assuming a 3-way tensor and that we do the clustering in the last mode, we can express the general *Boolean tensor clustering* (BTC) problem as follows:

**Problem 3 (BTC)** Given a binary  $n$ -by- $m$ -by- $l$  tensor  $\mathcal{X}$  and integers  $r_1$ ,  $r_2$ , and  $r_3$ , find a  $r_1$ -by- $r_2$ -by- $r_3$  binary tensor  $\mathcal{G}$  and matrices  $\mathbf{A} \in \{0, 1\}^{n \times r_1}$ ,  $\mathbf{B} \in \{0, 1\}^{m \times r_2}$ , and  $\mathbf{C} \in \{0, 1\}^{l \times r_3}$  such that  $\mathbf{C}$  is a cluster assignment matrix and that  $[[\mathcal{G}; \mathbf{A}, \mathbf{B}, \mathbf{C}]]_B$  is a good Boolean Tucker decomposition of  $\mathcal{X}$ .

If we let  $r_1 = n$  and  $r_2 = m$ , we obtain essentially a traditional (binary) clustering problem: Given a set of  $l$  binary matrices  $\{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_l\}$  (the frontal slices of  $\mathcal{X}$ ), cluster these matrices into  $r_3$  clusters, each represented by an  $n$ -by- $m$  binary matrix  $\mathbf{G}_i$  (the frontal slices of the core tensor  $\mathcal{G}$ ); the factor matrices  $\mathbf{A}$  and  $\mathbf{B}$  can be left as identity matrices. We call this problem *Unconstrained BTC*. Writing each of the involved matrices as an  $nm$ -dimensional vector, we obtain the following problem, called the *Discrete Basis Partitioning problem* (DBPP) in (Miettinen et al 2008):

**Problem 4 (DBPP, Miettinen et al 2008)** Given an  $l$ -by- $nm$  binary matrix  $\mathbf{X}$  and a positive integer  $r$ , find matrices  $\mathbf{C} \in \{0, 1\}^{l \times r}$  and  $\mathbf{G} \in \{0, 1\}^{r \times nm}$  such that  $\mathbf{C}$  is a cluster assignment matrix and  $\mathbf{C}$  and  $\mathbf{G}$  minimize  $|\mathbf{X} - \mathbf{C}\mathbf{G}|$ .

Miettinen et al (2008) show that DBPP is NP-hard and give a  $(10 + \varepsilon)$ -approximation algorithm that runs in polynomial time w.r.t.  $n$ ,  $m$ ,  $l$ , and  $r$ , while Jiang (2014) gives a 2-approximation algorithm that runs in time  $O(nml^r)$ .

#### 3.2 Boolean CP Clustering

Constraining  $r_1$ ,  $r_2$ , and  $\mathcal{G}$  yields to somewhat different looking problems, though. In what can be seen as the other extreme, we can restrict  $r_1 = r_2 = r_3 = r$  and  $\mathcal{G}$  (which now is  $r$ -by- $r$ -by- $r$ ) to having 1s exactly on the hyperdiagonal to obtain the *Boolean CP clustering* (BCPC) problem:

**Problem 5 (BCPC)** Given a binary  $n$ -by- $m$ -by- $l$  tensor  $\mathcal{X}$  and an integer  $r$ , find matrices  $\mathbf{A} \in \{0, 1\}^{n \times r}$ ,  $\mathbf{B} \in \{0, 1\}^{m \times r}$ , and  $\mathbf{C} \in \{0, 1\}^{l \times r}$  such that  $\mathbf{C}$  is a cluster assignment matrix and that  $[[\mathbf{A}, \mathbf{B}, \mathbf{C}]]_{\mathbf{B}}$  is a good Boolean CP decomposition of  $\mathcal{X}$ .

What BCPC does is perhaps easiest to understand if we use the unfolding rules (4) and write

$$\mathbf{X}_{(3)} \approx \mathbf{C} \circ (\mathbf{B} \odot \mathbf{A})^T, \quad (6)$$

where we can see that compared to the general BTC, we have restricted the type of centroids: each centroid must be a row of type  $(\mathbf{b} \otimes \mathbf{a})^T$ . This restriction plays a crucial role in the decomposition, as we shall see shortly. Notice also that using (2) we can rewrite (6) without the Boolean product:

$$\mathbf{X}_{(3)} \approx \mathbf{C}(\mathbf{B} \odot \mathbf{A})^T. \quad (7)$$

### 3.3 Clustering Perspective

We defined Unconstrained BTC and BCPC as tensor decomposition problems with additional constraints to enforce the clustering. Equivalently, we can define the problems as clustering problems with additional constraints to enforce the tensor structure. We provide these alternative definitions here, as they will make the connections to other clustering problems easier to see.

The Unconstrained BTC is straight forward to express as a clustering problem:

**Problem 6 (Unconstrained BTC, clustering version)** Given a set  $X$  of binary vectors,  $X = \{\mathbf{x}_1, \dots, \mathbf{x}_l\} \subset \{0, 1\}^{nm}$  and an integer  $r$ , partition  $X$  in  $r$  disjoint sets  $C_1, C_2, \dots, C_r$ , with each  $C_i$  associated with a centroid  $\mathbf{g}_i \in \{0, 1\}^{nm}$ , such that the centroids  $\mathbf{g}_i$  are good representations of vectors  $\mathbf{x}_j$  that are in their clusters.

The input vectors  $\mathbf{x}_j \in X$  are simply the rows of matrix  $\mathbf{X}_{(3)}$ , and the assignment of  $\mathbf{x}_j$ s to clusters  $C_i$  defines the matrix  $\mathbf{C} \in \{0, 1\}^{l \times r}$ :  $c_{ij} = 1$  if and only if vector  $\mathbf{x}_j$  is assigned to cluster  $C_i$ . Finally, treating vectors  $\mathbf{g}_1, \dots, \mathbf{g}_r$  as columns of an unfolded tensor  $\mathbf{G}_{(3)} \in \{0, 1\}^{nm \times r}$  gives us the  $m$ -by- $n$ -by- $r$  core tensor  $\mathcal{G}$  that we used earlier to describe Unconstrained BTC.

This formulation of Unconstrained BTC makes its connection to standard partition clustering problems obvious; in particular, if the quality of the result is measured as the sum of Hamming distances from input vectors to their centroids, it is clear that Unconstrained BTC is equivalent to the Binary  $k$ -Median clustering (Miettinen 2009). It is also clear from the formulation that Unconstrained BTC can easily suffer from the curse of dimensionality: a modestly-sized 100-by-100-by-100 tensor yields 100 vectors of dimensionality 10 000 to be clustered.

The BCPC problem imposes an additional constraint on the centroids: they must be of form  $\mathbf{g}_i = (\mathbf{b}_i \otimes \mathbf{a}_i)^T$  for some  $\mathbf{a} \in \{0, 1\}^n$  and  $\mathbf{b} \in \{0, 1\}^m$ .

**Problem 7 (BCPC, clustering version)** Given a set  $X$  of binary vectors,  $X = \{\mathbf{x}_1, \dots, \mathbf{x}_l\} \subset \{0, 1\}^{nm}$  and an integer  $r$ , partition  $X$  in  $r$  disjoint sets  $C_1, C_2, \dots, C_r$ , with each  $C_i$  associated with a centroid  $\mathbf{g}_i = (\mathbf{b}_i \otimes \mathbf{a}_i)^T \in \{0, 1\}^{nm}$ , such that the centroids  $\mathbf{g}_i$  are good representations of vectors  $\mathbf{x}_j$  that are in their clusters.

The vectors  $\mathbf{a}_1, \dots, \mathbf{a}_r$  yield the columns of the  $n$ -by- $r$  factor matrix  $\mathbf{A}$ , and the vectors  $\mathbf{b}_1, \dots, \mathbf{b}_r$  yield the columns of the  $m$ -by- $r$  factor matrix  $\mathbf{B}$  used in the previous section to define BCPC. While the additional constraints might not seem very intuitive in the clustering framework, they do increase the interpretability of the clusters. Also, they can help with the curse of dimensionality: Compared to the unconstrained version, the degree of freedom of the centroids, i.e. the number of possible centroids, is reduced to only those that admit a rank-1 decomposition. Conveniently, the centroids that admit the constraint can also be described in a more compact way than arbitrary centroids.

### 3.4 Extensions to $N$ -Way Tensors

All the above definitions are restricted to 3-way tensors. 3-way tensors are the most fundamental case, and they also allow relatively simple notation. Extending the definitions to  $N$ -way tensors, with  $N > 3$ , involves design decisions that are highly application-specific: for example, how many of the modes should be (co-)clustered and how high order should the centroids have. With 4-way tensors, for instance, we could either co-cluster two modes so that the centroids would still be matrices, or we could cluster just one mode and have the centroids as 3-way tensors.

If we cluster only the last mode of an  $N$ -way tensor  $\mathcal{X}$  and use  $\mathbf{C}$  to denote the cluster assignment matrix and  $\mathbf{A}^{(i)}$  to denote the factor for the  $i$ th mode ( $i = 1, \dots, N - 1$ ), Equation (7) is generalized into

$$\mathbf{X}_{(N)} \approx \mathbf{C}(\mathbf{A}^{(N-1)} \odot \mathbf{A}^{(N-2)} \odot \dots \odot \mathbf{A}^{(1)})^T. \quad (8)$$

If we keep the centroids as matrices and co-cluster the other modes, finding the centroids is similar to the 3-way setting, but the clustering must be changed to some type of tensor co-clustering (see, e.g. Zhao and Zaki 2005).

### 3.5 Similarity vs. Dissimilarity

So far we have avoided defining what we mean by “good” clustering. Arguably the most common approach is to say that any clustering is good if it minimizes the sum of distances (i.e. dissimilarities) between the original elements and their cluster’s representative (or centroid). An alternative is to maximize the similarity between the elements and their representatives: if the data and the centroids are binary, this problem is known as *Hypercube segmentation* (Kleinberg et al 2004). While maximizing similarity is obviously a dual of minimizing the dissimilarity – in the sense that an optimal solution to one is also an optimal solution to the other – the two behave differently when we aim at approximating the optimal solution. For example, for a fixed  $r$ , the best known approximation algorithm to DBPP is the aforementioned 2-approximation algorithm (Jiang 2014), while Alon and Sudakov (1999) gave a PTAS for the Hypercube segmentation problem.

The differences between the approximability, and in particular the reasons behind those differences, are important for data miners. Namely, it is not uncommon that our ability to minimize the dissimilarities is bounded by a family of problems where the optimal solution obtains extremely small dissimilarities, while the best polynomial-time algorithm has to do with slightly larger errors. These larger errors,

however, might still be small enough relative to the data size that we can ignore them. This is what essentially happens when we maximize the similarities.

For example, consider an instance of DBPP where the optimal solution causes ten errors. The best-known algorithm can guarantee to not cause more than twenty errors; assuming the data size is, say, 1000-by-500, making mistakes on twenty elements can still be considered excellent. On the other hand, should the optimal solution make an error on every fourth element, the best approximation could only guarantee to make errors in at most every second element. When considering the similarities instead of dissimilarities, these two setups look very different: In the first, we get within  $(500\,000 - 20)/(500\,000 - 10) = 0.99998$  of the best solution, while in the other, we are only within  $(0.5 \cdot 500\,000)/(500\,000 - 0.25 \cdot 500\,000) \approx 0.67$  of the optimal. Conversely, if the algorithm wants to be within, say, 0.9 of the optimal similarity in the latter setup, it must approximate the dissimilarity within a factor of 1.3 ( $((1-x) \cdot 500\,000)/(500\,000 - 0.25 \cdot 500\,000) = 0.9 \Rightarrow x = 0.325 \Rightarrow 0.325/0.25 \leq 1.3$ ). Hence similarity is less sensitive to small errors (relative to the data size) and more sensitive to large errors, than dissimilarity is. In many applications, this is exactly what we want.

None of this is to say that we should stop considering the error and concentrate only on the similarities. On the contrary: knowing the data size and the error, we can easily compute the similarity for the binary data, and especially when the error is relatively small, it is much easier to compare than the similarity scores. What we do argue, however, is that when analyzing the approximation ratio a data mining algorithm can obtain, similarity can give us a more interesting picture of the actual behavior of the algorithm. Hence, for the rest of the paper, we shall concentrate on the maximum-similarity variants of the problems; most notably, we concentrate on the maximum-similarity BCPC, denoted  $\text{BCPC}_{\max}$ .

#### 4 Solving the $\text{BCPC}_{\max}$

In this section, we present our main algorithm, **SaBoTeur**, for solving the BCPC problem. Our algorithm works in two parts: it first finds a solution to Unconstrained BTC, and then turns that into a valid solution to BCPC. The maximum-similarity Unconstrained BTC is equivalent to Hypercube Segmentation problem, explained in Section 4.1. The overall approach of **SaBoTeur** is explained in Section 4.2, and the problem of turning a solution of Unconstrained BTC to a valid solution of BCPC is further studied in Section 4.3. We present an intuitive iterative updates scheme in Section 4.4, and discuss about the implementation details in Section 4.5. Finally, we consider the problem of selecting the correct number of clusters in Section 4.6.

##### 4.1 The Hypercube Segmentation Problem

Given a tensor  $\mathcal{X}$ , for the optimal solution to  $\text{BCPC}_{\max}$ , we need matrices  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$  that maximize  $\text{sim}(\mathcal{X}_{(3)}, \mathbf{C}(\mathbf{B} \odot \mathbf{A})^T)$ . In the maximum-similarity unconstrained BTC we replace  $\mathbf{B} \odot \mathbf{A}$  with an arbitrary binary matrix  $\mathbf{G}$ . This latter problem is equal to the Hypercube segmentation problem defined by Kleinberg et al (2004):

**Algorithm 1** SaBoTeur algorithm for the BCPC**Input:** 3-way binary tensor  $\mathcal{X}$ , number of clusters  $r$ , number of samples  $s$ .**Output:** Binary factor matrices  $\mathbf{A}$  and  $\mathbf{B}$ , cluster assignment matrix  $\mathbf{C}$ .

---

```

1: function SaBoTeur( $\mathcal{X}, r, s$ )
2:   repeat
3:     Sample  $r$  rows of  $\mathbf{X}_{(3)}$  into matrix  $\mathbf{Y}$ 
4:     Find binary matrices  $\mathbf{A}$  and  $\mathbf{B}$  that maximize  $\text{sim}(\mathbf{Y}, (\mathbf{B} \odot \mathbf{A})^T)$ 
5:     Cluster  $\mathbf{C}$  by assigning each row of  $\mathbf{X}_{(3)}$  to its closest row of  $(\mathbf{B} \odot \mathbf{A})^T$ 
6:   until  $s$  resamples are done
7:   return best  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$ 

```

---

**Problem 8 (Hypercube Segmentation)** Given a set  $S$  of  $l$  vertices of the  $d$ -dimensional cube  $\{0, 1\}^d$ , find  $r$  vertices  $P_1, \dots, P_r \in \{0, 1\}^d$  and a partition of  $S$  into  $r$  segments to maximize  $\sum_{i=1}^r \sum_{c \in S} \text{sim}(P_i, c)$ .

The partition  $S$  defines the cluster assignment matrix  $\mathbf{C}$ , while the vertices  $P_i$  correspond to the rows of  $\mathbf{G}$  (c.f. Problem 6).

Alon and Sudakov (1999) gave a polynomial-time approximation scheme (PTAS) for the Hypercube segmentation problem. The PTAS obtains a similarity within  $(1 - \varepsilon)$  of the optimum with running time of  $e^{O((r^2/\varepsilon^2) \ln r)} nml$  for  $n$ -by- $m$ -by- $l$  data and  $r$  clusters. While technically linear in data size, the first term turns the running time unfeasible even for moderate values of  $r$  and  $\varepsilon$ . Earlier, Kleinberg et al (2004) proposed a sampling-based algorithm that obtains an approximation ratio of  $0.828 - \varepsilon$  with constant probability and running time  $O(nmlrs)$ , where  $s = (9/\varepsilon)^r \ln(1/\varepsilon)$  is the number of times the sampling is done. While the number of samples this algorithm requires is still exponential in  $r$ , in practice we can take  $s = \Theta(1)$  for a linear-time algorithm (with the cost of losing approximation guarantees, though). We will therefore follow a sampling approach similar to Kleinberg et al (2004) in SaBoTeur.

We will also note that both the PTAS of Alon and Sudakov (1999) and the sampling-based algorithm of Kleinberg et al (2004) can be used as-is to solve the maximum-similarity unconstrained BTC problem using the above connection between it and the Hypercube segmentation problem.

## 4.2 The SaBoTeur Algorithm

Our algorithm SaBoTeur (Sampling for Boolean Tensor clustering), Algorithm 1, considers only the unfolded tensor  $\mathbf{X}_{(3)}$ . In each iteration, it samples  $r$  rows of  $\mathbf{X}_{(3)}$  as the initial, unrestricted centroids. It then turns these unrestricted centroids into the restricted type in line 4, and then assigns each row of  $\mathbf{X}_{(3)}$  to its closest restricted centroid. The sampling is repeated multiple times, and in the end, the factors that gave the highest similarity are returned.

The algorithm is extremely simple and fast. The repeated sampling of rows of  $\mathbf{X}_{(3)}$  for (initial) centroids is similar to the aforementioned Hypercube segmentation algorithm of Kleinberg et al (2004), and we can use their analysis to justify why sampling the *rows* of the input matrix is justifiable (instead of finding arbitrary binary vectors). In particular, Kleinberg et al (2004) proved that among the rows of  $\mathbf{X}_{(3)}$  that are in the same optimal cluster, one is a good approximation of the (unrestricted) centroid of the cluster:

**Lemma 1 (Kleinberg et al 2004, Lemma 3.1)** *Let  $\mathbf{X}$  be an  $n$ -by- $m$  binary matrix and let*

$$\mathbf{y}^* = \arg \max_{\mathbf{y} \in \{0,1\}^m} \sum_{i=1}^n \text{sim}(\mathbf{x}_i, \mathbf{y}).$$

*Then there exists a row  $\mathbf{x}_j$  of  $\mathbf{X}$  such that*

$$\sum_{i=1}^n \text{sim}(\mathbf{x}_i, \mathbf{x}_j) \geq (2\sqrt{2} - 2) \sum_{i=1}^n \text{sim}(\mathbf{x}_i, \mathbf{y}^*). \quad (9)$$

That is, if we sample one row from each cluster, we have a high probability of inducing a close-optimal clustering, and with enough re-samples, we have a high probability to have at least one such sample. The crux of BCPChowever is the restriction of the centroids, and we will now turn our attention to the problem of enforcing this restriction in SaBoTeur.

### 4.3 Binary Rank-1 Matrix Decompositions

The next part of the SaBoTeur algorithm is to turn the unrestricted centroids into the restricted form (line 4). We start by showing that this problem is equivalent to finding the maximum-similarity binary rank-1 decomposition of a binary matrix:

**Problem 9 (Binary rank-1 decomposition)** Given an  $n$ -by- $m$  binary matrix  $\mathbf{X}$ , find an  $n$ -dimensional binary vector  $\mathbf{a}$  and an  $m$ -dimensional binary vector  $\mathbf{b}$  that maximize  $\text{sim}(\mathbf{X}, \mathbf{a} \boxtimes \mathbf{b})$ .

**Lemma 2** *Given an  $r$ -by- $nm$  binary matrix  $\mathbf{X}$ , finding  $n$ -by- $r$  and  $m$ -by- $r$  binary matrices  $\mathbf{A}$ ,  $\mathbf{B}$  that maximize  $\text{sim}(\mathbf{X}, (\mathbf{B} \odot \mathbf{A})^T)$  is equivalent to finding the most similar binary rank-1 approximation of each row  $\mathbf{x}$  of  $\mathbf{X}$ , where the rows are re-shaped as  $n$ -by- $m$  binary matrices.*

*Proof* If  $\mathbf{x}_i$  is the  $i$ th row of  $\mathbf{X}$  and  $\mathbf{z}_i$  is the corresponding row of  $(\mathbf{B} \odot \mathbf{A})^T$ , then  $\text{sim}(\mathbf{X}, (\mathbf{B} \odot \mathbf{A})^T) = \sum_{i=1}^k \text{sim}(\mathbf{x}_i, \mathbf{z}_i)$ , and hence we can solve the problem row-by-row. Let  $\mathbf{x} = (x_{1,1}, x_{2,1}, \dots, x_{n,1}, x_{1,2}, \dots, x_{n,m})$  be a row of  $\mathbf{X}$ . Re-write  $\mathbf{x}$  as an  $n$ -by- $m$  matrix  $\mathbf{Y}$ ,

$$\mathbf{Y} = \begin{pmatrix} x_{1,1} & x_{1,2} & \cdots & x_{1,m} \\ x_{2,1} & x_{2,2} & \cdots & x_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n,1} & x_{n,2} & \cdots & x_{n,m} \end{pmatrix}.$$

Consider the row of  $(\mathbf{B} \odot \mathbf{A})^T$  that corresponds to  $\mathbf{x}$ , and notice that it can be written as  $(\mathbf{b} \otimes \mathbf{a})^T$ , where  $\mathbf{a}$  and  $\mathbf{b}$  are the columns of  $\mathbf{A}$  and  $\mathbf{B}$  that correspond to  $\mathbf{x}$ . As  $(\mathbf{b} \otimes \mathbf{a})^T = (b_1 \mathbf{a}^T, b_2 \mathbf{a}^T, \dots, b_m \mathbf{a}^T)$ , re-writing it similarly as  $\mathbf{x}$  we obtain

$$\begin{pmatrix} a_1 b_1 & a_1 b_2 & \cdots & a_1 b_m \\ a_2 b_1 & a_2 b_2 & \cdots & a_2 b_m \\ \vdots & \vdots & \ddots & \vdots \\ a_n b_1 & a_n b_2 & \cdots & a_n b_m \end{pmatrix} = \mathbf{a} \mathbf{b}^T = \mathbf{a} \boxtimes \mathbf{b}.$$

Therefore,  $\text{sim}(\mathbf{x}, (\mathbf{b} \otimes \mathbf{a})^T) = \text{sim}(\mathbf{Y}, \mathbf{a} \boxtimes \mathbf{b})$ . □

**Algorithm 2** Randomized PTAS for maximum-similarity binary rank-1 decompositions**Input:** An  $n$ -by- $m$  binary matrix  $\mathbf{X}$ .**Output:** Binary vectors  $\mathbf{a}$  and  $\mathbf{b}$ .

---

```

1: function PTAS $_{\varepsilon}(\mathbf{X})$ 
2:   Sample  $l = \Theta(\varepsilon^{-2})$  rows of  $\mathbf{X}$  u.a.r. with replacements
3:   for all partitions of the sample into two sets do
4:     Let  $\mathbf{X}'$  contain the rows in the sample
5:     for both sets in the partition do
6:       Let  $\mathbf{a}$  be the incidence vector of the set
7:       Find  $\mathbf{b}^T$  maximizing  $\text{sim}(\mathbf{X}', \mathbf{a}\mathbf{b}^T)$ 
8:       Extend  $\mathbf{a}$  to the rows not in the sample maximizing  $\text{sim}(\mathbf{X}, \mathbf{a}\mathbf{b}^T)$ 
9:   return best vectors  $\mathbf{a}$  and  $\mathbf{b}$ 

```

---

**Algorithm 3** Approximation of maximum-similarity binary rank-1 decompositions**Input:** An  $n$ -by- $m$  binary matrix  $\mathbf{X}$ .**Output:** Binary vectors  $\mathbf{a}$  and  $\mathbf{b}$ .

---

```

1: function A( $\mathbf{X}$ )
2:   for all rows  $\mathbf{x}_i$  of  $\mathbf{X}$  do
3:     Let  $\mathbf{b} = \mathbf{x}_i$ 
4:     Find  $\mathbf{a}$  maximizing  $\text{sim}(\mathbf{X}, \mathbf{a}\mathbf{b}^T)$ 
5:   return best vectors  $\mathbf{a}$  and  $\mathbf{b}$ 

```

---

We start by showing that the maximum-similarity binary rank-1 decomposition admits a PTAS. To that end, we present a family of randomized algorithms that on expectation attain a similarity within  $(1 - \varepsilon)$  of the optimum. The family is similar to that of Alon and Sudakov's (1999) and can be analysed and de-randomized following the techniques presented by them. The family of algorithms (one for each  $\varepsilon$ ) is presented as Algorithm 2.

Lines 7 and 8 require us to solve one of the factor vectors when the other is given. To build  $\mathbf{b}$  (line 7), we simply take the column-wise majority element of those rows where  $\mathbf{a}$  is 1, and we extend  $\mathbf{a}$  similarly in line 8.

The analysis of Algorithm 2 follows closely the proofs by Alon and Sudakov (1999) and is omitted.

The running time of the PTAS algorithm becomes prohibitive even with moderate values of  $\varepsilon$ , and therefore we will not use it in **SaBoTeur**. Instead, we present a simple, deterministic algorithm that approximates the maximum similarity within 0.828, Algorithm 3. It is similar to the algorithm for Hypercube segmentation based on random sampling presented by Kleinberg et al (2004). The algorithm considers every row of  $\mathbf{X}$  as a potential vector  $\mathbf{b}$  and finds the best  $\mathbf{a}$  given  $\mathbf{b}$ . Using Lemma 1 it is straight forward to show that the algorithm achieves the claimed approximation ratio:

**Lemma 3** *Algorithm 3 approximates the optimum similarity within 0.828 in time  $O(nm \min\{n, m\})$ .*

*Proof* To prove the approximation ratio, let  $\mathbf{a}^*(\mathbf{b}^*)^T$  be the optimum decomposition. Consider the rows in which  $\mathbf{a}^*$  has 1. Per Lemma 1, selecting one of these rows, call it  $\mathbf{b}$ , gives us  $\text{sim}(\mathbf{X}, \mathbf{a}^*\mathbf{b}^T) \geq (2\sqrt{2} - 2)\text{sim}(\mathbf{X}, \mathbf{a}^*\mathbf{b}^T)$  (notice that  $\mathbf{a}^*\mathbf{b}^T$  agrees with the optimal solution in rows where  $\mathbf{a}^*$  is zero). Selecting  $\mathbf{a}$  that maximizes the similarity given  $\mathbf{b}$  can only improve the result, and the claim follows as we try every row of  $\mathbf{X}$ .

---

**Algorithm 4** SaBoTeur+itUp algorithm solves BCPC with  $k$ -means-like updates

---

**Input:** 3-way binary tensor  $\mathcal{X}$ , number of clusters  $r$ , number of samples  $s$ .

**Output:** Binary factor matrices  $\mathbf{A}$  and  $\mathbf{B}$ , cluster assignment matrix  $\mathbf{C}$ .

```

1: function SaBoTeur+itUp( $\mathcal{X}, r, s$ )
2:   repeat
3:     Sample  $r$  rows of  $\mathbf{X}_{(3)}$  into matrix  $\mathbf{Y}$ 
4:     repeat
5:       Find binary matrices  $\mathbf{A}$  and  $\mathbf{B}$  that maximize  $\text{sim}(\mathbf{Y}, (\mathbf{B} \odot \mathbf{A})^T)$ 
6:       Cluster  $\mathbf{C}$  by assigning each row of  $\mathbf{X}_{(3)}$  to its closest row of  $(\mathbf{B} \odot \mathbf{A})^T$ 
7:       Recompute centroids  $\mathbf{Y}$  by majority vote from cluster members
8:     until Similarity of  $[[\mathbf{A}, \mathbf{B}, \mathbf{C}]]_B$  to  $\mathcal{X}$  no longer improves
9:   until  $s$  resamples are done
10:  return best  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$ 

```

---

If  $n < m$ , the time complexity follows as for every candidate  $\mathbf{b}$  we have to make one sweep over the matrix. If  $m < n$ , we can operate on the transpose.  $\square$

#### 4.4 Iterative Updates

The SaBoTeur algorithm bears resemblance to the initialization phase of  $k$ -means style clustering algorithms. We propose a variation of SaBoTeur, SaBoTeur+itUp (Algorithm 4), where  $k$ -means style iterative updates are performed on the result of SaBoTeur until it does not improve further. In each iteration of the iterative updates phase, the new centroids are first determined by placing 1s in those positions where the majority of cluster members have a 1. Next, these new centroids are constrained to rank-1, and finally, the cluster assignment is recomputed.

This procedure clearly slows down the algorithm. Also, due to the constraint on the centroids the iterative updates might actually impair the result, unlike in the normal  $k$ -means algorithm that converges to a local optimum. However, as we shall see in Section 5, the results in practice improve slightly.

#### 4.5 Implementation Details

We implemented the SaBoTeur algorithm in C.<sup>1</sup> In the implementation, the tensor  $\mathcal{X}$  is represented as a bitmap. This allows for using fast vectorized instructions such as `xor` and `popcnt`. In addition this representation takes exactly one bit per entry (excluding necessary padding), being very memory efficient even compared to sparse tensor representations.

The second ingredient to the fast algorithm we present is parallelization. As every row of  $\mathbf{X}_{(3)}$  is handled independently, the algorithm is embarrassingly parallel. We use the OpenMP (Dagum and Menon 1998) and parallelize along the sampled cluster centroids. This means that the rank-1 decompositions of each centroid as well as the computation of the similarity in each of the columns of the Khatri–Rao product (line 4 of Algorithm 1) are computed in parallel.

---

<sup>1</sup>The code is available from <http://www.mpi-inf.mpg.de/~pmiettinen/btc/>

#### 4.6 Selecting the Number of Clusters

A common problem in data analysis is that many methods require the selection of the *model order* a priori; for example, most low-rank matrix factorization methods require the user to provide the target rank before the algorithm starts. Many clustering methods – ours included – assume the number of clusters as a parameter, although there are other methods that do not have this requirement.

When the user does not know what would be a proper number of clusters for the data, we propose the use of the *minimum description length principle* (MDL, Rissanen 1978) for automatically inferring it. In particular, we adopt the recent work of Miettinen and Vreeken (2014) on using MDL for Boolean matrix factorization (BMF) to our setting (we refer the reader to Miettinen and Vreeken (2014) for more information).

The intuition behind the MDL principle is that the best model for the data (best number of clusters, in our case) is the one that lets us compress the data most. We use the two-part (i.e. crude) MDL where the *description length* contains two parts: the description length of the model  $M$ ,  $L(M)$ , and that of the data  $D$  given the model  $M$ ,  $L(D | M)$ . The overall description length is simply  $L(M) + L(D | M)$ . The optimal model (in MDL’s sense) is the one where  $L(M) + L(D | M)$  is minimized.

In BCPC, the model consists of the two factor matrices  $\mathbf{A}$  and  $\mathbf{B}$  and the cluster assignment matrix  $\mathbf{C}$ . Given these, we can reconstruct the original tensor  $\mathcal{X}$  if we know the locations where  $[[\mathbf{A}, \mathbf{B}, \mathbf{C}]]_B$  differs from  $\mathcal{X}$ . Therefore,  $L(M)$  is the total length of encoding  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$ , while  $L(D | M)$  is the length of encoding the locations of the differences, i.e. tensor  $\mathcal{X} \oplus [[\mathbf{A}, \mathbf{B}, \mathbf{C}]]_B$ , where  $\oplus$  is the element-wise *exclusive-or*.

Encoding  $\mathbf{A}$  and  $\mathbf{B}$  is similar to BMF, so we can use the DtM encoding of Miettinen and Vreeken (2014) for encoding them (and the related dimensions of the data). To encode the matrix  $\mathbf{C}$ , we only need to store to which cluster each frontal slice is associated to, taking  $l \cdot \log_2(r)$  bits. This finalises the computation of  $L(M)$ .

To compute  $L(D | M)$ , we note that we can re-write  $\mathcal{X} \oplus [[\mathbf{A}, \mathbf{B}, \mathbf{C}]]_B = \mathbf{X}_{(3)} \oplus \mathbf{C} \circ (\mathbf{B} \circ \mathbf{A})^T$ . The computation of  $L(D | M)$  now becomes equivalent of computing it for Boolean matrix factorization with factor matrices  $\mathbf{C}$  and  $\mathbf{D} = (\mathbf{B} \circ \mathbf{A})^T$ . Hence, we can follow the Typed XOR DtM approach of Miettinen and Vreeken (2014) directly.

To sum up, in order to use MDL to select the number of clusters, we have to run **SaBoTeur** with different numbers of clusters, compute the description length for each result, and take the one that yields the smallest description length.

## 5 Experimental Evaluation

Our experimental evaluation of **SaBoTeur** focuses on two central aspects: First, we test whether **SaBoTeur** can compete with other solvers for the Boolean and continuous CP decomposition in terms of reconstruction error and speed. Second, we examine the clusterings obtained by **SaBoTeur** in comparison to a standard  $k$ -means-type approach. In particular, we assess the ability of **SaBoTeur** to classify unseen data after learning the centroids.

### 5.1 Other Methods and Evaluation Criteria

To the best of our knowledge, this paper presents the first algorithms for Boolean tensor clustering and hence we cannot compare directly to other methods. To assess the reconstruction error, we decided to compare **SaBoTeur** to continuous and Boolean tensor CP decompositions. We did not use other tensor clustering methods as they aim at optimizing significantly different targets (see Section 6 for more elaborate discussion). To evaluate the clustering quality, we compare **SaBoTeur** to a standard  $k$ -means-type approach explained subsequently in Section 5.2.5.

We used the **BCP\_ALS** (Miettinen 2011) and **Walk'n'Merge** (Erdős and Miettinen 2013b) algorithms for computing Boolean CP decompositions. **BCP\_ALS** is based on iteratively updating the factor matrices one at a time (similarly to the classical alternating least squares optimizations), while **Walk'n'Merge** is a recent algorithm for scalable Boolean tensor factorization in sparse binary tensors. We did not use **Walk'n'Merge** on synthetic data as **BCP\_ALS** is expected to perform better on smaller and denser tensors (Erdős and Miettinen 2013b), but we used it on some larger real-world tensors; **BCP\_ALS**, on the other hand, does not scale well to larger tensors and hence we had to omit it from most real-world experiments.

Of the continuous methods we used **ParCube** (Papalexakis et al 2012)<sup>2</sup> and **CP\_APR** (Chi and Kolda 2012) (implementation from the Matlab Tensor Toolbox v2.5<sup>3</sup>). **CP\_APR** is an alternating Poisson regression algorithm that is specifically developed for sparse (counting) data (which can be expected to follow the Poisson distribution), with the goal of returning sparse factors.

The other method, **ParCube**, uses sampling to find smaller sub-tensors. It then solves the CP decomposition in this sub-tensor, and merges the solutions back into one. We used a non-negative variant of **ParCube** that expects non-negative data, and returns non-negative factor matrices.

For synthetic data, we report the relative similarity, that is, the fraction of the elements where the data and the clustering agree. For real-world data, we report the error measured using squared Frobenius norm (i.e. the number of disagreements between the data and the clustering when both are binary). Comparing binary methods against continuous ones causes issues, however. Using the squared Frobenius can help the real-valued methods, as it scales all errors less than 1 down, but at the same time, small errors cumulate unlike with fully binary data. To alleviate this problem, we also rounded the reconstructed tensors from **CP\_APR** and **ParCube** to binary tensors. We tried different rounding thresholds between 0 and 1 and selected the one that gave the lowest (Boolean) reconstruction error. The rounded versions are denoted by **CP\_APR**<sub>0/1</sub> and **ParCube**<sub>0/1</sub>.

It is worth emphasizing that all of the methods we are comparing against are solving a relaxed version of the BCPC problem. Compared to BCPC, the Boolean CP factorization is not restricted to clustering the third mode while the normal CP factorization lifts even the requirement for binary factor matrices. Hence, a valid solution to BCPC is always also a valid solution to (Boolean and normal) CP factorization (notice, however, that the solutions for Boolean and normal CP factorization are not interchangeable). The methods we compare against should therefore always perform better than (or at least as good as) **SaBoTeur**.

<sup>2</sup><http://www.cs.cmu.edu/~epapalex/>

<sup>3</sup><http://www.sandia.gov/~tgkolda/TensorToolbox/>

## 5.2 Synthetic Experiments

To test the **SaBoTeur** algorithm in a controlled environment, we created synthetic data sets that measured the algorithm’s response to (1) different numbers of clusters, (2) different density of data, (3) different levels of additive noise, and (4) different levels of destructive noise. All tensors were of size 700-by-500-by-50. All data sets were created by first creating random binary factor matrices  $\mathbf{A}$ , and  $\mathbf{B}$ , and a random cluster assignment matrix  $\mathbf{C}$ . The outer product of factor matrices then serves as the ground-truth tensor  $\mathcal{X} = [[\mathbf{A}, \mathbf{B}, \mathbf{C}]]_B$ . The default number of clusters was  $r = 7$  and the default density of the tensor was  $d = 0.05$ . To obtain the desired density for the tensor  $\mathcal{X}$ , we uniformly at random set a fraction of  $\sqrt{d}$  elements of  $\mathbf{A}$  and of  $\mathbf{B}$  to 1. The outer product of these matrices and the random cluster assignment matrix  $\mathbf{C}$  then gives a tensor with the desired density.

Additive and destructive noise were applied to the tensor  $\mathcal{X}$ . Additive noise turns zeros into ones while destructive noise turns ones into zeros. The default noise level<sup>4</sup> for both types was 10%, yielding to a noised input tensor  $\tilde{\mathcal{X}}$ .

We varied each of the four features one at a time keeping the others in their default values, and created 5 random copies on each parameter combination. The results we report are mean values over these five random copies. In all experiments, the number of clusters (or factors) was set to the true number of clusters used to create the data. The number of re-samples in **SaBoTeur** was set to  $r = 20$  in all experiments.

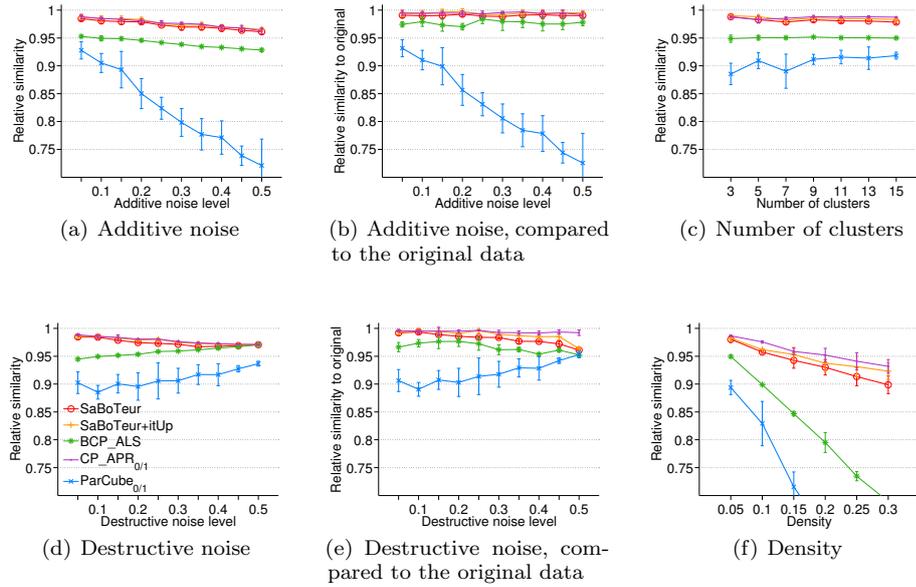
For these experiments, we compared the obtained reconstructions  $[[\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]]_B$  to both the original tensor  $\mathcal{X}$  and the noised tensor  $\tilde{\mathcal{X}}$  using the relative similarity  $1 - |\mathcal{X} - [[\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]]_B|/nml$  and  $1 - |\tilde{\mathcal{X}} - [[\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]]_B|/nml$ , respectively.

### 5.2.1 Varying the Noise

In the first experiment, we studied the effects different noise levels have to the reconstruction accuracy. First, we varied the level of additive noise from 5% to 50% (in steps of 5%). The results are in Figure 1(a), where we can see that **SaBoTeur** (with and without iterative updates) and **CP-APR**<sub>0/1</sub> all achieve very high similarity with a gentle downwards slope as the level of additive noise increases. **BCP-ALS** is slightly worse, but consistent, while the performance of **ParCube**<sub>0/1</sub> suffers significantly from increased noise levels. In order to test if the good performance of **SaBoTeur** means it is just modeling the noise, we also compared the similarity of the reconstruction to the original, noise-free tensor  $\mathcal{X}$  (Figure 1(b)). This did not change the results in meaningful ways, except that **BCP-ALS** improved somewhat.

When moving from additive to destructive noise (Figure 1(d)), **SaBoTeur**, **SaBoTeur+itUp**, and **CP-APR**<sub>0/1</sub> still stay as the best three methods, but as the destructive noise level increases, **BCP-ALS** approaches the three. That behaviour, however, is mostly driven by ‘modeling the noise’, as can be seen in Figure 1(e), which shows the similarity to the noise-free data. There, with highest levels of destructive noise, **SaBoTeur**’s performance suffers more than **CP-APR**<sub>0/1</sub>’s.

<sup>4</sup>The noise levels are reported w.r.t. number of non-zeros.



**Fig. 1** Results from synthetic experiments when varying different data characteristics. All  $y$ -axes show the relative similarity. All markers are mean values over 5 iterations and the width of the error bars is twice the standard deviation.

### 5.2.2 Varying the Number of Clusters

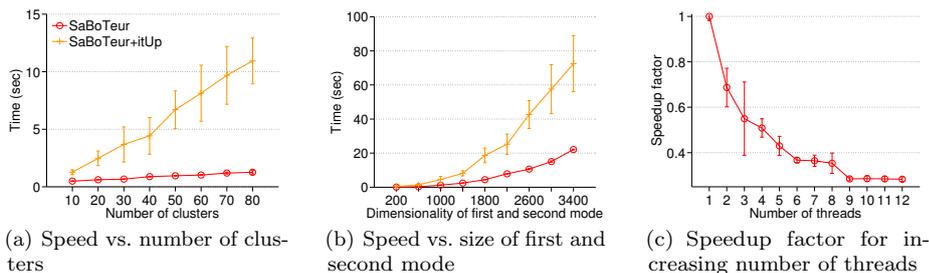
The number of clusters varied from 3 to 15 with steps of 2. The results are shown in Figure 1(c). None of the tested methods show any significant effect to the number of clusters, and the best three methods are still SaBoTeur, SaBoTeur+itUp, and CP\_APR<sub>0/1</sub>.

### 5.2.3 Varying the Density

The density of the tensor varied from 5% to 30% with steps of 5%. The results can be seen in Figure 1(f). All methods perform worse with denser data, with CP\_APR<sub>0/1</sub>, SaBoTeur+itUp, and SaBoTeur being the best three methods in that order. BCP\_ALS and ParCube's results quickly went below 75% similarity, ParCube<sub>0/1</sub> going as low as 55% with 30% density (we omit the worst results from the plots for clarity). Note that an increased density implies higher amounts of noise as the level of noise is defined with respect to the number of ones.

### 5.2.4 Scalability

We run the scalability tests on a dedicated machine with two Intel Xeon X5650 6-Core CPUs at 2.66GHz and 64GB of main memory. All reported times are wall-clock times. For these experiments, we created a new set of tensors with a default size of 800-by-800-by-500, a density of 5%, 10% of additive and destructive



**Fig. 2** Results from scalability tests. All times are wall-clock times. The speedup factor in (c) is provided with respect to the time a run with a single thread took. Markers are mean values over 5 iterations and the width of the error bars is twice the standard deviation.

noise, and 20 clusters by default. During the experiments, we varied the size, the number of clusters, and the number of threads used for the computation.

We also tried `ParCube`, `CP_APR`, and `BCP_ALS` with these data sets, but only `ParCube` was able to handle that big data sets without running out of memory. However, already with 200-by-200-by-500 data, the smallest we tried, it took 155.4 seconds. With 10 clusters, the least number we used, `ParCube` finished only after 1319.4 seconds. Therefore, we omitted these results from the figures.

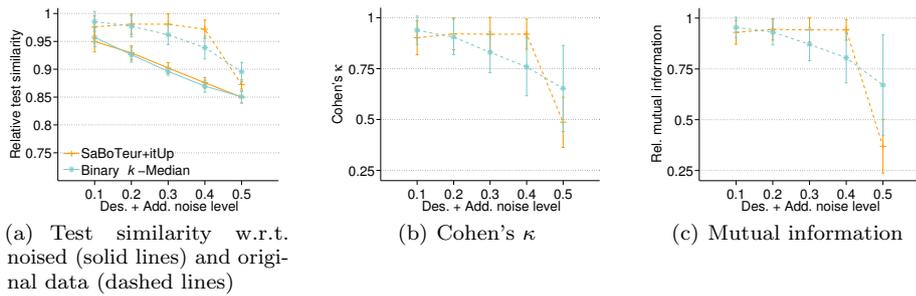
As can be seen in Figure 2(a), the `SaBoTeur` algorithm scales very well with the number of clusters. This is mostly due to efficient parallelization of the computing of the clusters (c.f. below). `SaBoTeur+itUp`, however, slows down with higher number of clusters. This is to be expected, given that more clusters require more update computations.

For the second experiment, we varied the dimensionality of the first and second mode between 200 and 3400 with steps of 400. The results can be seen in Figure 2(b). As we grew both modes simultaneously, the size of the tensor grows as a square (and, as the density was kept constant, so does the number of non-zeros). Given this, the running time of `SaBoTeur` grows as expected, or even slower, while `SaBoTeur+itUp` is again clearly slower.

In the last scalability experiment (Figure 2(c)), we tested how well `SaBoTeur` parallelizes. As the computer used had 12 cores, we set that as the maximum number of threads. Yet, after 9 threads, there were no more obvious benefits. We expect that the memory bus is the limiting factor here.

### 5.2.5 Generalization Tests

In our final test with synthetic data, we tested how well `SaBoTeur+itUp`'s clusterings generalize to yet-unseen data. Our hypothesis is that restricting the centroids to rank-1 matrices helps with the overfitting, and hence we compared `SaBoTeur+itUp` to an algorithm solving the Unconstrained Boolean tensor clustering (BTC), where the centroids are arbitrary binary matrices (Problem 6). As this problem is equivalent to clustering under binary constraints in the centroids (c.f. Section 3.3), the algorithm we used follows the standard  $k$ -means-type approach: first it samples  $k$  random points as its centroids, and then it alternates between assigning the frontal slices to the closest centroid and re-computing the centroid as the element-wise majority



**Fig. 3** Evaluation of the generalization quality using (a) reconstruction error on test data, and (b) Cohen’s  $\kappa$  statistic and (c) mutual information of cluster labels between the obtained and true labels in the test data. Markers are mean values over 50 different data sets and the width of the error bars is twice the standard deviation.

matrix. We call this algorithm **Binary  $k$ -Median**. Notice that **Binary  $k$ -Median** will always obtain at least as good reconstruction error on the training data as **SaBoTeur+itUp**. Notice further that an in-between approach would be to run **Binary  $k$ -Median** with a subsequent rank-1 decomposition on the centroids. We did not include this variation in the set of experiments but instead compared **SaBoTeur** to the more extreme counterpart with no additional constraint on the centroids.

We generated tensors of size 700-by-500-by-300 with 7 clusters, 15% of density and different levels of both additive and destructive noise, from 10% till 50%. We randomly selected 25% of the frontal slices as the test data and computed the clusterings on the remaining data. We then connected each frontal slice in the test set to its closest centroid.

The quality of the results was measured in two different ways. First, we used the reconstruction error in the test data (Figure 3(a)) and computed the overall similarity of expressing the testing data using the centroids. We also used the original (i.e. noise-free) frontal slices to assess whether we ‘model the noise’. The results show that with the noise-free test data (dashed lines), **SaBoTeur+itUp** is better than the **Binary  $k$ -Median** with 30–40% noise, and the results are reversed with 10% and 50% of noise. With noisy test data, the differences are less.

We also measured the similarity of the cluster labels w.r.t. the ground truth. We used two metrics for that, Cohen’s  $\kappa$  statistic (Figure 3(b)) and normalized mutual information (Figure 3(c)). Cohen’s  $\kappa$  takes values from  $[-1, 1]$ ,  $-1$  meaning complete disagreement and  $1$  meaning complete agreement. As Cohen’s  $\kappa$  requires the clusters’ labels to match, we paired the labels using bipartite maximum matching (we used the Hungarian method (Papadimitriou and Steiglitz 1998) to compute the matching). The mutual information was normalized by dividing it with the joint entropy; the resulting statistic takes values from  $[0, 1]$ , with  $1$  meaning perfect agreement. To assess the significance of the results we used one-sided Mann–Whitney U tests (i.e. Wilcoxon rank sum tests) for either side with a significance level of 5%.

By significance tests on all results from Cohen’s  $\kappa$  and on all results from the normalized mutual information, we find that **SaBoTeur+itUp** can recover clusterings overall significantly better than **Binary  $k$ -Median**. Looking at each level of noise separately, the two measures agree that **SaBoTeur+itUp** is significantly better at

30% and 40% of noise and worse with 10% and 50% of noise. With 20% of noise, the difference is not significant with either metric.

Perhaps the clearest sign of the benefits of the rank-1 centroids that distinguish **SaBoTeur+itUp** from **Binary  $k$ -Median** is the consistently smaller standard deviation obtained by **SaBoTeur+itUp**. This suggests that **SaBoTeur+itUp** is less susceptible to random variations in the data, which should improve the generalization quality. Indeed, our generalization experiments with the real-world data (Section 5.3.5) support this.

### 5.3 Real-World Data

We tested **SaBoTeur** also with multiple real-world data sets of varying characteristics. The purpose of these experiments is to verify our findings with synthetic data. To that end, we studied how well **SaBoTeur** (and the other methods) can reconstruct the data, how the methods scale with real-world data, and how **SaBoTeur** generalizes to unknown data. In addition, we also studied the sparsity of the factors, using MDL to select the number of clusters, and how interpretable the results **SaBoTeur** provides are.

#### 5.3.1 Data Sets

We used nine real-world data sets. The size and density for each data set are listed in Table 1. We cluster each data set along the dimension mentioned as tube in the table. The Delicious data<sup>5</sup> (Cantador et al 2011) contains user–bookmark–tag tuples from the Delicious social bookmarking system<sup>6</sup>. The Enron data<sup>7</sup> contains information about who sent an e-mail to whom (rows and columns) per months (tubes). The Facebook data set<sup>8</sup> (Viswanath et al 2009) contains information about who posted a message on whose wall per week. The Last.FM data<sup>5</sup> (Cantador et al 2011) consists of artist–user–tag tuples from the Last.FM online music system<sup>9</sup>. The MovieLens data set and the MAG data set are obtained from the same source<sup>5</sup> (Cantador et al 2011). While MovieLens is composed of user–movie–tag tuples, MAG consists of movies–actors–genres tuples. Note that the MAG data has a special structure, as the actors are connected to the genres only through the movies. The Resolver data contains entity–relation–entity tuples from the TextRunner open information extraction algorithm<sup>10</sup> (Yates and Etzioni 2009). We used the sample called ResolverL in Miettinen (2011). The TracePort data set<sup>11</sup> contains anonymized passive traffic traces (source and destination IP and port numbers) from 2009. The YAGO data set consist of subject–object–relation tuples from the semantic knowledge base YAGO<sup>12</sup> (Suchanek et al 2007).

<sup>5</sup><http://grouplens.org/datasets/hetrec-2011>

<sup>6</sup><http://www.delicious.com>

<sup>7</sup><http://www.cs.cmu.edu/~enron/>

<sup>8</sup><http://socialnetworks.mpi-sws.org/datasets.html>

<sup>9</sup><http://www.last.fm>

<sup>10</sup><http://www.cis.temple.edu/~yates/papers/jair-resolver.html>

<sup>11</sup>[http://www.caida.org/data/passive/passive\\_2009\\_dataset.xml](http://www.caida.org/data/passive/passive_2009_dataset.xml)

<sup>12</sup><http://www.mpi-inf.mpg.de/yago-naga/yago>

**Table 1** Content, size and density of the real-world data sets.

Data set	Rows	Columns	Tubes	Density ( $10^{-7}$ )
Delicious	1640 users	23584 bookmarks	7968 tags	8.23
Enron	146 senders	146 recipients	38 months	22752.86
Facebook	42390 users	39986 users	224 weeks	16.51
Last.FM	1892 users	12523 artists	9749 tags	8.07
MovieLens	2113 users	5908 movies	9079 tags	4.23
MAG	10197 movies	10673 actors	20 genres	1036.37
Resolver	343 entities	360 entities	200 relations	626.01
TracePort	10266 sources	8622 targets	501 ports	2.51
YAGO	190962 subjects	62428 objects	25 relations	67.35

As a preprocessing step we removed all the all-zero slices in every mode. In addition, for the Delicious data set, also all slices that had less than 6 entries were purged. For the MAG data, all actors that appeared in less than 5 movies were discarded.

### 5.3.2 Reconstruction Accuracy

As all real-world data sets are very sparse, we report the errors, not the similarity, for these experiments. The error is measured as the number of disagreements, when the reconstructed tensor is also binary, or as squared tensor Frobenius distance, when the reconstructed tensor is real-valued (not-rounded versions of `CP-APR` and `ParCube`). The results can be seen in Table 2. The results for `CP-APR0/1` and `ParCube0/1` for all data sets except Enron and Resolver are obtained by sampling: If  $|\mathcal{X}|$  is the number of non-zeros in data  $\mathcal{X}$ , we sampled  $200|\mathcal{X}|$  locations of 0s, and computed the error the methods make for every non-zero and for the sampled zero elements. The sampled results were then extrapolated to the size of the full tensor. We had to use sampling, as the factor matrices were too dense to allow reconstructing the full tensor. We will discuss more on density below.

The smallest reconstruction error is obtained by `ParCube0/1` in almost all experiments, YAGO being a notable exception. Remember, however, that `ParCube0/1` returns a non-negative tensor CP decomposition that is afterwards rounded to binary tensor, that is, it is neither clustering nor Boolean, and hence together with `CP-APR0/1` is expected to be better than `SaBoTeur`. Also, without the rounding, `ParCube` is often the worst method by a large margin. `CP-APR` benefits much less from the rounding, `CP-APR0/1` being comparable to `SaBoTeur`.

In those data sets where we were able to run `BGP-ALS`, its results were consistently worse than `SaBoTeur`'s results, despite it solving more relaxed problem. We were unable to get `Walk'n'Merge` to finish within a reasonable time with most data sets: it found rank-1 tensors sufficiently quickly, but took too much time to select the top ones from there (see Erdős and Miettinen (2013b) for explanation on how `Walk'n'Merge` finds a Boolean CP decomposition). When it did finish, however, it was slightly better than `SaBoTeur+itUp` or `SaBoTeur`.

Finally, the iterative updates of `SaBoTeur+itUp` consistently improved the results compared to `SaBoTeur`, but did so with significant increase to the running time. It seems that with most data sets, the iterative updates are not necessarily worth the extra wait.

**Table 2** Reconstruction errors on real data sets, measured as squared tensor Frobenius distance. ‘—’ means that the algorithm was not able to finish.

$r =$	Delicious				Enron				MAG	
	7	10	15	30	5	10	12	15	5	7
SaBoTeur	253608	253546	253447	253004	1811	1779	1769	1753	224322	223580
SaBoTeur+itUp	253153	253293	252738	252859	1793	1756	1750	1735	223869	223489
BCP_ALS	—	—	—	—	1850	1850	1850	1850	—	—
CP_APR	253558	253516	253433	253196	1718	1631	1598	1560	224996	224696
CP_APR <sub>0/1</sub>	253653	253653	253652	253639	1838	1817	1811	1781	225580	225155
ParCube	367521	373710	363581	418252	2137	2273	2352	2270	375290	322121
ParCube <sub>0/1</sub>	247129	246566	246016	241850	1802	1744	1751	1690	222133	221933
Walk’n’Merge	—	251034	—	—	—	—	1753	—	—	—

$r =$	MovieLens				Resolver				Facebook	
	5	15	20	30	5	10	15	30	15	20
SaBoTeur	56761	56712	56678	56571	1530	1509	1485	1455	626999	626776
SaBoTeur+itUp	47933	47479	47280	47316	1522	1503	1457	1443	—	—
BCP_ALS	—	—	—	—	1624	1624	1626	1632	—	—
CP_APR	47584	46927	46555	46287	1519	1509	1489	1457	626343	626202
CP_APR <sub>0/1</sub>	47178	44655	43614	42609	1545	1545	1538	1539	626945	626945
ParCube	252719	463025	454826	449534	1784	1762	1798	1939	750891	797534
ParCube <sub>0/1</sub>	41223	38098	37915	37393	1507	1471	1450	1370	619939	620894
Walk’n’Merge	—	46807	—	—	—	—	1534	—	—	—

$r =$	Last.FM				TracePort				YAGO
	8	15	20	30	5	10	15	30	7
SaBoTeur	195905	195735	195570	195278	11085	11004	10923	10776	1957737
SaBoTeur+itUp	186072	185713	185651	185399	10990	10961	10913	10673	—
BCP_ALS	—	—	—	—	—	—	—	—	—
CP_APR	184513	184038	183489	182659	11140	11114	11082	11027	—
CP_APR <sub>0/1</sub>	180766	177527	174491	170939	11110	11059	10893	10617	—
ParCube	2745259	2879381	2856941	2770361	19803	30793	31107	29796	619476763
ParCube <sub>0/1</sub>	133057	139207	119461	131322	10708	9913	10097	9647	4793390
Walk’n’Merge	—	—	—	—	—	—	10679	—	—

### 5.3.3 Sparsity of the Factors

An important question on the practical feasibility of the tensor decomposition and clustering algorithms is the density of the factor matrices. Too dense factor matrices increase the computational complexity and also the storage requirements. Furthermore, multiplying the dense factors together becomes prohibitively expensive, making it impossible to fully re-construct the tensor. This is the reason why, for example, we had to use sampling to compute the rounded representations of CP\_APR and ParCube.

We studied the sparsity of the factors using the Delicious data. In Table 3 we report the total densities of the matrices  $\mathbf{A}$  and  $\mathbf{B}$ , that is,  $(|\mathbf{A}| + |\mathbf{B}|)/r(n + m)$  for  $n$ -by- $r$  and  $m$ -by- $r$  factors.

**Table 3** Total density of the factor matrices  $\mathbf{A}$  and  $\mathbf{B}$  for different  $r$  on the Delicious data.

$r =$	7	10	15	30
SaBoTeur	0.00016	0.00031	0.00029	0.00029
CP_APR	0.20219	0.14786	0.10312	0.05396
ParCube	0.28186	0.21069	0.16658	0.09900
Walk'n'Merge	—	0.00022	—	—

**Table 4** Average wall-clock running times in seconds for real-world data sets. Averages are over 5 re-starts for MovieLens and 10 for Resolver. Both data sets used  $r = 15$  clusters.

	MovieLens	Resolver
SaBoTeur	45.6	0.02
SaBoTeur+itUp	303.6	0.11
CP_APR	196.6	22.00
ParCube	24.8	7.68
Walk'n'Merge	207.5	2.15

It is obvious that the Boolean methods, SaBoTeur and Walk'n'Merge, are orders of magnitude sparser than the continuous methods (CP\_APR and ParCube). This was to be expected as similar behaviour has already been observed with Boolean matrix factorizations (Miettinen 2010).

We did not compare the third-mode factor matrices for densities. For SaBoTeur, the clustering assignment matrix  $\mathbf{C}$  has density  $1/l$  that depends only on the number of frontal slices  $l$ ; obtaining density less than that requires having rows of  $\mathbf{C}$  that have no non-zeros.

#### 5.3.4 Scalability with Real-World Data

Table 4 shows the wall-clock times for MovieLens and Resolver with  $r = 15$ . MovieLens is one of the sparsest data sets, while Resolver is one of the densest ones, and hence these two show us how the relative speed of the methods changes as the density changes. The other data sets generally followed the behaviour we report here, adjusting to their density.

With the MovieLens data, ParCube is the fastest of the methods, followed by SaBoTeur. CP\_APR, Walk'n'Merge, and SaBoTeur+itUp are significantly slower than the two. With the denser Resolver data, however, the order is changed, with SaBoTeur and SaBoTeur+itUp being an order of magnitude faster than Walk'n'Merge, which still is faster than ParCube or CP\_APR. The density of the data does not affect the speed of SaBoTeur, while it does have a significant effect to CP\_APR and ParCube. SaBoTeur+itUp is less affected by density, and more affected by the size of the data, than the other methods.

#### 5.3.5 Generalization with Real-World Data

We repeated the generalization test we did with synthetic data with three real-world data sets – MovieLens, Resolver, and TracePort – keeping 85% of the frontal slices as training data while the remaining 15% we used for testing. The results can be seen in Table 5. We used  $r = 15$  clusters for MovieLens and TracePort, and 10 clusters

**Table 5** Training and testing errors on real-world data. The number of clusters was  $r = 15$  for MovieLens and TracePort and  $r = 10$  for Resolver.

		MovieLens	Resolver	TracePort
Train	SaBoTeur+itUp	40075	1282	9730
	Binary $k$ -Median	39351	1122	8843
Test	SaBoTeur+itUp	7258	225	1144
	Binary $k$ -Median	8256	274	1554

for the Resolver data. On each data set we conducted 10 repetitions of the training phase, and selected the one with the smallest training error.

As expected, Binary  $k$ -Median achieves lower training error, but larger testing error. This agrees with our synthetic experiments, and strongly indicates that using the rank-1 centroids helps to avoid overfitting.

### 5.3.6 Selecting the Number of Clusters

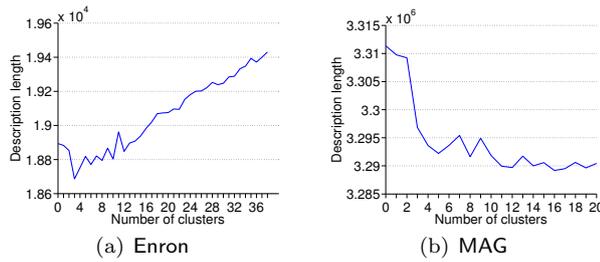
We tested the use of MDL to select the number of clusters with two real-world data sets, Enron and MAG. To that end, we ran SaBoTeur with every possible number of clusters (i.e. from having all slices in one cluster to having one cluster for each slice). Further, we computed the description length of representing the data with no clusters, i.e. having the full data explained in the  $L(D | M)$  part. The description lengths for the different numbers of clusters are presented in Figure 4. For Enron, the number of clusters that gave the smallest description length was 3, while for MAG it was 16.

Both Enron and MAG show non-smooth behaviour in the description length. This is likely partially due to SaBoTeur’s randomized behaviour and partially due to the fact that SaBoTeur does not consider the description length when it builds the clustering. Nevertheless, especially with Enron (Figure 4(a)), the selection of the best number of clusters – in the MDL’s sense – seems clear.

### 5.3.7 Interpretability of the Results

For the final experiment with the real-world data, we studied some of the results we obtained from SaBoTeur in order to see if they can provide insights to the data. We report results from four data sets: Enron, MAG, Delicious, and Last.FM. For Enron and MAG, we can interpret all three modes (email sender and recipient per month in Enron and movies, actors, genres in MAG); for the other two, we can only interpret two of the three modes (tags and URLs, and tags and artists, respectively), as the third mode corresponds to user IDs.

For Enron, the densest and smallest of the examined data sets, we mined 5 clusters. This data set is created from a set of email messages that was made public during the legal operation concerning the Enron corporation. The clusters SaBoTeur obtained can be associated to certain events in the timeline of this scandal or to the roles of certain people. One cluster we find is characterized by communication among the directors of Enron: In many months between September 2000 and December 2001, John Arnold (vice president), sends messages to Michael Maggi (director), to Andy Zipper (vice president), to Dutch Quigley (unknown



**Fig. 4** Description length vs. number of clusters.

position), and to John Griffith (unknown position). Another cluster has only November 2001 as member and describes communication of John Lavorato (CEO) to Kevin Presto (vice president), John Arnold (vice president), Richard Shapiro (vice president), Barry Tycholiz (vice president), Rick Buy (chief risk management officer), Stanley Horton (president), Bradley Mckay (employee), Kenneth Lay (CEO), Greg Whalley (president), and Mike Swerzbin (trader). This month is when the formal investigation started and Enron negotiated to be bought by Dynegy to avoid bankruptcy. Dynegy first agrees on a deal and withdraws its offer in the end of November 2001.

For MAG, we used 16 clusters (as suggested by MDL), and clustered the genres. As there are only 20 genres in the data, this means many clusters contained only one genre; this, however, is not surprising, as we would expect the sets of movies that fall in different genres be generally rather different. The results picked famous movies of the genre as the centroids: for example, the cluster of *animation* and *comedy* genres had the 1995 animation *Toy Story* appearing in its centroid. Another cluster, containing genres *adventure* and *fantasy*, had as its centroid Peter Jackson’s *Lord of the Rings* trilogy, and its main cast. Similarly to *Toy Story*, *Lord of the Rings* movies are arguably stereotypical adventure–fantasy movies.

The centroids for the MAG data contained very few movies (e.g. the three *Lord of the Rings* movies). This is due to the fact that the data sets are very sparse, and the algorithm can only pick up the most prominent elements in the centroids. With even sparser data sets, such as Delicious and Last.FM, the centroids get even sparser, perhaps containing just a singleton element from some mode. This is a common problem to many Boolean methods (see, e.g. Miettinen et al 2008). To alleviate it, we used a method proposed by Miettinen et al (2008), and weighted the reconstruction error. Normally, representing a 0 as a 1 and representing a 1 as a 0 both increase the error by 1, but for the next experiments, we set the weight so that representing a 1 with a 0 causes 10 times more error as representing 0 as a 1.

For Delicious and Last.FM, we mined 30 clusters. Even with the 10-fold weight, results for Delicious were rather sparse (indicating that the users’ behaviour was very heterogeneous). Some clusters were very informative, however, such as one with tags *software*, *apple*, *mac*, and *osx* that contained bookmarked pages on *HFS for Windows*, *iVPN*, *Object-Oriented Programming with Objective-C*, and *iBooks and ePub*, among others. For Last.FM, the clusters were generally more dense, as users tend to tag artists more homogeneously. For example, the cluster containing the tags *pop* and *rnb* contained the superstar artists such as *Beyoncé*, *Britney Spears*, *Katy Perry*,

and *Miley Cyrus* (and many more). But importantly, one could also find less-known artists: in a cluster containing tags *dance*, *trance*, *house*, *progressive trance*, *techno*, and *vocal trance*, we found artists such as *Infected Mushroom*, *T.M.Revolution*, *Dance Nation*, and *RuPaul*.

#### 5.4 Discussion

All the experiments show that **SaBoTeur** can hold its own even against methods that are theoretically superior: For reconstruction error (or similarity) with both synthetic (Sections 5.2.1–5.2.3) and real-world data (Section 5.3.2), **SaBoTeur** (and **SaBoTeur+itUp**) achieve results that are the best or close to the best results, notwithstanding that other methods, especially **CP-APR**<sub>0/1</sub> and **ParCube**<sub>0/1</sub>, benefit from significantly less constrained forms of decomposition, coupled with the post-hoc rounding. These relaxations also come with a price, as both methods create significantly denser factor matrices, and – as can be seen in Sections 5.2.4 and 5.3.4 – are unable to scale to denser data. Compared to the Boolean CP factorization methods, **BCP-ALS** and **Walk’n’Merge**, **SaBoTeur** is comparable (or sometimes clearly better) in reconstruction error and density of the factor matrices, while being significantly more scalable.

Our hypothesis that the rank-1 centroids help with overfitting was also confirmed (Sections 5.2.5 and 5.3.5): While **Binary *k*-Median** obtained lower training errors (as expected), **SaBoTeur** obtained better results with testing data in all real-world data sets and in most synthetic experiments (with **TracePort**, **SaBoTeur**’s testing error was more than 25% better than that of **Binary *k*-Median**).

Our experiments also show that MDL can be used to select the number of clusters automatically, even though we do think that more studies are needed to better understand the behaviour of the description length. Also, when studying the results, it seems obvious that **SaBoTeur** can return results that are easy to interpret and insightful; in part, this is due to the Boolean algebra, and in part, due to the clustering.

The last open question is whether one should use **SaBoTeur** or **SaBoTeur+itUp**. Judging by the results, **SaBoTeur+itUp** can provide tangible benefits over **SaBoTeur**, however, with a significant price in the running time. Thence, we think that it is best to start with **SaBoTeur**, and only if the user thinks she would benefit from more accurate results, move to **SaBoTeur+itUp**.

## 6 Related Work

Tensor clusterings have received some research interest in recent years. The existing work can be divided roughly into two separate approaches: on one hand Zhao and Zaki (2005), Jegelka et al (2009), and Papalexakis et al (2013) study the problem of clustering simultaneously all modes of a tensor (tensor co-clustering), and on the other hand, Huang et al (2008) and Liu et al (2010) (among others) study the problem where only one mode is clustered and the remaining modes are represented using a low-rank approximation. The latter form is closer to what we study in this paper, but the techniques used for the continuous methods do not apply to the binary case.

Normal tensor factorizations are well-studied, dating back to the late Twenties. The Tucker and CP decompositions were proposed in the Sixties (Tucker 1966) and Seventies (Carroll and Chang 1970; Harshman 1970), respectively. The topic has nevertheless attained growing interest in recent years, both in numerical linear algebra and computer science communities. For a comprehensive study of recent work, see Kolda and Bader (2009), and the recent work on scalable factorizations by Papalexakis et al (2012).

The first algorithm for Boolean CP factorization was presented by Leenen et al (1999), although without much analysis. Working in the framework of formal tri-concepts, Bělohávek et al (2012) studied the exact Boolean tensor decompositions, while Ignatov et al (2011) studied the problem of finding dense rank-1 subtensors from binary tensors. In data mining Miettinen (2011), studied computational complexity and sparsity of the Boolean tensor decompositions. Recently Erdős and Miettinen (2013b) proposed a scalable algorithm for Boolean CP and Tucker decompositions, and applied that algorithm for information extraction (Erdős and Miettinen 2013a). The tri-concepts are also related to closed  $n$ -ary relations, that is,  $n$ -ary extensions of closed itemsets (see, e.g. Cerf et al 2009, 2013). For more on these methods and their relation to Boolean CP factorization, see Miettinen (2011).

Miettinen et al (2008) presented the Discrete Basis Partitioning problem for clustering binary data and using binary centroids. They gave a  $(10 + \varepsilon)$  approximation algorithm based on the idea that any algorithm that can solve the so-called binary graph clustering (where centroids must be rows of the original data) within factor  $f$ , can solve the arbitrary binary centroid version within factor  $2f$ . Recently, Jiang (2014) gave a 2-approximation algorithm for the slightly more general case with  $k + 1$  clusters with one centroid restricted to be an all-zero vector.

The maximization dual of binary clustering, the hypercube segmentation problem, was studied by Kleinberg et al (1998, 2004) and they also gave three algorithms for the problem, obtaining an approximation ratio of  $2(\sqrt{2} - 1)$ . Later, Seppänen (2005) proved that this ratio is tight for this type of algorithms and Alon and Sudakov (1999) presented a PTAS for the problem.

Recently, Kim and Candan (2011) proposed the tensor-relational model, where relational information is represented as tensors. They defined the relational algebra over the (decompositions of) the tensors, noting that in the tensor-relational model, the tensor decomposition becomes the costliest operation. Hence, their subsequent work (Kim and Candan 2012, 2014) has concentrated on faster decompositions in the tensor-relational framework.

## 7 Conclusions and Future Work

We have studied the problem of clustering one mode of a 3-way binary tensor while simultaneously reducing the dimensionality in the two other modes. This problem bears close resemblance to the Boolean CP tensor decomposition, but the additional clustering constraint makes the problem significantly different. The main source of computational complexity, the consideration of overlapping factors in the tensor decomposition, does not play a role in BCPC. This lets us design algorithms with provable approximation guarantees better than what is known for the Boolean matrix and tensor decompositions.

We concentrated on 3-way tensors as they are the most common type of tensor data. The techniques we presented in this paper can be extended to higher-order tensors, though one has to first decide whether to cluster or decompose the further modes (or cluster some and decompose other). The exact resulting algorithms will depend on these decisions, and studying those algorithms more in-depth could yield interesting future research.

Our experiments show that the SaBoTeur algorithm obtains comparable reconstruction errors compared to the much less restricted tensor factorization methods while obtaining sparse factor matrices and overall best scalability. We also showed that restricting the cluster centroids to rank-1 matrices significantly helped the generalization to unseen data.

The idea of restricting the clustering centroids to a specific structure could be applied to other types of data, as well. Whether that will help with the curse of dimensionality remains an interesting topic for the future research, but based on the results we presented in this paper, we think this approach could have potential. On the other hand, one could also consider higher-rank centroids, moving from BCPC towards Unconstrained BTC. In a sense, the rank of the centroid can be seen as a regularization parameter, but the viability of this approach requires more research, both in Boolean and in continuous domains.

We also argue that the similarity – as opposed to distance or error – is often more meaningful metric for studying the (theoretical) performance of data mining algorithms due to its robustness towards small errors and more pronounced effects of large errors. To assess whether similarity is better than dissimilarity, more methods, existing and novel, should be analysed, and the results should be contrasted to those obtained when studying the dissimilarity and real-world performance. Our hypothesis is that good theoretical performance in terms of similarity correlates better with good real-world performance than good (or bad) performance correlates with the error.

Finally, our algorithms could be used to provide fast decompositions of Boolean tensors in the tensor-relational model.

## References

- Alon N, Sudakov B (1999) On two segmentation problems. *J Algorithm* 33:173–184
- Bělohávek R, Glodeanu C, Vychodil V (2012) Optimal Factorization of Three-Way Binary Data Using Triadic Concepts. *Order* 30(2):437–454
- Cantador I, Brusilovsky P, Kuffik T (2011) 2nd Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec '11). In: 5th ACM Conference on Recommender Systems (RecSys'11)
- Carroll JD, Chang JJ (1970) Analysis of individual differences in multidimensional scaling via an N-way generalization of “Eckart–Young” decomposition. *Psychometrika* 35(3):283–319
- Cerf L, Besson J, Robardet C, Boulicaut JF (2009) Closed patterns meet n-ary relations. *ACM Trans Knowl Discov Data* 3(1)
- Cerf L, Besson J, Nguyen KNT, Boulicaut JF (2013) Closed and noise-tolerant patterns in n-ary relations. *Data Min Knowl Discov* 26(3):574–619
- Chi EC, Kolda TG (2012) On Tensors, Sparsity, and Nonnegative Factorizations. *SIAM J Matrix Anal Appl* 33(4):1272–1299
- Dagum L, Menon R (1998) OpenMP: an industry standard API for shared-memory programming. *IEEE Comput Sci Eng Mag* 5(1):46–55
- Erdős D, Miettinen P (2013a) Discovering Facts with Boolean Tensor Tucker Decomposition. In: 22nd ACM International Conference on Information & Knowledge Management (CIKM '13), pp 1569–1572

- Erdős D, Miettinen P (2013b) Walk'n'Merge: A Scalable Algorithm for Boolean Tensor Factorization. In: 13th IEEE International Conference on Data Mining (ICDM '13), pp 1037–1042
- Harshman RA (1970) Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multimodal factor analysis. Tech. Rep. 16, UCLA Working Papers in Phonetics
- Huang H, Ding C, Luo D, Li T (2008) Simultaneous tensor subspace selection and clustering: The equivalence of high order SVD and k-means clustering. In: 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '08), pp 327–335
- Ignatov DI, Kuznetsov SO, Magizov RA, Zhukov LE (2011) From Triconcepts to Triclusters. In: Thirteenth International Conference on Rough Sets, Fuzzy Sets, Data Mining and Granular Computing (RSFDGrC '11), pp 257–264
- Jegelka S, Sra S, Banerjee A (2009) Approximation Algorithms for Tensor Clustering. In: International Conference on Algorithmic Learning Theory (ALT '09), pp 368–383
- Jiang P (2014) Pattern extraction and clustering for high-dimensional discrete data. PhD thesis, University of Illinois at Urbana-Champaign
- Kim M, Candan KS (2011) Approximate tensor decomposition within a tensor-relational algebraic framework. In: 20th ACM International Conference on Information & Knowledge Management (CIKM '11), pp 1737–1742
- Kim M, Candan KS (2012) Decomposition-by-normalization (DBN): Leveraging approximate functional dependencies for efficient tensor decomposition. In: 21st ACM International Conference on Information & Knowledge Management (CIKM '12), pp 355–364
- Kim M, Candan KS (2014) Pushing-Down Tensor Decompositions over Unions to Promote Reuse of Materialized Decompositions. In: European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD '14), pp 688–704
- Kleinberg J, Papadimitriou C, Raghavan P (1998) A Microeconomic View of Data Mining. *Data Min Knowl Discov* 2(4):311–324
- Kleinberg JM, Papadimitriou CH, Raghavan P (2004) Segmentation problems. *J ACM* 51(2):263–280
- Kolda TG, Bader BW (2009) Tensor decompositions and applications. *SIAM Rev* 51(3):455–500
- Leenen I, Van Mechelen I, De Boeck P, Rosenberg S (1999) INDCLAS: A three-way hierarchical classes model. *Psychometrika* 64(1):9–24
- Liu X, De Lathauwer L, Janssens F, De Moor B (2010) Hybrid Clustering of Multiple Information Sources via HOSVD. In: 7th International Conference on Advances in Neural Networks - Part II (ISNN '10), pp 337–345
- Miettinen P (2009) Matrix Decomposition Methods for Data Mining: Computational Complexity and Algorithms. PhD thesis, Department of Computer Science, University of Helsinki
- Miettinen P (2010) Sparse Boolean Matrix Factorizations. In: 10th IEEE International Conference on Data Mining (ICDM '10), pp 935–940
- Miettinen P (2011) Boolean Tensor Factorizations. In: 11th IEEE International Conference on Data Mining (ICDM '11), pp 447–456
- Miettinen P, Vreeken J (2014) MDL4BMF: Minimum description length for Boolean matrix factorization. *ACM Trans Knowl Discov Data* 8(4)
- Miettinen P, Mielikäinen T, Gionis A, Das G, Mannila H (2008) The Discrete Basis Problem. *IEEE Trans Knowl Data En* 20(10):1348–1362
- Papadimitriou CH, Steiglitz K (1998) *Combinatorial Optimization: Algorithms and Complexity*. Dover Publications, Mineola, New York
- Papalexakis EE, Faloutsos C, Sidiropoulos ND (2012) ParCube: Sparse Parallelizable Tensor Decompositions. In: European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD '12), pp 521–536
- Papalexakis EE, Sidiropoulos N, Bro R (2013) From K-Means to Higher-Way Co-Clustering: Multilinear Decomposition With Sparse Latent Factors. *IEEE Trans Sig Process* 61(2):493–506
- Rissanen J (1978) Modeling by shortest data description. *Automatica* 14(5):465–471
- Seppänen JK (2005) Upper bound for the approximation ratio of a class of hypercube segmentation algorithms. *Inform Process Lett* 93(3):139–141
- Suchanek FM, Kasneci G, Weikum G (2007) Yago: A core of semantic knowledge. In: 16th International Conference on World Wide Web (WWW '07), pp 697–706
- Tucker LR (1966) Some mathematical notes on three-mode factor analysis. *Psychometrika* 31(3):279–311

- 
- Viswanath B, Mislove A, Cha M, Gummadi KP (2009) On the evolution of user interaction in Facebook. In: 2nd ACM Workshop on Online Social Networks (WOSN '09), pp 37–42
- Yates A, Etzioni O (2009) Unsupervised methods for determining object and relation synonyms on the web. *J Artif Intell Res* 34
- Zhao L, Zaki MJ (2005) TRICLUSTER: an effective algorithm for mining coherent clusters in 3D microarray data. In: ACM SIGMOD International Conference on Management of Data (SIGMOD '05), pp 694–705