

# Higher-order Relation Schema Induction using Tensor Factorization with Back-off and Aggregation

**Madhav Nimishakavi**  
Indian Institute of Science  
Bangalore  
madhav@iisc.ac.in

**Manish Gupta**  
Microsoft  
Hyderabad  
manishg@microsoft.com

**Partha Talukdar**  
Indian Institute of Science  
Bangalore  
ppt@iisc.ac.in

## Abstract

Relation Schema Induction (RSI) is the problem of identifying type signatures of arguments of relations from unlabeled text. Most of the previous work in this area have focused only on *binary* RSI, i.e., inducing only the subject and object type signatures per relation. However, in practice, many relations are *high-order*, i.e., they have more than two arguments and inducing type signatures of all arguments is necessary. For example, in the sports domain, inducing a schema  $win(WinningPlayer, OpponentPlayer, Tournament, Location)$  is more informative than inducing just  $win(WinningPlayer, OpponentPlayer)$ . We refer to this problem as Higher-order Relation Schema Induction (HRSI). In this paper, we propose Tensor Factorization with Back-off and Aggregation (TFBA), a novel framework for the HRSI problem. To the best of our knowledge, this is the first attempt at inducing higher-order relation schemata from unlabeled text. Using the experimental analysis on three real world datasets, we show how TFBA helps in dealing with sparsity and induce higher order schemata.

## 1 Introduction

Building Knowledge Graphs (KGs) out of unstructured data is an area of active research. Research in this has resulted in the construction of several large scale KGs, such as NELL (Mitchell et al., 2015), Google Knowledge Vault (Dong et al., 2014) and YAGO (Suchanek et al., 2007). These KGs consist of millions of entities and beliefs involving those entities. Such KG construc-

tion methods are schema-guided as they require the list of input relations and their schemata (e.g.,  $playerPlaysSport(Player, Sport)$ ). In other words, knowledge of schemata is an important first step towards building such KGs.

While beliefs in such KGs are usually binary (i.e., involving two entities), many beliefs of interest go beyond two entities. For example, in the sports domain, one may be interested in beliefs of the form  $win(Roger Federer, Nadal, Wimbledon, London)$ , which is an instance of the high-order (or n-ary) relation  $win$  whose schema is given by  $win(WinningPlayer, OpponentPlayer, Tournament, Location)$ . We refer to the problem of inducing such relation schemata involving multiple arguments as Higher-order Relation Schema Induction (HRSI). In spite of its importance, HRSI is mostly unexplored.

Recently, tensor factorization-based methods have been proposed for *binary* relation schema induction (Nimishakavi et al., 2016), with gains in both speed and accuracy over previously proposed generative models. To the best of our knowledge, tensor factorization methods have not been used for HRSI. We address this gap in this paper.

Due to data sparsity, straightforward adaptation of tensor factorization from (Nimishakavi et al., 2016) to HRSI is not feasible, as we shall see in Section 3.1. We overcome this challenge in this paper, and make the following contributions.

- We propose Tensor Factorization with Back-off and Aggregation (TFBA), a novel tensor factorization-based method for Higher-order RSI (HRSI). In order to overcome data sparsity, TFBA *backs-off* and jointly factorizes multiple lower-order tensors derived from an extremely sparse higher-order tensor.
- As an aggregation step, we propose a constrained clique mining step which constructs

the higher-order schemata from multiple binary schemata.

- Through experiments on multiple real-world datasets, we show the effectiveness of TFBA for HRSI.

Source code of TFBA is available at <https://github.com/madhavcsa/TFBA>.

The remainder of the paper is organized as follows. We discuss related work in Section 2. In Section 3.1, we first motivate why a back-off strategy is needed for HRSI, rather than factorizing the higher-order tensor. Further, we discuss the proposed TFBA framework in Section 3.2. In Section 4, we demonstrate the effectiveness of the proposed approach using multiple real world datasets. We conclude with a brief summary in Section 5.

## 2 Related Work

In this section, we discuss related works in two broad areas: schema induction, and tensor and matrix factorizations.

**Schema Induction:** Most work on inducing schemata for relations has been in the binary setting (Mohamed et al., 2011; Movshovitz-Attias and Cohen, 2015; Nimishakavi et al., 2016). McDonald et al. (2005) and Peng et al. (2017) extract  $n$ -ary relations from Biomedical documents, but do not induce the schema, i.e., type signature of the  $n$ -ary relations. There has been significant amount of work on Semantic Role Labeling (Lang and Lapata, 2011; Titov and Khoddam, 2015; Roth and Lapata, 2016), which can be considered as  $n$ -ary relation extraction. However, we are interested in inducing the schemata, i.e., the type signature of these relations. Event Schema Induction is the problem of inducing schemata for events in the corpus (Balasubramanian et al., 2013; Chambers, 2013; Nguyen et al., 2015). Recently, a model for event representations is proposed in (Weber et al., 2018).

Cheung et al. (2013) propose a probabilistic model for inducing frames from text. Their notion of frame is closer to that of scripts (Schank and Abelson, 1977). Script learning is the process of automatically inferring sequence of events from text (Mooney and DeJong, 1985). There is a fair amount of recent work in statistical script learning (Pichotta and Mooney, 2016), (Pichotta and Mooney, 2014). While script learning deals with the sequence of events, we try to find the

schemata of relations at a corpus level. Ferraro and Durme (2016) propose a unified Bayesian model for scripts, frames and events. Their model tries to capture all levels of Minsky Frame structure (Minsky, 1974), however we work with the surface semantic frames.

**Tensor and Matrix Factorizations:** Matrix factorization and joint tensor-matrix factorizations have been used for the problem of predicting links in the Universal Schema setting (Riedel et al., 2013; Singh et al., 2015). Chen et al. (2015) use matrix factorizations for the problem of finding semantic slots for unsupervised spoken language understanding. Tensor factorization methods are also used in factorizing knowledge graphs (Chang et al., 2014; Nickel et al., 2012). Joint matrix and tensor factorization frameworks, where the matrix provides additional information, is proposed in (Acar et al., 2013) and (Wang et al., 2015). These models are based on PARAFAC (Harshman, 1970), a tensor factorization model which approximates the given tensor as a sum of rank-1 tensors. A boolean Tucker decomposition for discovering facts is proposed in (Erdos and Miettinen, 2013). In this paper, we use a modified version (Tucker2) of Tucker decomposition (Tucker, 1963).

RESCAL (Nickel et al., 2011) is a simplified Tucker model suitable for relational learning. Recently, SICTF (Nimishakavi et al., 2016), a variant of RESCAL with side information, is used for the problem of schema induction for binary relations. SICTF cannot be directly used to induce higher order schemata, as the higher-order tensors involved in inducing such schemata tend to be extremely sparse. TFBA overcomes these challenges to induce higher-order relation schemata by performing Non-Negative Tucker-style factorization of sparse tensor while utilizing a back-off strategy, as explained in the next section.

## 3 Higher Order Relation Schema Induction using Back-off Factorization

In this section, we start by discussing the approach of factorizing a higher-order tensor and provide the motivation for back-off strategy. Next, we discuss the proposed TFBA approach in detail. Please refer to Table 1 for notations used in this paper.

Notation	Definition
$\mathbb{R}_+$	Set of non-negative reals.
$\mathcal{X} \in \mathbb{R}_+^{n_1 \times n_2 \times \dots \times n_N}$	$N^{\text{th}}$ -order non-negative tensor.
$\mathcal{X}_{(i)}$	mode- $i$ matricization of tensor $\mathcal{X}$ . Please see (Kolda and Bader, 2009) for details.
$\mathbf{A} \in \mathbb{R}_+^{n \times r}$	Non-negative matrix of order $n \times r$ .
*	Hadamard product: $(\mathbf{A} * \mathbf{B})_{i,j} = \mathbf{A}_{i,j} \times \mathbf{B}_{i,j}$ .

Table 1: Notations used in the paper.

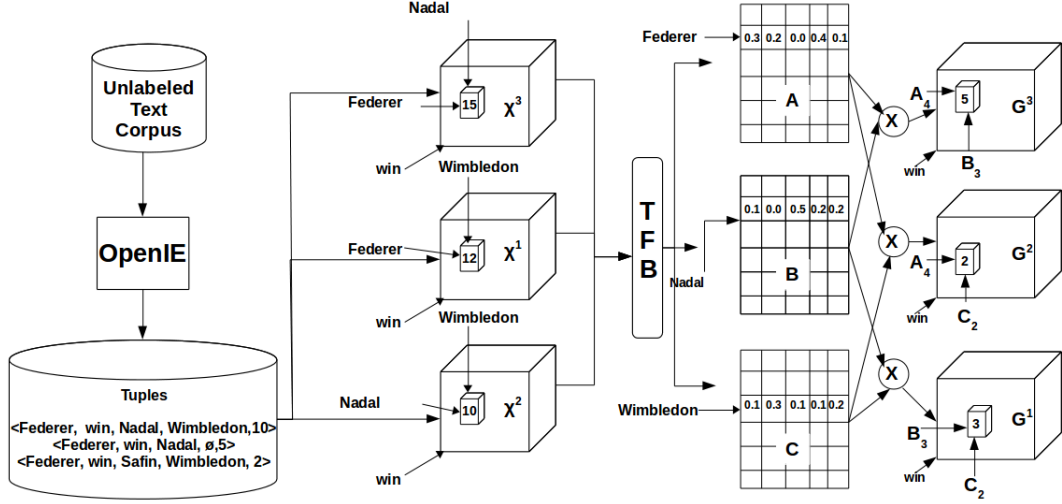


Figure 1: **Overview of Step 1 of TFBA.** Rather than factorizing the higher-order tensor  $\mathcal{X}$ , TFBA performs joint Tucker decomposition of multiple 3-mode tensors,  $\mathcal{X}^1$ ,  $\mathcal{X}^2$ , and  $\mathcal{X}^3$ , derived out of  $\mathcal{X}$ . This joint factorization is performed using shared latent factors  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$ . This results in binary schemata, each of which is stored as a cell in one of the core tensors  $\mathcal{G}^1$ ,  $\mathcal{G}^2$ , and  $\mathcal{G}^3$ . Please see Section 3.2.1 for details.

### 3.1 Factorizing a Higher-order Tensor

Given a text corpus, we use OpenIEv5 (Mausam, 2016) to extract tuples. Consider the following sentence “*Federer won against Nadal at Wimbledon.*”. Given this sentence, OpenIE extracts the 4-tuple (*Federer, won, against Nadal, at Wimbledon*). We lemmatize the relations in the tuples and only consider the noun phrases as arguments. Let  $\mathbb{T}$  represent the set of these 4-tuples. We can construct a 4-order tensor  $\mathcal{X} \in \mathbb{R}_+^{n_1 \times n_2 \times n_3 \times m}$  from  $\mathbb{T}$ . Here,  $n_1$  is the number of *subject* noun phrases (NPs),  $n_2$  is the number of *object* NPs,  $n_3$  is the number of *other* NPs, and  $m$  is the number of relations in  $\mathbb{T}$ . Values in the tensor correspond to the frequency of the tuples. In case of 5-tuples of the form (*subject, relation, object, other-1, other-2*), we split the 5-tuples into two 4-tuples of the form (*subject, relation, object, other-1*) and (*subject, relation, object, other-2*) and frequency of these 4-tuples is considered to be same as the original 5-tuple. Factorizing the tensor  $\mathcal{X}$  results in discovering latent categories of NPs, which help in in-

ducing the schemata. We propose the following approach to factorize  $\mathcal{X}$ .

$$\min_{\mathcal{G}, \mathbf{A}, \mathbf{B}, \mathbf{C}} \|\mathcal{X} - \mathcal{G} \times_1 \mathbf{A} \times_2 \mathbf{B} \times_3 \mathbf{C} \times_4 \mathbf{I}\|_F^2 + \lambda_a \|\mathbf{A}\|_F^2 + \lambda_b \|\mathbf{B}\|_F^2 + \lambda_c \|\mathbf{C}\|_F^2,$$

where,

$$\mathbf{A} \in \mathbb{R}_+^{n_1 \times r_1}, \mathbf{B} \in \mathbb{R}_+^{n_2 \times r_2}, \mathbf{C} \in \mathbb{R}_+^{n_3 \times r_3},$$

$$\mathcal{G} \in \mathbb{R}_+^{r_1 \times r_2 \times r_3 \times m}, \lambda_a \geq 0, \lambda_b \geq 0 \text{ and } \lambda_c \geq 0.$$

Here,  $\mathbf{I}$  is the identity matrix. Non-negative updates for the variables can be obtained following (Lee and Seung, 2000). Similar to (Nimishakavi et al., 2016), schemata induced will be of the form *relation*  $\langle \mathbf{A}_i, \mathbf{B}_j, \mathbf{C}_k \rangle$ . Here,  $\mathbf{P}_i$  represents the  $i^{\text{th}}$  column of a matrix  $\mathbf{P}$ .  $\mathbf{A}$  is the embedding matrix of subject NPs in  $\mathbb{T}$  (i.e., mode-1 of  $\mathcal{X}$ ),  $r_1$  is the embedding rank in mode-1 which is the number of latent categories of subject NPs. Similarly,  $\mathbf{B}$  and

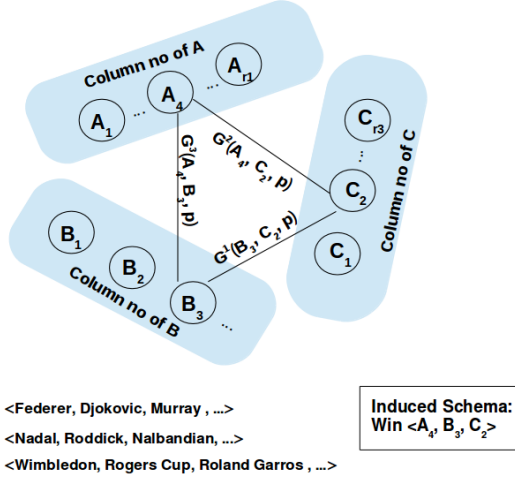


Figure 2: **Overview of Step 2 of TFBA.** Induction of higher-order schemata from the tri-partite graph formed from the columns of matrices  $A$ ,  $B$ , and  $C$ . Triangles in this graph (solid) represent a 3-ary schema,  $n$ -ary schemata for  $n > 3$  can be induced from the 3-ary schemata. Please refer to Section 3.2.2 for details.

$C$  are the embedding matrices of object NPs and other NPs respectively.  $r_2$  and  $r_3$  are the number of latent categories of object NPs and other NPs respectively.  $\mathcal{G}$  is the core tensor.  $\lambda_a$ ,  $\lambda_b$  and  $\lambda_c$  are the regularization weights.

However, the 4-order tensors are heavily sparse for all the datasets we consider in this work. The sparsity ratio of this 4-order tensor for all the datasets is of the order  $1e-7$ . As a result of the extreme sparsity, this approach fails to learn any schemata. Therefore, we propose a more successful back-off strategy for higher-order RSI in the next section.

### 3.2 TFBA: Proposed Framework

To alleviate the problem of sparsity, we construct three tensors  $\mathcal{X}^3$ ,  $\mathcal{X}^2$ , and  $\mathcal{X}^1$  from  $\mathbb{T}$  as follows:

- $\mathcal{X}^3 \in \mathbb{R}_+^{n_1 \times n_2 \times m}$  is constructed out of the tuples in  $\mathbb{T}$  by dropping the *other* argument and aggregating resulting tuples, i.e.,  $\mathcal{X}_{i,j,p}^3 = \sum_{k=1}^{n_3} \mathcal{X}_{i,j,k,p}$ . For example, 4-tuples  $\langle \langle \text{Federer}, \text{Win}, \text{Nadal}, \text{Wimbledon} \rangle, 10 \rangle$  and  $\langle \langle \text{Federer}, \text{Win}, \text{Nadal}, \text{Australian Open} \rangle, 5 \rangle$  will be aggregated to form a triple  $\langle \langle \text{Federer}, \text{Win}, \text{Nadal} \rangle, 15 \rangle$ .
- $\mathcal{X}^2 \in \mathbb{R}_+^{n_1 \times n_3 \times m}$  is constructed out of the tuples in  $\mathbb{T}$  by dropping the *object* argument

and aggregating resulting tuples i.e.,  $\mathcal{X}_{i,j,p}^2 = \sum_{k=1}^{n_2} \mathcal{X}_{i,k,j,p}$ .

- $\mathcal{X}^1 \in \mathbb{R}_+^{n_2 \times n_3 \times m}$  constructed out of the tuples in  $\mathbb{T}$  by dropping the *subject* argument and aggregating resulting tuples i.e.,  $\mathcal{X}_{i,j,p}^1 = \sum_{k=1}^{n_1} \mathcal{X}_{k,i,j,p}$ .

The proposed framework TFBA for inducing higher order schemata involves the following two steps.

- **Step 1:** In this step, TFBA factorizes multiple lower-order overlapping tensors,  $\mathcal{X}^1$ ,  $\mathcal{X}^2$ , and  $\mathcal{X}^3$ , derived from  $\mathcal{X}$  to induce binary schemata. This step is illustrated in Figure 1 and we discuss details in Section 3.2.1.
- **Step 2:** In this step, TFBA connects multiple binary schemata identified above to induce higher-order schemata. The method accomplishes this by solving a constrained clique problem. This step is illustrated in Figure 2 and we discuss the details in Section 3.2.2.

#### 3.2.1 Step 1: Back-off Tensor Factorization

A schematic overview of this step is shown in Figure 1. TFBA first preprocesses the corpus and extracts OpenIE tuple set  $\mathbb{T}$  out of it. The 4-mode tensor  $\mathcal{X}$  is constructed out of  $\mathbb{T}$ . Instead of performing factorization of the higher-order tensor  $\mathcal{X}$  as in Section 3.1, TFBA creates three tensors out of  $\mathcal{X}$ :  $\mathcal{X}_{n_2 \times n_3 \times m}^1$ ,  $\mathcal{X}_{n_1 \times n_3 \times m}^2$  and  $\mathcal{X}_{n_1 \times n_2 \times m}^3$ .

TFBA performs a coupled non-negative Tucker factorization of the input tensors  $\mathcal{X}^1$ ,  $\mathcal{X}^2$  and  $\mathcal{X}^3$  by solving the following optimization problem.

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}, \mathbf{C}} & f(\mathcal{X}^3, \mathcal{G}^3, \mathbf{A}, \mathbf{B}) + f(\mathcal{X}^2, \mathcal{G}^2, \mathbf{A}, \mathbf{C}) \\ & + f(\mathcal{X}^1, \mathcal{G}^1, \mathbf{B}, \mathbf{C}) \\ & + \lambda_a \|\mathbf{A}\|_F^2 + \lambda_b \|\mathbf{B}\|_F^2 + \lambda_c \|\mathbf{C}\|_F^2, \quad (1) \end{aligned}$$

where,

$$\begin{aligned} f(\mathcal{X}^i, \mathcal{G}^i, \mathbf{P}, \mathbf{Q}) &= \|\mathcal{X}^i - \mathcal{G}^i \times_1 \mathbf{P} \times_2 \mathbf{Q} \times_3 \mathbf{I}\|_F^2 \\ \mathbf{A} &\in \mathbb{R}_+^{n_1 \times r_1}, \mathbf{B} \in \mathbb{R}_+^{n_2 \times r_2}, \mathbf{C} \in \mathbb{R}_+^{n_3 \times r_3} \\ \mathcal{G}^1 &\in \mathbb{R}_+^{r_2 \times r_3 \times m}, \mathcal{G}^2 \in \mathbb{R}_+^{r_1 \times r_3 \times m}, \mathcal{G}^3 \in \mathbb{R}_+^{r_1 \times r_2 \times m}. \end{aligned}$$

We enforce non-negativity constraints on the matrices  $\mathbf{A}$ ,  $\mathbf{B}$ ,  $\mathbf{C}$  and the core tensors  $\mathcal{G}^i$  ( $i \in \{1, 2, 3\}$ ). Non-negativity is essential for learning interpretable latent factors (Murphy et al., 2012).

Each slice of the core tensor  $\mathcal{G}^3$  corresponds to one of the  $m$  relations. Each cell in a slice corresponds to an induced schema in terms of the latent factors from matrices  $\mathbf{A}$  and  $\mathbf{B}$ . In other words,  $\mathcal{G}_{i,j,k}^3$  is an induced binary schema for relation  $k$  involving induced categories represented by columns  $\mathbf{A}_i$  and  $\mathbf{B}_j$ . Cells in  $\mathcal{G}^1$  and  $\mathcal{G}^2$  may be interpreted accordingly.

We derive non-negative multiplicative updates for  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{C}$  following the NMF updating rules given in (Lee and Seung, 2000). For the update of  $\mathbf{A}$ , we consider the mode-1 matricization of first and the second term in Equation 1 along with the regularizer.

$$\mathbf{A} \leftarrow \mathbf{A} * \frac{\boldsymbol{\chi}_{(1)}^3 \mathcal{G}_{B_A}^\top + \boldsymbol{\chi}_{(1)}^2 \mathcal{G}_{C_A}^\top}{\mathbf{A}[\mathcal{G}_{B_A} \mathcal{G}_{B_A}^\top + \mathcal{G}_{C_A} \mathcal{G}_{C_A}^\top] + \lambda_a \mathbf{A}},$$

where,

$$\mathcal{G}_{B_A} = (\mathcal{G}^3 \times_2 \mathbf{B})_{(1)}, \quad \mathcal{G}_{C_A} = (\mathcal{G}^2 \times_2 \mathbf{C})_{(1)}.$$

In order to estimate  $\mathbf{B}$ , we consider mode-2 matricization of first term and mode-1 matricization of third term in Equation 1, along with the regularization term. We get the following update rule for  $\mathbf{B}$

$$\mathbf{B} \leftarrow \mathbf{B} * \frac{\boldsymbol{\chi}_{(2)}^3 \mathcal{G}_{A_B}^\top + \boldsymbol{\chi}_{(1)}^1 \mathcal{G}_{C_B}^\top}{\mathbf{B}[\mathcal{G}_{A_B} \mathcal{G}_{A_B}^\top + \mathcal{G}_{C_B} \mathcal{G}_{C_B}^\top] + \lambda_b \mathbf{B}},$$

where,

$$\mathcal{G}_{A_B} = (\mathcal{G}^3 \times_1 \mathbf{A})_{(2)}, \quad \mathcal{G}_{C_B} = (\mathcal{G}^1 \times_2 \mathbf{C})_{(1)}.$$

For updating  $\mathbf{C}$ , we consider mode-2 matricization of second and third terms in Equation 1 along with the regularization term, and we get

$$\mathbf{C} \leftarrow \mathbf{C} * \frac{\boldsymbol{\chi}_{(2)}^3 \mathcal{G}_{B_C}^\top + \boldsymbol{\chi}_{(2)}^2 \mathcal{G}_{A_C}^\top}{\mathbf{C}[\mathcal{G}_{A_C} \mathcal{G}_{A_C}^\top + \mathcal{G}_{B_C} \mathcal{G}_{B_C}^\top] + \lambda_c \mathbf{C}},$$

where,

$$\mathcal{G}_{A_C} = (\mathcal{G}^3 \times_1 \mathbf{B})_{(2)}, \quad \mathcal{G}_{B_C} = (\mathcal{G}^2 \times_1 \mathbf{A})_{(2)}.$$

Finally, we update the three core tensors in Equation 1 following (Kim and Choi, 2007) as follows,

$$\begin{aligned} \mathcal{G}^1 &\leftarrow \mathcal{G}^1 * \frac{\boldsymbol{\chi}^1 \times_1 \mathbf{B}^\top \times_2 \mathbf{C}^\top}{\mathcal{G}^1 \times_1 \mathbf{B}^\top \mathbf{B} \times_2 \mathbf{C}^\top \mathbf{C}}, \\ \mathcal{G}^2 &\leftarrow \mathcal{G}^2 * \frac{\boldsymbol{\chi}^2 \times_1 \mathbf{A}^\top \times_2 \mathbf{C}^\top}{\mathcal{G}^2 \times_1 \mathbf{A}^\top \mathbf{A} \times_2 \mathbf{C}^\top \mathbf{C}}, \\ \mathcal{G}^3 &\leftarrow \mathcal{G}^3 * \frac{\boldsymbol{\chi}^3 \times_1 \mathbf{A}^\top \times_2 \mathbf{B}^\top}{\mathcal{G}^3 \times_1 \mathbf{A}^\top \mathbf{A} \times_2 \mathbf{B}^\top \mathbf{B}}. \end{aligned}$$

In all the above updates,  $\frac{\mathbf{P}}{\mathbf{Q}}$  represents element-wise division and  $\mathbf{I}$  is the identity matrix.

**Initialization:** For initializing the component matrices  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$ , we first perform a non-negative Tucker2 Decomposition of the individual input tensors  $\boldsymbol{\chi}^1$ ,  $\boldsymbol{\chi}^2$ , and  $\boldsymbol{\chi}^3$ . Then compute the average of component matrices obtained from each individual decomposition for initialization. We initialize the core tensors  $\mathcal{G}^1$ ,  $\mathcal{G}^2$ , and  $\mathcal{G}^3$  with the core tensors obtained from the individual decompositions.

### 3.2.2 Step 2: Binary to Higher-Order Schema Induction

In this section, we describe how a higher-order schema is constructed from the factorization described in the previous sub-section. Each relation  $k$  has three representations given by the slices  $\mathcal{G}_k^1$ ,  $\mathcal{G}_k^2$  and  $\mathcal{G}_k^3$  from each core tensor. We need a principled way to produce a joint schema from these representations. For a relation, we select top- $n$  indices  $(i, j)$  with highest values from each matrix. The indices  $i$  and  $j$  from  $\mathcal{G}_k^3$  correspond to column numbers of  $\mathbf{A}$  and  $\mathbf{B}$  respectively, indices from  $\mathcal{G}_k^2$  correspond to columns from  $\mathbf{A}$  and  $\mathbf{C}$  and columns from  $\mathcal{G}_k^1$  correspond to columns from  $\mathbf{B}$  and  $\mathbf{C}$ .

We construct a tri-partite graph with the column numbers from each of the component matrices  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{C}$  as the vertices belonging to independent sets, the top- $n$  indices selected are the edges between these vertices. From this tri-partite graph, we find all the triangles which will give schema with three arguments for a relation, illustrated in Figure 2. We find higher order schemata, i.e., schemata with more than three arguments by merging two third order schemata with same column number from  $\mathbf{A}$  and  $\mathbf{B}$ . For example, if we find two schemata  $(\mathbf{A}_2, \mathbf{B}_4, \mathbf{C}_{10})$  and  $(\mathbf{A}_2, \mathbf{B}_4, \mathbf{C}_8)$  then we merge these two to give  $(\mathbf{A}_2, \mathbf{B}_4, \mathbf{C}_{10}, \mathbf{C}_8)$  as a higher order schema. This can be continued further for even higher order schemata. This process may be thought of as finding a **constrained**

**clique** over the tri-partite graph. Here the constraint is that in the maximal clique, there can only be one edge between sets corresponding to columns of  $\mathbf{A}$  and columns of  $\mathbf{B}$ .

The procedure above is inspired by (McDonald et al., 2005). However, we note that (McDonald et al., 2005) solved a different problem, viz., n-ary relation instance extraction, while our focus is on inducing schemata. Though we discuss the case of back-off from 4-order to 3-order, ideas presented above can be extended for even higher orders depending on the sparsity of the tensors.

## 4 Experiments

In this section, we evaluate the performance of TFBA for the task of HRSI. We also propose a baseline model for HRSI called HardClust.

**HardClust:** We propose a baseline model called the Hard Clustering Baseline (HardClust) for the task of higher order relation schema induction. This model induces schemata by grouping per-relation NP arguments from OpenIE extractions. In other words, for each relation, all the Noun Phrases (NPs) in first argument form a cluster that represents the subject of the relation, all the NPs in the second argument form a cluster that represents object and so on. Then from each cluster, the top most frequent NPs are chosen as the representative NPs for the argument type. We note that this method is only able to induce one schema per relation.

**Datasets:** We run our experiments on three datasets. The first dataset (Shootings) is a collection of 1,335 documents constructed from a publicly available database of mass shootings in the United States. The second is New York Times Sports (NYT Sports) dataset which is a collection of 20,940 sports documents from the period 2005 and 2007. And the third dataset (MUC) is a set of 1300 Latin American newswire documents about terrorism events. After performing the processing steps described in Section 3, we obtained 357,914 unique OpenIE extractions from the NYT Sports dataset, 10,847 from Shootings dataset, and 8,318 from the MUC dataset. However, in order to properly analyze and evaluate the model, we consider only the 50 most frequent relations in the datasets and their corresponding OpenIE extractions. This is done to avoid noisy OpenIE extractions to yield better data quality and to aid subsequent manual evaluation of the data. We construct input tensors

following the procedure described in Section 3.2. Details on the dimensions of tensors obtained are given in Table 2.

**Model Selection:** In order to select appropriate TFBA parameters, we perform a grid search over the space of hyper-parameters, and select the set of hyper-parameters that give best Average FIT score (AvgFIT).

$$\text{AvgFIT}(\mathcal{G}^1, \mathcal{G}^2, \mathcal{G}^3, \mathbf{A}, \mathbf{B}, \mathbf{C}, \mathcal{X}^1, \mathcal{X}^2, \mathcal{X}^3) = \frac{1}{3} \{ \text{FIT}(\mathcal{X}^1, \mathcal{G}^1, \mathbf{B}, \mathbf{C}) + \text{FIT}(\mathcal{X}^2, \mathcal{G}^2, \mathbf{A}, \mathbf{C}) + \text{FIT}(\mathcal{X}^3, \mathcal{G}^3, \mathbf{A}, \mathbf{B}) \},$$

where,

$$\text{FIT}(\mathcal{X}, \mathcal{G}, \mathbf{P}, \mathbf{Q}) = 1 - \frac{\|\mathcal{X} - \mathcal{G} \times_1 \mathbf{P} \times_2 \mathbf{Q}\|_F}{\|\mathcal{X}\|_F}.$$

We perform a grid search for the rank parameters between 5 and 20, for the regularization weights we perform a grid search over 0 and 1. Table 3 provides the details of hyper-parameters set for different datasets.

**Evaluation Protocol:** For TFBA, we follow the protocol mentioned in Section 3.2.2 for constructing higher order schemata. For every relation, we consider top 5 binary schemata from the factorization of each tensor. We construct a tri-partite graph, as explained in Section 3.2.2, and mine constrained maximal cliques from the tripartite graphs for schemata. Table 4 provides some qualitative examples of higher-order schemata induced by TFBA. Accuracy of the schemata induced by the model is evaluated by human evaluators. In our experiments, we use human judgments from three evaluators. For every relation, the first and second columns given in Table 4 are presented to the evaluators and they are asked to validate the schema. We present top 50 schemata based on the score of the constrained maximal clique induced by TFBA to the evaluators. This evaluation protocol was also used in (Movshovitz-Attias and Cohen, 2015) for evaluating ontology induction. All evaluations were blind, i.e., the evaluators were not aware of the model they were evaluating.

**Difficulty with Computing Recall:** Even though recall is a desirable measure, due to the lack of availability of gold higher-order schema annotated corpus, it is not possible to compute recall. Although the MUC dataset has gold annotations for some predefined list of events, it does not have annotations for the relations.

Dataset	$\mathcal{X}^1 shape$	$\mathcal{X}^2 shape$	$\mathcal{X}^3 shape$
Shootings	$3365 \times 1295 \times 50$	$2569 \times 1295 \times 50$	$2569 \times 3365 \times 50$
NYT Sports	$57,820 \times 20,356 \times 50$	$49,659 \times 20,356 \times 50$	$49,659 \times 57,820 \times 50$
MUC	$2825 \times 962 \times 50$	$2555 \times 962 \times 50$	$2555 \times 2825 \times 50$

Table 2: Details of dimensions of tensors constructed for each dataset used in the experiments.

Dataset	$(r_1, r_2, r_3)$	$(\lambda_a, \lambda_b, \lambda_c)$
Shootings	(10, 20, 15)	(0.3, 0.1, 0.7)
NYT Sports	(20, 15, 15)	(0.9, 0.5, 0.7)
MUC	(15, 12, 12)	(0.7, 0.7, 0.4)

Table 3: Details of hyper-parameters set for different datasets.

Experimental results comparing performance of various models for the task of HRSI are given in Table 5. We present evaluation results from three evaluators represented as E1, E2 and E3. As can be observed from Table 5, TFBA achieves better results than HardClust for the Shootings and NYT Sports datasets, however HardClust achieves better results for the MUC dataset. Percentage agreement of the evaluators for TFBA is 72%, 70% and 60% for Shootings, NYT Sports and MUC datasets respectively.

**HardClust Limitations:** Even though HardClust gives better induction for MUC corpus, this approach has some serious drawbacks. HardClust can only induce one schema per relation. This is a restrictive constraint as multiple senses can exist for a relation. For example, consider the schemata induced for the relation *shoot* as shown in Table 4. TFBA induces two senses for the relation, but HardClust can induce only one schema. For a set of 4-tuples, HardClust can only induce ternary schemata; the dimensionality of the schemata cannot be varied. Since the latent factors induced by HardClust are entirely based on frequency, the latent categories induced by HardClust are dominated by only a fixed set of noun phrases. For example, in NYT Sports dataset, subject category induced by HardClust for all the relations is  $\langle team, yankees, mets \rangle$ . In addition to inducing only one schema per relation, most of the times HardClust only induces a fixed set of categories. Whereas for TFBA, the number of categories depends on the rank of factorization, which is a user provided parameter, thus providing more flexibility to choose the latent categories.

#### 4.1 Using Event Schema Induction for HRSI

Event schema induction is defined as the task of learning high-level representations of events, like a tournament, and their entity roles, like winning-player etc, from unlabeled text. Even though the main focus of event schema induction is to induce the important roles of the events, as a side result most of the algorithms also provide schemata for the relations. In this section, we investigate the effectiveness of these schemata compared to the ones induced by TFBA.

Event schemata are represented as a set of  $(Actor, Rel, Actor)$  triples in (Balasubramanian et al., 2013). *Actors* represent groups of noun phrases and *Rel*s represent relations. From this style of representation, however, the n-ary schemata for relations cannot be induced. Event schemata generated in (Weber et al., 2018) are similar to that in (Balasubramanian et al., 2013). Event schema induction algorithm proposed in (Nguyen et al., 2015) doesn't induce schemata for relations, but rather induces the roles for the events. For this investigation we experiment with the following algorithm.

**Chambers-13** (Chambers, 2013): This model learns event templates from text documents. Each event template provides a distribution over slots, where slots are clusters of NPs. Each event template also provides a cluster of relations, which is most likely to appear in the context of the aforementioned slots. We evaluate the schemata of these relation clusters.

As can be observed from Table 5, the proposed TFBA performs much better than Chambers-13. HardClust also performs better than Chambers-13 on all the datasets. From this analysis we infer that there is a need for algorithms which induce higher-order schemata for relations, a gap we fill in this paper. Please note that the experimental results provided in (Chambers, 2013) for MUC dataset are for the task of event schema induction, but in this work we evaluate the relation schemata. Hence the results in (Chambers, 2013) and results in this paper are not comparable. Example

Relation Schema	NPs from the induced categories	Evaluator Judgment	(Human) Suggested Label
Shootings			
$leave(A_6, B_0, C_7)$	$A_6$ : shooting, shooting incident, double shooting $B_0$ : one person, two people, three people $C_7$ : dead, injured, on edge	valid	< shooting > < people > < injured >
$identify(A_1, B_1, C_5, C_6)$	$A_1$ : police, officers, huntsville police $B_1$ : man, victims, four victims $C_5$ : sunday, shooting staurday, wednesday afternoon $C_6$ : apartment, bedroom, building in the neighborhood	valid	< police > < victim(s) > < day/time > < place >
$shoot(A_7, B_6, C_1)$	$A_7$ : gunman, shooter, smith $B_6$ : freeman, slain woman, victims $C_1$ : friday, friday night, early monday morning	valid	< perpetrator > < victim > < time >
$shoot(A_4, B_2, C_{13})$	$A_4$ : <num>-year-old man, <num>-year-old george reavis, <num>-year-old brockton man $B_2$ : in the leg, in the head, in the neck $C_{13}$ : in macon, in chicago, in an alley	valid	< victim > < body part > < location >
$say(A_1, B_1, C_5)$	$A_1$ : police, officers, huntsville police $B_1$ : man, victims, four victims $C_5$ : sunday, shooting staurday, wednesday afternoon	invalid	-
NYT sports			
$spend(A_0, B_{16}, C_3)$	$A_0$ : yankees, mets, jets $B_{16}$ : \$ <num> million, \$ <num>, \$ <num> billion $C_3$ : <num>, year, last season	valid	< team > < money > < year >
$win(A_2, B_{10}, C_3)$	$A_2$ : red sox, team, yankees $B_{10}$ : world series, title, world cup $C_3$ : <num>, year, last season	valid	< team > < championship > < year >
$get(A_4, B_4, C_1)$	$A_4$ : umpire, mike cameron, andre agassi $B_4$ : ball, lives, grounder $C_1$ : back, forward, <num>-yard line	invalid	-
MUC			
$tell(A_7, B_2, C_0)$	$A_7$ : medardo gomez, jose azcona, gregorio roza chavez $B_2$ : media, reporters, newsmen $C_0$ : today, at <num>, tonight	valid	< politician > < media > < day/time >
$occur(A_9, B_5, C_{10})$	$A_9$ : bomb, blast, explosion $B_5$ : near san salvador, here in madrid, in the same office $C_{10}$ : at <num>, this time, simultaneously	valid	< bombing > < place > < time >
$suffer(A_5, B_4, C_4)$	$A_5$ : justice maria elena diaz, vargas escobar, judge sofia de roldan $B_4$ : casualties, car bomb, grenade $C_4$ : settlement of refugees, in san roman, now	invalid	-

Table 4: Examples of schemata induced by TFBA. Please note that some of them are 3-ary while others are 4-ary. For details about schema induction, please refer to Section 3.2.

	Shootings				NYT Sports				MUC			
	E1	E2	E3	Avg	E1	E2	E3	Avg	E1	E2	E3	Avg
HardClust	0.64	0.70	0.64	0.66	0.42	0.28	0.52	0.46	0.64	0.58	0.52	<b>0.58</b>
Chambers-13	0.32	0.42	0.28	0.34	0.08	0.02	0.04	0.07	0.28	0.34	0.30	0.30
TFBA	0.82	0.78	0.68	<b>0.76</b>	0.86	0.6	0.64	<b>0.70</b>	0.58	0.38	0.48	0.48

Table 5: Higher-order RSI accuracies of various methods on the three datasets. Induced schemata for each dataset and method are evaluated by three human evaluators, E1, E2, and E3. TFBA performs better than HardClust for Shootings and NYT Sports datasets. Even though HardClust achieves better accuracy on MUC dataset, it has several limitations, see Section 4 for more details. Chambers-13 solves a slightly different problem called event schema induction, for more details about the comparison with Chambers-13 see Section 4.1.

schemata induced by TFBA and (Chambers-13) are provided as part of the supplementary material.

## 5 Conclusion

Higher order Relation Schema Induction (HRSI) is an important first step towards building domain-specific Knowledge Graphs (KGs). In this paper, we proposed TFBA, a tensor factorization-based method for higher-order RSI. To the best of our knowledge, this is the first attempt at inducing higher-order (n-ary) schemata for relations from unlabeled text. Rather than factorizing a

severely sparse higher-order tensor directly, TFBA performs *back-off* and jointly factorizes multiple lower-order tensors derived out of the higher-order tensor. In the second step, TFBA solves a constrained clique problem to induce schemata out of multiple binary schemata. We are hopeful that the backoff-based factorization idea exploited in TFBA will be useful in other sparse factorization settings.



## Acknowledgment

We thank the anonymous reviewers for their insightful comments and suggestions. This research has been supported in part by the Ministry of Human Resource Development (Government of India), Accenture, and Google.

## References

- Evrin Acar, Morten Arendt Rasmussen, Francesco Savorani, Tormod Ns, and Rasmus Bro. 2013. Understanding data fusion within the framework of coupled matrix and tensor factorizations. *Chemometrics and Intelligent Laboratory Systems* 129:53–63.
- Niranjan Balasubramanian, Stephen Soderland, Mausam, and Oren Etzioni. 2013. Generating coherent event schemas at scale. In *EMNLP*.
- Nathanael Chambers. 2013. Event schema induction with a probabilistic entity-driven model. In *EMNLP*.
- Kai-Wei Chang, Wen tau Yih, Bishan Yang, and Christopher Meek. 2014. Typed tensor decomposition of knowledge bases for relation extraction. In *EMNLP*.
- Yun-Nung Chen, William Yang Wang, Anatole Gershan, and Alexander I. Rudnicky. 2015. Matrix factorization with knowledge graph propagation for unsupervised spoken language understanding. In *ACL*.
- Jackie Chi Kit Cheung, Hoifung Poon, and Lucy Vanderwende. 2013. Probabilistic frame induction. In *NAACL-HLT*.
- Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmman, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *KDD*.
- Dora Erdos and Pauli Miettinen. 2013. Discovering facts with boolean tensor tucker decomposition. In *CIKM*.
- Francis Ferraro and Benjamin Van Durme. 2016. A unified bayesian model of scripts, frames and language. In *AAAI*.
- R. A. Harshman. 1970. Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multi-modal factor analysis. *UCLA Working Papers in Phonetics* 16(1):84.
- Yong-Deok Kim and Seungjin Choi. 2007. Nonnegative tucker decomposition. In *CVPR*.
- Tamara G Kolda and Brett W Bader. 2009. Tensor decompositions and applications. *SIAM review* 51(3):455–500.
- Joel Lang and Mirella Lapata. 2011. Unsupervised semantic role induction via split-merge clustering. In *NAACL-HLT*.
- Daniel D. Lee and H. Sebastian Seung. 2000. Algorithms for non-negative matrix factorization. In *NIPS*.
- Mausam. 2016. Open information extraction systems and downstream applications. In *IJCAI*.
- Ryan McDonald, Fernando Pereira, Seth Kulick, Scott Winters, Yang Jin, and Pete White. 2005. Simple algorithms for complex relation extraction with applications to biomedical ie. In *ACL*.
- Marvin Minsky. 1974. A framework for representing knowledge. Technical report.
- T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. 2015. Never-ending learning. In *AAAI*.
- Thahir P. Mohamed, Jr. Estevam R. Hruschka, and Tom M. Mitchell. 2011. Discovering relations between noun categories. In *EMNLP*.
- Raymond Mooney and Gerald DeJong. 1985. Learning schemata for natural language processing. In *IJCAI*.
- Dana Movshovitz-Attias and William W. Cohen. 2015. Kb-lda: Jointly learning a knowledge base of hierarchy, relations, and facts. In *ACL*.
- Brian Murphy, Partha Talukdar, and Tom Mitchell. 2012. Learning effective and interpretable semantic models using non-negative sparse embedding. In *COLING*.
- Kiem-Hieu Nguyen, Xavier Tannier, Olivier Ferret, and Romaric Besançon. 2015. Generative event schema induction with entity disambiguation. In *ACL*.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A three-way model for collective learning on multi-relational data. In *ICML*.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2012. Factorizing yago: Scalable machine learning for linked data. In *WWW*.
- Madhav Nimishakavi, Uday Singh Saini, and Partha Talukdar. 2016. Relation schema induction using tensor factorization with side information. In *EMNLP*.
- Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova, and Wen-tau Yih. 2017. Cross-sentence n-ary relation extraction with graph lstms. *TACL* 5:101–115.

- Karl Pichotta and Raymond J. Mooney. 2014. Statistical script learning with multi-argument events. In *EACL*.
- Karl Pichotta and Raymond J. Mooney. 2016. Learning statistical scripts with lstm recurrent neural networks. In *AAAI*.
- Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M. Marlin. 2013. Relation extraction with matrix factorization and universal schemas. In *NAACL-HLT*.
- Michael Roth and Mirella Lapata. 2016. Neural semantic role labeling with dependency path embeddings. In *ACL*.
- R. Schank and R. Abelson. 1977. *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Sameer Singh, Tim Rocktäschel, and Sebastian Riedel. 2015. Towards Combined Matrix and Tensor Factorization for Universal Schema Relation Extraction. In *NAACL Workshop on Vector Space Modeling for NLP (VSM)*.
- Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *WWW*.
- Ivan Titov and Ehsan Khoddam. 2015. Unsupervised induction of semantic roles within a reconstruction-error minimization framework. In *NAACL-HLT*.
- L. R. Tucker. 1963. Implications of factor analysis of three-way matrices for measurement of change. In *Problems in measuring change.*, University of Wisconsin Press, Madison WI, pages 122–137.
- Yichen Wang, Robert Chen, Joydeep Ghosh, Joshua C. Denny, Abel N. Kho, You Chen, Bradley A. Malin, and Jimeng Sun. 2015. Rubik: Knowledge guided tensor factorization and completion for health data analytics. In *KDD*.
- Noah Weber, Niranjan Balasubramanian, and Nathanael Chambers. 2018. Event representations with tensor-based compositions. In *AAAI*.