

Towards Flexible Sparsity-Aware Modeling: Automatic Tensor Rank Learning Using the Generalized Hyperbolic Prior

Lei Cheng^{1b}, Zhongtao Chen, Qingjiang Shi^{1b}, Yik-Chung Wu^{1b}, *Senior Member, IEEE*,
and Sergios Theodoridis^{2b}, *Life Fellow, IEEE*

Abstract—Tensor rank learning for canonical polyadic decomposition (CPD) has long been deemed as an essential yet challenging problem. In particular, since the tensor rank controls the complexity of the CPD model, its inaccurate learning would cause overfitting to noise or underfitting to the signal sources, and even destroy the interpretability of model parameters. However, the optimal determination of a tensor rank is known to be a non-deterministic polynomial-time hard (NP-hard) task. Rather than exhaustively searching for the best tensor rank via trial-and-error experiments, Bayesian inference under the Gaussian-gamma prior was introduced in the context of probabilistic CPD modeling, and it was shown to be an effective strategy for automatic tensor rank determination. This triggered flourishing research on other structured tensor CPDs with automatic tensor rank learning. On the other side of the coin, these research works also reveal that the Gaussian-gamma model does not perform well for high-rank tensors and/or low signal-to-noise ratios (SNRs). To overcome these drawbacks, in this paper, we introduce a more advanced generalized hyperbolic (GH) prior to the probabilistic CPD model, which not only includes the Gaussian-gamma model as a special case, but also is more flexible to adapt to different levels of sparsity. Based on this novel probabilistic model, an algorithm is developed under the framework of variational inference, where each update is obtained in a closed-form. Extensive numerical results, using synthetic data and real-world datasets, demonstrate the significantly improved

performance of the proposed method in learning both low as well as high tensor ranks even for low SNR cases.

Index Terms—Automatic tensor rank learning, tensor CPD, generalized hyperbolic distribution, Bayesian learning, variational inference.

I. INTRODUCTION

IN THE Big Data era, tensor decomposition has become one of the most important tools in both theoretical studies of machine learning [1], [2] and a variety of real-world applications [3]–[10]. Among all the tensor decompositions, canonical polyadic decomposition (CPD) is the most fundamental format. It not only provides a faithful representation of a lot of real-world multidimensional data, see, e.g., [14]–[17], but it also allows a unique factor matrix recovery up to trivial scaling and permutation ambiguities [11]. This uniqueness bolsters the uncovering of the interpretable knowledge from the tensor data, which endows tensor CPD an advantage over other learning approaches (e.g., deep models), in various data analytic tasks including social group mining [14], drug discovery [15], biomedical data analytics [16] and functional Magnetic Resonance Imaging (fMRI) [17].

In tensor CPD, given an N dimensional (N-D) data tensor $\mathcal{Y} \in \mathbb{R}^{J_1 \times \dots \times J_N}$, a set of factor matrices $\{\mathbf{U}^{(n)} \in \mathbb{R}^{J_n \times R}\}$ are sought via solving the following problem [11]:

$$\min_{\{\mathbf{U}^{(n)}\}_{n=1}^N} \left\| \mathcal{Y} - \underbrace{\sum_{r=1}^R \mathbf{U}_{:,r}^{(1)} \circ \mathbf{U}_{:,r}^{(2)} \circ \dots \circ \mathbf{U}_{:,r}^{(N)}}_{\triangleq [\mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)}]} \right\|_F^2, \quad (1)$$

where the symbol \circ denotes vector outer product and the shorthand notation $[\cdot \cdot \cdot]$ is termed as the Kruskal operator. As illustrated in Fig. 1, the tensor CPD aims at decomposing an N-D tensor into a summation of R rank-1 tensors, with the r^{th} component constructed as the vector outer product of the r^{th} columns from all the factor matrices, i.e., $\{\mathbf{U}_{:,r}^{(n)}\}_{n=1}^N$. In problem (1), the number of columns R of each factor matrix, also known as tensor rank [11], determines the number of unknown model parameters and equivalently the model complexity. In practice, it needs to be carefully selected to achieve the best performance in both recovering the noise-free signals (e.g., image denoising [12]) and unveiling the underlying components (e.g., social group clustering [14]).

Manuscript received August 24, 2021; revised February 13, 2022; accepted March 25, 2022. Date of publication April 1, 2022; date of current version April 26, 2022. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Evrim Acar. This work was supported in part by the NSFC under Grant 62001309, in part by Science and Technology on Sonar Laboratory under Grant 6142109KF212204, in part by the National Key Research and Development Project under Grant 2017YFE0119300, in part by the NSFC under Grants 61671411, 61731018, and U1709219, and in part by the General Research Fund from the Hong Kong Research Grant Council under Grant 17207018. (*Corresponding author: Yik-Chung Wu.*)

Lei Cheng is with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310027, China, and also with Zhejiang Provincial Key Laboratory of Information Processing, Communication and Networking (IPCAN), Hangzhou 310027, China (e-mail: lei_cheng@zju.edu.cn).

Zhongtao Chen and Yik-Chung Wu are with the Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong (e-mail: ztchen@connect.hku.hk; ycwu@eee.hku.hk).

Qingjiang Shi is with the School of Software Engineering, Tongji University, Shanghai 201804, China, and also with the Shenzhen Research Institute of Big Data, Shenzhen 518172, China (e-mail: shiqj@tongji.edu.cn).

Sergios Theodoridis is with the National and Kapodistrian University of Athens, 157 72 Athens, Greece, and also with the Department of Electronic Systems, Aalborg University, 9220 Aalborg, Denmark (e-mail: stheodor@di.uoa.gr).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TSP.2022.3164200>, provided by the authors.

Digital Object Identifier 10.1109/TSP.2022.3164200




Fig. 1. Illustration of tensor CPD.

If the value of the tensor rank is known, problem (1) can be solved via nonlinear programming methods [18]. In particular, it has been found that problem (1) enjoys a nice block multi-convexity property, in the sense that after fixing all but one factor matrix, the problem is convex with respect to that matrix. This property motivates the use of block coordinate descent (BCD) methods (or alternating optimization methods) to devise fast and accurate algorithms for tensor CPD and its structured variants [19]–[21]. From a nonlinear programming perspective, these solutions need the knowledge of the tensor rank R , which, however, is unknown and, in general, it is non-deterministic polynomial-time hard (NP-hard) to obtain [11]. To acquire the optimal tensor rank (or equivalently the optimal model complexity), trial-and-error parameter tuning has been employed in previous works [12]–[16], which, unfortunately, is computationally costly.

To tackle the challenge of automatic tensor rank learning, a novel heuristic test method based on the core consistency diagnostic was proposed in [24], [25] and later efficiently implemented in [26], which, however, still needs to compute CPD multiple times. Recently, relying on a regularization scheme, model selection for block-term tensor decomposition, which includes CPD as its special case, was investigated in [27], in which hyper-parameter tuning is inevitable.

A breakthrough to the challenge of automatic tensor rank learning has been achieved under the framework of Bayesian modeling and inference. In the early work [28], sparsity-promoting priors (including Gaussian-gamma prior and Laplacian prior) were adopted, based on which the Tucker/CPD model parameters were estimated via the maximum-a-posterior (MAP) approach (also called “poor-man” Bayesian inference). Fully Bayesian treatments for probabilistic tensor CPD have been achieved in [29], [38] using sampling schemes, and the variational inference (VI) framework [22], [30], [31], which is more scalable to massive data [50]–[53].

Unlike [30], [31] which focus on count data, continuous tensor data was considered in [22] and the follow-up works [23], [32]–[37]. These two data types result in different likelihood functions (or cost functions) and prior distribution design principles. In particular, the likelihood functions in [30], [31] are poisson and binomial distributions, respectively, while [22], [23], [32]–[37] adopt Gaussian distributions (and its variants) as likelihood functions.

Since the inception of [22], the development of VI-based structured continuous-valued tensor decompositions with automatic tensor rank learning [23], [32]–[37] has been flourishing. Among these works, the key idea lies in the adoption of sparsity-promoting Gaussian-gamma prior and its variants for modeling the powers of columns in all the factor matrices, so that most columns in the factor matrices will be driven to




Fig. 2. Tensor rank learning results from probabilistic tensor CPD with Gaussian-gamma prior [22]. The vertical bars show the mean and the error bars indicate the standard derivation of tensor rank estimates. The black dashed lines show the true tensor rank.

zero during inference. Then, the number of remaining non-zero columns in each factor matrix is used to estimate the tensor rank. Extensive numerical studies using both synthetic data and real-world datasets have demonstrated the effectiveness of these methods [22], [23], [32]–[37].

In addition to learning the tensor rank of basic CPD, the Gaussian-gamma prior has recently become an important part in the probabilistic modeling of several advanced tensor models. For example, it was utilized as the cornerstone in [39] to infer the tensor tubal rank and multi-rank of tensor SVD, and cleverly integrated with the spike-and-slab prior in [40] to achieve simultaneous view-wise and feature-wise sparsity learning for multi-view tensor factorization, in which multiple tensors are jointly decomposed with shared latent factors.

While it may seem that tensor rank learning in CPD is solved, a closer inspection on the numerical results reveals that the performance of tensor rank learning deteriorates significantly when the noise power is large, and/or the true tensor rank is close to the dimension of the tensor data. As an illustration, consider three dimensional (3D) signal tensors¹ with dimension $30 \times 30 \times 30$ and the tensor ranks belonging to $\{6, 12, 24\}$. The observation tensor data are obtained by corrupting the signal tensors using additive white gaussian noises (AWGNs), with the noise power characterized by signal-to-noise ratio (SNR). With the tensor rank upper bound being 30, the probabilistic tensor CPD algorithm [22] was run on the observation tensor data. The tensor rank learning results from 100 Monte-Carlo trials are shown in Fig. 2. It is clear that when the SNR is 0 dB, the probabilistic tensor CPD algorithm with Gaussian-gamma model [22] either over-estimates or under-estimates the tensor rank. At a high SNR (i.e., 10 dB), although the method [22] estimates the correct rank value for low tensor ranks $\{6, 12\}$, it fails to recover the high tensor rank 24. In practice, over-estimation of a tensor rank will either lead to overfitting of the noise components or to generating uninterpretable “ghost” components. On the other hand, if the tensor rank is under-estimated, it is prone to missing important signal components. Therefore, there is a need for

¹Each signal tensor is generated via $\mathcal{X} = \llbracket \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \mathbf{U}^{(3)} \rrbracket \in \mathbb{R}^{30 \times 30 \times 30}$, where each element in the factor matrices $\{\mathbf{U}^{(n)}\}_{n=1}^3$ is independently drawn from $\mathcal{N}(0, 1)$.

further improving the accuracy of automatic tensor rank learning in CPD.

To achieve this goal via a principled approach, it is worthwhile to trace back the development of the Gaussian-gamma prior and the Bayesian framework from the early work of Mackay [41] and Tipping [42] on Bayesian neural networks and relevance vector machines (RVMs), in which important variables are automatically identified in the associated models. This hints us that further inspiration can be drawn from the early works [41], [42] and beyond [43]–[47]. In particular, to adapt to different levels of sparsity, which is measured by the ratio of zero-valued parameters to the total number of parameters, advanced sparsity-promoting priors including generalized-t distribution [43], normal-exponential gamma distribution [44], horseshoe distribution [45] and generalized hyperbolic distribution [46], [47] could be employed. Since these advanced priors are much more flexible in their functional forms compared to the Gaussian-gamma prior (and some of them even include the Gaussian-gamma model as their special cases), improved performance of variable selection was witnessed in linear regression models. We conjecture that this group of advanced sparsity-promoting priors would improve CPD rank selection compared to the Gaussian-gamma model.

In this paper, we take *the first step* towards developing “*Flexible Sparsity-Aware Modeling*” by introducing the generalized hyperbolic (GH) prior [46] into the research of probabilistic tensor CPD. The reason for choosing this prior is that not only it includes the widely-used Gaussian-gamma prior and Laplacian prior as its special cases, but also its mathematical form allows for an efficient expectation computation. Furthermore, the GH prior can be interpreted as a Gaussian scale mixture where the mixing distribution is the generalized inverse Gaussian (GIG) distribution [48]. This interpretation allows for a hierarchical construction of the probabilistic model with the conjugacy property within the exponential distribution family, based on which efficient variational inference (VI) algorithms [50]–[53] can be devised with closed-form update expressions.

By making full use of these advantages, we design a novel probabilistic tensor CPD model and the corresponding inference algorithm. Since the GH prior provides a more flexible sparsity-aware modeling than the Gaussian-gamma prior, and all the latent variables are updated during the learning process, the GH prior in essence has the potential to act as a better regularizer against the noise corruption, and to adapt to a wider range of sparsity levels. Numerical studies using both synthetic data and real-world datasets have demonstrated the improved performance of the proposed method over the Gaussian-gamma CPD in terms of tensor rank learning and factor matrix recovery, especially in the challenging high-rank and/or low-SNR regimes.

Note that the principle followed in this paper is a parametric way to seek flexible sparsity-aware modeling in the context of tensor CPD (specifically, investigating the advanced prior in the Gaussian scale mixture family). In parallel to this path, the pioneering work [38] proposed a non-parametric Bayesian CPD modeling based on a multiplicative gamma process (MGP) prior. Employing a Gibbs sampling method, the inference algorithm of [38] can deal with both continuous and binary tensor data. Due to the decaying effects of the length scales learnt through MGP [38], the inference algorithm is capable of learning low tensor rank, but it has the tendency to under-estimate the tensor rank when the ground-truth rank is high (as validated

in Section V), making it not very flexible in the high-rank regime.

The major contributions of this paper are summarized as follows:

- *Improved tensor rank learning capability:* This paper makes *the first* attempt to tackle the difficulty of tensor rank learning in high-rank and/or low-SNR regimes, under which previous works (especially those relying on Gaussian-gamma prior) *do not* achieve satisfactory performance. Therefore, this paper represents *the first step* to show that tensor rank learning performance can be *further improved* by employing a more advanced sparsity-promoting prior in the Gaussian scale mixture family.
- *Principled new prior design:* This paper exemplifies a principled design of a new parametric prior for tensor CPD, which can benefit future research of other Bayesian tensor decomposition formats. In particular, this paper clarifies why the GH prior should be chosen, and how this prior, which was originally devised for modeling scalars or vectors, can be adopted to model the rank-1 components of a tensor, thereby providing enhanced tensor rank learning capability.
- *Non-trivial inference algorithm derivation:* Given a new prior, the derivation of inference algorithm under VI framework is non-trivial, due to the complicated tensor algebra involved and the newly introduced nonlinearities from the GH prior. The derivations could be used as the reference or basis results for future related VI algorithm derivations.

The remainder of this paper is organized as follows. In Section II, the probabilistic tensor CPD using Gaussian-gamma prior is briefly reviewed. By leveraging the GH prior, we propose a new probabilistic model for tensor CPD in Section III. In Section IV, the framework of variational inference is utilized to derive an inference algorithm with closed-form update equations. In Sections V and VI, numerical results are presented to demonstrate the superior performance of the proposed method. Finally, conclusions and future directions are presented in Section VII.

Notation: Boldface lowercase and uppercase letters will be used for vectors and matrices, respectively. Tensors are written as calligraphic letters. $\mathbb{E}[\cdot]$ denotes the expectation of its argument. Superscript T denotes transpose, and the operator $\text{Tr}(\mathbf{A})$ denotes the trace of matrix \mathbf{A} . $\|\cdot\|_F$ represents the Frobenius norm of the argument. $\mathcal{N}(\mathbf{x}|\mathbf{u}, \mathbf{R})$ stands for the probability density function (pdf) of a Gaussian vector \mathbf{x} with mean \mathbf{u} and covariance matrix \mathbf{R} . The $N \times N$ diagonal matrix with diagonal elements a_1 through a_N is represented as $\text{diag}\{a_1, a_2, \dots, a_N\}$, while \mathbf{I}_M represents the $M \times M$ identity matrix. The $(i, j)^{\text{th}}$ element, the i^{th} row, and the j^{th} column of a matrix \mathbf{A} are represented by $\mathbf{A}_{i,j}$, $\mathbf{A}_{i,:}$, and $\mathbf{A}_{:,j}$, respectively. For easy reference, a table of symbol notations is also given in Appendix A.

II. REVIEW OF THE GAUSSIAN-GAMMA MODEL FOR CPD

In tensor CPD, as illustrated in Fig. 1, the l^{th} columns in all the factor matrices ($\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N$) constitute the building block of the model. Given an upper bound value L of the tensor rank, for each factor matrix, since there are $L - R$ columns being all zero, sparsity-promoting priors should be imposed on the

columns of each factor matrix to encode the information of over-parameterization.² In the pioneering works [22], [28], assuming statistical independence among the columns in $\{\mathbf{U}_{:,l}^{(n)}, \forall n, l\}$, a Gaussian-gamma prior was utilized to model them as

$$\begin{aligned} p(\{\mathbf{U}^{(n)}\}_{n=1}^N | \{\gamma_l\}_{l=1}^L) &= \prod_{l=1}^L p(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N | \gamma_l) \\ &= \prod_{l=1}^L \prod_{n=1}^N \mathcal{N}(\mathbf{U}_{:,l}^{(n)} | \mathbf{0}_{J_n \times 1}, \gamma_l^{-1} \mathbf{I}_{J_n}), \end{aligned} \quad (2)$$

$$\begin{aligned} p(\{\gamma_l\}_{l=1}^L | \{c_l^0, d_l^0\}_{l=1}^L) &= \prod_{l=1}^L p(\gamma_l | c_l^0, d_l^0) \\ &= \prod_{l=1}^L \text{gamma}(\gamma_l | c_l^0, d_l^0), \end{aligned} \quad (3)$$

where γ_l is the precision (i.e., the inverse of variance) of the l^{th} columns $\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N$, and $\{c_l^0, d_l^0\}$ are pre-determined hyper-parameters.

To see the sparsity-promoting property of the above Gaussian-gamma prior, we marginalize the precisions $\{\gamma_l\}_{l=1}^L$ to obtain the marginal probability density function (pdf) $p(\{\mathbf{U}^{(n)}\}_{n=1}^N)$ as follows:

$$\begin{aligned} p(\{\mathbf{U}^{(n)}\}_{n=1}^N) &= \prod_{l=1}^L p(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N) \\ &= \prod_{l=1}^L \int p(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N | \gamma_l) p(\gamma_l | c_l^0, d_l^0) d\gamma_l \\ &= \prod_{l=1}^L \left(\frac{1}{\pi}\right)^{\frac{\sum_{n=1}^N J_n}{2}} \frac{\Gamma(c_l^0 + \sum_{n=1}^N \frac{J_n}{2})}{2d_l^0 - c_l^0 \Gamma(c_l^0)} \\ &\quad \times \left(2d_l^0 + \|\text{vec}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N)\|_2^2\right)^{-c_l^0 - \sum_{n=1}^N \frac{J_n}{2}}, \end{aligned} \quad (4)$$

where $\Gamma(\cdot)$ denotes the gamma function and $\text{vec}(\cdot)$ denotes the vectorization³ of its argument. Equation (4) characterizes a multivariate student's t distribution with hyper-parameters $\{c_l^0, d_l^0\}_{l=1}^L$. To get insights from this marginal distribution, we illustrate its univariate case in Fig. 3 with different hyper-parameters. It is clear that each student's t pdf is strongly peaked at zero and with heavy tails. The prior with such features is known as sparsity-promoting prior, since the peak at zeros will inform the learning process to look for values around “zeros” while the heavy tails still allow the learning process to obtain components with large values [53].

The probabilistic CPD model is completed by specifying the likelihood function of \mathcal{Y} :

$$p(\mathcal{Y} | \{\mathbf{U}^{(n)}\}_{n=1}^N, \beta)$$

²The sparsity level of over-parameterized CPD model can be measured by $\frac{L-R}{L}$.

³The operation $\text{vec}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N)$ simply stacks all these columns into a long vector, i.e., $\text{vec}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N) = [\mathbf{U}_{:,l}^{(1)}; \mathbf{U}_{:,l}^{(2)}; \dots; \mathbf{U}_{:,l}^{(N)}] \in \mathbb{R}^{Z \times 1}$, with $Z = \sum_{n=1}^N J_n$.




Fig. 3. Univariate marginal probability density function in (4) with different values of hyper-parameters.




Fig. 4. Probabilistic CPD model with Gaussian-gamma prior.

$$\propto \exp\left(-\frac{\beta}{2} \|\mathcal{Y} - \llbracket \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)} \rrbracket\|_F^2\right). \quad (5)$$

Equation (5) assumes that the signal tensor $\llbracket \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)} \rrbracket$ is corrupted by AWGN tensor \mathcal{W} with each element having power β^{-1} . This is consistent with the least-squares (LS) problem in (1) if the AWGN power β^{-1} is known. However, in Bayesian modeling, β is modeled as another random variable. Since we have no prior information about the noise power, a non-informative prior $p(\beta) = \text{gamma}(\beta | \epsilon, \epsilon)$ with a very small ϵ (e.g., 10^{-6}) is usually employed.

By using the introduced prior distributions and likelihood function, a probabilistic model for tensor CPD was constructed, as illustrated in Fig. 4. Based on this model, a VI based algorithm was derived in [22] that can automatically drive most of the columns in each factor matrix to zero, by which the tensor rank is revealed. Inspired by the vanilla probabilistic CPD using the Gaussian-gamma prior, other structured and large-scale tensor CPDs with automatic tensor rank learning were further developed [23], [32]–[37] in recent years. The success of the previous works on automatic tensor rank learning [22], [23], [32]–[37] comes from the adoption of the sparsity-promoting Gaussian-gamma prior, while their performances are also limited by the rigid central and tail behaviors in modeling different levels of the sparsity. More specifically, Gaussian-gamma prior is effective only if the sparsity pattern of data matches well with that of Gaussian-gamma model. For example, when the tensor rank R is low, previous results have shown that a relatively

large upper bound value L (e.g., the maximal value of tensor dimensions [22], [23]) can give accurate tensor rank estimation. However, for a high tensor rank R , the upper bound value L selected for the low-rank case would be too small to render a sparsity pattern of columns, and thus it leads to performance degradation. Even though we can increase the value of L to large numbers, as it will be shown in Fig. 8(b)–(c), tensor rank learning accuracy using Gaussian-gamma prior is still not satisfactory, showing its lack of flexibility to adapt to different sparsity levels. Therefore, to further enhance the tensor rank learning capability, we explore the use of sparsity-promoting priors with more flexible central and tail behaviors.

III. NOVEL PROBABILISTIC MODELING: WHEN TENSOR CPD MEETS GENERALIZED HYPERBOLIC DISTRIBUTION

In particular, we focus on the GH prior, since it not only includes the Gaussian-gamma prior as a special case, but also it can be treated as the generalization of other widely-used sparsity-promoting distributions including the Laplacian distribution, normal-inverse chi-squared distribution, normal-inverse gamma distribution, variance-gamma distribution and McKay's Bessel distribution [46]. Therefore, it is expected that the functional flexibility of GH prior could lead to more flexibility in modeling different sparsity levels and thus more accurate learning for tensor rank.

Recall that the model building block is the l^{th} column group $\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N$. With the GH prior on each column group, we have a new prior distribution for factor matrices:

$$\begin{aligned} p(\{\mathbf{U}^{(n)}\}_{n=1}^N) &= \prod_{l=1}^L \text{GH}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N | a_l^0, b_l^0, \lambda_l^0) \\ &= \prod_{l=1}^L \frac{(a_l^0)^{\sum_{n=1}^N J_n} (b_l^0)^{-\frac{\lambda_l^0}{2}}}{(2\pi)^{\sum_{n=1}^N J_n} K_{\lambda_l^0}(\sqrt{a_l^0 b_l^0})} \\ &\quad \times \frac{K_{\lambda_l^0 - \sum_{n=1}^N J_n} \left(\sqrt{a_l^0 (b_l^0 + \|\text{vec}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N)\|_2^2)} \right)}{(b_l^0 + \|\text{vec}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N)\|_2^2)}, \end{aligned} \quad (6)$$

where $K(\cdot)$ is the modified Bessel function of the second kind, and $\text{GH}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N | a_l^0, b_l^0, \lambda_l^0)$ denotes the GH prior on the l^{th} column group $\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N$, in which the hyper-parameters $\{a_l^0, b_l^0, \lambda_l^0\}$ control the shape of the distribution. By setting $\{a_l^0, b_l^0, \lambda_l^0\}$ to specific values, the GH prior (6) reduces to other prevalent sparsity-promoting priors. Details on how GH prior reduces to student-t and Laplacian distributions are given in Appendix B.

To visualize the GH distribution and its special cases, the univariate GH pdfs with different hyper-parameters are illustrated in Fig. 5. It can be observed that the blue line is with a similar shape to those of the student-t distributions in Fig. 3, while the orange one resembles the shapes of Laplacian distributions [46], [47]. For other lines, they exhibit a wide range of the central and tail behaviors of the pdfs. This reveals the great functional flexibility of the GH prior in modeling different levels of sparsity.




Fig. 5. Univariate marginal probability density function in (6) with different values of hyper-parameters.




Fig. 6. Hierarchical construction of GH distribution.

On the other hand, the GH prior (6) can be expressed as a Gaussian scale mixture formulation [46], [47]:

$$\begin{aligned} p(\{\mathbf{U}^{(n)}\}_{n=1}^N) &= \prod_{l=1}^L \text{GH}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N | a_l^0, b_l^0, \lambda_l^0) \\ &= \prod_{l=1}^L \int \mathcal{N}(\text{vec}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N) | \mathbf{0}_{\sum_{n=1}^N J_n \times 1}, z_l \mathbf{I}_{\sum_{n=1}^N J_n}) \\ &\quad \times \text{GIG}(z_l | a_l^0, b_l^0, \lambda_l^0) dz_l, \end{aligned} \quad (7)$$

where z_l denotes the variance of the Gaussian distribution, and $\text{GIG}(z_l | a_l^0, b_l^0, \lambda_l^0)$ represents the generalized inverse Gaussian (GIG) pdf:

$$\begin{aligned} \text{GIG}(z_l | a_l^0, b_l^0, \lambda_l^0) &= \frac{\left(\frac{a_l^0}{b_l^0}\right)^{\frac{\lambda_l^0}{2}}}{2K_{\lambda_l^0}(\sqrt{a_l^0 b_l^0})} z_l^{\lambda_l^0 - 1} \exp\left(-\frac{1}{2}(a_l^0 z_l + b_l^0 z_l^{-1})\right). \end{aligned} \quad (8)$$

This Gaussian scale mixture formulation suggests that each GH distribution $\text{GH}(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N | a_l^0, b_l^0, \lambda_l^0)$ can be regarded as an infinite mixture of Gaussians with the mixing distribution being a GIG distribution. Besides revealing its inherent structure, the formulation (7) allows a hierarchical construction of each GH prior by introducing the latent variable z_l , as illustrated in Fig. 6. This gives us the following important conjugacy property [46].




Fig. 7. The probabilistic tensor CPD model with GH prior.

Property 1: For probability density functions (pdfs)

$$p\left(\left\{\mathbf{U}_{:,l}^{(n)}\right\}_{n=1}^N \mid z_l\right) = \mathcal{N}\left(\text{vec}\left(\left\{\mathbf{U}_{:,l}^{(n)}\right\}_{n=1}^N\right) \mid \mathbf{0}_{\sum_{n=1}^N J_n \times 1}, z_l \mathbf{I}_{\sum_{n=1}^N J_n}\right), \quad (9)$$

$$p(z_l) = \text{GIG}(z_l \mid a_l^0, b_l^0, \lambda_l^0), \quad (10)$$

pdf $p(z_l)$ is conjugate⁴ to $p(\{\mathbf{U}_{:,l}^{(n)}\}_{n=1}^N \mid z_l)$.

As will be seen later, the conjugacy property greatly facilitates the derivation of the Bayesian inference algorithm.

Finally, together with the likelihood function in (5), we propose a novel probabilistic model for tensor CPD using the hierarchical construction of the GH prior, as shown in Fig. 7. Denoting the model parameter set $\Theta = \{\{\mathbf{U}^{(n)}\}_{n=1}^N, \{z_l\}_{l=1}^L, \beta\}$, the proposed probabilistic tensor CPD model can be fully described by the joint pdf $p(\mathcal{Y}, \Theta)$ as

$$\begin{aligned} p(\mathcal{Y}, \Theta) &= p(\mathcal{Y} \mid \{\mathbf{U}^{(n)}\}_{n=1}^N, \beta) p(\{\mathbf{U}^{(n)}\}_{n=1}^N \mid \{z_l\}_{l=1}^L) \\ &\quad \times p(\{z_l\}_{l=1}^L) p(\beta) \\ &\propto \exp\left\{\frac{\prod_{n=1}^N J_n}{2} \ln \beta - \frac{\beta}{2} \|\mathcal{Y} - \llbracket \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)} \rrbracket\|_F^2\right. \\ &\quad \left. + \sum_{n=1}^N \left[\frac{J_n}{2} \sum_{l=1}^L \ln z_l^{-1} - \frac{1}{2} \text{Tr}\left(\mathbf{U}^{(n)} \mathbf{Z}^{-1} \mathbf{U}^{(n)T}\right) \right] \right. \\ &\quad \left. + \sum_{l=1}^L \left[\frac{\lambda_l^0}{2} \ln \frac{a_l^0}{b_l^0} - \ln \left[2K_{\lambda_l^0} \left(\sqrt{a_l^0 b_l^0} \right) \right] + (\lambda_l^0 - 1) \ln z_l \right. \right. \\ &\quad \left. \left. - \frac{1}{2} (a_l^0 z_l + b_l^0 z_l^{-1}) \right] + (\epsilon - 1) \ln \beta - \epsilon \beta \right\}, \quad (11) \end{aligned}$$

where $\mathbf{Z} = \text{diag}\{z_1, z_2, \dots, z_L\}$.

IV. INFERENCE ALGORITHM

A. General Philosophy of the Variational Inference

Given the probabilistic model $p(\mathcal{Y}, \Theta)$, the next task is to learn the model parameters in Θ from the tensor data \mathcal{Y} , in which the

⁴In Bayesian theory, a probability density function (pdf) $p(x)$ is said to be conjugate to a conditional pdf $p(y|x)$ if the resulting posterior pdf $p(x|y)$ is in the same distribution family as $p(x)$.

posterior distribution $p(\Theta \mid \mathcal{Y})$ is to be sought. However, for such a complicated probabilistic model (11), the multiple integrations in computing the posterior distribution $p(\Theta \mid \mathcal{Y})$ is not tractable. Fortunately, this challenge is not new, and similar obstacles have been faced in inferring other complicated Bayesian machine learning models such as Bayesian neural networks [54], [55], Bayesian structured matrix factorization [56], latent dirichlet allocation [57], and Gaussian mixture model [58]. It has been widely agreed that variational inference (VI), due to its efficiency in computations and the theoretical guarantee of convergence, is the major driving force for inferring complicated probabilistic models [52]. Rather than manipulating a huge number of samples from the probabilistic model, VI recasts the originally intractable multiple integration problem into the following functional optimization problem:

$$\begin{aligned} &\min_{Q(\Theta)} \text{KL}(Q(\Theta) \parallel p(\Theta \mid \mathcal{Y})) \\ &\triangleq -\mathbb{E}_{Q(\Theta)} \left\{ \ln \frac{p(\Theta \mid \mathcal{Y})}{Q(\Theta)} \right\} \\ &\text{s.t. } Q(\Theta) \in \mathcal{F}, \end{aligned} \quad (12)$$

where $\text{KL}(\cdot \parallel \cdot)$ denotes the Kullback-Leibler (KL) divergence between two arguments, and \mathcal{F} is a pre-selected family of pdfs. Its philosophy is to seek a tractable variational pdf $Q(\Theta)$ in \mathcal{F} that is the closest to the true posterior distribution $p(\Theta \mid \mathcal{Y})$ in terms of the KL divergence. Therefore, the art is to determine the family \mathcal{F} to balance the tractability of the algorithm and the accuracy of the posterior distribution learning. In this paper, we adopt the mean-field family due to its prevalence in Bayesian tensor research [22], [23], [32]–[37]. Other advanced choices could be found in [52].

Using the mean-field family, which restricts $Q(\Theta) = \prod_{k=1}^K Q(\Theta_k)$ where Θ is partitioned into mutually disjoint non-empty subsets Θ_k (i.e., Θ_k is a part of Θ with $\cup_{k=1}^K \Theta_k = \Theta$ and $\cap_{k=1}^K \Theta_k = \text{Oslash}$), the KL divergence minimization problem (12) becomes

$$\min_{\{Q(\Theta_k)\}_{k=1}^K} -\mathbb{E}_{\{Q(\Theta_k)\}_{k=1}^K} \left\{ \ln \frac{p(\Theta \mid \mathcal{Y})}{\prod_{k=1}^K Q(\Theta_k)} \right\}. \quad (13)$$

The factorable structure in (13) inspires the idea of block minimization in optimization theory [18]. In particular, after fixing variational pdfs $\{Q(\Theta_j)\}_{j \neq k}$ other than $Q(\Theta_k)$, the remaining problem is

$$\min_{Q(\Theta_k)} \int Q(\Theta_k) (-\mathbb{E}_{\prod_{j \neq k} Q(\Theta_j)} [\ln p(\Theta, \mathcal{Y})] + \ln Q(\Theta_k)) d\Theta_k, \quad (14)$$

and it has been shown that the optimal solution is [50], [53]:

$$Q^*(\Theta_k) = \frac{\exp\left(\mathbb{E}_{\prod_{j \neq k} Q(\Theta_j)} [\ln p(\Theta, \mathcal{Y})]\right)}{\int \exp\left(\mathbb{E}_{\prod_{j \neq k} Q(\Theta_j)} [\ln p(\Theta, \mathcal{Y})]\right) d\Theta_k}. \quad (15)$$

More discussions on the mean-field VI are provided in Appendix M.

B. Deriving Optimal Variational Pdfs

The optimal variational pdfs $\{Q^*(\Theta_k)\}_{k=1}^K$ can be obtained by substituting (11) into (15). Although straightforward as it may

TABLE I
OPTIMAL VARIATIONAL DENSITY FUNCTIONS

Variational pdfs	Remarks
$Q^*(\mathbf{U}^{(k)}) = \mathcal{MN}(\mathbf{U}^{(k)} \mathbf{M}^{(k)}, \mathbf{I}_{J_n}, \boldsymbol{\Sigma}^{(k)}), \forall k$	Matrix normal distribution with mean $\mathbf{M}^{(k)}$ and covariance matrix $\boldsymbol{\Sigma}^{(k)}$ given in (16) and (17), respectively.
$Q^*(z_l) = \text{GIG}(z_l a_l, b_l, \lambda_l), \forall l$	Generalized inverse Gaussian distribution with parameters $\{a_l, b_l, \lambda_l\}$ given in (18)–(20).
$Q^*(\beta) = \text{gamma}(\beta e, f)$	Gamma distribution with shape e and rate f given in (21), (22).

seem, the involvement of tensor algebras in (11) and the multiple integrations in the denominator of (15) make the derivation a challenge. On the other hand, since the proposed probabilistic model employs the GH prior, and is different from previous works using Gaussian-gamma prior [22], [23], [32]–[37], each optimal variational pdf $Q^*(\Theta_k)$ needs to be derived from first principles. To keep the main body of this paper concise, the lengthy derivations are moved to Appendix C, and we only present the optimal variational pdfs in Table I at the top of this page.

In particular, the optimal variational pdf $Q^*(\mathbf{U}^{(k)})$ was derived to be a matrix normal distribution [59] $\mathcal{MN}(\mathbf{U}^{(k)} | \mathbf{M}^{(k)}, \mathbf{I}_{J_n}, \boldsymbol{\Sigma}^{(k)})$ with the covariance matrix

$$\boldsymbol{\Sigma}^{(k)} = \left[\mathbb{E}[\beta] \mathbb{E} \left[\left(\underset{n=1, n \neq k}{\odot}^N \mathbf{U}^{(n)} \right)^T \times \left(\underset{n=1, n \neq k}{\odot}^N \mathbf{U}^{(n)} \right) \right] + \mathbb{E}[\mathbf{Z}^{-1}] \right]^{-1}, \quad (16)$$

and mean matrix

$$\mathbf{M}^{(k)} = \mathcal{Y}(k) \mathbb{E}[\beta] \left(\underset{n=1, n \neq k}{\odot}^N \mathbb{E}[\mathbf{U}^{(n)}] \right) \boldsymbol{\Sigma}^{(k)}. \quad (17)$$

In (16) and (17), $\mathcal{Y}(k)$ is a matrix obtained by unfolding the tensor \mathcal{Y} along its k^{th} dimension, and the multiple Khatri-Rao products $\underset{n=1, n \neq k}{\odot}^N \mathbf{A}^{(n)} = \mathbf{A}^{(N)} \odot \mathbf{A}^{(N-1)} \odot \dots \odot \mathbf{A}^{(k+1)} \odot \mathbf{A}^{(k-1)} \odot \dots \odot \mathbf{A}^{(1)}$. The expectations are taken with respect to the corresponding variational pdfs of the arguments. For the optimal variational pdf $Q(z_l)$, by using the conjugacy result in *Property 1*, it can be derived to be a GIG distribution $\text{GIG}(z_l | a_l, b_l, \lambda_l)$ with parameters

$$a_l = a_l^0, \quad (18)$$

$$b_l = b_l^0 + \sum_{n=1}^N \mathbb{E} \left[\left[\mathbf{U}_{:,l}^{(n)} \right]^T \mathbf{U}_{:,l}^{(n)} \right], \quad (19)$$

$$\lambda_l = \lambda_l^0 - \frac{1}{2} \sum_{n=1}^N J_n. \quad (20)$$

Finally, the optimal variational pdf $Q(\beta)$ was derived to be a gamma distribution $\text{gamma}(\beta | e, f)$ with parameters

$$e = \epsilon + \frac{1}{2} \prod_{n=1}^N J_n, \quad (21)$$

$$f = \epsilon + \frac{1}{2} \mathbb{E} \left[\left\| \mathcal{Y} - \llbracket \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket \right\|_F^2 \right]. \quad (22)$$

In (16)–(22), there are several expectations to be computed. They can be obtained either from the statistic literatures [59] or similar results in related works [22], [23], [32]–[37]. For easy reference, we listed the expectation results needed for (16)–(22) in Table II, where $\underset{n=1, n \neq k}{\otimes}^N \mathbf{A}^{(n)} = \mathbf{A}^{(N)} \otimes \mathbf{A}^{(N-1)} \otimes \dots \otimes \mathbf{A}^{(k+1)} \otimes \mathbf{A}^{(k-1)} \otimes \dots \otimes \mathbf{A}^{(1)}$ is the multiple Hadamard products.

C. Setting the Hyper-Parameters

From Table I, it can be found that the variational pdf $Q(\mathbf{U}^{(k)})$ and $\{Q(z_l)\}_{l=1}^L$ forms a more complicated Gaussian-GIG pdf pair than that in *Property 1*. Therefore, the shape of the variational pdf $Q(\mathbf{U}^{(k)})$, which determines both the factor matrix recovery and tensor rank learning, is affected by the variational pdf $\{Q(z_l)\}_{l=1}^L$. For each $Q(z_l)$, as seen in (18)–(20), its shape relies on the pre-selected hyper-parameters $\{a_l^0, b_l^0, \lambda_l^0\}$. In practice, we usually have no prior knowledge about the sparsity level before assessing the data, and a widely adopted approach is to make the prior non-informative.

In previous works using Gaussian-gamma prior [22], [23], [32]–[37], hyper-parameters are set equal to very small values in order to approach a non-informative prior. Although nearly zero hyper-parameters lead to an improper prior, the derived variational pdf is still proper since these parameters are updated using information from observations [22], [23], [32]–[37]. Therefore, in these works, the strategy of using non-informative prior is valid. On the other hand, for the employed GH prior, non-informative prior requires $\{a_l^0, b_l^0, \lambda_l^0\}$ all go to zero, which however would lead to an improper variational pdf $Q(z_l)$, since its parameter $a_l = a_l^0$ is fixed (as seen in (18)). This makes the expectation computation $\mathbb{E}[z_l]$ in Table II problematic. To tackle this issue, another viable approach is to optimize these hyper-parameters $\{a_l^0, b_l^0, \lambda_l^0\}$ so that they can be adapted during the procedure of model learning. However, as seen in (11), these three parameters are coupled together via the nonlinear modified Bessel function, and thus optimizing them jointly is prohibitively difficult. Therefore, in this paper, we propose to only optimize the most critical one, i.e., a_l^0 , since it directly determines the shape of $Q(z_l)$ and will not be updated in the learning procedure. For the other two parameters $\{b_l^0, \lambda_l^0\}$, as seen in (19) and (20), since they are updated with model learning results or tensor dimension, according to the Bayesian theory [49], their effect on the posterior distribution would become negligible when the observation tensor is large enough. This justifies the optimization of a_l^0 while not $\{b_l^0, \lambda_l^0\}$.

TABLE II
COMPUTATION RESULTS OF EXPECTATIONS

Expectations	Computation Results
$\mathbb{E} [\mathbf{U}^{(k)}], \forall k$	$\mathbf{M}^{(k)}, \forall k$
$\mathbb{E} [z_l], \forall l$	$\left(\frac{b_l}{a_l}\right)^{\frac{1}{2}} \frac{K_{\lambda_l+1}(\sqrt{a_l b_l})}{K_{\lambda_l}(\sqrt{a_l b_l})}$
$\mathbb{E} [z_l^{-1}], \forall l$	$\left(\frac{b_l}{a_l}\right)^{-\frac{1}{2}} \frac{K_{\lambda_l-1}(\sqrt{a_l b_l})}{K_{\lambda_l}(\sqrt{a_l b_l})}$
$\mathbb{E} [\beta]$	$\frac{e}{f}$
$\mathbb{E} \left[[\mathbf{U}_{:,l}^{(n)}]^T \mathbf{U}_{:,l}^{(n)} \right]$	$[\mathbf{M}_{:,l}^{(n)}]^T \mathbf{M}_{:,l}^{(n)} + J_n \boldsymbol{\Sigma}_{:,l}^{(n)}$
$\mathbb{E} \left[\left(\bigotimes_{n=1, n \neq k}^N \mathbf{U}^{(n)} \right)^T \left(\bigotimes_{n=1, n \neq k}^N \mathbf{U}^{(n)} \right) \right]$	$\bigotimes_{n=1, n \neq k}^N \left[[\mathbf{M}^{(n)}]^T \mathbf{M}^{(n)} + J_n \boldsymbol{\Sigma}^{(n)} \right]$
$\mathbb{E} \left[\ \mathcal{Y} - \llbracket \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket\ _F^2 \right]$	$\ \mathcal{Y}\ _F^2 + \text{Tr} \left(\bigotimes_{n=1}^N \left[\mathbf{M}^{(n)T} \mathbf{M}^{(n)} + J_n \boldsymbol{\Sigma}^{(n)} \right] \right) - 2\text{Tr} \left(\mathcal{Y}(1) \left(\bigotimes_{n=2}^N \mathbf{M}^{(n)} \right) \mathbf{M}^{(1)T} \right)$

For optimizing a_l^0 , following related works [50], [61], we introduce a conjugate hyper-prior $p(a_l^0) = \text{gamma}(a_l^0 | \kappa_{a_1}, \kappa_{a_2})$ to ensure the positiveness of a_l^0 during the optimization. To bypass the nonlinearity from the modified Bessel function, we set $b_l^0 \rightarrow 0$ so that $K_{\lambda_l^0}(\sqrt{a_l^0 b_l^0})$ becomes a constant. In the framework of VI, after fixing other variables, it has been derived in Appendix C that the hyper-parameter a_l^0 is updated via

$$a_l^0 = \frac{\kappa_{a_1} + \frac{\lambda_l^0}{2} - 1}{\kappa_{a_2} + \frac{\mathbb{E}[z_l]}{2}}. \quad (23)$$

Notice that it requires $\kappa_{a_1} > 1 - \lambda_l^0/2$ and $\kappa_{a_2} \geq 0$ to ensure the positiveness of a_l^0 .

D. Algorithm Summary and Insights

From (16)–(23), it can be seen that the statistics of each variational pdf rely on other variational pdfs. Therefore, they need to be updated in an alternating fashion, giving rise to an iterative algorithm summarized in Algorithm 1. To gain more insights from the proposed algorithm, discussions on its convergence property, computational complexity, and automatic tensor rank learning are presented in the following.

1) *Convergence Property*: While the proposed algorithm is devised with the novel probabilistic tensor CPD model using GH prior, its derivation follows the mean-field VI framework [50]–[53]. In particular, in each iteration, after fixing other variational pdfs, the problem that optimizes a single variational pdf has been shown to be convex and has a unique solution [50]–[53]. By treating each update step in mean-field VI as a block coordinate descent (BCD) step over the functional space, the limit point generated by the VI algorithm is at least a stationary point of the KL divergence [50]–[53].

2) *Automatic Tensor Rank Learning*: During the iterations, the mean of parameter z_l^{-1} (denoted by $m[z_l^{-1}]$) will be learnt using the updated parameters of other variational pdfs as seen in (26)–(29). Due to the sparsity-promoting nature of the GH prior, some of $m[z_l^{-1}]$ will take very large values, e.g., in the order of 10^6 . Since the inverse of $\{m[z_l^{-1}]\}_{l=1}^L$ contribute to the covariance matrix of each factor matrix (as seen in (24)) and then

rescale the columns in each factor matrix (as seen in (25)), a very large $m[z_l^{-1}]$ will shrink the l^{th} column of each factor matrix to all zero. Then, by enumerating how many non-zero columns in each factor matrix, the tensor rank can be automatically learnt.

In practice, to accelerate the learning algorithm, on-the-fly pruning is widely employed in Bayesian tensor research. In particular, in each iteration, if some of the columns in each factor matrix are found to be indistinguishable from all zeros, it indicates that these columns play no role in interpreting the data, and thus they can be safely pruned. This pruning procedure will not affect the convergence behavior of the algorithm, since each pruning is equivalent to restarting the algorithm for a reduced probabilistic model with the current variational pdfs acting as the initializations. More discussions on the pruning issue are provided in Appendix N.

3) *Computational Complexity*: For the proposed algorithm, in each iteration, the computational complexity is dominated by updating the factor matrices, costing $O(N \prod_{n=1}^N J_n L^2 + L^3 \sum_{n=1}^N J_n)$. Therefore, the computational complexity of the proposed algorithm is $O(q(N \prod_{n=1}^N J_n L^2 + L^3 \sum_{n=1}^N J_n))$ where q is the iteration number at convergence. The complexity is comparable to that of the inference algorithm using Gaussian-gamma prior [22].

V. NUMERICAL RESULTS ON SYNTHETIC DATA

In this and the next section, extensive numerical results are presented to assess the performance of the proposed algorithm using synthetic data and real-world datasets, respectively. For the proposed algorithm, its initialization follows the suggestions given in Appendix D unless stated otherwise. All experiments were conducted in Matlab R2015b with an Intel Core i7 CPU at 2.2 GHz.

A. Simulation Setup

We consider three-dimensional tensors $\mathcal{X} = \llbracket \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \mathbf{A}^{(3)} \rrbracket \in \mathbb{R}^{30 \times 30 \times 30}$ with different tensor ranks. Each element in the factor matrices $\{\mathbf{A}^{(n)}\}_{n=1}^3$ is independently drawn from a zero-mean Gaussian distribution with unit

Algorithm 1: Probabilistic Tensor CPD Using The GH Prior.

Initializations: Choose $L > R$ and initial values $\{\mathbf{M}^{(n)0}, [\boldsymbol{\Sigma}^{(n)0}]_{n=1}^N, \{m[z_l^{-1}]^0, a_l^0, b_l^0, \lambda_l^0\}_{l=1}^L, e^0, f^0$.
Choose $\kappa_{a_1} > -\lambda_l^0/2$ and $\kappa_{a_2} \geq 0$.

Iterations:

For the iteration $t + 1$ ($t \geq 0$),

For $k = 1, \dots, N$, update the parameters of $Q(\mathbf{U}^{(k)})^{t+1}$:

$$\begin{aligned} [\boldsymbol{\Sigma}^{(k)}]^{t+1} = & \left[\frac{c^t}{dt} \underset{n=1, n \neq k}{\otimes}^N \left(\left([\mathbf{M}^{(n)}] \right)^s \right)^T [\mathbf{M}^{(n)}] \right. \\ & \left. + J_n [\boldsymbol{\Sigma}^{(n)}] \right]^s + \text{diag} \{m[z_1^{-1}]^t, m[z_2^{-1}]^t, \dots, m[z_L^{-1}]^t\}^{-1}, \end{aligned} \quad (24)$$

$$[\mathbf{M}^{(k)}]^{t+1} = \mathcal{Y}^{(k)} \frac{c^t}{dt} \left(\underset{n=1, n \neq k}{\otimes}^N [\mathbf{M}^{(n)}] \right)^s [\boldsymbol{\Sigma}^{(n)}]^{t+1}, \quad (25)$$

where s denotes the most recent update index, i.e., $s = t + 1$ when $n < k$, and $s = t$ otherwise.

Update the parameters of $Q(z_l)^{t+1}$:

$$a_l^{t+1} = [a_l^0]^t, \quad (26)$$

$$\begin{aligned} b_l^{t+1} = & b_l^0 + \sum_{n=1}^N \left[\left([\mathbf{M}^{(n)}] \right)^{t+1} \right]^T [\mathbf{M}^{(n)}]^{t+1} \\ & + J_n [\boldsymbol{\Sigma}_{l,l}^{(n)}]^{t+1}, \end{aligned} \quad (27)$$

$$[\lambda_l]^{t+1} = \lambda_l^0 - \frac{1}{2} \sum_{n=1}^N J_n, \quad (28)$$

$$m[z_l^{-1}]^{t+1} = \left(\frac{b_l^{t+1}}{a_l^{t+1}} \right)^{-\frac{1}{2}} \frac{K_{[\lambda_l]^{t+1}-1} \left(\sqrt{a_l^{t+1} b_l^{t+1}} \right)}{K_{[\lambda_l]^{t+1}} \left(\sqrt{a_l^{t+1} b_l^{t+1}} \right)}, \quad (29)$$

$$m[z_l]^{t+1} = \left(\frac{b_l^{t+1}}{a_l^{t+1}} \right)^{\frac{1}{2}} \frac{K_{[\lambda_l]^{t+1}+1} \left(\sqrt{a_l^{t+1} b_l^{t+1}} \right)}{K_{[\lambda_l]^{t+1}} \left(\sqrt{a_l^{t+1} b_l^{t+1}} \right)}. \quad (30)$$

Update the parameters of $Q(\beta)^{t+1}$:

$$e^{t+1} = \epsilon + \frac{\prod_{n=1}^N J_n}{2}, \quad (31)$$

$$f^{t+1} = \epsilon + \frac{f^{t+1}}{2}, \quad (32)$$

where f^{t+1} is computed using the result in the last row of

Table II with $\{\mathbf{M}^{(n)}, \boldsymbol{\Sigma}^{(n)}\}$ being replaced by $\{[\mathbf{M}^{(n)}]^{t+1}, [\boldsymbol{\Sigma}^{(n)}]^{t+1}\}, \forall n$.

Update the hyper-parameter $[a_l^0]^{t+1}$:

$$[a_l^0]^{t+1} = \frac{\kappa_{a_1} + \frac{\lambda_l^0}{2} - 1}{\kappa_{a_2} + \frac{m[z_l]^{t+1}}{2}}. \quad (33)$$

Until Convergence

power. The observation model $\mathcal{Y} = \mathcal{X} + \mathcal{W}$, where each element of the noise tensor \mathcal{W} is independently drawn from a zero-mean Gaussian distribution with variance σ_w^2 . This data generation process follows that of [22]. The SNR is defined as $10 \log_{10}(\text{var}(\mathcal{X})/\sigma_w^2)$ [22], [23], where $\text{var}(\mathcal{X})$ is the variance⁵ of the noise-free tensor \mathcal{X} . All simulation results in this section are obtained by averaging 100 Monte-Carlo runs unless stated otherwise.

B. PCPD-GH versus PCPD-GG

Since the principle of this paper follows a parametric way of Bayesian modeling, we first compare the proposed algorithm using GH prior (labeled as PCPD-GH) with the benchmarking algorithm [22] (labeled as PCPD-GG).

1) *Tensor Rank Learning:* The performance of tensor rank learning is firstly evaluated. We regard the tensors as low-rank tensors when their ranks are smaller than or equal to half of the maximal tensor dimension, i.e., $R \leq \max\{J_n\}_{n=1}^N/2$. Similarly, high-rank tensors are those with $R > \max\{J_n\}_{n=1}^N/2$. In particular, in Fig. 8, we assess the tensor rank learning performances of the two algorithms for low-rank tensors with $R = \{3, 6, 9, 12, 15\}$ and high-rank tensors with $R = \{18, 21, 24, 27\}$ under SNR = 10 dB. In Fig. 8(a), the two algorithms are both with the tensor rank upper bound $\max\{J_n\}_{n=1}^N$. It can be seen that the PCPD-GH algorithm and the PCPD-GG algorithm achieve comparable performances in learning low tensor ranks. More specifically, the PCPD-GH algorithm achieves higher learning accuracies when $R = \{3, 6\}$ while the PCPD-GG method performs better when $R = \{9, 15\}$. However, when tackling high-rank tensors with $R > 15$, as seen in Fig. 8(a), both algorithms with tensor rank upper bound $\max\{J_n\}_{n=1}^N$ fail to work properly. The reason is that the upper bound value $\max\{J_n\}_{n=1}^N$ results in too small sparsity level $\frac{L-R}{L}$ to leverage the power of the sparsity-promoting priors in tensor rank learning. Therefore, the upper bound value should be set larger in case that the tensor rank is high. An immediate choice is $f \times \max\{J_n\}_{n=1}^N$ where $f = 1, 2, 3, \dots$. In Fig. 8(b) and (c), we assess the performances of tensor rank learning for the two methods using the upper bound $2 \max\{J_n\}_{n=1}^N$ and $5 \max\{J_n\}_{n=1}^N$, respectively. It can be seen that the PCPD-GG algorithm is very sensitive to the rank upper bound value, in the sense that its performance deteriorates significantly for low-rank tensors after employing the larger upper bounds. While PCPD-GG has an improved performance for high-rank tensors after adopting a larger upper bound, the chance of getting the correct rank is still very low. In contrast, the performance of the proposed PCPD-GH algorithm is stable for all cases and it achieves nearly 100% accuracies of tensor rank learning in a wide range of scenarios, showing its flexibility in adapting to different levels of sparsity. In Appendix L, further numerical results on the tensor rank learning accuracies versus different sparsity levels are presented, which show the better performance of the proposed algorithm in a wide range of sparsity levels.

To assess the tensor rank learning performance under different SNRs, in Fig. 9, the percentages of accurate tensor rank learning from the two methods are presented. We consider two scenarios: 1) low-rank tensor with $R = 6$ shown in Fig. 9(a) and 2) high-rank tensor with $R = 24$ shown in Fig. 9(b). For the proposed

⁵It means the empirical variance computed by treating all entries of the tensor as independent realizations of a same scalar random variable.




Fig. 8. Performance of tensor rank learning when the rank upper bound is (a) $\max\{J_n\}_{n=1}^N$, (b) $2 \max\{J_n\}_{n=1}^N$ and (c) $5 \max\{J_n\}_{n=1}^N$.

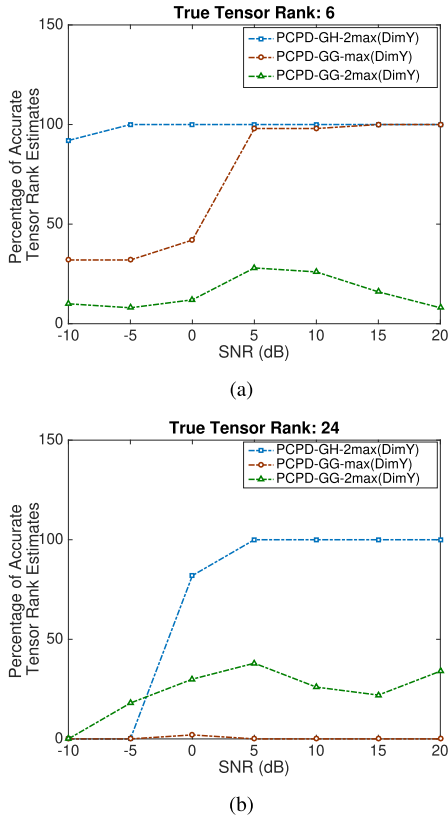


Fig. 9. Performance of tensor rank learning versus different SNRs: (a) low-rank tensors and (b) high-rank tensors.

PCPD-GH algorithm, due to its robustness to different rank upper bounds, $2 \max\{J_n\}_{n=1}^N$ is adopted as the upper bound value (labeled as PCPD-GH-2max(DimY)). For the PCPD-GG algorithm, both the upper bound value $\max\{J_n\}_{n=1}^N$ and $2 \max\{J_n\}_{n=1}^N$ are considered (labeled as PCPD-GG-max(DimY) and PCPD-GH-2max(DimY) respectively). From Fig. 9, it is clear that the performance of the PCPD-GG method, for all cases, highly relies on the choice of the rank upper bound value. In particular, when adopting $2 \max\{J_n\}_{n=1}^N$, its performance in tensor rank learning is not good (i.e., below than 50%) for both the low-rank tensor and the high-rank tensor cases. In contrast, when adopting $\max\{J_n\}_{n=1}^N$, its performance becomes much better for the low-rank cases. In Fig. 9(a), when SNR is larger than 5 dB, the PCPD-GG with upper bound value $\max\{J_n\}_{n=1}^N$ achieves nearly 100% accuracy, which is very

close to the accuracies of the PCPD-GH method. However, when the SNR is smaller than 5 dB, although the PCPD-GH method still achieves nearly 100% accuracies in tensor rank learning, the accuracies of the PCPD-GG method fall below 50%. For the high-rank case, as seen in Fig. 9(b), both the PCPD-GH and the PCPD-GG methods fail to recover the true tensor rank when SNR is smaller than 0 dB. However, when the SNR is larger than 0 dB, the accuracies of the PCPD-GH method are near 100% while those of the PCPD-GG at most achieve about 50% accuracy. Consequently, it can be concluded from Fig. 9 that the proposed PCPD-GH method achieves more stable and accurate tensor rank learning.

In summary, Figs. 8 and 9 show that the proposed method finds the correct tensor rank even if the initial tensor rank is *exceedingly over-estimated*, which is a *practically useful* feature since the rank is unknown in real-life cases.

2) *Insights From Learnt Length Scales*: To clearly show the substantial difference between the GG and GH prior, we compare the two algorithms (PCPD-GG and PCPD-GH) in terms of their learnt length scales. The length scale powers of GG and GH prior are denoted by $\{\gamma_l^{-1}\}_l$ and $\{z_l\}_l$ respectively. To assess the patterns of learnt length scales, we turn off the pruning procedure and let the two algorithms directly output $\{\gamma_l^{-1}\}_l$ and $\{z_l\}_l$ after convergence. Since in the different Monte-Carlo trial, the learnt length scale powers are possibly of different sparsity patterns, averaging them over Monte-Carlo trials is not informative. Instead, we present the learnt values of $\{\gamma_l^{-1}\}_l$ and $\{z_l\}_l$ in a single trial.

In particular, Fig. 10 shows the result for a typical low SNR and low rank case (SNR = -5 dB, R = 6) with rank upper bound being 60. From this figure, it can be seen that the learnt length scales of the two algorithms substantially *differ* from each other, in the sense that the number of learnt length scales (and the associated components) *with non-negligible magnitudes* are different. For example, in Fig. 10, PCPD-GG recovers 7 components with non-negligible magnitudes⁶ (the smallest one has value $4.3 \gg 0$), while PCPD-GH recovers 6 components. Note that the ground-truth rank is 6, and PCPD-GG produces a “ghost” component *with magnitude much larger than zero*. Additional simulation runs, and results of other simulation settings (e.g., high SNR and high rank case: SNR = 5 dB, R = 21) are included in Appendix P, from which similar conclusions can be drawn.

3) *Insights on Noise Precision Learning*: The learning of the noise precision β is crucial for reliable inference, since incorrect

⁶The magnitude of the l -th component is defined as $(\sum_{n=1}^3 [A_{:,l}^{(n)}]^T A_{:,l}^{(n)})^{\frac{1}{2}}$ [22].

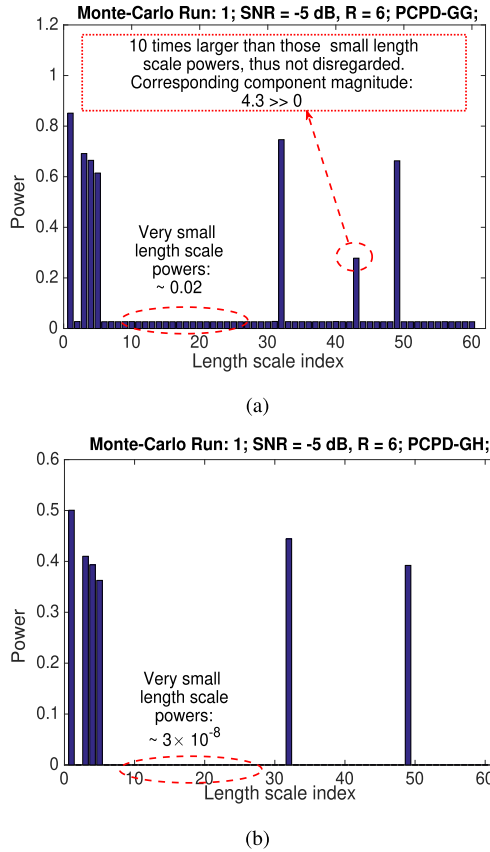


Fig. 10. (a) The powers of learnt length scales (i.e., $\{\gamma_l^{-1}\}_l$) for PCPD-GG; (b) The powers of learnt length scales (i.e., $\{z_l\}_l$) for PCPD-GH. It can be seen that PCPD-GG recovers 7 components with non-negligible magnitudes, while PCPD-GH recovers 6 components. The two algorithms are with the same upper bound value: 60.

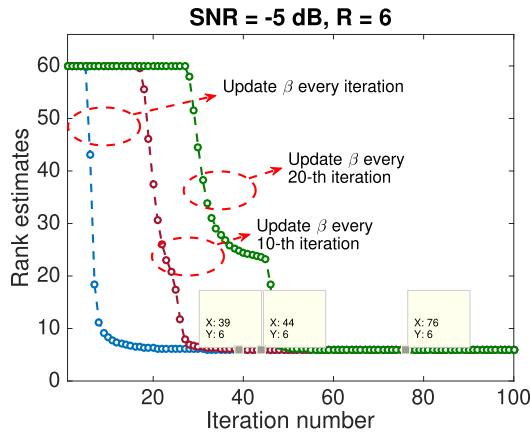


Fig. 11. Tensor rank estimates of PCPD-GH versus iteration number (averaged over 100 Monte-Carlo runs) with different noise precision learning speeds.

estimates will cause over-/under-regularization. To examine how the speed of noise learning affects the tensor rank (sparsity) learning when SNR is low (SNR = -5 dB), we turn on the pruning and present the rank learning results over iterations in three cases: i) Case I: update β every iteration; ii) Case II: update β every 10-th iteration; iii) Case III: update β every 20-th iteration. In Fig. 11, the rank estimates are averaged over

TABLE III
PERFORMANCE OF TENSOR RANK LEARNING UNDER DIFFERENT RANK UPPER BOUND VALUES. ALGORITHM: PCPD-MGP; SNR = 10 dB. THE SIMULATION SETTINGS ARE THE SAME AS THOSE OF FIG. 8

Rank Upper Bound	True Tensor Rank R					
	3	6	9	12	15	18-27
max(Dim Y)	100%	100%	100%	98%	90%	0%
2max(Dim Y)	100%	100%	100%	94%	64%	0%
5max(Dim Y)	100%	100%	100%	26%	24%	0%

100 Monte-Carlo runs. It can be seen that updating the noise precision β at earlier iterations will help will help the learning process to unveil the sparsity pattern more quickly.

Then, we investigate under which scenario slowing the noise precision learning will be helpful. We consider a very low SNR case, that is, SNR = -10 dB, and then evaluate the percentages of accurate rank learning over 100 Monte-Carlo runs. The results are: i) Case I: 76%; ii) Case II: 100%; iii) Case III: 100%. In other words, when the noise power is very large (e.g., SNR = -10 dB), slowing the noise precision learning will make the algorithm more robust to the noises.

Finally, if we fix the noise precision β and do not allow its update, PCPD-GH fails to identify the underlying sparsity pattern (tensor rank), see Figure Q1 in Appendix Q. Particularly, a small value of β (e.g., 0.01) leads to over-regularization, thus causing under-estimation of non-zero components; a large value of β (e.g., 100) causes under-regularization, thus inducing over-estimation of non-zero components. This shows the importance of modeling and updating of noise precision.

4) *Other Performance Metrics*: Additional results and discussions on the run time, tensor recovery root mean square error (RMSE), algorithm performance under factor matrix correlation, convergence behavior of the proposed algorithm in terms of evidence lower bound (ELBO) [51], and hyper-parameter learning of PCPD-GG are included in Appendix E, F, and G, J, R, respectively. The *key messages* of these simulation results are given as follows: i) PCPD-GH generally costs more run time than PCPD-GG; ii) Incorrect estimation of tensor rank degrades the tensor signal recovery; iii) PCPD-GH performs well under factor matrix correlation; iv) PCPD-GH monotonically increases the ELBO; v) The update of hyper-parameter of PCPD-GH does not help too much in improving rank estimation in the low SNR regime.

C. Comparisons With Nonparametric PCPD-MGP

After comparing to the parametric PCPD-GG, further comparisons are performed with the non-parametric Bayesian tensor CPD using MGP prior (labeled as PCPD-MGP).⁷ We implement PCPD-MGP in the variational inference framework. The initializations and hyper-parameters follow those used in [38].

1) *Tensor Rank Learning*: We first assess the performance of tensor rank learning over 100 Monte-Carlo trials, to facilitate the comparisons among the three algorithms (PCPD-MGP, PCPD-GG, PCPD-GH). The simulation settings follow those of Figs. 8 and 9.

In Table III, the percentages of accurate tensor rank estimates of PCPD-MGP under different rank upper bound values are presented. Comparing to Fig. 8, we can draw the following conclusions. i) When the tensor rank is low (e.g., R = 3, 6, 9)

⁷We appreciate Prof. Piyush Rai for sharing the code and data with us.

TABLE IV
PERFORMANCE OF TENSOR RANK LEARNING UNDER DIFFERENT SNRS.
ALGORITHM: PCPD-MGP; UPPER BOUND VALUE: $2\text{MAX}(\text{DIM}Y)$. THE
SIMULATION SETTINGS ARE THE SAME AS THOSE OF FIG. 9

True Tensor Rank	SNR (dB)						
	-10	-5	0	5	10	15	20
R = 6	2%	78%	90%	98%	100%	100%	100%
R = 24	0%	0%	0%	0%	0%	0%	0%

and the SNR is high (e.g., SNR = 10 dB), PCPD-MGP correctly learns the tensor rank over 100 Monte-Carlo trials under different rank upper bound values. Its performance is insensitive to the selection of the rank upper bound values, due to the decaying effects of the learnt length scales [38]. Therefore, in the low-rank and high-SNR scenario, the rank learning performance of PCPD-MGP is comparable to that of PCPD-GH, and much better than PCPD-GG (see Fig. 8). ii) When the tensor rank is high (e.g., R = 18, 21, 24, 27), PCPD-MGP fails to learn the correct tensor rank under different rank upper bound values, since the decaying effects of learnt length scales tend to *under-estimate* the tensor rank. On the contrary, the PCPD-GH method always accurately estimates the high tensor ranks (see Fig. 8).

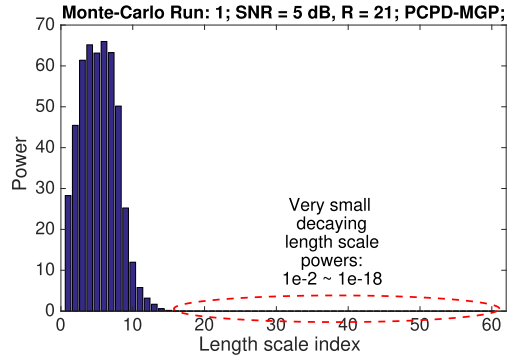
In Table IV, we present the rank estimation performance of PCPD-MGP under different SNRs, with the same settings as those of Fig. 9. It can be seen that when the tensor rank is low (e.g., R = 6), PCPD-MGP shows good performance when SNR is larger than -5 dB. It outperforms the PCPD-GG method, but it is still not as good as the proposed PCPD-GH. At SNR = -10 dB, PCPD-MGP fails to correctly estimate the tensor rank, while PCPD-GH shows good performance in a wider range of SNRs (from -10 dB to 20 dB). Furthermore, when the tensor rank becomes large (e.g., R = 24), PCPD-MGP fails to learn the underlying true tensor rank, making it inferior to PCPD-GH (and even PCPD-GG) in the high-rank regime.

2) *Insights From Learnt Length Scales*: To reveal more insights, we present the learnt length scales (without pruning) of PCPD-MGP. Following the notations in [38], the learnt length scale powers of PCPD-MGP are denoted by $\{\tau_l^{-1}\}_l$.

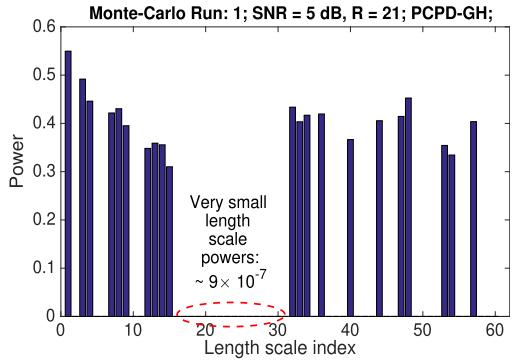
Two typical cases are considered: i) Case I: SNR = 5 dB, R = 21, corresponding to high rank and high SNR; ii) Case II: SNR = -10 dB, R = 6, corresponding to very low SNR and low rank. For each case, we present the learnt length scales in a single trial in Figs. 12 and 13 respectively. From these figures, we have the following observations. i) Due to the *decaying effect* of the MGP prior, the learnt length scale power τ_l^{-1} quickly decreases as l becomes larger. This drives PCPD-MGP to *fail* to recover the sparsity-pattern of the high-rank CPD (e.g., R = 21), in which a large number of length scale powers should be much larger than zero, see Fig. 12(a). In contrast, without the decaying effect, the proposed PCPD-GH successfully identifies the 21 non-negligible components, as seen in Fig. 12(b). ii) When the SNR is very low (e.g., SNR = -10 dB) and the rank is low (e.g., R = 6), the sparsity pattern of length scales learnt by PCPD-MGP is *not as accurate as* that obtained by PCPD-GH (see Fig. 13). Additional simulation runs, and results of more simulation settings are included in Appendix S, from which similar conclusions can be drawn.

VI. (SEMI-) REAL-WORLD DATASET EXPERIMENTS

In this section, experimental results on several (semi-) real-world datasets are presented. Since the scope of this paper is to



(a)



(b)

Fig. 12. (a) The powers of learnt length scales (i.e., $\{\tau_l^{-1}\}_l$) for PCPD-MGP; (b) The powers of learnt length scales (i.e., $\{z_l\}_l$) for PCPD-GH. It can be seen that PCPD-MGP recovers 15 components with non-negligible magnitudes, while PCPD-GH recovers 21 components. The two algorithms are with the same upper bound value: 60. SNR = 5 dB, R = 21.

TABLE V
FIT VALUES AND ESTIMATED TENSOR RANKS OF FLUORESCENCE DATA UNDER
DIFFERENT SNRS (WITH RANK UPPER BOUND $\max\{J_n\}_{n=1}^N$)

SNR (dB)	PCPD-GG		PCPD-GH	
	Fit Value	Estimated Tensor Rank	Fit Value	Estimated Tensor Rank
-10	71.8109	4	72.6401	3
-5	83.9269	4	84.3424	3
0	90.6007	4	90.8433	3
5	94.2554	4	94.3554	3
10	96.0907	3	96.0951	3
15	96.8412	3	96.8431	3
20	97.1197	3	97.1204	3

investigate flexible sparsity-aware modeling for tensor CPD in the parametric Gaussian scale mixture family, the comparisons are performed between PCPD-GG and PCPD-GH.

A. Semi-Real Experiment: Fluorescence Data Analytics

Tensor CPD is an important tool in fluorescence data analytics, with the aim to reveal the underlying signal components.

TABLE VI
TENSOR RANK LEARNING FOR CPD OF TONGUE DATASET AND ENRON EMAIL DATASET

Dataset	Benchmarking Tensor Rank	PCPD-GG				PCPD-GH			
		max(DimY): 13		2max(DimY): 26		max(DimY): 13		2max(DimY): 26	
Tongue	2	2		3		2	2	2	2
ENRON	5	max(DimY): 184	max(DimY)+20: 204	max(DimY)+40: 224	2max(DimY): 368	max(DimY): 184	max(DimY)+20: 204	max(DimY)+40: 224	2max(DimY): 368
		5	5	6	11	5	5	5	9

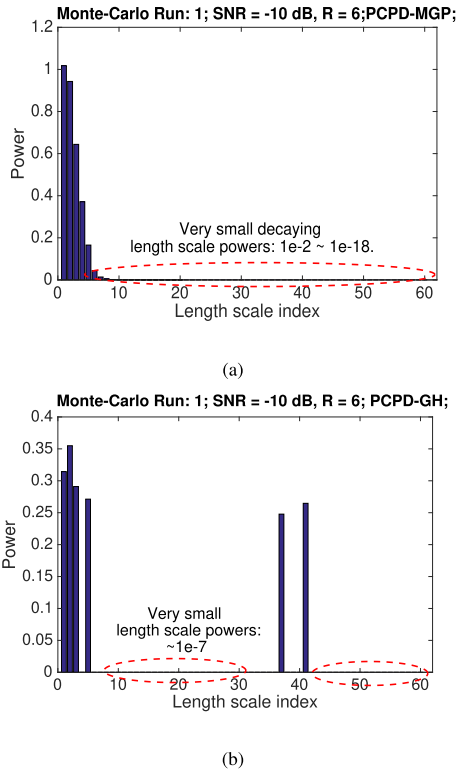


Fig. 13. (a) The powers of learnt length scales (i.e., $\{\tau_l^{-1}\}_l$) for PCPD-MGP; (b) The powers of learnt length scales (i.e., $\{z_l\}_l$) for PCPD-GH. It can be seen that PCPD-MGP recovers 5 components with non-negligible magnitudes, while PCPD-GH recovers 6 components. The two algorithms are with the same upper bound value: 60. SNR = -10 dB, R = 6.

TABLE VII
SNR OUTPUTS AND ESTIMATED TENSOR RANKS OF HSI DATA UNDER DIFFERENT RANK UPPER BOUNDS

Algorithm		PCPD-GG		PCPD-GH	
Dataset	Rank Upper Bound	SNR Output (dB)	Estimated Tensor Rank	SNR Output (dB)	Estimated Tensor Rank
Salinas-A	$\max\{J_n\}_{n=1}^N$	43.7374	137	44.0519	143
	$2\max\{J_n\}_{n=1}^N$	46.7221	257	46.7846	260
Indian Pines	$\max\{J_n\}_{n=1}^N$	30.4207	169	30.5541	178
	$2\max\{J_n\}_{n=1}^N$	31.9047	317	32.0612	335

We consider the popular amino acids fluorescence data⁸ \mathcal{X} with size $5 \times 201 \times 61$ [24], which consists of five laboratory-made samples. Each sample contains different amounts of tyrosine,

⁸<http://www.models.life.ku.dk>

tryptophan and phenylalanine dissolved in phosphate buffered water. Since there are three different types of amino acid, when adopting the CPD model, the optimal tensor rank should be 3. In particular, with the optimal tensor rank 3, the clean spectra for the three types of amino acid, which are recovered by the alternative least-squares (ALS) algorithm [24], are presented in Figure H1 (in Appendix H) as the benchmark.

In practice, it is impossible to know how many components are present in the data in advance, and this calls for automatic tensor rank learning. In this subsection, we assess both the rank learning performance and the noise mitigation performance for the two algorithms (i.e., PCPD-GH and PCPD-GG) under different levels of noise sources. In particular, the Fit value [10], which is defined as $(1 - \frac{\|\hat{\mathcal{X}} - \mathcal{X}\|_F}{\|\mathcal{X}\|_F}) \times 100\%$, is adopted, where $\hat{\mathcal{X}}$ represents the reconstructed fluorescence tensor data from the algorithm. In Table V and H1 (see Appendix H), the performances of the two algorithms are presented assuming different upper bound values of tensor rank. It can be observed that with different upper bound values, the proposed PCPD-GH algorithm always gives the correct tensor rank estimates, even when the SNR is smaller than 0 dB. On the other hand, the PCPD-GG method is quite sensitive to the choice of the upper bound value. Its performance with upper bound $2\max\{J_n\}_{n=1}^N$ becomes much worse than that with $\max\{J_n\}_{n=1}^N$ in tensor rank learning. Even with the upper bound being equal to $\max\{J_n\}_{n=1}^N$, PCPD-GG fails to recover the optimal tensor rank 3 in the low SNR region (i.e., SNR ≤ 5 dB). With the over-estimated tensor rank, the reconstructed fluorescence tensor data $\hat{\mathcal{X}}$ will be overfitted to the noise sources, leading to lower Fit values. As a result, the Fit values of the proposed method are generally higher than those of the PCPD-GG method under different SNRs.

In this application, since the tensor rank denotes the number of underlying components inside the data, its incorrect estimation will not only lead to overfitting to the noise, but also will cause “ghost” components that cannot be interpreted. The figures for recovered spectra and the “ghost” component under incorrect rank estimation are given in Appendix H. In particular, in Figure H3 (of Appendix H), we present the learnt length scale powers and the associated component magnitudes from PCPD-GG and PCPD-GH. It can be seen that PCPD-GG recovers four components with non-negligible magnitudes. The smallest magnitude of the learnt four components is 10.6, which is much larger than zero. In practical data analysis, disregarding a learnt latent component with magnitude 10.6 is not reasonable. Since the “ghost” component has a relatively large magnitude, it degrades the performance of the tensor signal recovery. From Table V, when the SNR is low (e.g., -10 dB), the Fit value of PCPD-GH is higher by 0.8, compared to PCPD-GG. Note that in the high SNR regime (no “ghost” component), their Fit value difference is about 0.001. Therefore, the “ghost” component largely degrades the tensor signal recovery.

B. Real-World Dataset Experiments

In this subsection, four real-world datasets, namely the ENRON Email dataset,⁹ the Tongue dataset [63], the Salinas-A hyperspectral image (HSI) dataset,¹⁰ and the Indian Pines HSI dataset¹¹ are used for evaluation. The ENRON Email dataset (with size $184 \times 184 \times 44$) contains e-mail communication records between 184 people within 44 months, and the Tongue dataset (with size $5 \times 10 \times 13$) was obtained from X-rays of five different speakers during their pronunciation of the vowels. The Salinas-A dataset (with size $83 \times 86 \times 204$) was collected by the AVIRIS¹² sensor over Salinas Valley, California; and the Indian Pines dataset (with size $145 \times 145 \times 200$) was collected by AVIRIS sensor over the Indian Pines test site in North-western Indiana.

For the ENRON Email dataset and the Tongue dataset, it was reported in previous research, see, e.g., [14], [63], that interpretable decomposition results can be obtained in the context of each data analysis task, when the tensor rank is 5 and 2, respectively. In particular, rank-5 tensor CPD can give interpretable clustering results for the ENRON Email dataset [14]; and rank-2 tensor CPD is good for extracting the two underlying components from the Tongue dataset [63]. Using these as benchmarking tensor ranks, the performance of the PCPD-GH method and the PCPD-GG method was reported in Table VI, from which it can be observed that the proposed PCPD-GH method is more robust to different tensor rank upper bounds, and outperforms the PCPD-GG method, in the sense that the estimated tensor ranks of PCPD-GH match benchmarking results in more scenarios than PCPD-GG.

On the other hand, since hyperspectral image (HSI) data are naturally three dimensional (two spatial dimensions and one spectral dimension), tensor CPD is a good tool to analyze such data. However, due to the radiometric noise, photon effects and calibration errors, it is crucial to mitigate these corruptions before putting the HSI data into use. Since each HSI is rich in details, previous works using searching-based methods [12], [13] revealed that the tensor rank in HSI data is usually larger than half of the maximal tensor dimension. This corresponds to the *high tensor rank* scenario considered in this paper. In these two real-world datasets, different bands of HSIs were corrupted by different levels of noises. Some of the HSIs are quite clean while some of them are quite noisy. For such types of real-world data, since no ground-truth is available, a no-reference quality assessment score is usually adopted [12], [13]. Following [12], the SNR output, which is defined as $10 \log_{10} \|\hat{\mathcal{X}}\|_F^2 / \|\mathcal{X} - \hat{\mathcal{X}}\|_F^2$ is utilized as the denoising performance measure, where $\hat{\mathcal{X}}$ is the restored tensor data and \mathcal{X} is the original HSI data. In Table VII, the SNR outputs of the two methods using different rank upper bound values are presented, from which it can be seen that the proposed PCPD-GH method gives higher SNR outputs than PCPD-GG. Samples of denoised HSIs are depicted in Appendix I.

Remark: In HSI denoising, the state-of-the-art performance¹³ is usually achieved via the integration of a tensor method and

deep learning. In this work, we never claim that the proposed method gives the best performance in this specific task, but provides a more advanced solution for the CPD, which can be utilized as a building component for future HSI machine learning model design.

VII. CONCLUSION

In this paper, we investigated the automatic tensor rank learning problem for canonical polyadic decomposition (CPD) models under the framework of parametric Bayesian modeling and variational inference. By noticing that the performance of tensor rank learning is related to the flexibility of the prior distribution, we introduced the generalized hyperbolic (GH) prior to the probabilistic modeling of the CPD problem, based on which an inference algorithm is further developed. Extensive numerical results on synthetic data showed that the proposed method exhibits excellent performance in learning both low and high tensor ranks, even when the noise power is large. This advantage is further evidenced by real-world data analysis.

This paper exemplified how tensor rank learning performance can be enhanced via employing more advanced prior distributions in the Gaussian scale mixture family. Besides GH prior, there are many other advanced priors (including generalized-t distribution [43], normal-exponential gamma distribution [44]) worth investigating. Besides exploiting the variants of the GH prior, informative structures such as non-negativeness [37] and orthogonality [35] can also be incorporated into the newly proposed probabilistic CPD model. Finally, recent deep models utilize tensor decomposition formats to represent its multi-dimensional parameters [64]–[66], and thus there is a surging need to automatically learn the ranks of these decompositions for overfitting avoidance. Early attempts based on Gaussian-gamma prior were reported in [66]. Given that the GH model performs better than the Gaussian-gamma prior model, it is viable to extend these related works [64]–[66] by using the prior model proposed in this paper. In the nonparametric front, replacing the gamma distribution of MGP [38] by a more advanced scale distribution in Gaussian scale mixture family might further improve its flexibility in sparsity modeling and learning. More discussions on future theoretical research are provided in Appendix O.

REFERENCES

- [1] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky, "Tensor decompositions for learning latent variable models," *J. Mach. Learn. Res.*, vol. 15, pp. 2773–2832, 2004.
- [2] N. Cohen, O. Sharir, and A. Shashua, "On the expressive power of deep learning: A tensor analysis," in *Proc. Conf. Learn. Theory*, 2016, pp. 698–728.
- [3] V. Hore *et al.*, "Tensor decomposition for multiple-tissue gene expression experiments," *Nature Genet.*, vol. 48, no. 9, pp. 1094–1100, 2016.
- [4] Y. H. Taguchi, "Drug candidate identification based on gene expression of treated cells using tensor decomposition-based unsupervised feature extraction for large-scale data," *BMC Bioinf.*, vol. 19, no. 13, 2019, Art. no. 388.
- [5] S. D. Bachman, B. Fox-Kemper, and F. O. Bryan, "A diagnosis of anisotropic eddy diffusion from a high-resolution global ocean model," *J. Adv. Model. Earth Syst.*, vol. 12, no. 2, 2020, Art. no. e2019MS001904.
- [6] Y. Cheng, G. Li, N. Wong, H. B. Chen, and H. Yu., "DEEPEYE: A deeply tensor-compressed neural network for video comprehension on terminal devices," *ACM Trans. Embedded Comput. Syst.*, vol. 19, no. 3, pp. 1–25, 2020.

⁹The original source of the data is from [62], and we greatly appreciate Prof. Vagelis Papalexakis for sharing the data with us.

¹⁰http://www.ehu.es/ccwintco/index.php/Hyperspectral_Remote_Sensing_Scenes#Salinas-A_scene.

¹¹http://www.ehu.es/ccwintco/index.php/Hyperspectral_Remote_Sensing_Scenes#Indian_Pines.

¹²An acronym for the Airborne Visible InfraRed Imaging Spectrometer.

¹³<https://paperswithcode.com/task/hyperspectral-image-classification>.

- [7] E. E. Papalexakis, C. Faloutsos, and N. D. Sidiropoulos, "Tensors for data mining and data fusion: Models, applications, and scalable algorithms," *ACM Trans. Intell. Syst. Technol.*, vol. 8, no. 2, pp. 1–44, 2016.
- [8] N. D. Sidiropoulos, L. D. Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos, "Tensor decomposition for signal processing and machine learning," *IEEE Trans. Signal Process.*, vol. 65, no. 13, pp. 3551–3582, Jul. 2017.
- [9] Z. Bai, Y. Li, M. Woźniak, M. Zhou, and D. Li, "DecomVQANet: Decomposing visual question answering deep network via tensor decomposition and regression," *Pattern Recognit.*, vol. 110, 2020, Art. no. 107538.
- [10] A. Cichocki, R. Zdunek, A. H. Phan, and S. I. Amari, *Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-Way Data Analysis and Blind Source Separation*, Hoboken, NJ, USA: Wiley, 2009.
- [11] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Rev.*, vol. 51, no. 3, pp. 455–500, Aug. 2009.
- [12] X. Liu, S. Bourennane, and C. Fossati, "Denoising of hyperspectral images using the PARAFAC model and statistical performance analysis," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 10, pp. 3717–3724, Oct. 2012.
- [13] B. Rasti, P. Scheunders, P. Ghamsi, G. Licciardi, and J. Chanussot, "Noise reduction in hyperspectral imagery: Overview and application," *Remote Sens.*, vol. 10, no. 3, 2018, Art. no. 482.
- [14] E. E. Papalexakis, N. D. Sidiropoulos, and R. Bro, "From k-means to higher-way co-clustering: Multilinear decomposition with sparse latent factors," *IEEE Trans. Signal Process.*, vol. 61, no. 2, pp. 493–506, Jan. 2013.
- [15] H. Chen and J. Li, "DrugCom: Synergistic discovery of drug combinations using tensor decomposition," in *Proc. IEEE Int. Conf. Data Mining*, 2018, pp. 899–904.
- [16] R. Bro, "Review on multiway analysis in chemistry 2000-2005," *Crit. Rev. Anal. Chem.*, vol. 35, pp. 279–293, Jan. 2006.
- [17] C. Chatzichristos, E. Kofidis, M. Morante, and S. Theodoridis, "Blind fMRI source unmixing via higher-order tensor decompositions," *J. Neurosci. Methods*, vol. 315, pp. 17–47, 2019.
- [18] D. P. Bertsekas, *Nonlinear Programming*, 3rd ed. Belmont, MA, USA: Athena Scientific, 2016.
- [19] Y. Xu and W. Yin, "A block coordinate descent method for regularized multi-convex optimization with applications to nonnegative tensor factorization and completion," *SIAM J. Imag. Sci.*, vol. 6, no. 3, pp. 1758–1789, 2013.
- [20] A. P. Liavas, G. Kostoulas, G. Lourakis, K. Huang, and N. D. Sidiropoulos, "Nesterov-based alternating optimization for nonnegative tensor factorization: Algorithm and parallel implementations," *IEEE Trans. Signal Process.*, vol. 66, no. 4, pp. 944–953, Feb. 2018.
- [21] X. Fu, S. Ibrahim, H.-T. Wai, C. Gao, and K. Huang, "Block-randomized stochastic proximal optimization for low-rank tensor factorization," *IEEE Trans. Signal Process.*, vol. 68, pp. 2170–2185, Mar. 2020.
- [22] Q. Zhao, L. Zhang, and A. Cichocki, "Bayesian CP factorization of incomplete tensors with automatic rank determination," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 9, pp. 1751–1763, Sep. 2015.
- [23] Q. Zhao, G. Zhou, L. Zhang, A. Cichocki, and S. I. Amari, "Bayesian robust tensor factorization for incomplete multiway data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 27, no. 4, pp. 736–748, Apr. 2016.
- [24] R. Bro, "Multi-way analysis in the food industry-models, algorithms, and applications," *Proc. ICSLP 2000*, 1998.
- [25] R. Bro and H. Kiers, "A new efficient method for determining the number of components in PARAFAC models," *J. Chemom.*, vol. 17, pp. 274–286, 2003.
- [26] E. E. Papalexakis and C. Faloutsos, "Fast efficient and scalable core consistency diagnostic for the parafac decomposition for big sparse tensors," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Brisbane, QLD, 2015, pp. 5441–5445.
- [27] A. A. Rontogiannis, E. Kofidis, and P. V. Giampouras, "Block-term tensor decomposition: Model selection and computation," in *Proc. 28th Eur. Signal Process. Conf.*, Amsterdam, 2020, pp. 1976–1980.
- [28] M. Mørup and L. K. Hansen, "Automatic relevance determination for multi-way models," *J. Chemom.*, vol. 23, no. 7/8, pp. 352–363, 2009.
- [29] M. N. Schmidt and S. Mohamed, "Probabilistic nonnegative tensor factorization using Markov chain monte carlo," in *Proc. 17th Eur. Signal Process. Conf.*, 2009, pp. 1918–1922.
- [30] B. Ermis and A. T. Cemgil, "A Bayesian tensor factorization model via variational inference for link prediction," 2014, *arXiv:1409.8276*.
- [31] P. Rai, C. Hu, M. Harding, and L. Carin, "Scalable probabilistic tensor factorization for binary and count data," in *Proc. Int. Joint Conf. Artif. Intell.*, 2015, pp. 3770–3776.
- [32] Z. Zhang and C. Hawkins, "Variational Bayesian inference for robust streaming tensor factorization and completion," in *Proc. IEEE Int. Conf. Data Mining*, 2018, pp. 1446–1451.
- [33] J. Luan and Z. Zhang, "Prediction of multi-dimensional spatial variation data via Bayesian tensor completion," *IEEE Trans. CAD Integr. Circuits Syst.*, vol. 39, no. 2, pp. 547–551, Feb. 2020.
- [34] Q. Zhao, L. Zhang, and A. Cichocki, "Bayesian sparse Tucker models for dimension reduction and tensor completion," 2015. [Online]. Available: <http://arxiv.org/abs/1505.02343>
- [35] L. Cheng, Y.-C. Wu, and H. V. Poor, "Probabilistic tensor canonical polyadic decomposition with orthogonal factors," *IEEE Trans. Signal Process.*, vol. 65, no. 3, pp. 663–676, Feb. 2017.
- [36] L. Cheng, Y.-C. Wu, and H. V. Poor, "Scaling probabilistic tensor canonical polyadic decomposition to massive data," *IEEE Trans. Signal Process.*, vol. 66, no. 21, pp. 5534–5548, Nov. 2018.
- [37] L. Cheng, X. Tong, S. Wang, Y.-C. Wu, and H. V. Poor, "Learning nonnegative factors from tensor data: Probabilistic modeling and inference algorithm," *IEEE Trans. Signal Process.*, vol. 68, pp. 1792–1806, Feb. 2020.
- [38] P. Rai, Y. Wang, S. Guo, G. Chen, D. Dunson, and L. Carin, "Scalable bayesian low-rank decomposition of incomplete multiway tensors," in *Proc. 31st Int. Conf. Mach. Learn.*, 2014, pp. 1800–1808.
- [39] Y. Zhou and Y. M. Cheung, "Bayesian low-tubal-rank robust tensor factorization with multi-rank determination," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 1, pp. 62–76, Jan. 2020.
- [40] S. Khan and S. Kaski, "Bayesian multi-view tensor factorization," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discov. Databases*, 2014, pp. 656–671.
- [41] D. J. MacKay, "Probable networks and plausible predictions—a review of practical Bayesian methods for supervised neural networks," *Comput. Neural Syst.*, vol. 6, no. 3, pp. 469–505, 1995.
- [42] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, Jun. 2001.
- [43] A. Armagan, D. B. Dunson, and J. Lee, "Generalized double Pareto shrinkage," *Statistica Sinica*, vol. 23, no. 1, 2013, Art. no. 119.
- [44] J. Griffin and P. Brown, "Inference with normal-gamma prior distributions in regression problems," *Bayesian Anal.*, vol. 5, no. 1, pp. 171–188, 2010.
- [45] C. Carvahlo, N. Polson, and J. Scott, "The horseshoe estimator for sparse signals," *Biometrika*, vol. 97, no. 2, 2010, Art. no. 465.
- [46] L. Thabane and M. Safiul Haq, "On the matrix-variate generalized hyperbolic distribution and its Bayesian applications," *Statistics*, vol. 38, no. 6, pp. 511–526, 2004.
- [47] S. D. Babacan, S. Nakajima, and M. N. Do, "Bayesian group-sparse modeling and variational inference," *IEEE Trans. Signal Process.*, vol. 62, no. 11, pp. 2906–2921, Jun. 2014.
- [48] B. Jørgensen, *Statistical Properties of the Generalized Inverse Gaussian Distribution*, Berlin, Germany: Springer Science & Business Media, 2012.
- [49] J. M. Bernardo and A. F. Smith, *Bayesian Theory*. Hoboken, NJ, USA: Wiley, 2009.
- [50] M. J. Beal, *Variational Algorithms for Approximate Bayesian Inference*. London, U.K.: Univ. London, Univ. College London, 2003.
- [51] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Found. Trends Mach. Learn.*, vol. 1, no. 102, pp. 1–305, Jan. 2008.
- [52] C. Zhang, J. Butepage, H. Kjellstrom, and S. Mandt, "Advances in variational inference," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 8, pp. 2008–2026, Aug. 2019.
- [53] S. Theodoridis, *Machine Learning: A Bayesian and Optimization Perspective*, 2nd ed., Cambridge, MA, USA: Academic Press, 2020.
- [54] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose Bayesian inference algorithm," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2378–2386.
- [55] K. Panousis, S. Chatzis, and S. Theodoridis, "Nonparametric Bayesian deep networks with local competition," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 4980–4988.
- [56] Y. Weng, L. Wu, and W. Hong, "Bayesian inference via variational approximation for collaborative filtering," *Neural Process. Lett.*, vol. 49, no. 3, pp. 1041–1054, 2019.
- [57] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, 2003.
- [58] Z. Ghahramani and M. J. Beal, "Variational inference for Bayesian mixtures of factor analysers," in *Proc. Adv. Neural Inf. Process. Syst.*, 2000, pp. 449–455.

- [59] A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions*, Boca Raton, FL, USA: CRC Press, 1999.
- [60] W. Deng, W. Yin, and Y. Zhang, "Group sparse optimization by alternating direction method," *Wavelets Sparsity XV*, vol. 8858, Sep. 2013, Art. no. 88580R.
- [61] Z. Ghahramani and M. J. Beal, "Variational inference for Bayesian mixtures of factor analysers," in *Proc. Adv. Neural Inf. Process. Syst.*, 2000, pp. 449–455.
- [62] B. W. Bader, R. A. Harshman, and T. G. Kolda, "Temporal analysis of social network using three-way DEDICOM," Sandia National Labs, TR SAND 2006–2161, 2006.
- [63] R. A. Harshman, P. Ladefoged, and L. Goldstein, "Factor analysis of tongue shapes," *J. Acoust. Soc. Amer.*, vol. 62, pp. 693–713, 1977.
- [64] Z. He, S. Gao, L. Xiao, D. Liu, H. He, and D. Barber, "Wider and deeper, cheaper and faster: Tensorized LSTMs for sequence learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 1–11.
- [65] Y.-D. Kim, E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin, "Compression of deep convolutional neural networks for fast and low power mobile applications," 2015, *arXiv:1511.06530*.
- [66] C. Hawkins, X. Liu, and Z. Zhang, "Towards compact neural networks via end-to-end training: A Bayesian tensor approach with automatic rank determination," 2020, *arXiv:2010.08689*.
- [67] N. Vervliet, O. Debals, L. Sorber, M. Van Barel, and L. De Lathauwer, "Tensorlab 3.0—Numerical optimization strategies for large-scale constrained and coupled matrix/tensor factorization," in *Proc. 50th Asilomar Conf. Signals, Syst. Comput.*, Mar. 2016, pp. 1733–1738. [Online]. Available: <https://www.tensorlab.net/>
- [68] P. J. Green, "Reversible jump Markov chain Monte Carlo computation and bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.
- [69] D. P. Wipf and S. Nagarajan, "A new view of automatic relevance determination," in *Proc. Adv. Neural Inf. Process. Syst.*, 2008, pp. 1625–1632.
- [70] S. Nakajima, M. Sugiyama, S. D. Babacan, and R. Tomioka, "Global analytic solution of fully-observed variational bayesian matrix factorization," *J. Mach. Learn. Res.*, vol. 14, pp. 1–37, 2013.
- [71] G. S. Ranganath and B. David, "Black box variational inference," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2014, pp. 814–822.



analytics, and interpretable machine learning for information systems.



Zhongtao Chen received the B.Eng. degree from The Chinese University of Hong Kong, Shenzhen, China, in 2021. He is currently working toward the Ph.D. degree with The University of Hong Kong, Hong Kong. His research interests include signal processing and machine learning using Bayesian methods.



Qingjiang Shi received the Ph.D. degree in electronic engineering from Shanghai Jiao Tong University, Shanghai, China, in 2011. From September 2009 to September 2010, he visited Prof. Z.-Q. (Tom) Luo's Research Group, University of Minnesota, Minneapolis, MN, USA. In 2011, he was a Research Scientist with Bell Laboratories China. From 2012, He was with the School of Information and Science Technology, Zhejiang Sci-Tech University, Hangzhou, China. From February 2016 to March 2017, he was a Research Fellow with Iowa State University, Ames, IA, USA. From March 2018, he is currently a Full Professor with the School of Software Engineering, Tongji University, Shanghai, China. He is also with the Shenzhen Research Institute of Big Data. He has authored or coauthored more than 70 IEEE journals and filed about 30 national patents. His research interests include algorithm design and analysis with applications in machine learning, signal processing and wireless networks. Dr. Shi was an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING. He was the recipient of the Huawei Outstanding Technical Achievement Award in 2021, the Huawei Technical Cooperation Achievement Transformation Award (2nd Prize) in 2022, the Golden Medal at the 46th International Exhibition of Inventions of Geneva in 2018, the First Prize of Science and Technology Award from China Institute of Communications in 2017, the National Excellent Doctoral Dissertation Nomination Award in 2013, the Shanghai Excellent Doctoral Dissertation Award in 2012, and the Best Paper Award from the IEEE PIMRC'09 conference.



Yik-Chung Wu (Senior Member, IEEE) received the B.Eng. (EEE) and M.Phil. degrees from The University of Hong Kong (HKU), Hong Kong, in 1998 and 2001, respectively, and the Ph.D. degree from Texas A&M University, College Station, TX, USA, in 2005. From 2005 to 2006, he was with Thomson Corporate Research, Princeton, NJ, USA, as a Member of Technical Staff. Since 2006, he has been with HKU, where he is currently an Associate Professor. He was a Visiting Scholar with Princeton University, Princeton, NJ, USA, between 2015 and 2017. His research interests include signal processing, machine learning, and communication systems. He was the Editor of IEEE COMMUNICATIONS LETTERS and IEEE TRANSACTIONS ON COMMUNICATIONS. He is currently the Editor of IEEE TRANSACTIONS ON SIGNAL PROCESSING and *Journal of Communications and Networks*.



Sergios Theodoridis (Life Fellow, IEEE) received Ph.D. from Department of Electronics and Electrical Engineering, University of Birmingham, U.K, in 1978. He is currently Professor Emeritus with the Department of Informatics and Telecommunications, National and Kapodistrian University of Athens and he is a Distinguished Professor with Aalborg University, Aalborg, Denmark. His research interests include the cross section of signal processing and machine learning. He is the author of the book *Machine Learning: A Bayesian and Optimization Perspective*, Academic Press, second edition, 2020, the co-author of the best-selling book *Pattern Recognition*, Academic Press, fourth edition 2009, and co-author of the book *Introduction to Pattern Recognition: A MATLAB Approach*, Academic Press, 2010. He is the coauthor of seven papers that have received Best Paper Awards including the 2014 *IEEE Signal Processing Magazine* Best Paper award and the 2009 IEEE Computational Intelligence Society Transactions on Neural Networks Outstanding Paper Award. He was the recipient of the 2021 *IEEE Signal Processing* Norbert Wiener Award, the 2017 EURASIP Athanasios Papoulis Award, the 2014 IEEE Signal Processing Society Carl Friedrich Gauss Education Award and the 2014 EURASIP Meritorious Service Award. He was the Vice President IEEE Signal Processing Society, Editor-in-Chief of the IEEE TRANSACTIONS ON SIGNAL PROCESSING and as President of the *European Association for Signal Processing* (EURASIP). He is currently the Chair for the IEEE Signal Processing Society awards committee. He is Fellow of IET, a Corresponding Fellow of the Royal Society of Edinburgh (RSE), and a Fellow of EURASIP.