# Bayesian CP Factorization of Incomplete Tensors with Automatic Rank Determination 

Qibin Zhao, Member, IEEE, Liqing Zhang, Member, IEEE, and Andrzej Cichocki Fellow, IEEE


#### Abstract

CANDECOMP/PARAFAC (CP) tensor factorization of incomplete data is a powerful technique for tensor completion through explicitly capturing the multilinear latent factors. The existing CP algorithms require the tensor rank to be manually specified, however, the determination of tensor rank remains a challenging problem especially for $C P$ rank. In addition, existing approaches do not take into account uncertainty information of latent factors, as well as missing entries. To address these issues, we formulate CP factorization using a hierarchical probabilistic model and employ a fully Bayesian treatment by incorporating a sparsity-inducing prior over multiple latent factors and the appropriate hyperpriors over all hyperparameters, resulting in automatic rank determination. To learn the model, we develop an efficient deterministic Bayesian inference algorithm, which scales linearly with data size. Our method is characterized as a tuning parameter-free approach, which can effectively infer underlying multilinear factors with a low-rank constraint, while also providing predictive distributions over missing entries. Extensive simulations on synthetic data illustrate the intrinsic capability of our method to recover the ground-truth of $C P$ rank and prevent the overfitting problem, even when a large amount of entries are missing. Moreover, the results from real-world applications, including image inpainting and facial image synthesis, demonstrate that our method outperforms state-of-the-art approaches for both tensor factorization and tensor completion in terms of predictive performance.


Index Terms-Tensor factorization, tensor completion, tensor rank determination, Bayesian inference, image synthesis

## 1 Introduction

TENSORS (i.e., multiway arrays) provide an effective and faithful representation of the structural properties of data, in particular, when multidimensional data or a data ensemble affected by multiple factors are involved. For instance, a video sequence can be represented by a third-order tensor with dimensionality of height $\times$ width $\times$ time; an image ensemble measured under multiple conditions can be represented by a higher order tensor with dimensionality of pixel $\times$ person $\times$ pose $\times$ illumination. Tensor factorization enables us to explicitly take into account the structure information by effectively capturing the multilinear interactions among multiple latent factors. Therefore, its theory and algorithms have been an active area of study during the past decade (see e.g., [1], [2]), and have been successfully applied to various application fields, such as face recognition, social network analysis, image and video completion, and brain signal processing. The two most popular tensor factorization frameworks are Tucker [3] and CANDECOMP/PARAFAC (CP), also known as canonical polyadic decomposition (CPD) [4], [5], [6].

Most existing tensor factorization methods assume

- Q. Zhao is at the Laboratory for Advanced Brain Signal Processing, RIKEN Brain Science Institute, Japan and the Department of Computer Science and Engineering, Shanghai Jiao Tong University, China.
- L. Zhang is at the MOE-Microsoft Laboratory for Intelligent Computing and Intelligent Systems and the Department of Computer Science and Engineering, Shanghai Jiao Tong University, China.
- A. Cichocki is at the Laboratory for Advanced Brain Signal Processing, RIKEN Brain Science Institute, Japan and the Systems Research Institute at the Polish Academy of Science, Warsaw, Poland.
that the tensor is complete, whereas the problem of missing data can arise in a variety of real-world applications. This issue has attracted a great deal of research interest in tensor completion in recent years. The objective of tensor factorization of incomplete data is to capture the underlying multilinear factors from only partially observed entries, which can in turn predict the missing entries. In [7], CP factorization with missing data was formulated as a weighted least squares problem, termed CP weighted optimization (CPWOPT). A structured CPD using nonlinear least squares (CPNLS) was proposed in [8]. In [9], geometric nonlinear conjugate gradient (geomCG) based on Riemannian optimization on the manifold of tensors were presented. However, as the number of missing entries increases, tensor factorization schemes tend to overfit the model because of an incorrectly specified tensor rank, resulting in severe deterioration of their predictive performance. In contrast, another popular technique is to exploit a low-rank assumption for recovering the missing entries; the rank minimization can be formulated as a convex optimization problem on a nuclear norm. This technique has been extended to higher order tensors by defining the nuclear norm of a tensor, yielding the tensor completion [10]. Some variants were also proposed, such as a framework based on convex optimization and spectral regularization, which uses an inexact splitting method [11], and fast composite splitting algorithms (FCSA) [12]. To improve the efficiency, the Douglas-Rachford splitting technique [13], nonlinear Gauss-Seidal method [14] were also investigated. Recently, the nuclear norm based optimization was also applied to a supervised tensor dimensionality reduction method [15]. An alternative
method for tensor completion is to employ adaptive sampling schemes [16]. However, the nuclear norm of a tensor is defined straightforwardly by a weighted sum of the nuclear norm of mode- $n$ matricizations, which is related to multilinear rank rather than $C P$ rank. In addition, these completion-based methods cannot explicitly capture the underlying factors. Hence, a simultaneous tensor decomposition and completion (STDC) method was introduced in which a rank minimization technique was combined with Tucker decomposition [17]. To improve completion accuracy, auxiliary information was also exploited in [17], [18], which strongly depends on the specific application. It is also noteworthy that the rank minimization based on convex optimization of the nuclear norm is affected by tuning parameters, which may tend to over- or under-estimate the true tensor rank.

It is important to emphasize that our knowledge about the properties of $C P$ rank, defined by the minimum number of rank-one terms in CP decomposition, is surprisingly limited. There is no straightforward algorithm to compute the rank even for a given specific tensor, and the problem has been shown to be NP-complete [19]. The lower and upper bound of tensor rank was studied in [20], [21]. The ill-posedness of the best low-rank approximation of a tensor was investigated in [22]. In fact, determining or even bounding the rank of an arbitrary tensor is quite difficult in contrast to the matrix rank [23], and this difficulty would be significantly exacerbated in the presence of missing data.

Probabilistic models for matrix/tensor factorization have attracted much interest in collaborative filtering and matrix/tensor completion. Probabilistic matrix factorization was proposed in [24], and its fully Bayesian treatment using Markow chain Monte Carlo (MCMC) inference was shown in [25] and using variational Bayesian inference in [26], [27]. Further extensions of nonparametric and robust variants were presented in [28], [29]. The probabilistic frameworks of tensor factorization were presented in [30], [31], [32]. Other variants include extensions of the exponential family model [33] and the nonparametric Bayesian model [34]. However, the tensor rank or model complexity are often given by a tuning parameter selected by either maximum likelihood or cross-validations, which are computationally expensive and inaccurate. Another important issue is that the inference of factor matrices is performed by either point estimation, which is prone to overfitting, or MCMC inference, which tends to converge very slowly.

To address these issues, we propose a fully Bayesian probabilistic tensor factorization model according to the CP factorization framework. Our objective is to infer the underlying multilinear factors from a noisy incomplete tensor and the predictive distribution of missing entries, while the rank of the true latent tensor can be determined automatically and implicitly. To achieve this, we specify a sparsity-inducing hierarchical prior over multiple factor matrices with individual hyperparameters associated to each latent dimension, such that the
number of components in factor matrices is constrained to be minimum. All the model parameters, including noise precision, are considered to be latent variables over which the corresponding priors are placed. Due to complex interactions among multiple factors and fully Bayesian treatment, learning the model is analytically intractable. Thus, we resort to the variational Bayesian inference and derive a deterministic solution to approximate the posteriors of all the model parameters and hyperparameters. Our method is characterized as a tuning parameter-free approach that can effectively avoid parameter selections. The extensive experiments and comparisons on synthetic data illustrate the advantages of our approach in terms of rank determination, predictive capability, and robustness to overfitting. Moreover, several real-word applications, including image completion, restoration, and synthesis, demonstrate that our method outperforms state-of-the-art approaches, including both tensor factorization and tensor completion, in terms of the predictive performance.

The rest of this paper is organized as follows. In Section 2, preliminary multilinear operations and notations are presented. In Section 3, we introduce the probabilistic CP model specification and the model learning via Bayesian inference. A variant of our method using mixture priors is proposed in Section 4 In Section 5, we present the comprehensive experimental results for both synthetic data and real-world applications, followed by our conclusion in Section 6

## 2 Preliminaries and Notations

The order of a tensor is the number of dimensions, also known as ways or modes. Vectors (first-order tensors) are denoted by boldface lowercase letters, e.g., a. Matrices (second-order tensors) are denoted by boldface capital letters, e.g., A. Higher-order tensors (order $\geq 3$ ) are denoted by boldface calligraphic letters, e.g., $\mathcal{A}$. Given an $N$ th-order tensor $\mathcal{X} \in \mathbb{R}^{I_{1} \times I_{2} \times \cdots \times I_{N}}$, its $\left(i_{1}, i_{2}, \ldots, i_{N}\right)$ th entry is denoted by $\mathcal{X}_{i_{1} i_{2} \ldots i_{N}}$, where the indices typically range from 1 to their capital version, e.g., $i_{n}=1,2, \ldots, I_{n}, \forall n \in[1, N]$.

The inner product of two tensors is defined by $\langle\boldsymbol{\mathcal { A }}, \boldsymbol{\mathcal { B }}\rangle=$ $\sum_{i_{1}, i_{2}, \ldots, i_{N}} \mathcal{A}_{i_{1} i_{2} \ldots i_{N}} \mathcal{B}_{i_{1} i_{2} \ldots i_{N}}$, and the squared Frobenius norm by $\|\mathcal{A}\|_{F}^{2}=\langle\mathcal{A}, \mathcal{A}\rangle$. As an extension to $N$ variables, the generalized inner product of a set of vectors, matrices, or tensors is defined as a sum of element-wise products. For example, given $\left\{\mathbf{A}^{(n)} \mid n=1, \ldots, N\right\}$, we define

$$
\begin{equation*}
\left\langle\mathbf{A}^{(1)}, \cdots, \mathbf{A}^{(N)}\right\rangle=\sum_{i, j} \prod_{n} A_{i j}^{(n)} \tag{1}
\end{equation*}
$$

The Hadamard product is an entrywise product of two vectors, matrices, or tensors of the same sizes. For instance, given two matrices, $\mathbf{A} \in \mathbb{R}^{I \times J}$ and $\mathbf{B} \in \mathbb{R}^{I \times J}$, their Hadamard product is a matrix of size $I \times J$ and is denoted by $\mathbf{A} \circledast \mathbf{B}$. Without loss of generality, the Hadamard product of a set of matrices can be simply
denoted by

$$
\begin{equation*}
\underset{n}{\circledast} \mathbf{A}^{(n)}=\mathbf{A}^{(1)} \circledast \mathbf{A}^{(2)} \circledast \cdots \circledast \mathbf{A}^{(N)} . \tag{2}
\end{equation*}
$$

The Kronecker product [1] of matrices $\mathbf{A} \in \mathbb{R}^{I \times J}$ and $\mathbf{B} \in$ $\mathbb{R}^{K \times L}$ is a matrix of size $I K \times J L$, denoted by $\mathbf{A} \otimes \mathbf{B}$. The Khatri-Rao product of matrices $\mathbf{A} \in \mathbb{R}^{I \times K}$ and $\mathbf{B} \in \mathbb{R}^{J \times K}$ is a matrix of size $I J \times K$, defined by a columnwise Kronecker product and denoted by $\mathbf{A} \odot \mathbf{B}$. In particular, the Khatri-Rao product of a set of matrices in reverse order is defined by

$$
\begin{equation*}
\bigodot_{n} \mathbf{A}^{(n)}=\mathbf{A}^{(N)} \odot \mathbf{A}^{(N-1)} \odot \cdots \odot \mathbf{A}^{(1)} \tag{3}
\end{equation*}
$$

while the Khatri-Rao product of a set of matrices, except the $n$th matrix, denoted by $\mathbf{A}^{(\backslash n)}$, is

$$
\begin{equation*}
\bigodot_{k \neq n} \mathbf{A}^{(k)}=\mathbf{A}^{(N)} \odot \cdots \odot \mathbf{A}^{(n+1)} \odot \mathbf{A}^{(n-1)} \odot \cdots \odot \mathbf{A}^{(1)} \tag{4}
\end{equation*}
$$

## 3 BAYESIAN TENSOR FACTORIZATION

### 3.1 Probabilistic Model and Priors

Let $\mathcal{Y}$ be an incomplete $N$ th-order tensor of size $I_{1} \times$ $I_{2} \times \cdots \times I_{N}$ with missing entries. The element $\mathcal{Y}_{i_{1} i_{2} \ldots i_{N}}$ is observed if $\left(i_{1}, i_{2}, \cdots, i_{N}\right) \in \Omega$, where $\Omega$ denotes a set of indices. For simplicity, we also define a binary tensor $\mathcal{O}$ of the same size as $\mathcal{Y}$ as an indicator of observed entries. We assume $\mathcal{Y}$ is a noisy observation of true latent tensor $\mathcal{X}$, that is, $\mathcal{Y}=\mathcal{X}+\varepsilon$, where the noise term is assumed to be an i.i.d. Gaussian distribution, i.e., $\varepsilon \sim \prod_{i_{1}, \ldots, i_{N}} \mathcal{N}\left(0, \tau^{-1}\right)$, and the latent tensor $\mathcal{X}$ can be exactly represented by a CP model, given by

$$
\begin{equation*}
\boldsymbol{\mathcal { X }}=\sum_{r=1}^{R} \mathbf{a}_{r}^{(1)} \circ \cdots \circ \mathbf{a}_{r}^{(N)}=\llbracket \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)} \rrbracket \tag{5}
\end{equation*}
$$

where $\circ$ denotes the outer product of vectors and $\llbracket \cdots \rrbracket$ is a shorthand notation, also termed as the Kruskal operator. CP factorization can be interpreted as a sum of $R$ rank-one tensors, while the smallest integer $R$ is defined as $C P \operatorname{rank}$ [1]. $\left\{\mathbf{A}^{(n)}\right\}_{n=1}^{N}$ are a set of factor matrices where mode- $n$ factor matrix $\mathbf{A}^{(n)} \in \mathbb{R}^{I_{n} \times R}$ can be denoted by row-wise or column-wise vectors
$\mathbf{A}^{(n)}=\left[\mathbf{a}_{1}^{(n)}, \ldots, \mathbf{a}_{i_{n}}^{(n)}, \ldots, \mathbf{a}_{I_{n}}^{(n)}\right]^{T}=\left[\mathbf{a}_{\cdot 1}^{(n)}, \ldots, \mathbf{a}_{\cdot r}^{(n)}, \ldots, \mathbf{a}_{\cdot R}^{(n)}\right]$.
The CP generative model, together with noise assumption, directly give rise to the observation model, which is factorized over observed tensor elements

$$
\begin{align*}
& p\left(\mathcal{Y}_{\Omega} \mid\left\{\mathbf{A}^{(n)}\right\}_{n=1}^{N}, \tau\right)=\prod_{i_{1}=1}^{I_{1}} \cdots \prod_{i_{N}=1}^{I_{N}} \\
& \mathcal{N}\left(\mathcal{Y}_{i_{1} i_{2} \ldots i_{N}} \mid\left\langle\mathbf{a}_{i_{1}}^{(1)}, \mathbf{a}_{i_{2}}^{(2)}, \cdots, \mathbf{a}_{i_{N}}^{(N)}\right\rangle, \tau^{-1}\right)^{\mathcal{O}_{i_{1} \cdots i_{N}}} \tag{6}
\end{align*}
$$

where the parameter $\tau$ denotes the noise precision, and $\left\langle\mathbf{a}_{i_{1}}^{(1)}, \mathbf{a}_{i_{2}}^{(2)}, \cdots, \mathbf{a}_{i_{N}}^{(N)}\right\rangle=\sum_{r} \prod_{n} a_{i_{n} r}^{(n)}$ denotes a generalized inner-product of $N$ vectors. The likelihood model


Fig. 1. Probabilistic graphical model of Bayesian CP factorization of an $N$ th-order tensor.
in (6) indicates that $\mathcal{Y}_{i_{1} \cdots i_{N}}$ is generated by multiple $R$ dimensional latent vectors $\left\{\mathbf{a}_{i_{n}}^{(n)} \mid n=1, \ldots, N\right\}$, where each latent vector $\mathbf{a}_{i_{n}}^{(n)}$ contributes to a set of observations, i.e., a subtensor whose mode- $n$ index is $i_{n}$. The essential difference between matrix and tensor factorization is that the inner product of $N \geq 3$ vectors allows us to model the multilinear interaction structure, which however leads to many more difficulties in model learning.

In general, the effective dimensionality of the latent space, i.e., $\operatorname{Rank}_{C P}(\boldsymbol{\mathcal { X }})=R$, is a tuning parameter whose selection is quite challenging and computational costly. Therefore, we seek an elegant automatic model selection, which can not only infer the rank of the latent tensor $\mathcal{X}$, but also effectively avoid overfitting. To achieve this, a set of continuous hyperparameters are employed to control the variance related to each dimensionality of the latent space, respectively. Since the minimum $R$ is desired in the sense of low rank approximation, a sparsity-inducing prior is specified over these hyperparameters, resulting in it being possible to achieve automatic rank determination as a part of the Baybesian inference process. This technique is related to automatic relevance determination (ARD) [35] or sparse Bayesian learning [36]. However, unlike the traditional methods that place the ARD prior over either latent variables or weight parameters, such as Bayesian principle component analysis [37], our method considers all model parameters as latent variables over which a sparsityinducing prior is placed with shared hyperparameters.

More specifically, we place a prior distribution over the latent factors, governed by hyperparameters $\boldsymbol{\lambda}=$ $\left[\lambda_{1}, \ldots, \lambda_{R}\right.$ ] where each $\lambda_{r}$ controls $r$ th component in $\mathbf{A}^{(n)}$, which is

$$
\begin{equation*}
p\left(\mathbf{A}^{(n)} \mid \boldsymbol{\lambda}\right)=\prod_{i_{n}=1}^{I_{n}} \mathcal{N}\left(\mathbf{a}_{i_{n}}^{(n)} \mid \mathbf{0}, \boldsymbol{\Lambda}^{-1}\right), \forall n \in[1, N] \tag{7}
\end{equation*}
$$

where $\boldsymbol{\Lambda}=\operatorname{diag}(\boldsymbol{\lambda})$ denotes the inverse covariance matrix, also known as the precision matrix, and is shared by latent factor matrices in all modes. We can further define a hyperprior over $\boldsymbol{\lambda}$, which is factorized over
latent dimensions

$$
\begin{equation*}
p(\boldsymbol{\lambda})=\prod_{r=1}^{R} \mathrm{Ga}\left(\lambda_{r} \mid c_{0}^{r}, d_{0}^{r}\right) \tag{8}
\end{equation*}
$$

where $\operatorname{Ga}(x \mid a, b)=\frac{b^{a} x^{a-1} e^{-b x}}{\Gamma(a)}$ denotes a Gamma distribution and $\Gamma(a)$ is the Gamma function.

Since the sparsity is enforced in the latent dimensions, the initialization point of the dimensionality of latent space (i.e., R) is usually set to its maximum possible value, while the effective dimensionality can be inferred automatically under a Bayesian inference framework. It should be noted that since the priors are shared across $N$ latent matrices, our framework can learn the same sparsity pattern for them, yielding the minimum number of rank-one terms. Therefore, our model can effectively infer the rank of tensor while performing the tensor factorization, which can be treated as a Bayesian low-rank tensor factorization.

To complete the model with a fully Bayesian treatment, we also place a hyperprior over the noise precision $\tau$, that is,

$$
\begin{equation*}
p(\tau)=\mathrm{Ga}\left(\tau \mid a_{0}, b_{0}\right) \tag{9}
\end{equation*}
$$

For simplicity of notation, all unknowns including latent variables and hyperparameters are collected and denoted together by $\Theta=\left\{\mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)}, \boldsymbol{\lambda}, \tau\right\}$. The probabilistic graph model is illustrated in Fig. 1. from which we can easily write the joint distribution of the model as

$$
p\left(\boldsymbol{Y}_{\Omega}, \Theta\right)=p\left(\mathcal{Y}_{\Omega} \mid\left\{\mathbf{A}^{(n)}\right\}_{n=1}^{N}, \tau\right) \prod_{n=1}^{N} p\left(\mathbf{A}^{(n)} \mid \boldsymbol{\lambda}\right) p(\boldsymbol{\lambda}) p(\tau)
$$

By combining the likelihood in (6), the priors of model parameters in (7), and the hyperpriors in (8) and (9), the logarithm of the joint distribution is given by (see Sec. 1 of Appendix for details)

$$
\begin{align*}
& \ell(\Theta)=- \frac{\tau}{2}\left\|\mathcal{O} \circledast\left(\mathcal{Y}-\llbracket \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)} \rrbracket\right)\right\|_{F}^{2} \\
&-\frac{1}{2} \operatorname{Tr}\left(\boldsymbol{\Lambda} \sum_{n} \mathbf{A}^{(n) T} \mathbf{A}^{(n)}\right)+\left(\frac{M}{2}+a_{0}-1\right) \ln \tau \\
& \quad+\sum_{r}\left[\left(\frac{\sum_{n} I_{n}}{2}+\left(c_{0}^{r}-1\right)\right) \ln \lambda_{r}\right] \\
&-\sum_{r} d_{0}^{r} \lambda_{r}-b_{0} \tau+\mathrm{const} \tag{10}
\end{align*}
$$

where $M=\sum_{i_{1}, \ldots, i_{N}} \mathcal{O}_{i_{1} \ldots i_{N}}$ denotes the total number of observations. Without loss of generality, we can perform maximum a posteriori (MAP) estimation of $\Theta$ by maximizing [10, which is, to some extent, equivalent to optimizing a squared error function with regularizations imposed on the factor matrices and additional constraints imposed on the regularization parameters.

However, our objective is to develop a method that, in contrast to the point estimation, computes the full
posterior distribution of all variables in $\Theta$ given the observed data, that is,

$$
\begin{equation*}
p\left(\Theta \mid \mathcal{Y}_{\Omega}\right)=\frac{p\left(\Theta, \mathcal{Y}_{\Omega}\right)}{\int p\left(\Theta, \mathcal{Y}_{\Omega}\right) d \Theta} \tag{11}
\end{equation*}
$$

Based on the posterior distribution of $\Theta$, the predictive distribution over missing entries, denoted by $\mathcal{Y}_{\backslash \Omega}$, can be inferred by

$$
\begin{equation*}
p\left(\mathcal{Y}_{\backslash \Omega} \mid \mathcal{Y}_{\Omega}\right)=\int p\left(\mathcal{Y}_{\backslash \Omega} \mid \Theta\right) p\left(\Theta \mid \mathcal{Y}_{\Omega}\right) \mathrm{d} \Theta \tag{12}
\end{equation*}
$$

### 3.2 Model Learning via Bayesian Inference

An exact Bayesian inference in (11) and (12) would integrate over all latent variables as well as hyperparameters, which is obviously analytically intractable. In this section, we describe the development of a deterministic approximate inference under variational Bayesian (VB) framework [38], [39] to learn the probabilistic CP factorization model.

We therefore seek a distribution $q(\Theta)$ to approximate the true posterior distribution $p\left(\Theta \mid \mathcal{Y}_{\Omega}\right)$ by minimizing the KL divergence, that is,

$$
\begin{align*}
\mathrm{KL}\left(q(\Theta) \| p\left(\Theta \mid \mathcal{Y}_{\Omega}\right)\right) & =\int q(\Theta) \ln \left\{\frac{q(\Theta)}{p\left(\Theta \mid \mathcal{Y}_{\Omega}\right)}\right\} d \Theta \\
=\ln p\left(\mathcal{Y}_{\Omega}\right) & -\int q(\Theta) \ln \left\{\frac{p\left(\mathcal{Y}_{\Omega}, \Theta\right)}{q(\Theta)}\right\} d \Theta \tag{13}
\end{align*}
$$

where $\ln p\left(\mathcal{Y}_{\Omega}\right)$ represents the model evidence, and its lower bound is defined by $\mathcal{L}(q)=\int q(\Theta) \ln \left\{\frac{p\left(\boldsymbol{\mathcal { Y }}_{\Omega}, \Theta\right)}{q(\Theta)}\right\} d \Theta$. Since the model evidence is a constant, the maximum of the lower bound occurs when the KL divergence vanishes, which implies that $q(\Theta)=p\left(\Theta \mid \mathcal{Y}_{\Omega}\right)$.

For the initial derivation, it will be assumed that the variational distribution is factorized w.r.t. each variable $\Theta_{j}$ and therefore can be written as

$$
\begin{equation*}
q(\Theta)=q_{\lambda}(\boldsymbol{\lambda}) q_{\tau}(\tau) \prod_{n=1}^{N} q_{n}\left(\mathbf{A}^{(n)}\right) \tag{14}
\end{equation*}
$$

It should be noted that this is the only assumption about the distribution, while the particular functional forms of the individual factors $q_{j}\left(\Theta_{j}\right)$ can be explicitly derived in turn. The optimised form of the $j$ th factor based on the maximization of $\mathcal{L}(q)$ is given by

$$
\begin{equation*}
\ln q_{j}\left(\Theta_{j}\right)=\mathbb{E}_{q_{\left(\Theta \backslash \Theta_{j}\right)}}\left[\ln p\left(\mathcal{Y}_{\Omega}, \Theta\right)\right]+\text { const } \tag{15}
\end{equation*}
$$

where $\mathbb{E}_{q_{\left(\Theta \backslash \Theta_{j}\right)}}[\cdot]$ denotes an expectation w.r.t. the $q$ distributions over all variables except $\Theta_{j}$. Since the distributions of all parameters are drawn from the exponential family and are conjugate w.r.t. the distributions of their parents (see Fig. 11, we can derive the closed-form posterior update rules for each parameter in $\Theta$ by using (15.

### 3.2.1 Posterior distribution of factor matrices

As can be seen from the graphical model shown in Fig. 1. the inference of mode- $n$ factor matrix $\mathbf{A}^{(n)}$ can be performed by receiving the messages from observed data and its co-parents, including other factors $\mathbf{A}^{(k)}, k \neq n$ and the hyperparameter $\tau$, which are expressed by the likelihood term (6), and incorporating the messages from its parents, which are expressed by the prior term (7). By applying (15), it has been shown that their posteriors can be factorized as independent distributions of their rows, which are also Gaussian (see Sec. 2 of Appendix for details), given by

$$
\begin{equation*}
q_{n}\left(\mathbf{A}^{(n)}\right)=\prod_{i_{n}=1}^{I_{n}} \mathcal{N}\left(\mathbf{a}_{i_{n}}^{(n)} \mid \tilde{\mathbf{a}}_{i_{n}}^{(n)}, \mathbf{V}_{i_{n}}^{(n)}\right), \forall n \in[1, N] \tag{16}
\end{equation*}
$$

where the posterior parameters can be updated by

$$
\begin{align*}
&{\tilde{\mathbf{a}} i_{n}}_{(n)}=\mathbb{E}_{q}[\tau] \mathbf{V}_{i_{n}}^{(n)} \mathbb{E}_{q}\left[\mathbf{A}_{i_{n}}^{(\backslash n) T}\right] \operatorname{vec}\left(\mathcal{Y}_{\mathbb{I}\left(\mathcal{O}_{i_{n}}=1\right)}\right) \\
& \mathbf{V}_{i_{n}}^{(n)}=\left(\mathbb{E}_{q}[\tau] \mathbb{E}_{q}\left[\mathbf{A}_{i_{n}}^{(\backslash n) T} \mathbf{A}_{i_{n}}^{(\backslash n)}\right]+\mathbb{E}_{q}[\boldsymbol{\Lambda}]\right)^{-1}, \tag{17}
\end{align*}
$$

where $\mathcal{Y}_{\mathbb{I}\left(\mathcal{O}_{i_{n}}=1\right)}$ is a sample function denoting a subset of the observed entries $\mathcal{Y}_{\Omega}$, whose mode- $n$ index is $i_{n}$, i.e., the observed entries associated to the latent factor $\mathbf{a}_{i_{n}}^{\left(n_{n}^{\prime}\right)}$. The most complex term in 17, is related to

$$
\begin{equation*}
\mathbf{A}_{i_{n}}^{(\backslash n) T}=\left(\bigodot_{k \neq n} \mathbf{A}^{(k)}\right)_{\mathbb{I}\left(\mathcal{O}_{i_{n}}=1\right)}^{T}, \tag{18}
\end{equation*}
$$

where $\left(\bigodot_{k \neq n} \mathbf{A}^{(k)}\right)^{T}$ is of size $R \times \prod_{k \neq n} I_{k}$, and each column is computed by $\circledast_{k \neq n} \mathbf{a}_{i_{k}}^{(k)}$ with varying mode$k$ index $i_{k}$. The symbol $(\cdot)_{\mathbb{I}\left(\mathcal{O}_{i_{n}}=1\right)}$ denotes a subset of columns sampled according to the subtensor $\operatorname{vec}\left(\mathcal{O}_{\left.\ldots i_{n} \ldots\right)}=1\right.$. Hence, $\mathbb{E}_{q}\left[\mathbf{A}_{i_{n}}^{(\backslash n) T} \mathbf{A}_{i_{n}}^{(\backslash n)}\right]$ denotes the posterior covariance matrix of the Khatri-Rao product of latent factors in all modes except the $n$ th-mode, and is computed by only the columns corresponding to the observed entries whose mode- $n$ index is $i_{n}$. In order to evaluate this posterior covariance matrix, first we need to introduce the following results.
Theorem 3.1. Given a set of independent random matrices $\left\{\mathbf{A}^{(n)} \in \mathbb{R}^{I_{n} \times R} \mid n=1, \ldots, N\right\}$, we assume that $\forall n, \forall i_{n}$, the row vectors $\left\{\mathbf{a}_{i_{n}}^{(n)}\right\}$ are independent, then

$$
\begin{equation*}
\mathbb{E}\left[\left(\bigodot_{n} \mathbf{A}^{(n)}\right)^{T}\left(\bigodot_{n} \mathbf{A}^{(n)}\right)\right]=\sum_{i_{1}, \ldots, i_{N}} \underset{n}{\circledast}\left(\mathbb{E}\left[\mathbf{a}_{i_{n}}^{(n)} \mathbf{a}_{i_{n}}^{(n) T}\right]\right) \tag{19}
\end{equation*}
$$

where $\mathbb{E}\left[\mathbf{a}_{i_{n}}^{(n)} \mathbf{a}_{i_{n}}^{(n) T}\right]=\mathbb{E}\left[\mathbf{a}_{i_{n}}^{(n)}\right] \mathbb{E}\left[\mathbf{a}_{i_{n}}^{(n) T}\right]+\operatorname{Var}\left(\mathbf{a}_{i_{n}}^{(n)}\right)$.
Proof: See Sec. 3 of Appendix for details.
For simplicity, we attempt to compute (19) by multilinear operations. Let $\forall n, \mathbf{B}^{(n)}$ of size $I_{n} \times R^{2}$ denote an expectation of a quadratic form related to $\mathbf{A}^{(n)}$ by defining the $i_{n}$ th-row vector as

$$
\begin{equation*}
\mathbf{b}_{i_{n}}^{(n)}=\operatorname{vec}\left(\mathbb{E}_{q}\left[\mathbf{a}_{i_{n}}^{(n)} \mathbf{a}_{i_{n}}^{(n) T}\right]\right)=\operatorname{vec}\left(\tilde{\mathbf{a}}_{i_{n}}^{(n)} \tilde{\mathbf{a}}_{i_{n}}^{(n) T}+\mathbf{V}_{i_{n}}^{(n)}\right), \tag{20}
\end{equation*}
$$

then we have
$\operatorname{vec}\left(\sum_{i_{1}, \ldots, i_{N}} \not{ }_{n}\left(\mathbb{E}\left[\mathbf{a}_{i_{n}}^{(n)} \mathbf{a}_{i_{n}}^{(n) T}\right]\right)\right)=\left(\bigodot_{n} \mathbf{B}^{(n)}\right)^{T} \mathbf{1}_{\Pi_{n} I_{n}}$,
where $\mathbf{1}_{\Pi_{n} I_{n}}$ denotes a vector of length $\prod_{n} I_{n}$ and all elements are equal to one.

According to Theorem 3.1 and the computation form in 21, the term $\mathbb{E}_{q}\left[\mathbf{A}_{i_{n}}^{(\backslash n) T} \mathbf{A}_{i_{n}} \backslash n\right)$ in 17 can be evaluated efficiently by

$$
\begin{equation*}
\operatorname{vec}\left(\mathbb{E}_{q}\left[\mathbf{A}_{i_{n}}^{(\backslash n) T} \mathbf{A}_{i_{n}}^{(\backslash n)}\right]\right)=\left(\bigodot_{k \neq n} \mathbf{B}^{(k)}\right)^{T} \operatorname{vec}\left(\mathcal{O}_{\ldots i_{n} \ldots}\right) \tag{22}
\end{equation*}
$$

where $\mathcal{O} \ldots i_{n} \ldots$ denotes a subtensor by fixing model $-n$ index to $i_{n}$. It should be noted that the Khatri-Rao product is computed by all mode factors except the $n$th mode, while the sum is performed according to the indices of observations, implying that only factors that interact with $\mathbf{a}_{i_{n}}^{(n)}$ are taken into account. Another complicated part in (17) can also be simplified by multilinear operations, i.e.,

$$
\begin{align*}
& \mathbb{E}_{q}\left[\mathbf{A}_{i_{n}}^{(\backslash n) T}\right] \operatorname{vec}\left(\mathcal{Y}_{\mathbb{I}\left(\mathcal{O}_{i_{n}}=1\right)}\right) \\
& \quad=\left(\bigodot_{k \neq n} \mathbb{E}_{q}\left[\mathbf{A}^{(k)}\right]\right)^{T} \operatorname{vec}\left\{(\mathcal{O} \circledast \mathcal{Y}) \ldots i_{n} \ldots\right\} \tag{23}
\end{align*}
$$

Finally, the variational posterior approximation of factor matrices can be updated by (17). On the basis of the approximated posterior, the posterior moments, including $\forall n, \forall i_{n}, \mathbb{E}_{q}\left[\mathbf{a}_{i_{n}}^{(n)}\right], \operatorname{Var}\left(\mathbf{a}_{i_{n}}^{(n)}\right), \mathbb{E}_{q}\left[\mathbf{A}^{(n)}\right]$, and $\mathbb{E}_{q}\left[\mathbf{a}_{i_{n}}^{(n)} \mathbf{a}_{i_{n}}^{(n) T}\right], \mathbb{E}_{q}\left[\mathbf{a}_{i_{n}}^{(n) T} \mathbf{a}_{i_{n}}^{(n)}\right]$, can be easily evaluated, which are required by the inference of other hyperparameters in $\Theta$.

An intuitive interpretation of (17) is given as follows. The posterior covariance $\mathbf{V}_{i_{n}}^{(n)}$ is updated by combining the prior information $\mathbb{E}_{q}[\boldsymbol{\Lambda}]$ and the posterior information from other factor matrices computed by (22), while the tradeoff between these two terms is controlled by $\mathbb{E}_{q}[\tau]$ that is related to the quality of model fitting. In other words, the better fitness of the current model leads to more information from other factors than from prior information. The posterior mean $\tilde{\mathbf{a}}_{i_{n}}^{(n)}$ is updated firstly by linear combination of all other factors, expressed by (23), where the coefficients are observed values. This implies that the larger observation leads to more similarity of its corresponding latent factors. Then, $\tilde{\mathbf{a}}_{i_{n}}^{(n)}$ is rotated by $\mathbf{V}_{i_{n}}^{(n)}$ to obtain the property of sparsity and is scaled according to the model fitness $\mathbb{E}_{q}[\tau]$.

### 3.2.2 Posterior distribution of hyperparameters $\boldsymbol{\lambda}$

It should be noted that, instead of point estimation via optimizations, learning the posterior of $\boldsymbol{\lambda}$ is crucial for automatic rank determination. As seen in Fig. 1 the inference of $\boldsymbol{\lambda}$ can be performed by receiving messages from $N$ factor matrices and incorporating the messages
from its hyperprior. By applying (15), we can identify the posteriors of $\lambda_{r}, \forall r \in[1, R]$ as an independent Gamma distribution (see Sec. 4 of Appendix for details),

$$
\begin{equation*}
q_{\boldsymbol{\lambda}}(\boldsymbol{\lambda})=\prod_{r=1}^{R} \mathrm{Ga}\left(\lambda_{r} \mid c_{M}^{r}, d_{M}^{r}\right) \tag{24}
\end{equation*}
$$

where $c_{M}^{r}, d_{M}^{r}$ denote the posterior parameters learned from $M$ observations and can be updated by

$$
\begin{align*}
& c_{M}^{r}=c_{0}^{r}+\frac{1}{2} \sum_{n=1}^{N} I_{n}  \tag{25}\\
& d_{M}^{r}=d_{0}^{r}+\frac{1}{2} \sum_{n=1}^{N} \mathbb{E}_{q}\left[\mathbf{a}_{\cdot r}^{(n) T} \mathbf{a}_{\cdot r}^{(n)}\right] .
\end{align*}
$$

The expectation of the inner product of the $r$ th component in mode- $n$ matrix w.r.t. $q$ distribution can be evaluated using the posterior parameters in (16), i.e.,

$$
\begin{equation*}
\mathbb{E}_{q}\left[\mathbf{a}_{\cdot r}^{(n) T} \mathbf{a}_{\cdot r}^{(n)}\right]=\tilde{\mathbf{a}}_{\cdot r}^{(n) T} \tilde{\mathbf{a}}_{\cdot r}^{(n)}+\sum_{i_{n}}\left(\mathbf{V}_{i_{n}}^{(n)}\right)_{r r} \tag{26}
\end{equation*}
$$

By combining (25) and 26, we can further simplify the computation of $\mathbf{d}_{M}=\left[d_{M}^{T}, \ldots d_{M}^{R}\right]^{T}$ as

$$
\begin{equation*}
\mathbf{d}_{M}=\sum_{n=1}^{N}\left\{\operatorname{diag}\left(\tilde{\mathbf{A}}^{(n) T} \tilde{\mathbf{A}}^{(n)}+\sum_{i_{n}} \mathbf{V}_{i_{n}}^{(n)}\right)\right\} \tag{27}
\end{equation*}
$$

where $\tilde{\mathbf{A}}=\mathbb{E}_{q}\left[\mathbf{A}^{(n)}\right]$. Hence, the posterior expectation can be obtained by $\mathbb{E}_{q}[\boldsymbol{\lambda}]=\left[c_{M}^{1} / d_{M}^{1}, \ldots, c_{M}^{R} / d_{M}^{R}\right]^{T}$, and thus, $\mathbb{E}_{q}[\boldsymbol{\Lambda}]=\operatorname{diag}\left(\mathbb{E}_{q}[\boldsymbol{\lambda}]\right)$.

An intuitive interpretation of (25) is that $\lambda_{r}$ is updated by the sum of squared $L_{2}$-norm of $r$ th component, expressed by 26 , from $N$ factor matrices. Therefore, the smaller of $\left\|\mathbf{a}_{\cdot r}\right\|_{2}^{2}$ leads to larger $\mathbb{E}_{q}\left[\lambda_{r}\right]$ and updated priors of factor matrices, which in turn enforces more strongly the $r$ th component to be zero.

### 3.2.3 Posterior distribution of hyperparameter $\tau$

The inference of the noise precision $\tau$ can be performed by receiving the messages from observed data and its coparents, including $N$ factor matrices, and incorporating the messages from its hyperprior. By applying (15), the variational posterior is a Gamma distribution (see Sec. 5 of Appendix for details), given by

$$
\begin{equation*}
q_{\tau}(\tau)=\mathrm{Ga}\left(\tau \mid a_{M}, b_{M}\right) \tag{28}
\end{equation*}
$$

where the posterior parameters can be updated by

$$
\begin{align*}
a_{M} & =a_{0}+\frac{1}{2} \sum_{i_{1}, \ldots, i_{N}} \mathcal{O}_{i_{1} \ldots i_{N}} \\
b_{M} & =b_{0}+\frac{1}{2} \mathbb{E}_{q}\left[\left\|\mathcal{O} \circledast\left(\mathcal{Y}-\llbracket \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)} \rrbracket\right)\right\|_{F}^{2}\right] \tag{29}
\end{align*}
$$

However, the posterior expectation of model error in the above expression cannot be computed straightforwardly, and therefore, we need to introduce the following results.

Theorem 3.2. Assume a set of independent $R$-dimensional random vectors $\left\{\mathbf{x}^{(n)} \mid n=1, \ldots, N\right\}$, then

$$
\begin{equation*}
\mathbb{E}\left[\left\langle\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(N)}\right\rangle^{2}\right]=\left\langle\mathbb{E}\left[\mathbf{x}^{(1)} \mathbf{x}^{(1) T}\right], \ldots, \mathbb{E}\left[\mathbf{x}^{(N)} \mathbf{x}^{(N) T}\right]\right\rangle \tag{30}
\end{equation*}
$$

where the left term denotes the expectation of the squared inner product of $N$ vectors, and the right term denotes the inner product of $N$ matrices, where each matrix of size $R \times R$ denotes an expectation of the outer product of the $n$th vector, respectively.

Proof: See Sec. 6 of Appendix for details.
Theorem 3.3. Given a set of independent random matrices $\left\{\mathbf{A}^{(n)} \mid n=1, \ldots, N\right\}$, we assume that $\forall n, \forall i_{n}$, the row vectors $\left\{\mathbf{a}_{i_{n}}^{(n)}\right\}$ are independent, then

$$
\begin{align*}
\mathbb{E} & {\left[\left\|\llbracket \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)} \rrbracket\right\|_{F}^{2}\right] } \\
& =\sum_{i_{1}, \ldots, i_{N}}\left\langle\mathbb{E}\left[\mathbf{a}_{i_{1}}^{(1)} \mathbf{a}_{i_{1}}^{(1) T}\right], \ldots, \mathbb{E}\left[\mathbf{a}_{i_{N}}^{(N)} \mathbf{a}_{i_{N}}^{(N) T}\right]\right\rangle \tag{31}
\end{align*}
$$

Let $\mathbf{B}^{(n)}$ denote the expectation of a quadratic form related to $\mathbf{A}^{(n)}$ with $i_{n}$ th-row vector $\mathbf{b}_{i_{n}}^{(n)}=\operatorname{vec}\left(\mathbb{E}\left[\mathbf{a}_{i_{n}}^{(n)} \mathbf{a}_{i_{n}}^{(n) T}\right]\right)$; thus, (31) can be computed by

$$
\mathbb{E}\left[\left\|\llbracket \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)} \rrbracket\right\|_{F}^{2}\right]=\mathbf{1}_{\prod_{n} I_{n}}^{T}\left(\bigodot_{n} \mathbf{B}^{(n)}\right) \mathbf{1}_{R^{2}}
$$

Proof: See Sec. 7 of Appendix for details.
From Theorems 3.2 and 3.3 , the posterior expectation term in (29) can be evaluated explicitly. Due to the missing entries in $\mathcal{Y}$, the evaluation form is finally written as (see Sec. 8 of Appendix for details)

$$
\begin{align*}
& \mathbb{E}_{q} {\left[\left\|\mathcal{O} \circledast\left(\mathcal{Y}-\llbracket \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)} \rrbracket\right)\right\|_{F}^{2}\right] } \\
&=\left\|\mathcal{Y}_{\Omega}\right\|_{F}^{2}-2 \operatorname{vec}^{T}\left(\mathcal{Y}_{\Omega}\right) \operatorname{vec}\left(\llbracket \tilde{\mathbf{A}}^{(1)}, \ldots, \tilde{\mathbf{A}}^{(N)} \rrbracket_{\Omega}\right) \\
&+\operatorname{vec}^{T}(\mathcal{O})\left(\bigodot_{n} \mathbf{B}^{(n)}\right) \mathbf{1}_{R^{2}}, \tag{32}
\end{align*}
$$

where $\tilde{\mathbf{A}}^{(n)}=\mathbb{E}_{q}\left[\mathbf{A}^{(n)}\right]$ and $\mathbf{B}^{(n)}$ is computed by 20 . Hence, the posterior approximation of $\tau$ can be obtained by (29) together with the posterior expectation $\mathbb{E}_{q}[\tau]=$ $a_{M} / b_{M}$.

An intuitive interpretation of 29 is straightforward. $a_{M}$ is related to the number of observations and $b_{M}$ is related to the residual of model fitting measured by the squared Frobenius norm on observed entries.

### 3.2.4 Lower bound of model evidence

The inference framework presented in the previous section can essentially maximize the lower bound of model evidence that is defined in 13). Since the lower bound should not decrease at each iteration, it can be used to test for convergence. The lower bound of the logmarginal likelihood is computed by

$$
\begin{equation*}
\mathcal{L}(q)=\mathbb{E}_{q(\Theta)}\left[\ln p\left(\mathcal{Y}_{\Omega}, \Theta\right)\right]+H(q(\Theta)) \tag{33}
\end{equation*}
$$

where the first term denotes the posterior expectation of joint distribution, and the second term denotes the entropy of posterior distributions.

Various terms in the lower bound are evaluated and derived by taking parametric forms of $q$ distribution, giving the following results (see Sec. 9 of Appendix for details)

$$
\begin{align*}
& \mathcal{L}(q)= \\
& -\frac{a_{M}}{2 b_{M}} \mathbb{E}_{q}\left[\left\|\mathcal{O} \circledast\left(\mathcal{Y}-\llbracket \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)} \rrbracket\right)\right\|_{F}^{2}\right] \\
& -\frac{1}{2} \operatorname{Tr}\left\{\tilde{\mathbf{\Lambda}} \sum_{n}\left(\tilde{\mathbf{A}}^{(n) T} \tilde{\mathbf{A}}^{(n)}+\sum_{i_{n}} \mathbf{V}_{i_{n}}^{(n)}\right)\right\} \\
& +\frac{1}{2} \sum_{n} \sum_{i_{n}}\left\{\ln \left|\mathbf{V}_{i_{n}}^{(n)}\right|\right\}+\sum_{r}\left\{\ln \Gamma\left(c_{M}^{r}\right)\right\}  \tag{34}\\
& +\sum_{r}\left\{c_{M}^{r}\left(1-\ln d_{M}^{r}-\frac{d_{0}^{r}}{d_{M}^{r}}\right)\right\}+\ln \Gamma\left(a_{M}\right) \\
& +a_{M}\left(1-\ln b_{M}-\frac{b_{0}}{b_{M}}\right)+\mathrm{const} .
\end{align*}
$$

The posterior expectation of model error denoted by $\mathbb{E}_{q}[\cdot]$ can be computed using (32).

An intuitive interpretation of 34 is as follows. The first term is related to model residual; the second term is related to the weighted sum of squared $L_{2}$-norm of each component in factor matrices, while the uncertainty information is also considered; the rest terms are related to negative KL divergence between the posterior and prior distributions of hyperparameters.

### 3.2.5 Initialization of model parameters

The variational Bayesian inference is guaranteed to converge only to a local minimum. To avoid getting stuck in poor local solutions, it is important to choose an initialization point. In our model, the top level hyperparameters including $\mathbf{c}_{0}, \mathbf{d}_{0}, a_{0}, b_{0}$ are set to $10^{-6}$, resulting in a noninformative prior. Thus, we have $\mathbb{E}[\boldsymbol{\Lambda}]=\mathbf{I}$ and $\mathbb{E}[\tau]=1$. For the factor matrices, $\left\{\mathbb{E}\left[\mathbf{A}^{(n)}\right]\right\}_{n=1}^{N}$ can be initialized by two different strategies, one is randomly drawn from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ for $\mathbf{a}_{i_{n}}^{(n)}, \forall i_{n} \in\left[1, I_{n}\right], \forall n \in[1, N]$. The other is set to $\mathbf{A}^{(n)}=\mathbf{U}^{(n)} \boldsymbol{\Sigma}^{(n)^{\frac{1}{2}}}$, where $\mathbf{U}^{(n)}$ denotes the left singular vectors and $\boldsymbol{\Sigma}^{(n)}$ denotes the diagonal singular values matrix, obtained by SVD of mode- $n$ matricization of tensor $\mathcal{Y}$. The covariance matrix $\mathbf{V}^{(n)}$ is simply set to $\mathbf{I}$. The tensor rank $R$ is usually initialized by the weak upper bound on its maximum rank, i.e., $R \leq \min _{n} P_{n}$, where $P_{n}=\prod_{i \neq n} I_{i}$. In practice, we can also manually define the initialization value of $R$ for computational efficiency.

### 3.2.6 Interpretaion of automatic rank determination

The entire procedure of model inference is summarized in Algorithm 1 . It should be noted that tensor rank is determined automatically and implicitly. More specifically, updating $\boldsymbol{\lambda}$ in each iteration results in a new prior over $\left\{\mathbf{A}^{(n)}\right\}$, and then, $\left\{\mathbf{A}^{(n)}\right\}$ can be updated using this

```
Algorithm 1 Fully Bayesian CP Factorization (FBCP)
    Input: an \(N\) th-order incomplete tensor \(\mathcal{Y}_{\Omega}\) and an
    indicator tensor \(\mathcal{O}\).
    Initialization: \(\tilde{\mathbf{A}}^{(n)}, \mathbf{V}_{i_{n}}^{(n)}, \forall i_{n} \in\left[1, I_{n}\right], \forall n \in[1, N]\),
    \(a_{0}, b_{0}, \mathbf{c}_{0}, \mathbf{d}_{0}\), and \(\tau=a_{0} / b_{0}, \lambda_{r}=c_{0}^{r} / d_{0}^{r}, \forall r \in[1, R]\).
    repeat
        for \(n=1\) to \(N\) do
            Update the posterior \(q\left(\mathbf{A}^{(n)}\right)\) using 17);
        end for
        Update the posterior \(q(\boldsymbol{\lambda})\) using (25);
        Update the posterior \(q(\tau)\) using (29);
        Evaluate the lower bound using (34);
        Reduce rank \(R\) by eliminating zero-components of
        \(\left\{\mathbf{A}^{(n)}\right\}\) (an optional procedure);
    until convergence.
    Computation of predictive distributions using (35).
```

new prior in the subsequent iteration, which in turn affects $\boldsymbol{\lambda}$. Hence, if the posterior mean of $\lambda_{r}$ becomes very large, the $r$ th components in $\left\{\mathbf{A}^{(n)}\right\}, \forall n \in[1, N]$ are forced to be zero because of their prior information, and the tensor rank can be obtained by simply counting the number of non-zero components in the factor matrices. For implementation of the algorithm, we can keep the size of $\left\{\mathbf{A}^{(n)}\right\}$ unchanged during iterations; an alternative method is to eliminate the zero-components of $\left\{\mathbf{A}^{(n)}\right\}$ after each iteration.

### 3.3 Predictive Distribution

The predictive distributions over missing entries, given observed entries, can be approximated by using variational posterior distribution, that is,

$$
\begin{align*}
& p\left(\mathcal{Y}_{i_{1} \ldots i_{N}} \mid \mathcal{Y}_{\Omega}\right)=\int p\left(\mathcal{Y}_{i_{1} \ldots i_{N}} \mid \Theta\right) p\left(\Theta \mid \mathcal{Y}_{\Omega}\right) \mathrm{d} \Theta \\
& \simeq \iint p\left(\mathcal{Y}_{i_{1} \ldots i_{N}} \mid\left\{\mathbf{a}_{i_{n}}^{(n)}\right\}, \tau^{-1}\right) q\left(\left\{\mathbf{a}_{i_{n}}^{(n)}\right\}\right) q(\tau) \mathrm{d}\left\{\mathbf{a}_{i_{n}}^{(n)}\right\} \mathrm{d} \tau . \tag{35}
\end{align*}
$$

We can now approximate these integrations, yielding a Student's t-distribution $\mathcal{Y}_{i_{1} \ldots i_{N}} \mid \mathcal{Y}_{\Omega} \sim$ $\mathcal{T}\left(\tilde{\mathcal{Y}}_{i_{1} \ldots i_{N}}, \mathcal{S}_{i_{1} \ldots i_{N}}, \nu_{y}\right)$ (see Sec. 10 of Appendix for details) with its parameters given by

$$
\begin{aligned}
& \tilde{\mathcal{Y}}_{i_{1} \ldots i_{N}}=\left\langle\tilde{\mathbf{a}}_{i_{1}}^{(1)}, \cdots, \tilde{\mathbf{a}}_{i_{N}}^{(n)}\right\rangle, \quad \nu_{y}=2 a_{M}, \\
& \mathcal{S}_{i_{1} \ldots i_{N}}=\left\{\frac{b_{M}}{a_{M}}+\sum_{n}\left\{\left(\underset{k \neq n}{\circledast} \tilde{\mathbf{a}}_{i_{k}}^{(k)}\right)^{T} \mathbf{V}_{i_{n}}^{(n)}\left(\underset{k \neq n}{\circledast} \tilde{\mathbf{a}}_{i_{k}}^{(k)}\right)\right\}\right\}^{-1} .
\end{aligned}
$$

Thus, the predictive variance can be obtained by $\operatorname{Var}\left(\mathcal{Y}_{i_{1} \ldots i_{N}}\right)=\frac{\nu_{y}}{\nu_{y}-2} \mathcal{S}_{i_{1} \ldots i_{N}}^{-1}$.

### 3.4 Computational Complexity

The computation cost of the $N$ factor matrices in 17 is $O\left(N R^{2} M+R^{3} \sum_{n} I_{n}\right)$, where $N$ is the order of the tensor, $M$ denotes the number of observations, i.e., the input data size. $R$ is the number of latent components in
each $\mathbf{A}^{(n)}$, i.e., model complexity or tensor rank, and is generally much smaller than the data size, i.e., $R \ll M$. Hence, it has linear complexity w.r.t. the data size and polynomial complexity w.r.t. the model complexity. It should be noted that, because of the automatic model selection, the excessive latent components are pruned out in the first few iterations such that $R$ reduces rapidly in practice. The computation cost of the hyperparameter $\boldsymbol{\lambda}$ in 25 is $O\left(R^{2} \sum_{n} I_{n}\right)$, which is dominated by the model complexity, while the computation cost of noise precision $\tau$ in 29 is $O\left(R^{2} M\right)$. Therefore, the overall complexity of our algorithm is $O\left(N R^{2} M+R^{3}\right)$, which scales linearly with the data size but polynomially with the model complexity.

### 3.5 Discussion of Advantages

The advantages of our method are discussed as follows:

- The automatic determination of CP rank enables us to obtain an optimal low-rank tensor approximation, even from a highly noisy and incomplete tensor.
- Our method is characterized as a tuning parameterfree approach and all model parameters can be inferred from the observed data, which avoids the computational expensive parameter selection procedure. In contrast, the existing tensor factorization methods require a predefined rank, while the tensor completion methods based on nuclear norm require several tuning parameters.
- The uncertainty information over both latent factors and predictions of missing entries can be inferred by our method, while most existing tensor factorization and completion methods provide only the point estimations.
- An efficient and deterministic Bayesian inference is developed for model learning, which empirically shows a fast convergence.


## 4 Mixture Factor Priors

The low-rank assumption is powerful in general cases, however if the tensor data does not satisfy an intrinsic low-rank structure and a large amount of entries are missing, it may yield an oversimplified model. In this section, we present a variant of Bayesian CP factorization model which can take into account the local similarity in addition to the low-rank assumption.

We specify a Gaussian mixture prior over factor matrices such that the prior distribution in (7) can be rewritten as $\forall i_{n} \in\left[1, I_{n}\right], \forall n \in[1, N]$,
$p\left(\mathbf{a}_{i_{n}}^{(n)} \mid \boldsymbol{\lambda},\left\{\mathbf{a}_{k}^{(n)}\right\}\right)=w_{i_{n}, i_{n}} \mathcal{N}\left(\mathbf{0}, \boldsymbol{\Lambda}^{-1}\right)+\sum_{k \neq i_{n}} w_{i_{n}, k} \mathcal{N}\left(\mathbf{a}_{k}^{(n)}, \beta_{k}^{-1} \mathbf{I}\right)$, where $\sum_{k} w_{i_{n}, k}=1$. This indicates that the $i_{n}$ th row vector $a_{i_{n}}^{(n)}$ is similar to $k$ th row vectors with the probability of $w_{i_{n}, k}$. Based on our assumption that the adjacent rows are highly correlated, we can define the mixture coefficients by $w_{i, j}=z_{i} \exp \left(-|i-j|^{2}\right)$ where $z_{i}=$ $1 / \sum_{j} \exp \left(-|i-j|^{2}\right)$ is used to ensure the sum of mixture
coefficients to be 1 . For model learning, we can easily verify that the posterior distribution is also a mixture distribution. For simplicity, we set $\forall k, \beta_{k}=0$, thus the posterior mean of factor matrix can be updated firstly by 17 and then applying $\mathbb{E}_{q}\left[\mathbf{A}^{(n)}\right] \leftarrow \mathbf{W} \mathbb{E}_{q}\left[\mathbf{A}^{(n)}\right]$, while the posterior covariance $\left\{\mathbf{V}_{i_{n}}^{(n)}\right\}_{i_{n}=1}^{I_{n}}$ keep unchanged. Furthermore, the inference of all other variables do not need any changes.

## 5 EXPERIMENTAL Results

We conducted extensive experiments using both synthetic data and real-world applications, and compared our fully Bayesian CP factorization (FBCP) ${ }^{1}$ with seven state-of-the-art methods. Tensor factorization based scheme includes CPWOPT [7] and CPNLS [8], [40], while the completion based scheme includes HaLRTC and FaLRTC [10], FCSA [12], hard-completion (HardC.) [11], geomCG |9] and STDC [17]. Our objective when using synthetic data was to validate our method from several aspects: i) capability of rank determination; ii) reconstruction performance given a complete tensor; iii) predictive performance over missing entries given an incomplete tensor. Two real-world applications including image inpainting and facial image synthesis were used for demonstration. All experiments were performed by a PC (Intel Xeon(R) 3.3GHz, 64GB memory).

### 5.1 Validation on Synthetic Data

The synthetic tensor data is generated by the following procedure. $N$ factor matrices $\left\{\mathbf{A}^{(n)}\right\}_{n=1}^{N}$ are drawn from a standard normal distribution, i.e., $\forall n, \forall i_{n}, \mathbf{a}_{i_{n}}^{(n)} \sim$ $\mathcal{N}\left(\mathbf{0}, \mathbf{I}_{R}\right)$. Then, the true tensor is constructed by $\mathcal{X}=$ $\llbracket \mathbf{A}^{1}, \ldots, \mathbf{A}^{(N)} \rrbracket$, and an observed tensor by $\mathcal{Y}=\mathcal{X}+\varepsilon$, where $\varepsilon \sim \prod_{i_{1}, \ldots, i_{N}} \mathcal{N}\left(0, \sigma_{\varepsilon}^{2}\right)$ denotes i.i.d. additive noise. The missing entries, chosen uniformly, are marked by an indicator tensor $\mathcal{O}$.

### 5.1.1 A toy example

To illustrate our model, we provide two demo videos in the supplemental materials. A true latent tensor $\mathcal{X}$ is of size $10 \times 10 \times 10$ with $C P$ rank $R=5$, the noise parameter was $\sigma_{\varepsilon}^{2}=0.001$, and $40 \%$ of entries were missing. Then, we applied our method with the initial rank being set to 10. As shown in Fig. 2, three factor matrices are inferred in which five components are effectively pruned out, resulting in correct estimation of tensor rank. The lower bound of model evidence increases monotonically, which indicates the effectiveness and convergence of our algorithm. Finally, the posterior of noise precision $\tau \approx 1000$ implies the method's capability of denoising and the estimation of $\sigma_{\varepsilon}^{2} \approx 0.001, \mathrm{SNR}=10 \log \frac{\sigma_{\mathcal{x}}^{2}}{\tau-1}$.

1. Matlab codes are available at http://www.bsp.brain.riken.jp/ $\sim$ qibin/homepage/BayesTensorFactorization.html


$\begin{array}{lllll}2 & 4 & 6 & 8 & 10\end{array}$
Latent components


Fig. 2. A toy example illustrating FBCP applied on an incomplete tensor. The top row shows Hinton diagram of factor matrices, where the color and size of each square represent the sign and magnitude of the value, respectively. The bottom row shows the posterior of $\lambda$, the lower bound of model evidence, and the posterior of $\tau$ from left to right.

### 5.1.2 Automatic determination of tensor rank

To evaluate the automatic determination of tensor rank (i.e., $C P r a n k$ ), extensive simulations were performed under varying experimental conditions related to tensor size, tensor rank, noise level, missing ratio, and the initialization method of factor matrices (e.g., SVD or random sample). Each result is evaluated by 50 repetitions corresponding to 50 different tensors generated under the same criterion. There are four groups of experiments. (A) Given complete tensors of size $20 \times 20 \times 20$ with $R=5$, the evaluations were performed under varying noise levels and by two different initializations (see Fig. 3(a)). (B) Given incomplete tensors of size $20 \times 20 \times 20$ with $R=5$ and $\mathrm{SNR}=20 \mathrm{~dB}$, the evaluations were performed under five different missing ratios, and by different initializations (see Fig. 3(b)). (C) Given incomplete tensors with $R=5$ and $\mathrm{SNR}=0 \mathrm{~dB}$, the evaluations were performed under varying missing ratios and two different tensor sizes (see Fig. 3(c)). (D) Given incomplete tensors of size $20 \times 20 \times 20$ with SNR=20 dB, the evaluations were performed under varying missing ratios and two different true ranks (see Fig. 3(d)).

From Fig. 3. we observe that SVD initialization is slightly better than random initialization in terms of the determination of tensor rank. If the tensor is complete, our model can detect the true tensor rank with $100 \%$ accuracy when $S N R \geq 10 \mathrm{~dB}$. Although the accuracy decreased to $70 \%$ under a high noise level of 0 dB , the error deviation is only $\pm 1$. On the other hand, if the tensor is incomplete and almost free of noise, the detection rate is $100 \%$, when missing ratio is 0.7 , and is $44 \%$ with an error deviation of only $\pm 1$, even under a high missing ratio of 0.9 . As both missing data and high noise level are


Fig. 3. Determination of tensor rank under varying experimental conditions. Each vertical bar shows the mean and standard deviation of estimations from 50 repetitions, while the accuracy of detections is shown on the top of the corresponding bar. The red and blue horizontal dash dotted lines indicate the true tensor rank.
presented, our model can achieve $90 \%$ accuracy under the condition of $\mathrm{SNR}=0 \mathrm{~dB}$ and 0.5 missing ratio. It should be noted that, when the data size is larger, such as $50 \times 50 \times 50$, our model can achieve $90 \%$ accuracy, even when $\mathrm{SNR}=0 \mathrm{~dB}$ and the missing ratio is 0.9 . If the true rank is larger, such as $R=15$, the model can correctly recover the rank from a complete tensor, but fails to do so when the missing ratio is larger than 0.5 .

We can conclude from these results that the determination of the tensor rank depends primarily on the number of observed entries and the true tensor rank. In general, more observations are necessary if the tensor rank is larger; however, when high-level noise occurs, the excessive number of observations may not be helpful for rank determination.

### 5.1.3 Predictive performance

In this experiment, we considered incomplete tensors of size $20 \times 20 \times 20$ generated by the true rank $R=$ 5 and $\mathrm{SNR}=30 \mathrm{~dB}$ under varying missing ratios. The initial rank was set to 10 . The relative standard error $R S E=\frac{\|\mathcal{X}-\mathcal{X}\|_{F}}{\|\mathcal{X}\|_{F}}$, where $\hat{\mathcal{X}}$ denotes the estimation of the true tensor $\mathcal{X}$, was used to evaluate the performance. To ensure statistically consistent results, the performance is evaluated by 50 repetitions for each condition. As shown in Fig. 4. our method significantly outperforms other algorithms under all missing ratios. Factorization-based methods, including CPWOPT, and CPNLS show a better performance than completion-based methods when the missing ratio is relatively small, while they perform


Fig. 4. Predictive performance when $\mathrm{SNR}=30 \mathrm{~dB}$.
worse than completion methods when the missing ratio is large, e.g., 0.9. FaLRTC, FCSA, and HardC. achieve similar performances, because they are all based on nuclear norm optimization. geomCG achieves a performance comparable with that of CWOPT and CPNLS when data is complete, while it fails as the missing ratio becomes high. This is because geomCG requires a large number of observations and precisely defined rank. It should be noted that STDC outperforms all algorithms except FBCP as the missing ratio becomes extremely high. These results demonstrate that FBCP, as a tensor factorization method, can be also effective for tensor completion, even when an extremely sparse tensor is presented.

We also conducted two additional experiments. One is the reconstruction from a complete tensor, the other is the tensor completion when the noise level is high, i.e., $\mathrm{SNR}=0 \mathrm{~dB}$. The results of these two experiments are presented in Appendix (see Sec. 11, 12).

### 5.2 Image Inpainting



Fig. 5. Ground-truth of eight benchmark images.
In this section, the applications of image inpainting based on several benchmark images, shown in Fig. 5. are used to evaluate and compare the performance of different methods. The colorful image can be represented by a third-order tensor of size $200 \times 200 \times 3$. We conducted various experiments under four groups of conditions.

TABLE 1
Performance (RSEs) evaluated on missing pixels. "NF", " N " indicate noise free or noisy image. " T ", " S " indicate the text corruption or scrabbled image.

| Method | Facade |  | Lenna |  | Non-random |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | NF | N | NF | N | T | S |
| FBCP | 0.13 | 0.17 | 0.17 | 0.20 | 0.13 | 0.14 |
| CPWOPT | 0.33 | 0.41 | 0.25 | 0.41 | 0.18 | 0.18 |
| CPNLS | 0.84 | 0.84 | 0.62 | 0.73 | 0.22 | 0.30 |
| HaLRTC | 0.15 | 0.21 | 0.19 | 0.21 | 0.29 | 0.28 |
| FaLRTC | 0.16 | 0.21 | 0.19 | 0.22 | 0.29 | 0.29 |
| FCSA | 0.19 | 0.21 | 0.19 | 0.21 | 0.28 | 0.28 |
| HardC. | 0.19 | 0.25 | 0.20 | 0.23 | 0.31 | 0.30 |
| STDC | 0.14 | 0.22 | 0.11 | 0.20 | 0.15 | 0.16 |

(A) Structural image with uniformly random missing pixels. A building facade image with $95 \%$ missing pixels under two noise conditions, i.e., noise free and $\mathrm{SNR}=5 \mathrm{~dB}$, were considered as observations. (B) Natural image with uniformly random missing pixels. The Lena image of size $300 \times 300$ with $90 \%$ missing pixels under two noise conditions, i.e., noise free and $\mathrm{SNR}=10 \mathrm{~dB}$, were considered. (C) Non-random missing pixels. We conducted two experiments for image restoration from a corrupted image: 1) The Lenna image corrupted by superimposed text was used as an observed imag\& ${ }^{2}$. In practice, the location of text pixels are difficult to detect exactly; we can simply indicate missing entries by a value larger than 200 to ensure that the text pixels are completely missing. 2) The scrabbled Lenna image was used as an observed image and pixels with values larger than 200 can be marked as missing. (D) Object removal. Given an image and a mask covering the object area, our goal was to complete the image without that object. The algorithm settings of compared methods are described as follows. For factorization-based methods, the initial rank was set to 50 in cases of (A) and (B) due to the high missing ratios, and 100 in cases of (C) and (D). For completionbased methods, the tuning parameters were chosen by multiple runs and performance evaluated on the groundtruth of missing pixels.

The visual effects of image inpainting are shown in Fig. 6, and the predictive performances are shown in Table 1 where case (D) is not available due to the lack of ground-truth. In case (A), we observe that FBCP outperforms other methods for a structural image under an extremely high missing ratio and the superiority is more significant when an additive noise is involved. In case (B), observe that STDC obtains the best performance followed by FBCP that is better than other methods. However, STDC is severely degraded when noise is involved, and obtains the same performance as FBCP, while its visual quality is still much better than others. These indicate that the additional smooth constraints in STDC are suitable for natural image. In case (C), FBCP is superior to all other methods, followed by

[^0]

Fig. 6. Visual effects of image inpainting. Seven examples shown from top to bottom are (1) facade image with $95 \%$ missing; (2) facade image with $95 \%$ missing and an additive noise; (3) lena image with $90 \%$ missing; (4) lena image with $90 \%$ missing and an additive noise; (5) lena image with superimposed text; (6) scribbled lena image; (7) an image of ocean with an object.

STDC. The completion-based methods obtain relatively smoother effects than factorization-based methods, but the global color of the image is not recovered well, resulting in a poor predictive performance. In case (D), FBCP obtains the most clean image by removing the object completely while the ghost effects appear in all other methods. HaLRTC, FaLRTC and FCSA outperform CPWOPT, CPNLS and STDC.

From these results we can conclude that the completion-based methods generally outperforms factorization-based methods for image completion. However, FBCP significantly improves ability of factorization-based scheme by automatic model selection and robustness to overfitting, resulting in potential applications for various image inpainting problems. The necessary number of observed entries mainly depends on the rank of true image. For instance, a structural image with an intrinsic low-rank need very fewer observations than a natural image. However only $10 \%$ observed pixels from lena image are not sufficient
to recover the whole image, which is caused by the low-rank assumption. This property is common for all these algorithms except STDC, because STDC employs an auxiliary information as additional constraints. The advantage of STDC has been shown for lena image, while its disadvantages are that the auxiliary information must be well designed for each specific application, which make it difficult to be applied to other types of data. In addition, STDC degrades in presence of the non-random missing pixels or noise. Moreover, the performance of HaLRTC and STDC are sensitive to tuning parameters that must be carefully selected for each specific condition. Therefore, a crucial drawback of completion-based scheme lies in the tuning parameters whose selection is quite challenging when the ground-truth of missing data is unknown.
Next, we perform image completion extensively on eight images in Fig. 5 with randomly missing pixels. Since most of these images are natural images on which the low-rank approximation cannot recover the missing

TABLE 2
The averaged recovery performance (RSE, PSNR, SSIM) and runtime (seconds) on eight images with missing rates of $70 \%, 80 \%, 90 \%$ and $95 \%$. For methods that need to tune parameters, both the runtime with the best tuning parameter and the overall runtime are reported.

|  |  | FBCP | FBCP-MP | CPWOPT | STDC | HaLRTC | FaLRTC | FCSA | HardC. |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $70 \%$ | RSE | 0.1209 | $\mathbf{0 . 0 9 8 6}$ | 0.1493 | 0.1003 | 0.1205 | 0.1205 | 0.1406 | 0.1254 |
|  | PSNR | 25.13 | 26.78 | 23.35 | $\mathbf{2 6 . 7 9}$ | 25.12 | 25.12 | 23.64 | 24.71 |
|  | SSIM | 0.7546 | $\mathbf{0 . 8 5 3 1}$ | 0.6417 | 0.8245 | 0.7830 | 0.7831 | 0.7437 | 0.7730 |
|  | Runtime | 83 | 251 | $1807 / 2908$ | $32 / 292$ | $18 / 139$ | $46 / 232$ | $\mathbf{9}$ | 36 |
| $80 \%$ | RSE | 0.1423 | $\mathbf{0 . 1 0 8 4}$ | 0.1700 | 0.1095 | 0.1479 | 0.1479 | 0.1675 | 0.1548 |
|  | PSNR | 23.09 | 25.36 | 21.52 | $\mathbf{2 5 . 4 0}$ | 22.72 | 22.72 | 21.53 | 22.32 |
|  | SSIM | 0.6515 | $\mathbf{0 . 7 9 4 1}$ | 0.5567 | 0.7781 | 0.6716 | 0.6716 | 0.6410 | 0.6579 |
|  | Runtime | 76 | 196 | $590 / 2316$ | $30 / 328$ | $25 / 122$ | $57 / 282$ | $\mathbf{9}$ | 39 |
| $90 \%$ | RSE | 0.1878 | $\mathbf{0 . 1 2 9 5}$ | 0.2372 | 0.1316 | 0.1992 | 0.1995 | 0.2342 | 0.2121 |
|  | PSNR | 20.12 | $\mathbf{2 3 . 2 6}$ | 18.08 | 23.21 | 19.62 | 19.61 | 18.09 | 19.11 |
|  | SSIM | 0.4842 | $\mathbf{0 . 6 9 5 6}$ | 0.3628 | 0.6950 | 0.5005 | 0.4998 | 0.4477 | 0.4790 |
|  | Runtime | 69 | 169 | $390 / 1475$ | $32 / 378$ | $21 / 127$ | $61 / 307$ | $\mathbf{9}$ | 32 |
| $95 \%$ | RSE | 0.2420 | $\mathbf{0 . 1 5 6 6}$ | 0.3231 | 0.1600 | 0.2549 | 0.2564 | 0.2777 | 0.2903 |
|  | PSNR | 17.76 | $\mathbf{2 1 . 3 4}$ | 15.35 | 21.18 | 17.25 | 17.19 | 16.39 | 16.12 |
|  | SSIM | 0.3455 | $\mathbf{0 . 6 0 3 1}$ | 0.2539 | 0.5810 | 0.3676 | 0.3649 | 0.3535 | 0.3369 |
|  | Runtime | 66 | 133 | $201 / 881$ | $35 / 400$ | $24 / 137$ | $63 / 313$ | $\mathbf{8}$ | 32 |

pixels well, we apply the fully Bayesian CP with mixture priors (FBCP-MP) for comparison with FBCP and other related methods. For FBCP and FBCP-MP, the same initialization of $R=100$ was applied, while CPWOPT was performed by using the optimal ranks obtained from FBCP and FBCP-MP to show the best performance. The parameter selection for other methods was same with previous experiments. The size of all images is $256 \times 256 \times 3$. Table 2 shows quantitative results in terms of recovery performance and runtime. Observe that FBCP-MP improves the performance of FBCP significantly and achieves the best recovery performance, especially in the case of high missing rate. The time costs of FBCP and FBCP-MP are comparable with completionbased methods and significantly lower than other tensor factorization method. STDC obtains the comparable performance with FBCP-MP, however the parameters must be manually tuned for the specific condition. More detailed results on each image are shown visually and quantitatively in the supplemental materials. These results demonstrate the effectiveness of mixture priors and advantages when the local similarity is taken into account in addition to the low-rank assumption.

### 5.3 Facial Image Synthesis

For recognition of face images captured from surveillance videos, the ideal solution is to create a robust classifier that is invariant to some factors, such as pose and illumination. Hence, there arises the question whether we can generate novel facial images under multiple conditions given images under other conditions. Tensors are highly suitable for modeling a multifactor image ensemble, and therefore, we introduce a novel application
of facial image synthesis that utilizes tensor factorization approaches.


Fig. 7. Facial images under multiple conditions where some images are fully missing.

We used the dataset of 3D Basel Face Model [41], which contains an ensemble of facial images of 10 people, each rendered in 9 different poses under 3 different illuminations. All 270 facial images were decimated and cropped to $68 \times 68$ pixels, and were then represented by a fourth-order tensor of size $4624 \times 10 \times 9 \times 3$. As shown in Fig. 7. some images were fully missing. Since some methods are either computationally intractable or not applicable to $N \geq 4$ order tensor, five algorithms were finally applied on this dataset under different missing ratios. The initial rank was set to 100 in factorization based methods, while the parameters of completion based methods were well tuned based on the groundtruth of missing images.


Fig. 8. The ground-truth of 49 missing facial images and the synthetic images by different methods.

TABLE 3
RSEs on observed images ( $O$ ) and missing images (M). The cases of $36,49,64$ and 81 missing images were tested.

| Method | 36/270 |  | 49/270 |  | 64/270 |  | 81/270 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | O | M | O | M | O | M | O | M |
| FBCP | 0.05 | 0.10 | 0.05 | 0.10 | 0.05 | 0.15 | 0.05 | 0.20 |
| CPWOPT | 0.50 | 0.65 | 0.55 | 0.61 | 0.57 | 0.59 | 0.62 | 0.73 |
| HaLRTC | N/A | 0.28 | N/A | 0.30 | N/A | 0.31 | N/A | 0.34 |
| FaLRTC | N/A | 0.28 | N/A | 0.30 | N/A | 0.31 | N/A | 0.34 |
| HardC. | 0.37 | 0.37 | 0.36 | 0.40 | 0.36 | 0.40 | 0.36 | 0.40 |

As shown in Fig. 8, the visual effects of image synthesis produced by FBCP are significantly superior to those produced by other methods. Although both CPWOPT and HaLRTC produce images that are too smooth and blurred, HaLRTC obtains much better visual quality than CPWOPT. The detailed performances are compared in Table 3, where RSE w.r.t. observed entries reflects the performance of model fitting, and RSE w.r.t. missing entries particularly reflects the predictive ability. Note that $\mathrm{RSE}=\mathrm{N} / \mathrm{A}$ implies that HaLRTC and FaLRTC donot model the observed entries. The inferred rank by FBCP is within the range of $[98,100]$ that are close to the initialization. Observe that completion based methods including HaLRTC, FaLRTC and HardC. achieve better performance than CPWOPT. However, FBCP demonstrates the possibility that factorization-based scheme can significantly outperform completion-based methods, especially in terms of performance on missing images.

## 6 Conclusion

In this paper, we proposed a fully Bayesian CP factorization which can naturally handle incomplete and noisy tensor data. By employing hierarchical priors over all unknown parameters, we derived a deterministic model inference under a fully Bayesian treatment. The most significant advantage is automatic determination of $C P$ rank. Moreover, as a tuning parameter-free approach, our method avoids the parameter selection problem and can also effectively prevent overfitting. In addition, we proposed a variant of our method by using mixture priors, which shows advantages on natural images with
a large amount of missing pixels. Empirical results validate the effectiveness in terms of discovering the ground-truth of tensor rank and imputing missing values for an extremely sparse tensor. Several real-world applications, such as image completion and image synthesis, demonstrated the superiority of our method over state-of-the-art techniques. Due to several interesting properties, our method would be attractive for many potential applications.

## References

[1] T. Kolda and B. Bader, "Tensor Decompositions and Applications," SIAM Review, vol. 51, no. 3, pp. 455-500, 2009.
[2] A. Cichocki, R. Zdunek, A. H. Phan, and S. I. Amari, Nonnegative Matrix and Tensor Factorizations. John Wiley \& Sons, 2009.
[3] L. R. Tucker, "Some mathematical notes on three-mode factor analysis," Psychometrika, vol. 31, no. 3, pp. 279-311, 1966.
[4] R. A. Harshman, "Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multimodal factor analysis," UCLA Working Papers in Phonetics, vol. 16, pp. 1-84, 1970.
[5] R. Bro, "PARAFAC. Tutorial and applications," Chemometrics and intelligent laboratory systems, vol. 38, no. 2, pp. 149-171, 1997.
[6] M. Sørensen, L. De Lathauwer, P. COMON, S. ICART, and L. DENEIRE, "Canonical polyadic decomposition with orthogonality constraints," SIAM Journal on Matrix Analysis and Applications, no. accepted, 2012.
[7] E. Acar, D. M. Dunlavy, T. G. Kolda, and M. Mørup, "Scalable tensor factorizations for incomplete data," Chemometrics and Intelligent Laboratory Systems, vol. 106, no. 1, pp. 41-56, 2011.
[8] L. Sorber, M. Van Barel, and L. De Lathauwer, "Optimizationbased algorithms for tensor decompositions: Canonical polyadic decomposition, decomposition in rank-(l_r,1_r,1) terms, and a new generalization," SIAM Journal on Optimization, vol. 23, no. 2, pp. 695-720, 2013.
[9] D. Kressner, M. Steinlechner, and B. Vandereycken, "Low-rank tensor completion by riemannian optimization," BIT Numerical Mathematics, pp. 1-22, 2013.
[10] J. Liu, P. Musialski, P. Wonka, and J. Ye, "Tensor completion for estimating missing values in visual data," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 35, no. 1, pp. 208220, 2013.
[11] M. Signoretto, Q. T. Dinh, L. De Lathauwer, and J. A. Suykens, "Learning with tensors: a framework based on convex optimization and spectral regularization," Machine Learning, pp. 1-49, 2013.
[12] J. Huang, S. Zhang, H. Li, and D. Metaxas, "Composite splitting algorithms for convex optimization," Computer Vision and Image Understanding, vol. 115, no. 12, pp. 1610-1622, 2011.
[13] S. Gandy, B. Recht, and I. Yamada, "Tensor completion and low-n-rank tensor recovery via convex optimization," Inverse Problems, vol. 27, no. 2, p. 025010, 2011.
[14] H. Tan, B. Cheng, W. Wang, Y.-J. Zhang, and B. Ran, "Tensor completion via a multi-linear low-n-rank factorization model," Neurocomputing, vol. 133, pp. 161-169, 2014.
[15] G. Zhong and M. Cheriet, "Large margin low rank tensor analysis," Neural computation, vol. 26, no. 4, pp. 761-780, 2014.
[16] A. Krishnamurthy and A. Singh, "Low-rank matrix and tensor completion via adaptive sampling," in Advances in Neural Information Processing Systems, 2013, pp. 836-844.
[17] Y. Chen, C. Hsu, and H. Liao, "Simultaneous tensor decomposition and completion using factor priors," Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. PP, no. 99, pp. 1-1, 2013.
[18] A. Narita, K. Hayashi, R. Tomioka, and H. Kashima, "Tensor factorization using auxiliary information," Data Mining and Knowledge Discovery, vol. 25, no. 2, pp. 298-324, 2012.
[19] J. Håstad, "Tensor rank is NP-complete," Journal of Algorithms, vol. 11, no. 4, pp. 644-654, 1990.
[20] B. Alexeev, M. A. Forbes, and J. Tsimerman, "Tensor rank: Some lower and upper bounds," in Computational Complexity (CCC), 2011 IEEE 26th Annual Conference on. IEEE, 2011, pp. 283-291.
[21] P. Bürgisser and C. Ikenmeyer, "Geometric complexity theory and tensor rank," in Proceedings of the 43 rd annual ACM symposium on Theory of computing. ACM, 2011, pp. 509-518.
[22] V. De Silva and L.-H. Lim, "Tensor rank and the ill-posedness of the best low-rank approximation problem," SIAM Journal on Matrix Analysis and Applications, vol. 30, no. 3, pp. 1084-1127, 2008.
[23] E. S. Allman, P. D. Jarvis, J. A. Rhodes, and J. G. Sumner, "Tensor rank, invariants, inequalities, and applications," SIAM Journal on Matrix Analysis and Applications, vol. 34, no. 3, pp. 1014-1045, 2013.
[24] R. Salakhutdinov and A. Mnih, "Probabilistic matrix factorization," Advances in neural information processing systems, vol. 20, pp. 1257-1264, 2008.
[25] _, "Bayesian probabilistic matrix factorization using Markov chain Monte Carlo," in Proceedings of the 25th international conference on Machine learning. ACM, 2008, pp. 880-887.
[26] S. D. Babacan, M. Luessi, R. Molina, and A. K. Katsaggelos, "Sparse Bayesian methods for low-rank matrix estimation," IEEE Transactions on Signal Processing, vol. 60, no. 8, pp. 3964-3977, 2012.
[27] Y. J. Lim and Y. W. Teh, "Variational Bayesian approach to movie rating prediction," in Proceedings of KDD Cup and Workshop, 2007, pp. 15-21.
[28] N. D. Lawrence and R. Urtasun, "Non-linear matrix factorization with Gaussian processes," in Proceedings of the 26th Annual International Conference on Machine Learning. ACM, 2009, pp. 601-608.
[29] B. Lakshminarayanan, G. Bouchard, and C. Archambeau, "Robust Bayesian matrix factorisation," in Proc. Intl. Conf. on Artificial Intelligence and Statistics (AISTATS), 2011.
[30] W. Chu and Z. Ghahramani, "Probabilistic models for incomplete multi-dimensional arrays," in JMLR Workshop and Conference Proceedings, vol. 5, 2009, pp. 89-96.
[31] S. Gao, L. Denoyer, P. Gallinari, and J. GUO, "Probabilistic latent tensor factorization model for link pattern prediction in multirelational networks," The Journal of China Universities of Posts and Telecommunications, vol. 19, pp. 172-181, 2012.
[32] L. Xiong, X. Chen, T.-K. Huang, J. Schneider, and J. G. Carbonell, "Temporal collaborative filtering with Bayesian probabilistic tensor factorization," in Proceedings of SIAM Data Mining, vol. 2010, 2010.
[33] K. Hayashi, T. Takenouchi, T. Shibata, Y. Kamiya, D. Kato, K. Kunieda, K. Yamada, and K. Ikeda, "Exponential family tensor factorization for missing-values prediction and anomaly detection," in Data Mining (ICDM), 2010 IEEE 10th International Conference on. IEEE, 2010, pp. 216-225.
[34] Z. Xu, F. Yan, and A. Qi, "Infinite Tucker decomposition: Nonparametric Bayesian models for multiway data analysis," in Proceedings of the 29th International Conference on Machine Learning (ICML-12), 2012, pp. 1023-1030.
[35] D. J. MacKay, "Bayesian methods for backpropagation networks," in Models of Neural Networks III. Springer, 1996, pp. 211-254.
[36] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," The Journal of Machine Learning Research, vol. 1, pp. 211244, 2001.
[37] C. M. Bishop, "Bayesian PCA," Advances in neural information processing systems, pp. 382-388, 1999.
[38] J. M. Winn and C. M. Bishop, "Variational message passing," Journal of Machine Learning Research, vol. 6, pp. 661-694, 2005.
[39] C. M. Bishop et al., Pattern recognition and machine learning. springer New York, 2006.
[40] L. Sorber, M. V. Barel, and L. D. Lathauwer, "Tensorlab v1.0," Available online, URL: http://esat. kuleuven. be/sista/tensorlab, 2013. [Online]. Available: http://esat.kuleuven.be/sista/tensorlab/
[41] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and I. Vetter, "A 3D face model for pose and illumination invariant face recognition," in Advanced Video and Signal Based Surveillance, Sixth IEEE International Conference on. IEEE, 2009, pp. 296-301.


Qibin Zhao received the Ph.D. degree in engineering from Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China, in 2009. He is currently a research scientist at Laboratory for Advanced Brain Signal Processing in RIKEN Brain Science Institute, Japan. His research interests include machine learning, tensor factorization, computer vision and brain computer interface.


Liqing Zhang received the Ph.D. degree from Zhongshan University, Guangzhou, China, in 1988. He is now a Professor with Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China. His current research interests cover computational theory for cortical networks, brain-computer interface, perception and cognition computing model, statistical learning and inference. He has published more than 170 papers in international journals and conferences.


Andrzej Cichocki received the Ph.D. and Dr.Sc. (Habilitation) degrees, all in electrical engineering, from Warsaw University of Technology (Poland). He is currently the senior team leader of the laboratory for Advanced Brain Signal Processing, at RIKEN Brain Science Institute (JAPAN). He is co-author of more than 250 technical papers and 4 monographs (two of them translated to Chinese). He is Associate Editor of Journal of Neuroscience Methods and IEEE Transaction on Signal Processing.


[^0]:    2. A demo video is available in the supplemental materials.
