

Tensor Ring Decomposition with Rank Minimization on Latent Space: An Efficient Approach for Tensor Completion

Longhao Yuan,^{1,2} Chao Li,² Danilo Mandic,⁵ Jianting Cao,^{1,2,4,†} Qibin Zhao^{2,3,†}

¹Graduate School of Engineering, Saitama Institute of Technology, Japan

²Tensor Learning Unit, RIKEN Center for Advanced Intelligence Project (AIP), Japan

³School of Automation, Guangdong University of Technology, China

⁴School of Computer Science and Technology, Hangzhou Dianzi University, China

⁵Department of Electrical and Electronic Engineering, Imperial College London, United Kingdom

Abstract

In tensor completion tasks, the traditional low-rank tensor decomposition models suffer from the laborious model selection problem due to their high model sensitivity. In particular, for tensor ring (TR) decomposition, the number of model possibilities grows exponentially with the tensor order, which makes it rather challenging to find the optimal TR decomposition. In this paper, by exploiting the low-rank structure of the TR latent space, we propose a novel tensor completion method which is robust to model selection. In contrast to imposing the low-rank constraint on the data space, we introduce nuclear norm regularization on the latent TR factors, resulting in the optimization step using singular value decomposition (SVD) being performed at a much smaller scale. By leveraging the alternating direction method of multipliers (ADMM) scheme, the latent TR factors with optimal rank and the recovered tensor can be obtained simultaneously. Our proposed algorithm is shown to effectively alleviate the burden of TR-rank selection, thereby greatly reducing the computational cost. The extensive experimental results on both synthetic and real-world data demonstrate the superior performance and efficiency of the proposed approach against the state-of-the-art algorithms.

Introduction

Tensor decompositions aim to find latent factors in tensor-valued data (i.e., the generalization of multi-dimensional arrays), thereby casting large-scale and intractable tensor problems into a multilinear tensor latent space of low-dimensionality (very few degrees of freedom designated by the rank). The latent factors within tensor decomposition can be considered as the latent features of data, which makes them an ideal set of bases to predict missing entries when the acquired data is incomplete. The specific forms and operations among latent factors determine the type of tensor decomposition. The most classical and successful tensor decomposition models are the Tucker decomposition

(TKD) and the CANDECOMP/PARAFAC (CP) decomposition (Kolda and Bader 2009). More recently, the matrix product state/tensor-train (MPS/TT) decomposition has become very attractive, owing to its super-compression and computational efficiency properties (Oseledets 2011). Currently, a generalization of TT decomposition, termed the tensor ring (TR) decomposition, has been studied across scientific disciplines (Zhao et al. 2016a; 2018). These tensor decomposition models have found application in various fields such as machine learning (Wang et al. 2018; Novikov et al. 2015; Anandkumar et al. 2014; Kanagawa et al. 2016), signal processing (Cong et al. 2015), image/video completion (Liu et al. 2013; Zhao et al. 2016b), compressed sensing (Gandy, Recht, and Yamada 2011), to name but a few. Tensor completion is one of the most important applications of tensor decompositions, with the goal to recover an incomplete tensor from partially observed entries. The theoretical lynchpin in tensor completion problems is the tensor low-rank assumption, and the methods can mainly be categorized into two types: (i) tensor-decomposition-based approach and (ii) rank-minimization-based approach.

Tensor decomposition based methods find latent factors of tensor by the incomplete tensor, and then the latent factors are used to predict the missing entries. Many completion algorithms have been proposed based on alternating least squares (ALS) method (Grasedyck, Kluge, and Kramer 2015; Wang, Aggarwal, and Aeron 2017), gradient-based method (Yuan, Zhao, and Cao 2017; Acar et al. 2011), to mention but a few. Though ALS and gradient-based algorithms are free from burdensome hyper-parameter tuning, the performance of these algorithms is rather sensitive to model selection, i.e., rank selection of the tensor decomposition. Moreover, since the optimal rank is generally data-dependent, it is very challenging to specify the optimal rank beforehand. This is especially the case for Tucker, TT, and TR decompositions, for which the rank is defined as a vector; it is therefore impossible to find the optimal ranks by cross-validation due to the immense possibilities.

Rank minimization based methods employ convex surrogates to minimize the tensor rank. One of the most commonly-used surrogates is the nuclear norm (a.k.a. Schatten norm, or trace norm), which is defined as the sum of singular values of a matrix and it is the most popular convex surrogate for rank regularization. Based on differ-

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This work was supported by JSPS KAKENHI (Grant No. 17K00326, 15H04002, 18K04178), JST CREST (Grant No. JPMJCR1784), and the National Natural Science Foundation of China (Grant No. 61773129).

The code is available at: <https://github.com/yuanlonghao/TRLRF>

†Corresponding authors: qibin.zhao@riken.jp, cao@sit.ac.jp

ent definitions of tensor rank, various nuclear norm regularized algorithms have been proposed (Liu et al. 2013; Imaizumi, Maehara, and Hayashi 2017; Liu et al. 2014; 2015). Rank minimization based methods do not need to specify the rank of the employed tensor decompositions beforehand, and the rank of the recovered tensor will be automatically learned from the limited observations. However, these algorithms face multiple large-scale singular value decomposition (SVD) operations on the 2D unfoldings of the tensor when employing the nuclear norm and numerous hyper-parameter tuning, which in turn leads to high computational cost and low efficiency.

To address the problems of high sensitivity to rank selection and low computational efficiency which are inherent in traditional tensor completion methods, in this paper, we propose a new algorithm named tensor ring low-rank factors (TRLRF) which effectively alleviates the burden of rank selection and reduces the computational cost. By virtue of employing both nuclear norm regularization and tensor decomposition, our model provides performance stability and high computational efficiency. The proposed TRLRF is efficiently solved by the ADMM algorithm and it simultaneously achieves both the underlying tensor decomposition and completion based on TR decomposition. Our main contributions in this paper are:

- A theoretical relationship between the multilinear tensor rank and the rank of TR factors is established, which allows the low-rank constraint to be performed implicitly on TR latent space. This has led to fast SVD calculation on small size factors.
- The nuclear norm is further imposed to regularize the TR-ranks, which enables our algorithm to always obtain a stable solution, even if the TR-rank is inappropriately given. This highlights rank-robustness of the proposed TRLRF algorithm.
- An efficient algorithm based on ADMM is developed to optimize the proposed model, so as to obtain the TR-factors and the recovered tensor simultaneously.

Preliminaries and Related Works

Notations

The notations in (Kolda and Bader 2009) are adopted in this paper. A scalar is denoted by a standard lowercase letter or an uppercase letter, e.g., $x, X \in \mathbb{R}$, and a vector is denoted by a boldface lowercase letter, e.g., $\mathbf{x} \in \mathbb{R}^I$. A matrix is denoted by a boldface capital letter, e.g., $\mathbf{X} \in \mathbb{R}^{I \times J}$. A tensor of order $N \geq 3$ is denoted by calligraphic letters, e.g., $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$. The set $\{\mathcal{X}^{(n)}\}_{n=1}^N := \{\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \dots, \mathcal{X}^{(N)}\}$ denotes a tensor sequence, with $\mathcal{X}^{(n)}$ being the n -th tensor of the sequence. Where appropriate, a tensor sequence can also be written as $[\mathcal{X}]$. The representations of matrix sequences and vector sequences are designated in the same way. An element of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ of index (i_1, i_2, \dots, i_N) is denoted by $\mathcal{X}(i_1, i_2, \dots, i_N)$ or $x_{i_1 i_2 \dots i_N}$. The inner product of two tensors \mathcal{X}, \mathcal{Y} with the same size $\mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is defined

as $\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{i_1} \sum_{i_2} \dots \sum_{i_N} x_{i_1 i_2 \dots i_N} y_{i_1 i_2 \dots i_N}$. Furthermore, the Frobenius norm of \mathcal{X} is defined by $\|\mathcal{X}\|_F = \sqrt{\langle \mathcal{X}, \mathcal{X} \rangle}$.

We employ two types of tensor unfolding (matricization) operations in this paper. The standard mode- n unfolding (Kolda and Bader 2009) of tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is denoted by $\mathbf{X}_{(n)} \in \mathbb{R}^{I_n \times I_1 \dots I_{n-1} I_{n+1} \dots I_N}$. Another mode- n unfolding of tensor \mathcal{X} which is often used in TR operations (Zhao et al. 2016a) is denoted by $\mathbf{X}_{\langle n \rangle} \in \mathbb{R}^{I_n \times I_{n+1} \dots I_N I_1 \dots I_{n-1}}$. Furthermore, the inverse operation of unfolding is matrix folding (tensorization), which transforms matrices to higher-order tensors. In this paper, we only define the folding operation for the first type of mode- n unfolding as $\text{fold}_n(\cdot)$, i.e., for a tensor \mathcal{X} , we have $\text{fold}_n(\mathbf{X}_{(n)}) = \mathcal{X}$.

Tensor ring decomposition

The tensor ring (TR) decomposition is a more general decomposition model than the tensor-train (TT) decomposition. It represents a tensor of higher-order by circular multilinear products over a sequence of low-order latent core tensors, i.e., TR factors. For $n = 1, \dots, N$, the TR factors are denoted by $\mathcal{G}^{(n)} \in \mathbb{R}^{R_n \times I_n \times R_{n+1}}$ and each consists of two rank-modes (i.e., mode-1 and mode-3) and one dimension-mode (i.e., mode-2). The syntax $\{R_1, R_2, \dots, R_{N+1}\}$ denotes the TR-rank which controls the model complexity of TR decomposition. The TR decomposition applies trace operations and all of the TR factors are set to be 3rd-order; thus the TR decomposition relaxes the rank constraint on the first and last core of TT to $R_1 = R_{N+1}$. Moreover, TR decomposition linearly scales to the order of the tensor, and in this way it overcomes the ‘curse of dimensionality’. In this case, TR can be considered as a linear combination of TTs and hence offers a powerful and generalized representation ability. The element-wise relation of TR decomposition and the generated tensor is given by:

$$\mathcal{X}(i_1, i_2, \dots, i_N) = \text{Trace} \left\{ \prod_{n=1}^N \mathbf{G}_{i_n}^{(n)} \right\}, \quad (1)$$

where $\text{Trace}\{\cdot\}$ is the matrix trace operation, $\mathbf{G}_{i_n}^{(n)} \in \mathbb{R}^{R_n \times R_{n+1}}$ is the i_n th mode-2 slice matrix of $\mathcal{G}^{(n)}$, which can also be denoted by $\mathcal{G}^{(n)}(:, i_n, :)$ according to the Matlab notation.

Tensor completion

Completion by TR decomposition Tensor decomposition based algorithms do not directly employ the rank constraint to the object tensor. Instead, they try to find the low-rank representation (i.e., tensor decompositions) of the incomplete data from the observed entries. The obtained latent factors of the tensor decomposition are used to predict the missing entries. For model formulation, the tensor completion problem is set as a weighted least squares (WLS) model. Based on different tensor decompositions, various tensor completion algorithms have been proposed, e.g., weighted CP (Acar et al. 2011), weighted Tucker (Filipović and Jukić

2015), TRWOPT (Yuan et al. 2018) and TRALS (Wang, Aggarwal, and Aeron 2017). To the best of our knowledge, there are two proposed TR-based tensor completion algorithms: the TRALS and TRWOPT. They apply the same optimization model which is formulated as:

$$\min_{[\mathcal{G}]} \|P_{\Omega}(\mathcal{T} - \Psi([\mathcal{G}]))\|_F^2, \quad (2)$$

where the optimization objective is the TR factors, $[\mathcal{G}]$, $P_{\Omega}(\mathcal{T})$ denotes all the observed entries w.r.t. the set of indices of observed entries represented by Ω , and $\Psi([\mathcal{G}])$ denotes the approximated tensor generated by $[\mathcal{G}]$. Every element of $\Psi([\mathcal{G}])$ is calculated by equation (1). The two algorithms are both based on the model in (2). However, TRALS applies alternative least squares (ALS) method and TRWOPT uses a gradient-based algorithm to solve the model, respectively. They perform well for both low-order and high-order tensors due to the high representation ability and flexibility of TR decomposition. However, these algorithms are shown to suffer from high sensitivity to rank selection, which would lead to high computational cost.

Completion by nuclear norm regularization The model of rank minimization-based tensor completion can be formulated as:

$$\min_{\mathcal{X}} \text{Rank}(\mathcal{X}) + \frac{\lambda}{2} \|P_{\Omega}(\mathcal{T} - \mathcal{X})\|_F^2, \quad (3)$$

where \mathcal{X} is the recovered low-rank tensor, and $\text{Rank}(\cdot)$ is a rank regularizer. The model can therefore find the low-rank structure of the data and approximate the recovered tensor. Because determining the tensor rank is an NP-hard problem (Hillar and Lim 2013; Kolda and Bader 2009), work in (Liu et al. 2013) and (Signoretto et al. 2014) extends the concept of low-rank matrix completion and defines tensor rank as a sum of the rank of mode- n unfolding of the object tensor. Moreover, the convex surrogate named nuclear norm is applied to the tensor low-rank model and it simultaneously regularizes all the mode- n unfoldings of the object tensor. In this way, the model in (3) can be reformulated as:

$$\min_{\mathcal{X}} \sum_{n=1}^N \|\mathbf{X}_{(n)}\|_* + \frac{\lambda}{2} \|P_{\Omega}(\mathcal{T} - \mathcal{X})\|_F^2, \quad (4)$$

where $\|\cdot\|_*$ denotes the nuclear norm regularization in the form of a sum of the singular values of the matrix. Usually, the model is solved by ADMM algorithms and it is shown to have fast convergence and good performance when data size is small. However, when dealing with large-scale data, the multiple SVD operations in the optimization step will be intractable due to high computational cost.

Tensor Ring Low-rank Factors

To solve the issues traditional tensor completion methods have, we impose low-rankness on each of the TR factors and so that our basic tensor completion model is formulated as follow:

$$\min_{[\mathcal{G}], \mathcal{X}} \sum_{n=1}^N \|\mathcal{G}^{(n)}\|_* + \frac{\lambda}{2} \|\mathcal{X} - \Psi([\mathcal{G}])\|_F^2, \quad (5)$$

s.t. $P_{\Omega}(\mathcal{X}) = P_{\Omega}(\mathcal{T})$.

To solve (5), we first need to deduce the relation of the tensor rank and the corresponding core tensor rank, which can be explained by the following theorem.

Theorem 1. *Given an N -th order tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ which can be represented by equation (1), then the following inequality holds for all $n = 1, \dots, N$:*

$$\text{Rank}(\mathbf{G}_{(2)}^{(n)}) \geq \text{Rank}(\mathbf{X}_{(n)}). \quad (6)$$

Proof. For the n -th core tensor $\mathcal{G}^{(n)}$, according to the work in (Zhao et al. 2016a), we have:

$$\mathbf{X}_{\langle n \rangle} = \mathbf{G}_{(2)}^{(n)} (\mathbf{G}_{\langle 2 \rangle}^{(\neq n)})^{\top}, \quad (7)$$

where $\mathcal{G}^{(\neq n)} \in \mathbb{R}^{R_{n+1} \times \prod_{i=1, i \neq n}^N I_i \times R_n}$ is a subchain tensor generated by merging all but the n -th core tensor. Hence, the relation of the rank satisfies:

$$\begin{aligned} \text{Rank}(\mathbf{X}_{\langle n \rangle}) &\leq \min\{\text{Rank}(\mathbf{G}_{(2)}^{(n)}), \text{Rank}(\mathbf{G}_{\langle 2 \rangle}^{(\neq n)})\} \\ &\leq \text{Rank}(\mathbf{G}_{(2)}^{(n)}). \end{aligned} \quad (8)$$

The proof is completed by

$$\text{Rank}(\mathbf{X}_{\langle n \rangle}) = \text{Rank}(\mathbf{X}_{(n)}) \leq \text{Rank}(\mathbf{G}_{(2)}^{(n)}). \quad (9)$$

□

This theorem proves the relation between the tensor rank and the rank of the TR factors. The rank of mode- n unfolding of the tensor \mathcal{X} is upper bounded by the rank of the dimension-mode unfolding of the corresponding core tensor $\mathcal{G}^{(n)}$, which allows us to impose a low-rank constraint on $\mathcal{G}^{(n)}$. By the new surrogate, our model (5) is reformulated by:

$$\begin{aligned} \min_{[\mathcal{G}], \mathcal{X}} \sum_{n=1}^N \|\mathbf{G}_{(2)}^{(n)}\|_* + \frac{\lambda}{2} \|\mathcal{X} - \Psi([\mathcal{G}])\|_F^2 \\ \text{s.t. } P_{\Omega}(\mathcal{X}) = P_{\Omega}(\mathcal{T}). \end{aligned} \quad (10)$$

The above model imposes nuclear norm regularization on the dimension-mode unfoldings of the TR factors, which can largely decrease the computational complexity compared to the algorithms which are based on model (4). Moreover, we consider to give low-rank constraints on the two rank-modes of the TR factors, i.e., the unfoldings of the TR factors along mode-1 and mode-3, which can be expressed by $\sum_{n=1}^N \|\mathbf{G}_{(1)}^{(n)}\|_* + \sum_{n=1}^N \|\mathbf{G}_{(3)}^{(n)}\|_*$. When the model is optimized, nuclear norms of the rank-mode unfoldings and the fitting error of the approximated tensor are minimized simultaneously, resulting in the initial TR-rank becoming the upper bound of the real TR-rank of the tensor, thus equipping our model with robustness to rank selection. The tensor ring low-rank factors (TRLRF) model can be finally expressed as:

$$\min_{[\mathcal{G}], \mathcal{X}} \sum_{n=1}^N \sum_{i=1}^3 \|\mathbf{G}_{(i)}^{(n)}\|_* + \frac{\lambda}{2} \|\mathcal{X} - \Psi([\mathcal{G}])\|_F^2 \quad (11)$$

s.t. $P_{\Omega}(\mathcal{X}) = P_{\Omega}(\mathcal{T})$.

Our TRLRF model has two distinctive advantages. Firstly, the low-rank assumption is placed on tensor factors instead of on the original tensor, this greatly reduces the computational complexity of the SVD operation. Secondly, low-rankness of tensor factors can enhance the robustness to rank selection, which can alleviate the burden of searching for optimal TR-rank and reduce the computational cost in the implementation.

Solving scheme

To solve the model in (11), we apply the alternating direction method of multipliers (ADMM) which is efficient and widely used (Boyd et al. 2011). Moreover, because the variables of TRLRF model are inter-dependent, we impose auxiliary variables to simplify the optimization. Thus, the TRLRF model can be rewritten as

$$\begin{aligned} \min_{[\mathcal{M}], [\mathcal{G}], \mathcal{X}} \quad & \sum_{n=1}^N \sum_{i=1}^3 \|\mathbf{M}_{(i)}^{(n,i)}\|_* + \frac{\lambda}{2} \|\mathcal{X} - \Psi([\mathcal{G}])\|_F^2, \\ \text{s.t.} \quad & \mathbf{M}_{(i)}^{(n,i)} = \mathbf{G}_{(i)}^{(n)}, n = 1, \dots, N, i = 1, 2, 3, \\ & P_{\Omega}(\mathcal{X}) = P_{\Omega}(\mathcal{T}), \end{aligned} \quad (12)$$

where $[\mathcal{M}] := \{\mathcal{M}^{(n,i)}\}_{n=1, i=1}^{N,3}$ are the auxiliary variables of $[\mathcal{G}]$. By merging the equal constraints of the auxiliary variables into the Lagrangian equation, the augmented Lagrangian function of TRLRF model becomes

$$\begin{aligned} L([\mathcal{G}], \mathcal{X}, [\mathcal{M}], [\mathcal{Y}]) \\ = \sum_{n=1}^N \sum_{i=1}^3 (\|\mathbf{M}_{(i)}^{(n,i)}\|_* + \langle \mathcal{Y}^{(n,i)}, \mathcal{M}^{(n,i)} - \mathcal{G}^{(n)} \rangle \\ + \frac{\mu}{2} \|\mathcal{M}^{(n,i)} - \mathcal{G}^{(n)}\|_F^2) + \frac{\lambda}{2} \|\mathcal{X} - \Psi([\mathcal{G}])\|_F^2, \\ \text{s.t.} \quad P_{\Omega}(\mathcal{X}) = P_{\Omega}(\mathcal{T}), \end{aligned} \quad (13)$$

where $[\mathcal{Y}] := \{\mathcal{Y}^{(n,i)}\}_{n=1, i=1}^{N,3}$ are the Lagrangian multipliers, and $\mu > 0$ is a penalty parameter. For $n = 1, \dots, N$, $i = 1, 2, 3$, $\mathcal{G}^{(n)}$, $\mathcal{M}^{(n,i)}$ and $\mathcal{Y}^{(n,i)}$ are each independent, so we can update them by the updating scheme below.

Update of $\mathcal{G}^{(n)}$. By using (13), the augmented Lagrangian function w.r.t. $\mathcal{G}^{(n)}$ can be simplified as

$$\begin{aligned} L(\mathcal{G}^{(n)}) = \sum_{i=1}^3 \frac{\mu}{2} \|\mathcal{M}^{(n,i)} - \mathcal{G}^{(n)} + \frac{1}{\mu} \mathcal{Y}^{(n,i)}\|_F^2 \\ + \frac{\lambda}{2} \|\mathcal{X} - \Psi([\mathcal{G}])\|_F^2 + C_{\mathcal{G}}, \end{aligned} \quad (14)$$

where the constant $C_{\mathcal{G}}$ consists of other parts of the Lagrangian function which is irrelevant to updating $\mathcal{G}^{(n)}$. This is a least squares problem, so for $n = 1, \dots, N$, $\mathcal{G}^{(n)}$ can be updated by

$$\begin{aligned} \mathcal{G}_+^{(n)} = \text{fold}_2 \left(\left(\sum_{i=1}^3 (\mu \mathbf{M}_{(2)}^{(n,i)} + \mathbf{Y}_{(2)}^{(n,i)}) \right) \right. \\ \left. + \lambda \mathbf{X}_{\langle n \rangle} \mathbf{G}_{\langle 2 \rangle}^{(\neq n)} \right) (\lambda \mathbf{G}_{\langle 2 \rangle}^{(\neq n), T} \mathbf{G}_{\langle 2 \rangle}^{(\neq n)} + 3\mu \mathbf{I})^{-1}, \end{aligned} \quad (15)$$

where $\mathbf{I} \in \mathbb{R}^{R_n^2 \times R_n^2}$ denotes the identity matrix.

Update of $\mathcal{M}^{(n,i)}$. For $i = 1, 2, 3$, the augmented Lagrangian functions w.r.t. $[\mathcal{M}]$ is expressed as

$$\begin{aligned} L(\mathcal{M}^{(n,i)}) = \frac{\mu}{2} \|\mathcal{M}^{(n,i)} - \mathcal{G}^{(n)} + \frac{1}{\mu} \mathcal{Y}^{(n,i)}\|_F^2 \\ + \|\mathbf{M}_{(i)}^{(n,i)}\|_* + C_{\mathcal{M}}. \end{aligned} \quad (16)$$

The above formulation has a closed-form (Cai, Candès, and Shen 2010), which is given by

$$\mathcal{M}_+^{(n,i)} = \text{fold}_i \left(D_{\frac{1}{\mu}} \left(\mathbf{G}_{(i)}^{(n)} - \frac{1}{\mu} \mathbf{Y}_{(i)}^{(n,i)} \right) \right), \quad (17)$$

where $D_{\beta}(\cdot)$ is the singular value thresholding (SVT) operation, e.g., if $\mathbf{U}\mathbf{S}\mathbf{V}^T$ is the singular value decomposition of matrix \mathbf{A} , then $D_{\beta}(\mathbf{A}) = \mathbf{U} \text{max}\{\mathbf{S} - \beta \mathbf{I}, 0\} \mathbf{V}^T$.

Update of \mathcal{X} . The augmented Lagrangian functions w.r.t. \mathcal{X} is given by

$$\begin{aligned} L(\mathcal{X}) = \frac{\lambda}{2} \|\mathcal{X} - \Psi([\mathcal{G}])\|_F^2 + C_{\mathcal{X}}, \\ \text{s.t.} \quad P_{\Omega}(\mathcal{X}) = P_{\Omega}(\mathcal{T}), \end{aligned} \quad (18)$$

which is equivalent to the tensor decomposition based model in (2). The expression for \mathcal{X} is updated by inputting the observed values in the corresponding entries, and by approximating the missing entries by updated TR factors $[\mathcal{G}]$ for every iteration, i.e.,

$$\mathcal{X}_+ = P_{\Omega}(\mathcal{T}) + P_{\Omega}(\Psi([\mathcal{G}])), \quad (19)$$

where $\bar{\Omega}$ is the set of indices of missing entries which is a complement to Ω .

Update of $\mathcal{Y}^{(n,i)}$. For $n = 1, \dots, N$ and $i = 1, 2, 3$, the Lagrangian multiplier $\mathcal{Y}^{(n,i)}$ is updated as

$$\mathcal{Y}_+^{(n,i)} = \mathcal{Y}^{(n,i)} + \mu (\mathcal{M}^{(n,i)} - \mathcal{G}^{(n)}). \quad (20)$$

In addition, the penalty term of the Lagrangian functions L is restricted by μ which is also updated for every iteration by $\mu_+ = \text{max}\{\rho\mu, \mu_{max}\}$, where $1 < \rho < 1.5$ is a tuning hyper parameter.

The ADMM based solving scheme is updated iteratively based on the above equations. Moreover, we consider to set two optimization stopping conditions: (i) maximum number of iterations k_{max} and (ii) the difference between two iterations (i.e., $\|\mathcal{X} - \mathcal{X}_{last}\|_F / \|\mathcal{X}\|_F$) which is thresholded by the tolerance tol . The implementation process and hyper-parameter selection of TRLRF is summarized in Algorithm 1. It should be noted that our TRLRF model is non-convex, so the convergence to the global minimum cannot be theoretically guaranteed. However, the convergence of our algorithm can be verified empirically (see experiment details in Figure 1). Moreover, the extensive experimental results in the next section also illustrate the stability and effectiveness of TRLRF.

Computational complexity

We analyze the computational complexity of our TRLRF algorithm as follows. For a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$, the

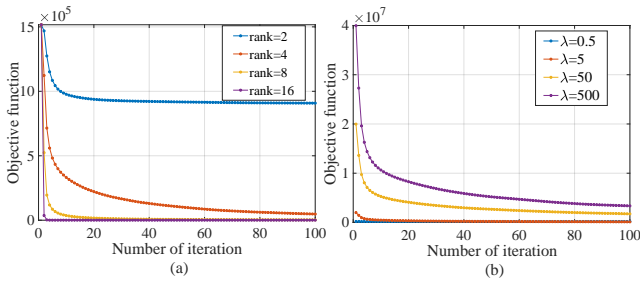


Figure 1: Illustration of convergence for TRLRF under different hyper-parameter choices. A synthetic tensor with TR structure (size $7 \times 8 \times 7 \times 8$ with TR-rank $\{4,4,4,4\}$, missing rate 0.5) is tested. The experiment records the change of the objective function values along the number of iterations. Each independent experiment is conducted 100 times and the average results are shown in the graphs. Panels (a) and (b) represent the convergence curve when TR-rank and λ are changed respectively.

Algorithm 1. Tensor ring low-rank factors (TRLRF)

- 1: **Input:** $P_{\Omega}(\mathcal{T})$, initial TR-rank $\{R_n\}_{n=1}^N$.
- 2: **Initialization:** For $n = 1, \dots, N, i = 1, 2, 3$, random sample $\mathcal{G}^{(n)}$ by distribution $N \sim (0, 1)$, $\mathcal{Y}^{(n,i)} = 0, \mathcal{M}^{(n,i)} = 0, \lambda = 5, \mu^0 = 1, \mu_{max} = 10^2, \rho = 1.01, tol = 10^{-6}, k = 0, k_{max} = 300$.
- 3: **For** $k = 1$ to k_{max} **do**
- 4: $\mathcal{X}_{last} = \mathcal{X}$.
- 5: Update $\{\mathcal{G}^{(n)}\}_{n=1}^N$ by (15).
- 6: Update $\{\mathcal{M}^{(n,i)}\}_{n=1, i=1}^{N,3}$ by (17).
- 7: Update \mathcal{X} by (19).
- 8: Update $\{\mathcal{Y}^{(n,i)}\}_{n=1, i=1}^{N,3}$ by (20).
- 9: $\mu = \max(\rho\mu, \mu_{max})$
- 6: **If** $\|\mathcal{X} - \mathcal{X}_{last}\|_F / \|\mathcal{X}\|_F < tol$, **break**
- 7: **End for**
- 8: **Output:** completed tensor \mathcal{X} and TR factors $[\mathcal{G}]$.

TR-rank is set as $R_1 = R_2 = \dots = R_N = R$, then the computational complexity of updating $[\mathcal{M}]$ represents mainly the cost of SVD operation, which is $\mathcal{O}(\sum_{n=1}^N 2I_n R^3 + I_n^2 R^2)$. The computational complexities incurred calculating $\mathcal{G}_{\langle 2 \rangle}^{(\neq n)}$ and updating $[\mathcal{G}]$ are $\mathcal{O}(NR^3 \prod_{i=1, i \neq n}^N I_i)$ and $\mathcal{O}(NR^2 \prod_{i=1}^N I_i + NR^6)$, respectively. If we assume $I_1 = I_2 = \dots = I_N = I$, then overall complexity of our proposed algorithm can be written as $\mathcal{O}(NR^2 I^N + NR^6)$.

Compared to HaLRTC and TRALS which are the representative of the nuclear-norm-based and the tensor decomposition based algorithms, the computational complexity of HaLRTC is $\mathcal{O}(NI^{N+1})$. Since TRALS is based on ALS method and TR decomposition, its computational complexity is $\mathcal{O}(PNR^4 I^N + NR^6)$, where P denotes the observation rate. We can see that the computational complexity of our TRLRF is similar to that of the two related algorithms. However, the desirable characteristic of rank selection ro-

bustness of our algorithm can help relieve the workload for model selection in practice, and thus the computational cost can be reduced. Moreover, though the computational complexity of TRLRF is of high power in R , due to the high representation ability and flexibility of TR decomposition, the TR-rank is always set as a small value. In addition, from experiments, we find out that our algorithm is capable of working efficiently for high-order tensors so that we can tensorize the data to a higher-order tensor and choose a small TR-rank to reduce the computational complexity.

Experimental Results

Synthetic data

We first conducted experiments to testify the rank robustness of our algorithm by comparing TRALS, TRWOPT, and our TRLRF. To verify the performance of the three algorithms, we tested two tensors of size $20 \times 20 \times 20 \times 20$ and $7 \times 8 \times 7 \times 8 \times 7 \times 8$. The tensors were generated by TR factors of TR-ranks $\{6, 6, 6, 6\}$ and $\{4, 4, 4, 4, 4, 4\}$ respectively. The values of the TR factors were drawn from i.i.d. Gaussian distribution $N \sim (0, 0.5)$. The observed entries of the tensors were randomly removed by a missing rate of 0.5, where the missing rate is calculated by $1 - M/\text{num}(\mathcal{T}_{real})$ and M is the number of sampled entries (i.e., observed entries). We recorded the completion performance of the three algorithms by selecting different TR-ranks. The evaluation index was RSE which is defined by $RSE = \|\mathcal{T}_{real} - \mathcal{X}\|_F / \|\mathcal{T}_{real}\|_F$, where \mathcal{T}_{real} is the known tensor with full observations and \mathcal{X} is the recovered tensor calculated by each tensor completion algorithm. The hyper-parameters of our TRLRF were set according to Algorithm 1. All the hyper-parameters of TRALS and TRWOPT are set according to the recommended settings in the corresponding papers to get the best results.

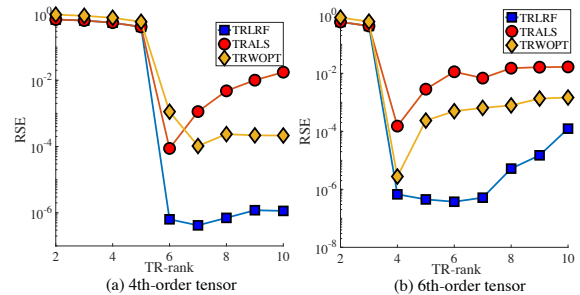


Figure 2: Completion performance of three TR-based algorithms in the synthetic data experiment. The RSE values of different selected TR-ranks are recorded. The missing rate of the two target tensors equals 0.5 and the real TR-ranks are 6 and 4 respectively.

Figure 2 shows the final RSE results which represent the average values of 100 independent experiments for each case. From the figure, we can see that all the three algorithms had their lowest RSE values when the real TR-ranks of the tensors were chosen and the best performance was obtained from our TRLRF. Moreover, when the TR-rank increased,

the performance of TRLRF remained stable while the performance of the other two compared algorithms fell drastically. This indicates that imposing low-rankness assumption on the TR factors can bring robustness to rank selection, which largely alleviates the model selection problem in the experiments.

Benchmark images inpainting



Figure 3: The eight benchmark images. The first image is named “Lena” and is used in the next two experiments.

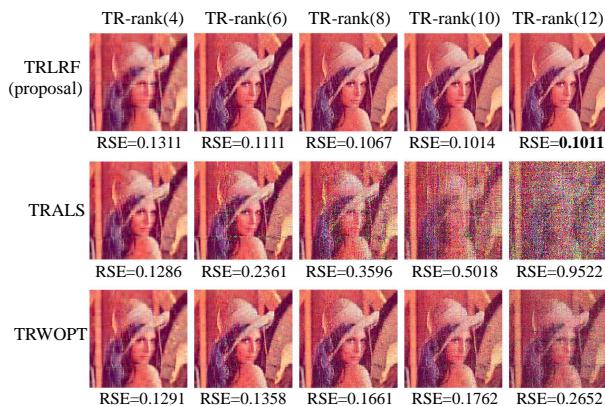


Figure 4: Visual completion results of the TRLRF (proposed), TRALS, and TRWOPT on image “Lena” with different TR-ranks, when the missing rate is 0.8. The selected TR-ranks are 4, 6, 8, 10, 12 respectively, from the first column to the last column. The RSE results are noted under each picture.

In this section, we tested our TRLRF against the state-of-the-art algorithms on eight benchmark images which are shown in Figure 3. The size of each RGB image was $256 \times 256 \times 3$ which can be considered as a 3rd-order tensor. For the first experiment, we continued to verify the TR-rank robustness of TRLRF on the image named “Lena”. Figure 4 shows the completion results of TRLRF, TRALS, and TRWOPT when different TR-ranks for each algorithm are selected. The missing rate of the image was set as 0.8, which is the case that the TR decompositions are prone to overfitting. From the figure, we can see that our TRLRF gives better results than the other two TR-based algorithms in each case and the highest performance was obtained when the TR rank was set as 12. When TR-rank increases, the completion performance of TRALS and TRLRF decreases due to redundant

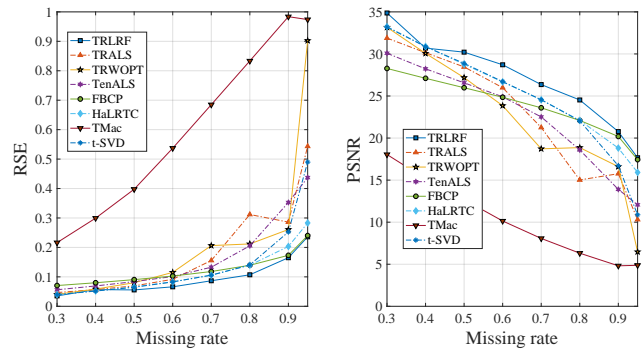


Figure 5: Average completion performance of the eight considered algorithms, under different data missing rates.

model complexity and overfitting of the algorithms, while our TRLRF shows better results even the selected TR-rank is larger than the desired TR-rank.

In the next experiment, we compared our TRLRF to the two TR-based algorithm, TRALS and TRWOPT, and the other state-of-the-art algorithms, i.e., TenALS (Jain and Oh 2014), FBCP (Zhao, Zhang, and Cichocki 2015), HaLRTC (Liu et al. 2013), TMac (Xu et al. 2013) and t-SVD (Zhang et al. 2014). We tested these algorithms on all the eight benchmark images and for different missing rates: 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9 and 0.95. The relative square error (RSE) and peak signal-to-noise ratio (PSNR) were adopted for the evaluation of the completion performance. For RGB image data, PSNR is defined as $PSNR = 10 \log_{10}(255^2/MSE)$ where MSE is calculated by $MSE = \|\mathcal{T}_{real} - \mathcal{X}\|_F^2 / \text{num}(\mathcal{T}_{real})$, and $\text{num}(\cdot)$ denotes the number of element of the fully observed tensor.

For the three TR-based algorithms, we assumed the TR-ranks were equal for every core tensor (i.e., $R_1 = R_2 = \dots = R_N$). The best completion results for each algorithm were obtained by selecting the best TR-ranks for the TR-based algorithms by a cross-validation method. Actually, finding the best TR-rank to obtain the best completion results is very tedious. However, this is much easier for our proposed algorithm because the performance of TRLRF is fairly stable even though the TR-rank is selected from a wide large. For the other five compared algorithms, we tuned the hyper-parameters according to the suggestions of each paper to obtain the best completion results. Finally, we show the average performance of the eight images for each algorithm under different missing rates by line graphs. Figure 5 shows the RSE and PSNR results of each algorithm. The smaller RSE value and the larger PSNR value indicate better performance. Our TRLRF performed the best among all the considered algorithms in most cases. When the missing rate increased, the completion results of all the algorithms decreased, especially when the missing rate was near 0.9. The performance of most algorithm fell drastically when the missing rate was 0.95. However, the performance of TRLRF, HaLRTC, and FBCP remained stable and the best performance was obtained from our TRLRF.

Table 1: HSI completion results (RSE) under three different tensor orders with different rank selections

	TRLRF	TRALS	TRWOPT	TMac	TenALS	t-SVD	FBCP	HaLRTC
3rd-order, high-rank($R_n = 12$)	0.06548	0.07049	0.06695	0.1662	0.3448	0.4223	0.2363	0.1254
3rd-order, low-rank($R_n = 8$)	0.1166	0.1245	0.1249	0.2963	0.3312	-	-	-
5th-order, high-rank($R_n = 22$)	0.1035	0.1392	0.1200	0.8064	-	0.9504	0.3833	0.3944
5th-order, low-rank($R_n = 18$)	0.1062	0.1122	0.1072	0.7411	-	-	-	-
8th-order, high-rank($R_n = 24$)	0.1190	0.1319	0.1637	0.9487	-	0.9443	0.4021	0.9099
8th-order, low-rank($R_n = 20$)	0.1421	0.1581	0.1767	0.9488	-	0.9450	0.4135	0.9097

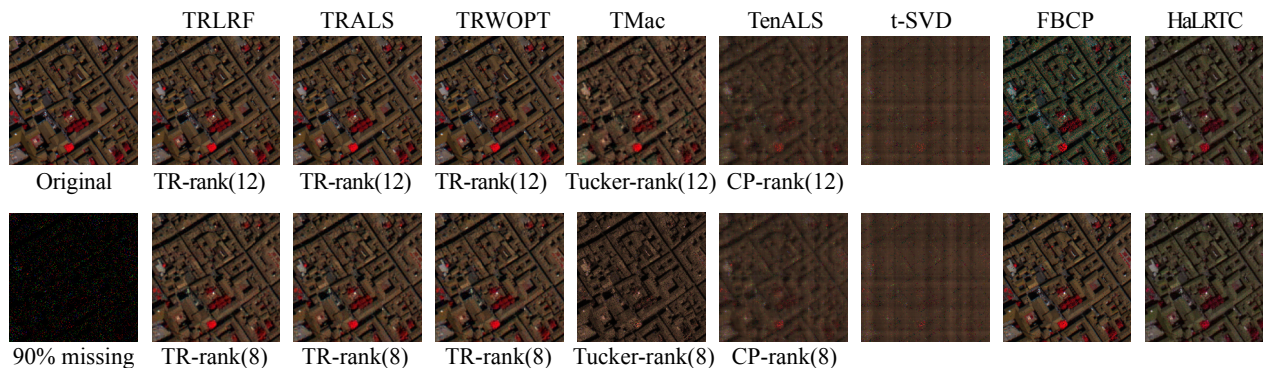


Figure 6: Completion results under the 0.9 missing rate HSI data. The channels 80, 34, 9 are picked to show the visual results. The rank selection of TRLRF, TRALS, TRWOPT, TMac and TenALS are given under the corresponding images.

Hyperspectral image

A hyperspectral image (HSI) of size $200 \times 200 \times 80$ which records an area of the urban landscape was tested in this section¹. In order to test the performance of TRLRF on higher-order tensors, the HSI data was reshaped to higher-order tensors, which is an easy way to find more low-rank features of the data. We compared our TRLRF to the other seven tensor completion algorithms in 3rd-order tensor ($200 \times 200 \times 80$), 5th-order tensor ($10 \times 20 \times 10 \times 20 \times 80$) and 8th-order tensor ($8 \times 5 \times 5 \times 8 \times 5 \times 5 \times 8 \times 10$) cases. The higher-order tensors were generated from original HSI data by directly reshaping it to the specified size and order.

The experiment aims to verify the completion performance of the eight algorithms under different model selection, whereby the experiment variables are the tensor order and tensor rank. The missing rates of all the cases are set as 0.9. All the tuning parameters of every algorithm were set according to the statement in the previous experiments. Besides, for the experiments which need to set rank manually, we chose two different tensor ranks: high-rank and low-rank for algorithms. It should be noted that the CP-rank of TenALS and the Tucker-rank of TMac were set to the same values as TR-rank. The completion performance of RSE and visual results are listed in Table 1 and shown in Figure 6. The results of FBCP, HaLRTC and t-SVD were not affected by tensor rank, so the cases of the same order with different rank are left blank in Table 1. The TenALS could not deal with tensor more than three order, so the high-order

tensor cases for TenALS are also left blank. As shown in Table 1, our TRLRF gives the best recovery performance for the HSI image. In the 3rd-order cases, the best performance was obtained when the TR-rank was 12, however, when the rank was set to 8, the performance of TRLRF, TRALS, TRWOPT, TMac, and TenALS failed because of the underfitting of the selected models. For 5th-order cases, when the rank increased from 18 to 22, the performance of TRLRF kept steady while the performance of TRALS, TRWOPT, and TMac decreased. This is because the high-rank makes the models overfit while our TRLRF performs without any issues, owing to its inherent TR-rank robustness. In the 8th-order tensor cases, similar properties can be obtained and our TRLRF also performed the best.

Conclusion

We propose a tensor completion method based on TR decomposition. By exploiting the rank relationship between the tensor and the TR latent space, we employ low-rank constraints on the TR factors. To solve the model, we propose an ADMM solving scheme. Our method can effectively solve the model selection problem which is common in traditional tensor completion methods. The extensive experiments on both synthetic data and real-world data have demonstrated that our algorithm outperforms the state-of-the-art algorithms. Furthermore, it is expected that the idea of imposing rank minimization constraint on tensor latent space can be extended to various tensor decomposition models in order to develop more efficient and robust algorithms.

¹http://www.ehu.es/ccwintco/index.php/Hyperspectral_Remote_Sensing_Scenes

References

- Acar, E.; Dunlavy, D. M.; Kolda, T. G.; and Mørup, M. 2011. Scalable tensor factorizations for incomplete data. *Chemometrics and Intelligent Laboratory Systems* 106(1):41–56.
- Anandkumar, A.; Ge, R.; Hsu, D.; Kakade, S. M.; and Telgarsky, M. 2014. Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research* 15(1):2773–2832.
- Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; Eckstein, J.; et al. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3(1):1–122.
- Cai, J.-F.; Candès, E. J.; and Shen, Z. 2010. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization* 20(4):1956–1982.
- Cong, F.; Lin, Q.-H.; Kuang, L.-D.; Gong, X.-F.; Astikainen, P.; and Ristaniemi, T. 2015. Tensor decomposition of EEG signals: a brief review. *Journal of Neuroscience Methods* 248:59–69.
- Filipović, M., and Jukić, A. 2015. Tucker factorization with missing data with application to low- n -rank tensor completion. *Multidimensional Systems and Signal Processing* 26(3):677–692.
- Gandy, S.; Recht, B.; and Yamada, I. 2011. Tensor completion and low- n -rank tensor recovery via convex optimization. *Inverse Problems* 27(2):025010.
- Grasedyck, L.; Kluge, M.; and Kramer, S. 2015. Variants of alternating least squares tensor completion in the tensor train format. *SIAM Journal on Scientific Computing* 37(5):A2424–A2450.
- Hillar, C. J., and Lim, L.-H. 2013. Most tensor problems are np-hard. *Journal of the ACM (JACM)* 60(6):45.
- Imaizumi, M.; Maehara, T.; and Hayashi, K. 2017. On tensor train rank minimization: Statistical efficiency and scalable algorithm. In *Advances in Neural Information Processing Systems*, 3933–3942.
- Jain, P., and Oh, S. 2014. Provable tensor factorization with missing data. In *Advances in Neural Information Processing Systems*, 1431–1439.
- Kanagawa, H.; Suzuki, T.; Kobayashi, H.; Shimizu, N.; and Tagami, Y. 2016. Gaussian process nonparametric tensor estimator and its minimax optimality. In *International Conference on Machine Learning*, 1632–1641.
- Kolda, T. G., and Bader, B. W. 2009. Tensor decompositions and applications. *SIAM Review* 51(3):455–500.
- Liu, J.; Musialski, P.; Wonka, P.; and Ye, J. 2013. Tensor completion for estimating missing values in visual data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(1):208–220.
- Liu, Y.; Shang, F.; Fan, W.; Cheng, J.; and Cheng, H. 2014. Generalized higher-order orthogonal iteration for tensor decomposition and completion. In *Advances in Neural Information Processing Systems*, 1763–1771.
- Liu, Y.; Shang, F.; Jiao, L.; Cheng, J.; and Cheng, H. 2015. Trace norm regularized CANDECOMP/PARAFAC decomposition with missing data. *IEEE Transactions on Cybernetics* 45(11):2437–2448.
- Novikov, A.; Podoprikin, D.; Osokin, A.; and Vetrov, D. P. 2015. Tensorizing neural networks. In *Advances in Neural Information Processing Systems*, 442–450.
- Oseledets, I. V. 2011. Tensor-train decomposition. *SIAM Journal on Scientific Computing* 33(5):2295–2317.
- Signoretto, M.; Dinh, Q. T.; De Lathauwer, L.; and Suykens, J. A. 2014. Learning with tensors: a framework based on convex optimization and spectral regularization. *Machine Learning* 94(3):303–351.
- Wang, W.; Aggarwal, V.; and Aeron, S. 2017. Efficient low rank tensor ring completion. In *IEEE International Conference on Computer Vision (ICCV)*, 5698–5706. IEEE.
- Wang, W.; Sun, Y.; Eriksson, B.; Wang, W.; and Aggarwal, V. 2018. Wide compression: Tensor ring nets. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 9329–9338. IEEE.
- Xu, Y.; Hao, R.; Yin, W.; and Su, Z. 2013. Parallel matrix factorization for low-rank tensor completion. *arXiv preprint arXiv:1312.1254*.
- Yuan, L.; Cao, J.; Wu, Q.; and Zhao, Q. 2018. Higher-dimension tensor completion via low-rank tensor ring decomposition. *arXiv preprint arXiv:1807.01589*.
- Yuan, L.; Zhao, Q.; and Cao, J. 2017. Completion of high order tensor data with missing entries via tensor-train decomposition. In *International Conference on Neural Information Processing*, 222–229. Springer.
- Zhang, Z.; Ely, G.; Aeron, S.; Hao, N.; and Kilmer, M. 2014. Novel methods for multilinear data completion and de-noising based on tensor-SVD. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3842–3849.
- Zhao, Q.; Zhou, G.; Xie, S.; Zhang, L.; and Cichocki, A. 2016a. Tensor ring decomposition. *arXiv preprint arXiv:1606.05535*.
- Zhao, Q.; Zhou, G.; Zhang, L.; Cichocki, A.; and Amari, S.-I. 2016b. Bayesian robust tensor factorization for incomplete multiway data. *IEEE Transactions on Neural Networks and Learning Systems* 27(4):736–748.
- Zhao, Q.; Sugiyama, M.; Yuan, L.; and Cichocki, A. 2018. Learning efficient tensor representations with ring structure networks. In *Sixth International Conference on Learning Representations (ICLR workshop)*.
- Zhao, Q.; Zhang, L.; and Cichocki, A. 2015. Bayesian CP factorization of incomplete tensors with automatic rank determination. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37(9):1751–1763.