The Annals of Statistics
2022, Vol. 50, No. 1, 1–29
https://doi.org/10.1214/21-AOS2061

AN OPTIMAL STATISTICAL AND COMPUTATIONAL FRAMEWORK FOR GENERALIZED TENSOR ESTIMATION

BY RUNGANG HAN[1,*], REBECCA WILLETT[2] AND ANRU R. ZHANG[1,†]

[1]*Department of Statistics, University of Wisconsin-Madison,* [*]*rhan32@stat.wisc.edu;* [†]*anruzhang@stat.wisc.edu*
[2]*Departments of Statistics and Computer Science, University of Chicago, willett@uchicago.edu*

This paper describes a flexible framework for generalized low-rank tensor estimation problems that includes many important instances arising from applications in computational imaging, genomics, and network analysis. The proposed estimator consists of finding a low-rank tensor fit to the data under generalized parametric models. To overcome the difficulty of nonconvexity in these problems, we introduce a unified approach of projected gradient descent that adapts to the underlying low-rank structure. Under mild conditions on the loss function, we establish both an upper bound on statistical error and the linear rate of computational convergence through a general deterministic analysis. Then we further consider a suite of generalized tensor estimation problems, including sub-Gaussian tensor PCA, tensor regression, and Poisson and binomial tensor PCA. We prove that the proposed algorithm achieves the minimax optimal rate of convergence in estimation error. Finally, we demonstrate the superiority of the proposed framework via extensive experiments on both simulated and real data.

1. Introduction. In recent years, the analysis of tensors or high-order arrays has emerged as an active topic in statistics, applied mathematics, machine learning, and data science. Datasets in the form of tensors arise from various scientific applications (Kroonenberg (2008)), such as collaborative filtering (Bi, Qu and Shen (2018), Shah (2018)), neuroimaging analysis (Li et al. (2018), Zhou, Li and Zhu (2013)), hyperspectral imaging (Li and Li (2010)), longitudinal data analysis (Hoff (2015)), and more. In many of these problems, although the tensor of interest is high-dimensional in the sense that the ambient dimension of the dataset is substantially greater than the sample size, there is often hidden low-dimensional structures in the tensor that can be exploited to facilitate the data analysis. In particular, the low-rank condition renders convenient decomposable structure and has been proposed and widely used in the analysis of tensor data (Kolda and Bader (2009), Kroonenberg (2008)). However, leveraging these hidden low-rank structures in estimation and inference can pose great statistical and computational challenges in real practice.

1.1. *Generalized tensor estimation.* In this paper, we consider a statistical and optimization framework for generalized tensor estimation. Suppose we observe a random sample D drawn from some distribution parametrized by an unknown low-rank tensor parameter $\mathcal{X}^* \in \mathbb{R}^{p_1 \times p_2 \times p_3}$. A straightforward idea to estimate $\mathcal{X}^*$ is via optimization:

$$\hat{\mathcal{X}} = \underset{\mathcal{X} \text{ is low-rank}}{\arg\min} \quad L(\mathcal{X}; D). \tag{1.1}$$

Here, $L(\mathcal{X}; D)$ can be taken as the negative log-likelihood function (then $\hat{\mathcal{X}}$ becomes the maximum likelihood estimator (MLE)) or any more general loss function. We can even

Received February 2020; revised February 2021.
MSC2020 subject classifications. Primary 62H12, 62H25; secondary 62C20.
Key words and phrases. Generalize tensor estimation, gradient descent, image denoising, low-rank tensor, minimax optimality, nonconvex optimization.

broaden the scope of this framework to a deterministic setting: suppose we observe D that is "associated" with an unknown tensor parameter $\mathcal{X}^* \in \mathbb{R}^{p_1 \times p_2 \times p_3}$; to estimate $\mathcal{X}^*$, we try to minimize the loss function $L(\mathcal{X}; D)$ that is specified by the problem scenario. This general framework includes many important instances arising in real applications. For example:

- *Computational imaging*. Photon-limited imaging appears in signal processing (Salmon et al. (2014)), material science (Yankovich et al. (2016)), astronomy (Timmermann and Nowak (1999), Willett and Nowak (2007)), and often involves arrays with nonnegative photon counts contaminated by substantial noise. Data from photon-limited imaging are often in the form of tensors (e.g., stacks of spectral images in which each image corresponds to a different wavelength of light). How to denoise these images is often crucial for the subsequent analysis. To this end, Poisson tensor PCA serves as a prototypical model for tensor photon-limited imaging analysis; see Sections 4.3 and 7.2 for more details.
- *Analysis of multilayer network data*. In network analysis, one often observes multiple snapshots of static or dynamic networks (Arroyo et al. (2021), Lei, Chen and Lynch (2020), Pensky and Zhang (2019), Sewell and Chen (2015)). How to perform an integrative analysis for the network structure using multilayer network data has become an important problem in practice. By stacking adjacency matrices from multiple snapshots to an adjacency tensor, the hidden community structure of network can be transformed to the low-rankness of adjacency tensor, and the generalized tensor learning framework can provide a new perspective on the analysis of multilayer network data.
- *Biological sequencing data analysis*. Tensor data also commonly appear in biological sequencing data analysis (Faust et al. (2012), Flores et al. (2014), Wang, Fischer and Song (2017)). The identification of significant triclusters or modules, that is, coexpressions of different genes or coexistence of different microbes, often has significant biological meanings (Henriques and Madeira (2019)). From a statistical perspective, these modules often correspond to low-rank tensor structure, so the generalized tensor learning framework could be naturally applied.
- *Online-click through Prediction.* Online click-through data analysis in e-commerce has become an increasingly important tool in building the online recommendation system (McMahan et al. (2013), Shan et al. (2016), Sun and Li (2017)). There are three major entities: users, items, and time, and the data can be organized as a tensor, where each entry represents the click times of one user on a specific category of items in a time period (e.g., noon or evening). Then generalized tensor estimation could be applied to study the implicit features of users and items for better prediction of user behaviors.

Additional applications include neuroimaging analysis (Zhou, Li and Zhu (2013)), collaborative filtering (Yu, Gupta and Kolar (2020)), mortality rate analysis (Wilmoth and Shkolnikov (2006)), and more. We also elucidate specific model setups and real data examples in detail later in Sections 4 and 7.2, respectively.

The central tasks of generalized tensor estimation problems include two elements. From a statistical perspective, it is important to investigate how well one can estimate the target tensor parameter $\mathcal{X}^*$ and the optimal rates of estimation error. From an optimization perspective, it is crucial to develop a computationally efficient procedure for estimating $\mathcal{X}^*$ with provable theoretical guarantees. To estimate the low-rank tensor parameter $\mathcal{X}^*$, a straightforward idea is to perform the rank constrained minimization on the loss function $L(\mathcal{X}; D)$ in (1.1). Since the low-rank constraint is highly nonconvex, the direct implementation of (1.1) is computationally infeasible in practice. If $\mathcal{X}^*$ is a sparse vector or low-rank matrix, common substitutions often involve convex regularization methods, such as M-estimators with an ℓ_1 penalty or matrix nuclear norm penalty for estimating sparse or low-rank structure (Fazel (2002), Tibshirani (1996)). These methods enjoy great empirical and theoretical success for

vector/matrix estimators, but it is unclear whether they can achieve good performance on generalized tensor estimation problems. First, different from the matrix nuclear norm, tensor nuclear norm is generally NP-hard to even approximate (Friedland and Lim (2018)), so that the tensor nuclear norm regularization approach can be computationally intractable. Second, other computationally feasible convex regularization methods, such as the overlapped nuclear norm minimization (Tomioka and Suzuki (2013), Tomioka et al. (2011)), may be statistical sub-optimal based on the theory of simultaneously structured model estimation (Oymak et al. (2015)).

In contrast, we focus on a unified nonconvex approach for generalized tensor estimation problems in this paper. Our central idea is to decompose the low-rank tensor into $\mathcal{X} = [\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]$ (see Section 2.1 for explanations of tensor algebra) and reformulate the original problem to

$$(\hat{\mathcal{S}}, \hat{\mathbf{U}}_1, \hat{\mathbf{U}}_2, \hat{\mathbf{U}}_3) = \mathop{\arg\min}_{\mathcal{S}, \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3} \left\{ L\big([\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]; D\big) + \frac{a}{2} \sum_{k=1}^{3} \|\mathbf{U}_k^\top \mathbf{U}_k - b^2 \mathbf{I}_{r_k}\|_{\mathrm{F}}^2 \right\}, \tag{1.2}$$

which can be efficiently solved by (projected) gradient descent on all components. The resulting $\hat{\mathcal{X}} = [\![\hat{\mathcal{S}}; \hat{\mathbf{U}}_1, \hat{\mathbf{U}}_2, \hat{\mathbf{U}}_3]\!]$ naturally admits a low-rank structure. The auxiliary regularizers $\|\mathbf{U}_k^\top \mathbf{U}_k - b^2 \mathbf{I}_{r_k}\|_{\mathrm{F}}^2$ in (1.2) can keep $\hat{\mathbf{U}}_k$ from being singular. It is actually easy to check that (1.1) and (1.2) are exactly equivalent.

We provide strong theoretical guarantees for the proposed procedure on generalized tensor estimation problems. In particular, we establish the linear rate of local convergence for gradient descent methods under a general deterministic setting with the *Restricted Correlated Gradient* condition (see Section 3.1 for details). An informal statement of the result is given below,

$$\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 \lesssim \xi^2 + (1-c)^t \|\mathcal{X}^{(0)} - \mathcal{X}^*\|_{\mathrm{F}}^2 \quad \text{for all } t \geq 1 \tag{1.3}$$

with high probability. Here, we use ξ^2 to characterize the statistical noise and its definition and interpretation will be given in Section 3.2. Then for specific statistical models, including sub-Gaussian tensor PCA, tensor regression, Poisson tensor PCA, and binomial tensor PCA, based on the general result (1.3), we prove that the proposed algorithm achieves the minimax optimal rate of convergence in estimation error. Specifically for the low-rank tensor regression problem, Table 1 illustrates the advantage of our method through a comparison with existing ones.

Finally, we apply the proposed framework to synthetic and real data examples, including photon-limited 4D-STEM (scanning transmission electron microscopy) imaging data and click-through e-commerce data. The comparison of performance with existing methods illustrates the merit of our proposed procedure.

1.2. *Related literature*. This work is related to a broad range of literature on tensor analysis. For example, tensor decomposition/SVD/PCA focuses on the extraction of low-rank structures from noisy tensor observations (Anandkumar et al. (2014), Chen (2019), Hopkins, Shi and Steurer (2015), Johndrow, Bhattacharya and Dunson (2017), Lesieur, Krzakala and Zdeborová (2017), Montanari, Reichman and Zeitouni (2017), Richard and Montanari (2014)). Correspondingly, a number of methods have been proposed and analyzed under either deterministic or random Gaussian noise, such as the maximum likelihood estimation (Richard and Montanari (2014)), (truncated) power iterations (Anandkumar et al. (2014), Sun et al. (2017)), higher-order SVD (De Lathauwer, De Moor and Vandewalle (2000a)), higher-order orthogonal iteration (HOOI) (De Lathauwer, De Moor and Vandewalle (2000b), Zhang and Xia (2018)), STAT-SVD (Zhang and Han (2019)).

TABLE 1

Comparison of different tensor regression methods when the rank is known. For simplicity, we assume $r_1 = r_2 = r_3 = r$, $p_1 = p_2 = p_3 = p$ *and* $\sigma^2 \ll \|\mathcal{X}^*\|_{\mathrm{F}}^2$. *Here, the sample complexity* is the minimal sample size required to achieve the corresponding estimation error*

Algorithm	Sample complexity*	Estimation Error Upper Bound	Recovery (noiseless)
Our Method	$p^{3/2}r$	$\sigma^2 pr/n$	Exact
Tucker reg. (Zhou, Li and Zhu (2013))	N.A.	N.A.	Exact
Nonconvex-PGD. (Chen, Raskutti and Yuan (2019))	$p^2 r$	$\sigma^2 p^2 r/n$	Exact
Nuclear Norm Min. (Raskutti, Yuan and Chen (2019))	N.A.	$\sigma^2 pr^2/n$	Exact
Schatten-1 Norm Min. (Tomioka and Suzuki (2013))	$p^2 r$	$\sigma^2 p^2 r/n$	Exact
ISLET (Zhang et al. (2020))	$p^{3/2}r$	$\sigma^2 pr/n$	Inexact
Iterative Hard Thresholding[a] (Rauhut, Schneider and Stojanac (2017))	pr	σ^2	Exact

[a]The analysis in Rauhut, Schneider and Stojanac (2017) relies on an assumption that the projection on low-rank tensor manifold can be approximately done by high-order SVD. It is, however, unclear whether this assumption holds in general.

Since non-Gaussian-valued tensor data also commonly appear in practice, Chi and Kolda (2012), Hong, Kolda and Duersch (2020), Signoretto et al. (2011) considered the generalized tensor decomposition and introduced computational efficient algorithms. However, the theoretical guarantees for these procedures and the statistical performances of the generalized tensor decomposition still remain open.

Our proposed framework includes the topic of tensor recovery and tensor regression. Various methods, such as the convex regularization (Raskutti, Yuan and Chen (2019), Tomioka and Suzuki (2013)), alternating minimization (Zhou, Li and Zhu (2013)), hard thresholding iteration (Chen, Raskutti and Yuan (2019), Rauhut, Schneider and Stojanac (2017, 2015)), importance-sketching (Zhang et al. (2020)) were introduced and studied. A more detailed comparison of these methods is summarized in Table 1.

In addition, high-order interaction pursuits (Hao, Zhang and Cheng (2020)), tensor completion (Cai et al. (2019), Liu et al. (2013), Montanari and Sun (2018), Xia and Yuan (2019), Xia, Yuan and Zhang (2021), Yuan and Zhang (2016), Zhang (2019)), and tensor block models (Chi et al. (2020), Lei, Chen and Lynch (2020), Wang and Zeng (2019)) are important topics in tensor analysis that have attracted enormous attention recently. Departing from the existing results, this paper, to the best of our knowledge, is the first to give a unified treatment for a broad range of tensor estimation problems with both statistical optimality and computational efficiency.

This work is also related to a substantial body of literature on low-rank matrix recovery, where the goal is to estimate a low-rank matrix based on a limited number of observations. Specific examples of this topic include matrix completion (Candes and Plan (2010), Candès and Recht (2009)), phase retrieval (Cai, Li and Ma (2016), Candès, Li and Soltanolkotabi (2015)), blind deconvolution (Ahmed, Recht and Romberg (2014)), low-rank matrix trace regression (Chen and Chi (2018), Fan, Gong and Zhu (2019), Keshavan, Montanari and Oh (2010), Koltchinskii, Lounici and Tsybakov (2011)), and many others. A common approach for low-rank matrix recovery is via explicit low-rank factorization: one can decompose the target p_1-by-p_2 rank-r matrix $\mathbf{X}$ into $\mathbf{X} = \mathbf{U}\mathbf{V}^\top$, where $\mathbf{U} \in \mathbb{R}^{p_1 \times r}$, $\mathbf{V} \in \mathbb{R}^{p_2 \times r}$, then mini-

mize the loss function $L(\mathbf{U}\mathbf{V}^\top)$ with respect to both $\mathbf{U}$ and $\mathbf{V}$ (Wen, Yin and Zhang (2012)). Previously, (Zhao, Wang and Liu (2015)) considered the noiseless setting of trace regression and proved that under good initialization, the first order alternating optimization on $\mathbf{U}$ and $\mathbf{V}$ achieves exact recovery. (Park et al. (2018), Tu et al. (2016)) established the local convergence of gradient descent for strongly convex and smooth loss function L. The readers are referred to a recent survey paper (Chi, Lu and Chen (2019)) on the applications and optimization landmarks of the nonconvex factorized optimization. Despite significant developments in low-rank matrix recovery and nonconvex optimization, they cannot be directly generalized to tensor estimation problems for many reasons. First, many basic matrix concepts or methods cannot be directly generalized to high-order ones (Hillar and Lim (2013)). Naive generalization of matrix concepts (e.g., operator norm, singular values, eigenvalues) are possible but often computationally NP-hard. Second, tensors have more complicated algebraic structure than matrices. As what we will illustrate later, one has to simultaneously handle all arm matrices (i.e., $\mathbf{U}_1$, $\mathbf{U}_2$, and $\mathbf{U}_3$) and the core tensor (i.e., $\mathcal{S}$) with distinct dimensions in the theoretical error contraction analysis. To this end, we develop new technical tools on tensor algebra and perturbation results (e.g., E.2, Lemmas E.3 in the Appendix). More technical issues of generalized tensor estimation will be addressed in Section 3.3.

The projected gradient schemes, which apply gradient descent on the parameter tensor $\mathcal{X}$ followed by the low-rank tensor retraction/projection operators, form another important class of methods in the literature (Rauhut, Schneider and Stojanac (2015, 2017), Chen, Raskutti and Yuan (2019)):

$$\mathcal{X}^{(t+1)} = \mathcal{P}\big(\mathcal{X}^{(t)} - \eta \nabla L(\mathcal{X}^{(t)})\big).$$

Different from the low-rank projection for matrices, the exact low-rank tensor projection (i.e., the best rank-(r_1, r_2, r_3) approximation: $\mathcal{P}(\mathcal{X}) = \arg\min_{\mathcal{T} \text{ is rank-}(r_1,r_2,r_3)} \|\mathcal{T} - \mathcal{X}\|_{\mathrm{F}}$) is NP-hard in general (Hillar and Lim (2013)) and less practical. Several inexact but efficient projection methods were developed and studied to overcome this issue. In particular, Rauhut, Schneider and Stojanac (2015) proposed a polynomial-time computable projected gradient scheme that converges linearly to the true tensor parameter for the noiseless tensor completion problem, given the initialization is sufficiently close to the solution. On the other hand, it is not clear if such schemes with inexact projection operators can achieve optimal statistical rate in the noisy setting (Chen, Raskutti and Yuan (2019)). In contrast, the proposed method in this paper is both computationally efficient and statistically optimal in a variety of settings with provable guarantees.

1.3. *Organization of the paper.* The rest of the article is organized as follows. After a brief introduction of the notation and preliminaries in Section 2.1, we introduce the general problem formulation in Section 2.2. A deterministic error and local convergence analysis of the projected gradient descent algorithm for order-3 tensor estimation is discussed in Section 3. Then we apply the results on a variety of generalized tensor estimation problems in Section 4, including sub-Gaussian tensor PCA, tensor regression, Poisson tensor PCA, and binomial tensor PCA. We develop the upper and minimax matching lower bounds in each of these scenarios. In Section 5, we propose a data-driven rank selection method with theoretical guarantee. The extension to general order-d tensor estimation is discussed in Section 6. Simulation and real data analysis are presented in Section 7. All proofs of technical results and more implementation details of algorithms are collected in the Supplementary Material (Han, Willett and Zhang (2021)).

2. Generalized tensor estimation model.

2.1. *Notation and preliminaries.* The following notation and preliminaries are used throughout this paper. The lowercase letters, for example, x, y, u, v, are used to denote scalars or vectors. For any $a, b \in \mathbb{R}$, let $a \wedge b$ and $a \vee b$ be the minimum and maximum of a and b, respectively. We use $C, C_0, C_1, \ldots$ and $c, c_0, c_1, \ldots$ to represent generic large and small positive constants respectively. The actual values of these generic symbols may differ from line to line.

We use bold uppercase letters $\mathbf{A}$, $\mathbf{B}$ to denote matrices. Let $\mathbb{O}_{p,r}$ be the collection of all p-by-r matrices with orthonormal columns: $\mathbb{O}_{p,r} = \{\mathbf{U} \in \mathbb{R}^{p\times r} : \mathbf{U}^\top\mathbf{U} = \mathbf{I}_r\}$, where $\mathbf{I}_r$ is the r-by-r identity matrix. For any matrix $\mathbf{A} \in \mathbb{R}^{p_1\times p_2}$, let $\sigma_1(\mathbf{A}) \geq \cdots \geq \sigma_{p_1\wedge p_2}(\mathbf{A}) \ldots \geq 0$ be its singular values in descending order. We also define $\mathrm{SVD}_r(\mathbf{A}) \in \mathbb{O}_{p,r}$ to be the matrix comprised of the top r left singular vectors of $\mathbf{A}$. For any matrix $\mathbf{A}$, let $\mathbf{A}_{ij}$, $\mathbf{A}_{i\cdot}$, and $\mathbf{A}_{\cdot j}$ be the entry on the ith row and jth column, the ith row, and the jth column of $\mathbf{A}$, respectively. The inner product of two matrices with the same dimension is defined as $\langle \mathbf{A}, \mathbf{B}\rangle = \mathrm{tr}(\mathbf{A}^\top\mathbf{B})$, where $\mathrm{tr}(\cdot)$ is the trace operator. We use $\|\mathbf{A}\| = \sigma_1(\mathbf{A})$ to denote the spectral norm of $\mathbf{A}$, use $\|\mathbf{A}\|_{\mathrm{F}} = \sqrt{\sum_{i,j}\mathbf{A}_{ij}^2} = \sqrt{\sum_{k=1}^{p_1\wedge p_2}\sigma_k^2}$ to denote the Frobenius norm of $\mathbf{A}$, and use $\|\mathbf{A}\|_* = \sum_{k=1}^{p_1\wedge p_2}\sigma_k$ to denote the nuclear norm of $\mathbf{A}$. The $l_{2,\infty}$ norm of $\mathbf{A}$ is defined as the largest row-wise l_2 norm of $\mathbf{A}$: $\|\mathbf{A}\|_{2,\infty} = \max_i \|\mathbf{A}_{i\cdot}\|_2$. For any matrix $\mathbf{A} = [a_1, \ldots, a_J] \in \mathbb{R}^{I\times J}$ and $\mathbf{B} \in \mathbb{R}^{K\times L}$, the *Kronecker product* is defined as the (IK)-by-(JL) matrix $\mathbf{A}\otimes\mathbf{B} = [a_1\otimes\mathbf{B}\cdots a_J\otimes\mathbf{B}]$.

In addition, we use calligraphic letters, for example, $\mathcal{S}, \mathcal{X}, \mathcal{Y}$, to denote higher-order tensors. To simplify the presentation, we mainly focuses on order-3 tensors in this paper while all results for higher-order tensors can be carried out similarly. For tensor $\mathcal{S} \in \mathbb{R}^{r_1\times r_2\times r_3}$ and matrix $\mathbf{U}_1 \in \mathbb{R}^{p_1\times r_1}$, the mode-1 tensor-matrix product is defined as:

$$\mathcal{S}\times_1\mathbf{U}_1 \in \mathbb{R}^{p_1\times r_2\times r_3}, \qquad (\mathcal{S}\times\mathbf{U}_1)_{i_1i_2i_3} = \sum_{j=1}^{r_1}\mathcal{S}_{ji_2i_3}(\mathbf{U}_1)_{i_1j}.$$

For any $\mathbf{U}_2 \in \mathbb{R}^{p_2\times r_2}$, $\mathbf{U}_3 \in \mathbb{R}^{p_3\times r_3}$, the tensor-matrix products $\mathcal{S}\times_2\mathbf{U}_2$ and $\mathcal{S}\times_3\mathbf{U}_3$ are defined in a similarly way. Importantly, multiplication along different directions is commutative invariant: $(\mathcal{S}\times_{k_1}\mathbf{U}_{k_1})\times_{k_2}\mathbf{U}_{k_2} = (\mathcal{S}\times_{k_2}\mathbf{U}_{k_2})\times_{k_1}\mathbf{U}_{k_1}$ for any $k_1 \neq k_2$. We simply denote

$$\mathcal{S}\times_1\mathbf{U}_1\times_2\mathbf{U}_2\times_3\mathbf{U}_3 = [\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!],$$

as this formula commonly appears in the analysis. We also introduce the matricization operator that transforms tensors to matrices: for $\mathcal{X} \in \mathbb{R}^{p_1\times p_2\times p_3}$, define

$$\begin{aligned}
&\mathcal{M}_1(\mathcal{X}) \in \mathbb{R}^{p_1\times p_2p_3} \quad \text{where } \big[\mathcal{M}_1(\mathcal{X})\big]_{i_1,i_2+p_2(i_3-1)} = \mathcal{X}_{i_1i_2i_3},\\
&\mathcal{M}_2(\mathcal{X}) \in \mathbb{R}^{p_2\times p_1p_3} \quad \text{where } \big[\mathcal{M}_2(\mathcal{X})\big]_{i_2,i_3+p_3(i_1-1)} = \mathcal{X}_{i_1i_2i_3},\\
&\mathcal{M}_3(\mathcal{X}) \in \mathbb{R}^{p_3\times p_1p_2} \quad \text{where } \big[\mathcal{M}_3(\mathcal{X})\big]_{i_3,i_1+p_1(i_2-1)} = \mathcal{X}_{i_1i_2i_3},
\end{aligned}$$

and $\mathcal{M}_k^{-1} : \mathbb{R}^{p_k\times(p_{k+1}p_{k+2})} \to \mathbb{R}^{p_1\times p_2\times p_3}$ as the inverse operator of $\mathcal{M}_k$ where $k+1$ and $k+2$ are computed modulo 3. Essentially, $\mathcal{M}_k$ "flattens" all but the kth directions of any tensor. The following identity that relates the matrix-tensor product and matricization plays an important role in our analysis:

$$\mathcal{M}_k(\mathcal{S}\times_1\mathbf{U}_1\times_2\mathbf{U}_2\times_3\mathbf{U}_3) = \mathbf{U}_k\mathcal{M}_k(\mathcal{S})(\mathbf{U}_{k+2}\otimes\mathbf{U}_{k+1})^\top, \quad k = 1, 2, 3.$$

Here again, $k+1$ and $k+2$ are computed modulo 3. The inner product of two tensors with the same dimension is defined as $\langle\mathcal{X}, \mathcal{Y}\rangle = \sum_{ijk}\mathcal{X}_{ijk}\mathcal{Y}_{ijk}$. The Frobenius norm of a tensor $\mathcal{X}$ is defined as $\|\mathcal{X}\|_{\mathrm{F}} = \sqrt{\sum_{i,j,k}\mathcal{X}_{ijk}^2}$. For any smooth tensor-variate function $f : \mathbb{R}^{p_1\times p_2\times p_3} \to$

$\mathbb{R}$, let $\nabla f : \mathbb{R}^{p_1\times p_2\times p_3} \to \mathbb{R}^{p_1\times p_2\times p_3}$ be the gradient function such that $(\nabla f(\mathcal{X}))_{ijk} = \frac{\partial f}{\partial \mathcal{X}_{ijk}}$. We simply write this as ∇f when there is no confusion. Finally, the readers are also referred to Kolda and Bader (2009) for a comprehensive discussions on tensor algebra. The focus of this paper is on the following low-Tucker-rank tensors.

DEFINITION 2.1 (Low Tucker rank). We say $\mathcal{X}^* \in \mathbb{R}^{p_1\times p_2\times p_3}$ is Tucker rank-(r_1, r_2, r_3) if and only if $\mathcal{X}^*$ can be decomposed as

$$\mathcal{X}^* = \mathcal{S}^* \times_1 \mathbf{U}_1^* \times_2 \mathbf{U}_2^* \times_3 \mathbf{U}_3^* =: [\![\mathcal{S}^*; \mathbf{U}_1^*, \mathbf{U}_2^*, \mathbf{U}_3^*]\!]$$

for some $\mathcal{S}^* \in \mathbb{R}^{r_1\times r_2\times r_3}$ and $\mathbf{U}_k^* \in \mathbb{R}^{p_k\times r_k}$, $k = 1, 2, 3$.

In addition, $\mathcal{X}^*$ is Tucker rank-(r_1, r_2, r_3) if and only if $\operatorname{rank}(\mathcal{M}_k(\mathcal{X}^*)) \le r_k$ for $k = 1, 2, 3$. For convenience of presentation, we denote $\bar{p} = \max\{p_1, p_2, p_3\}$, $\bar{r} = \max\{r_1, r_2, r_3\}$, $\underline{p} = \min\{p_1, p_2, p_3\}$, $\underline{r} = \min\{r_1, r_2, r_3\}$, $p_{-k} = p_1p_2p_3/p_k$ and $r_{-k} = r_1r_2r_3/r_{-k}$.

2.2. *Generalized tensor estimation.* Suppose we observe a dataset D associated with an unknown parameter $\mathcal{X}^*$. Here, $\mathcal{X}^*$ is a p_1-by-p_2-by-p_3 rank-(r_1, r_2, r_3) tensor and $r_k \ll p_k$. For example, D can be a random sample drawn from some distribution parametrized by $\mathcal{X}^*$ or a deterministic perturbation of $\mathcal{X}^*$. The central goal of this paper is to have an efficient and accurate estimation of $\mathcal{X}^*$.

Let $L(\mathcal{X}; D)$ be an empirical loss function known a priori, such as the negative log-likelihood function from the generating distribution or more general objective function. Then the following rank constrained optimization provides a straightforward way to estimate $\mathcal{X}^*$:

$$\min_{\mathcal{X}} L(\mathcal{X}; D) \quad \text{subject to} \quad \operatorname{rank}\big(\mathcal{M}_k(\mathcal{X})\big) \le r_k, \quad k = 1, 2, 3. \tag{2.1}$$

As mentioned earlier, this framework includes many instances arising from applications in various fields. Due to the connection between low Tucker rank and the decomposition discussed in Section (2.1), it is natural to consider the following minimization problem:

$$\begin{aligned} \hat{\mathcal{X}} &= \hat{\mathcal{S}} \times_1 \hat{\mathbf{U}}_1 \times_2 \hat{\mathbf{U}}_2 \times_3 \hat{\mathbf{U}}_3, \\ (\hat{\mathcal{S}}, \hat{\mathbf{U}}_1, \hat{\mathbf{U}}_2, \hat{\mathbf{U}}_3) &= \operatorname*{arg\,min}_{\mathcal{S}, \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3} L(\mathcal{S} \times_1 \mathbf{U}_1 \times_2 \mathbf{U}_2 \times_3 \mathbf{U}_3; D), \end{aligned} \tag{2.2}$$

and consider a gradient-based optimization algorithm to estimate $\mathcal{X}^*$. Let $\nabla L : \mathbb{R}^{p_1\times p_2\times p_3} \to \mathbb{R}^{p_1\times p_2\times p_3}$ be the gradient of loss function. The following lemma gives the partial gradients of L on $\mathbf{U}_k$ and $\mathcal{S}$. The proof is provided in the Supplementary Material (Appendix E.1).

LEMMA 2.1 (Partial gradients of loss).

$$\begin{aligned} \nabla_{\mathbf{U}_1} L([\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]) &= \mathcal{M}_1(\nabla L)(\mathbf{U}_3 \otimes \mathbf{U}_2)\mathcal{M}_1(\mathcal{S})^\top, \\ \nabla_{\mathbf{U}_2} L([\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]) &= \mathcal{M}_2(\nabla L)(\mathbf{U}_1 \otimes \mathbf{U}_3)\mathcal{M}_2(\mathcal{S})^\top, \\ \nabla_{\mathbf{U}_3} L([\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]) &= \mathcal{M}_3(\nabla L)(\mathbf{U}_2 \otimes \mathbf{U}_1)\mathcal{M}_3(\mathcal{S})^\top, \\ \nabla_{\mathcal{S}} L([\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]) &= \nabla L \times_1 \mathbf{U}_1^\top \times_2 \mathbf{U}_2^\top \times_3 \mathbf{U}_3^\top. \end{aligned} \tag{2.3}$$

Here, ∇L is short for $\nabla L([\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!])$.

As mentioned earlier, we consider optimizing the following objective function:

$$F(\mathcal{S}, \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3) = L([\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]; D) + \frac{a}{2}\sum_{k=1}^{3} \|\mathbf{U}_k^\top \mathbf{U}_k - b^2\mathbf{I}_{r_k}\|_{\mathrm{F}}^2, \tag{2.4}$$

where $a \geq 0, b > 0$ are tuning parameters to be discussed later. By adding regularizers $\|\mathbf{U}_k^\top \mathbf{U}_k - b^2 \mathbf{I}_{r_k}\|_{\mathrm{F}}^2$, we can prevent $\mathbf{U}_k$ from being singular throughout gradient descent, while do not alter the minimizer. This can be summarized as the following proposition, whose proof is provided in Appendix E.2.

PROPOSITION 2.1. *Suppose* $(\hat{\mathcal{S}}, \hat{\mathbf{U}}_1, \hat{\mathbf{U}}_2, \hat{\mathbf{U}}_2) = \arg\min F(\mathcal{S}, \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3)$ *for* F *defined in* (2.4). *Then*

$$\hat{\mathcal{X}} = [\![\hat{\mathcal{S}}, \hat{\mathbf{U}}_1, \hat{\mathbf{U}}_2, \hat{\mathbf{U}}_2]\!] = \mathop{\arg\min}_{\mathcal{X}:\mathrm{rank}(\mathcal{X}) \leq (r_1, r_2, r_3)} L(\mathcal{X}; D).$$

Similar regularizers have been widely used on nonconvex low-rank matrix optimization (Park et al. (2018), Tu et al. (2016)) and more technical interpretations are provided in Section 3.3.

3. Projected gradient descent. In this section, we study the local convergence of the projected gradient descent under a general deterministic framework.

3.1. *Restricted correlated gradient condition.* We first introduce the regularity condition on the loss function L and set $\mathcal{C}$.

DEFINITION 3.1 (Restricted Correlated Gradient (RCG)). Let f be a real-valued function. We say f satisfies RCG$(\alpha, \beta, \mathcal{C})$ condition for $\alpha, \beta > 0$ and the set $\mathcal{C}$ if

$$\langle \nabla f(x) - \nabla f(x^*), x - x^* \rangle \geq \alpha \|x - x^*\|_2^2 + \beta \|\nabla f(x) - \nabla f(x^*)\|_2^2 \tag{3.1}$$

for any $x \in \mathcal{C}$. Here, x^* is some fixed target parameter.

Our later analysis will be based on the assumption that L satisfies the RCG condition on to-be-specified sets of tensors with $x^* = \mathcal{X}^*$ being the true parameter tensor.

REMARK 3.1 (Interpretation of the RCG condition). The RCG condition is similar to the "regularity condition" appearing in recent nonconvex optimization literature (Candès, Li and Soltanolkotabi (2015), Chen and Candès (2015), Chi, Lu and Chen (2019), Yonel and Yazici (2020)):

$$\langle \nabla f(x), x - x^\# \rangle \geq \alpha \|x - x^\#\|_2^2 + \beta \|\nabla f(x)\|_2^2, \tag{3.2}$$

where $f(\cdot)$ is the objective function in their context and $x^\#$ is the minimizer of $f(\cdot)$. The RCG condition can be seen as a generalization of (3.2): in the deterministic case without statistical noise, the target x^* usually becomes an exact stationery point of $f(\cdot)$ and (3.1) reduces to (3.2). In addition, it is worthy noting that RCG condition does not require the function f to be convex since x^* is only a fixed target parameter in the requirement (3.2) (also see Figure 1 in Chi, Lu and Chen (2019), for an example).

3.2. *Theoretical analysis.* We now consider a general setting that the loss function L satisfies the RCG condition in a constrained domain:

$$\mathcal{C} = \{\mathcal{X} \in \mathbb{R}^{p_1 \times p_2 \times p_3} : \mathcal{X} = [\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!], \mathbf{U}_k \in \mathcal{C}_k, \mathcal{S} \in \mathcal{C}_{\mathcal{S}}\}, \tag{3.3}$$

where the true parameter tensor $\mathcal{X}^*$ is feasible—that is, $\mathcal{X}^* \in \mathcal{C}$. Here, $\mathcal{C}_k$ and $\mathcal{C}_{\mathcal{S}}$ are some convex and rotation invariant sets: for any $\mathbf{U}_k \in \mathcal{C}_k$, $\mathcal{S} \in \mathcal{C}_{\mathcal{S}}$, we have $\mathbf{U}_k \mathbf{R}_k \in \mathcal{C}_k$ and $[\![\mathcal{S}; \mathbf{R}_1, \mathbf{R}_2, \mathbf{R}_3]\!] \in \mathcal{C}_{\mathcal{S}}$ for arbitrary orthogonal matrices $\mathbf{R}_k \in \mathbb{O}_{r_k}$. Some specific problems of this general setting will be discussed in Section 4.

Algorithm 1 Projected Gradient Descent

Require: Initialization $(\mathcal{S}^{(0)}, \mathbf{U}_1^{(0)}, \mathbf{U}_2^{(0)}, \mathbf{U}_3^{(0)})$, constraint sets $\{\mathcal{C}_k\}_{k=1}^3$, $\mathcal{C}_\mathcal{S}$, tuning parameters $a, b > 0$, step size η.
  for all $t = 0$ to $T-1$ **do**
    for all $k = 1, 2, 3$ **do**
      $\tilde{\mathbf{U}}_k^{(t+1)} = \mathbf{U}_k^{(t)} - \eta(\nabla_{\mathbf{U}_k} L(\mathcal{S}^{(t)}, \mathbf{U}_1^{(t)}, \mathbf{U}_2^{(t)}, \mathbf{U}_3^{(t)}) + a\mathbf{U}_k^{(t)}(\mathbf{U}_k^{(t)\top}\mathbf{U}_k^{(t)} - b^2\mathbf{I}))$
      $\mathbf{U}_k^{(t+1)} = \mathcal{P}_{\mathcal{C}_k}(\tilde{\mathbf{U}}_k^{(t+1)})$, where $\mathcal{P}_{\mathcal{C}_k}(\cdot)$ is the projection onto $\mathcal{C}_k$.
    end for
    $\tilde{\mathcal{S}}^{(t+1)} = \mathcal{S}^{(t)} - \eta\nabla_\mathcal{S} L(\mathcal{S}^{(t)}, \mathbf{U}_1^{(t)}, \mathbf{U}_2^{(t)}, \mathbf{U}_3^{(t)})$
    $\mathcal{S}^{(t+1)} = \mathcal{P}_{\mathcal{C}_\mathcal{S}}(\tilde{\mathcal{S}}^{(t+1)})$, where $\mathcal{P}_{\mathcal{C}_\mathcal{S}}(\cdot)$ is the projection onto $\mathcal{C}_\mathcal{S}$.
  end for
  return $\mathcal{X}^{(T)} = \mathcal{S}^{(T)} \times_1 \mathbf{U}_1^{(T)} \times_2 \mathbf{U}_2^{(T)} \times_3 \mathbf{U}_3^{(T)}$

When L and $\mathcal{X}^*$ satisfy the condition above, we introduce the projected gradient descent in Algorithm 1. In addition to the vanilla gradient descent, the proposed Algorithm 1 includes multiple projection steps to ensure that $\mathcal{X}^*$ is in the regularized domain $\mathcal{C}$ throughout the iterations.

Suppose the true parameter $\mathcal{X}^*$ is of Tucker rank-(r_1, r_2, r_3). We also introduce the following value to quantify how different the $\mathcal{X}^*$ is from being a stationary point of $L(\mathcal{X}; D)$:

$$\xi := \sup_{\substack{\mathcal{T}\in\mathbb{R}^{p_1\times p_2\times p_3}, \|\mathcal{T}\|_{\mathrm{F}}\le 1 \\ \mathrm{rank}(\mathcal{T})\le(r_1,r_2,r_3)}} |\langle \nabla L(\mathcal{X}^*), \mathcal{T}\rangle|. \tag{3.4}$$

Intuitively speaking, ξ measures the amplitude of $\nabla L(\mathcal{X}^*)$ projected onto the manifold of low-rank tensors. In many statistical models, ξ essentially characterizes the amplitude of statistical noise. Specifically in the noiseless setting, $\mathcal{X}^*$ is exactly a stationary point of L, then $\nabla L(\mathcal{X}^*) = 0$, $\xi = 0$. In various probabilistic settings, a suitable L often satisfies $\mathbb{E}\nabla L(\mathcal{X}^*) = 0$; then ξ reflects the reduction of variance of $\nabla L(\mathcal{X}^*)$ after projection onto the low-rank tensor manifold. We also define

$$\overline{\lambda} := \max\{\|\mathcal{M}_1(\mathcal{X}^*)\|, \|\mathcal{M}_2(\mathcal{X}^*)\|, \|\mathcal{M}_3(\mathcal{X}^*)\|\},$$
$$\underline{\lambda} := \min\{\sigma_{r_1}(\mathcal{M}_1(\mathcal{X}^*)), \sigma_{r_2}(\mathcal{M}_2(\mathcal{X}^*)), \sigma_{r_3}(\mathcal{M}_3(\mathcal{X}^*))\},$$

and $\kappa = \overline{\lambda}/\underline{\lambda}$ can be regarded as a tensor condition number, as similarly defined for matrices. It is note worthy that the curvature of Tucker rank-(r_1, r_2, r_3) tensor manifold on $\mathcal{X}^*$ can be bounded by $\underline{\lambda}^{-1}$ Lubich et al. (2013), Lemma 4.5.

We are now in position to establish a deterministic upper bound on the estimation error and a linear rate of convergence for the proposed Algorithm 1 when a warm initialization is provided. Specific initialization algorithms for different applications will be discussed in Section 4.

THEOREM 3.1 (Local convergence). *Suppose L satisfies* RCG$(\alpha, \beta, \mathcal{C})$ *for $\mathcal{C}$ defined in (3.3) and $b \asymp \overline{\lambda}^{1/4}$, $a = \frac{4\alpha b^4}{3\kappa^2}$ in (2.4). Assume $\mathcal{X}^* = [\![\mathcal{S}^*; \mathbf{U}_1^*, \mathbf{U}_2^*, \mathbf{U}_3^*]\!]$ such that $\mathbf{U}_k^{*\top}\mathbf{U}_k^* = b^2\mathbf{I}_{r_k}, \mathbf{U}_k^* \in \mathcal{C}_k, k = 1, 2, 3$, and $\mathcal{S}^* \in \mathcal{C}_\mathcal{S}$. Suppose the initialization $\mathcal{X}^{(0)} = [\![\mathcal{S}^{(0)}; \mathbf{U}_1^{(0)}, \mathbf{U}_2^{(0)}, \mathbf{U}_3^{(0)}]\!]$ satisfies $\|\mathcal{X}^{(0)} - \mathcal{X}^*\|_{\mathrm{F}}^2 \le c\frac{\alpha\beta\underline{\lambda}^2}{\kappa^2}$ for some small constant $c > 0$, $\mathbf{U}_k^{(0)}\mathbf{U}_k^{(0)} = b^2\mathbf{I}_{r_k}, \mathbf{U}_k^{(0)} \in \mathcal{C}_k, k = 1, 2, 3$ and $\mathcal{S}^{(0)} \in \mathcal{C}_\mathcal{S}$. Also, the signal-noise-ratio satisfies $\underline{\lambda}^2 \ge C_0\frac{\kappa^4}{\alpha^3\beta}\xi^2$ for some universal constant C_0. Then there exists a constant $c > 0$ such that if*

$\eta = \frac{\eta_0 \beta}{b^6}$ *for* $\eta_0 \le c$, *we have*

$$\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 \le C\left(\frac{\kappa^4}{\alpha^2}\xi^2 + \kappa^2\left(1 - \frac{2\rho\alpha\beta\eta_0}{\kappa^2}\right)^t \|\mathcal{X}^{(0)} - \mathcal{X}^*\|_{\mathrm{F}}^2\right).$$

In addition, the following corollary provides a theoretical guarantee for the estimation loss of the proposed Algorithm 1 after a logarithmic number of iterations.

COROLLARY 3.1. *Suppose the conditions of Theorem 3.1 hold and* α, β, κ *are constants. Then after at most* $T = \Omega(\log(\|\mathcal{X}^{(0)} - \mathcal{X}^*\|_{\mathrm{F}}/\xi))$ *iterations and for a constant* C *that only relies on* $\alpha, \beta, \kappa > 0$, *we have*

$$\|\mathcal{X}^{(T)} - \mathcal{X}^*\|_{\mathrm{F}}^2 \le C\xi^2.$$

REMARK 3.2. When $\nabla L(\mathcal{X}^*) = 0$, that is, there is no statistical noise or perturbation, we have $\xi = 0$. In this case, Theorem 3.1 and Corollary 3.1 imply that the proposed algorithm converges to the true target parameter $\mathcal{X}^*$ at a linear rate:

$$\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 \le C\left(1 - \frac{2\rho\alpha\beta\eta_0}{\kappa^2}\right)^t \|\mathcal{X}^{(0)} - \mathcal{X}^*\|_{\mathrm{F}}^2.$$

When $\nabla L(\mathcal{X}^*) \neq 0$, we have $\xi > 0$ and $\mathcal{X}^*$ is not an exact stationary point of the loss function L. Then the estimation error $\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2$ is naturally not expected to go to zero, which matches the upper bounds of Theorem 3.1 and Corollary 3.1. In a statistical model where noise or perturbation is in presence, the upper bound on the estimation error can be determined by evaluating ξ under the specific random environment and these bounds are often minimax-optimal. See Section 4 for more detail.

REMARK 3.3. If $\mathcal{C}_{\mathcal{S}}$ and $\mathcal{C}_k$ are unbounded domains, then $\mathcal{C}$ is the set of all rank-(r_1, r_2, r_3) tensors, $\mathcal{P}_{\mathcal{C}_{\mathcal{S}}}, \mathcal{P}_{\mathcal{C}_k}$ are identity operators, and the proposed Algorithm 1 essentially becomes the vanilla gradient descent. When $\mathcal{C}_{\mathcal{S}}$ and $\mathcal{C}_k$ are nontrivial convex subsets, the projection steps ensure that $\mathcal{X}^{(t)} = [\![\mathcal{S}^{(t)}; \mathbf{U}_1^{(t)}, \mathbf{U}_2^{(t)}, \mathbf{U}_3^{(t)}]\!] \in \mathcal{C}$ and the RCG condition can be applied throughout the iterations. In fact, we found that the projection steps can be omitted in many numerical cases even if L does not satisfy the RCG condition for the full set of low-rank tensors, such as the forthcoming Poisson and binomial tensor PCA. See Sections 4 and 7 for more discussions.

3.3. *Proof sketch of main results.* We briefly discuss the idea for the proof of Theorem 3.1 here. The complete proof is provided in Appendix C. A key step in our analysis is to establish an error contraction inequality to characterize the estimation error of $\mathcal{X}^{(t+1)}$ based on the one of $\mathcal{X}^{(t)}$. Since the proposed nonconvex gradient descent is performed on $\mathcal{S}^{(t)}, \mathbf{U}_1^{(t)}, \mathbf{U}_2^{(t)}, \mathbf{U}_3^{(t)}$ jointly in lieu of $\mathcal{X}^{(t)}$ directly, it becomes technically difficult to develop a direct link between $\|\mathcal{X}^{(t+1)} - \mathcal{X}^*\|_{\mathrm{F}}^2$ and $\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2$. To overcome this difficulty, a "lifting" scheme was proposed and widely used in the recent literature on low-rank asymmetric matrix optimization (Tu et al. (2016), Zhu et al. (2017), Park et al. (2018)): one can factorize any rank-r matrix estimator $\mathbf{A}^{(t)}$ and the target matrix parameter $\mathbf{A}^*$ into $\mathbf{A}^{(t)} = \mathbf{U}^{(t)}(\mathbf{V}^{(t)})^\top, \mathbf{A}^* = \mathbf{U}^*(\mathbf{V}^*)^\top$, where $\mathbf{U}^{(t)}, \mathbf{V}^{(t)}$ (or $\mathbf{U}^*, \mathbf{V}^*$) both have r columns and share the same singular values. Then, one can stack them into one matrix

$$\mathbf{W}^{(t)} = \begin{bmatrix} \mathbf{U}^{(t)} \\ \mathbf{V}^{(t)} \end{bmatrix}, \qquad \mathbf{W}^* = \begin{bmatrix} \mathbf{U}^* \\ \mathbf{V}^* \end{bmatrix}.$$

By establishing the equivalence between $\min_{\mathbf{R}\in\mathbb{O}_r}\|\mathbf{W}^{(t)}-\mathbf{W}^*\mathbf{R}\|_{\mathrm{F}}^2$ and $\|\mathbf{A}^{(t)}-\mathbf{A}^*\|_{\mathrm{F}}^2$, and analyzing on $\min_{\mathbf{R}\in\mathbb{O}_r}\|\mathbf{W}^{(t)}-\mathbf{W}^*\mathbf{R}\|_{\mathrm{F}}^2$, a local convergence of $\mathbf{A}^{(t)}$ to $\mathbf{A}^*$ can be established. However, the "lifting" scheme is not applicable to the tensor problem here since $\mathbf{U}_1,\mathbf{U}_2,\mathbf{U}_3,\mathcal{S}$ have distinct shapes and cannot be simply stacked together. To overcome this technical issue in the generalized tensor estimation problems, we propose to assess the following criterion:

$$E^{(t)}=\min_{\substack{\mathbf{R}_k\in\mathbb{O}_{p_k,r_k}\\ k=1,2,3}}\left\{\sum_{k=1}^{3}\|\mathbf{U}_k^{(t)}-\mathbf{U}_k^*\mathbf{R}_k\|_{\mathrm{F}}^2+\|\mathcal{S}^{(t)}-[\![\mathcal{S}^*;\mathbf{R}_1^\top,\mathbf{R}_2^\top,\mathbf{R}_3^\top]\!]\|_{\mathrm{F}}^2\right\}. \tag{3.5}$$

Intuitively, $E^{(t)}$ measures the difference between a pair of tensor components $(\mathcal{S}^*,\mathbf{U}_1,\mathbf{U}_2,\mathbf{U}_3)$ and $(\mathcal{S}^{(t)},\mathbf{U}_1^{(t)},\mathbf{U}_2^{(t)},\mathbf{U}_3^{(t)})$ under rotation. The introduction of $E^{(t)}$ enables a convenient error contraction analysis as being an additive form of tensor components. In particular, the following lemma exhibits that $E^{(t)}$ is equivalent to the estimation error $\|\mathcal{X}^{(t)}-\mathcal{X}^*\|_{\mathrm{F}}^2$ under regularity conditions.

LEMMA 3.1 (An informal version of Lemma E.2).

$$cE^{(t)}\le\|\mathcal{X}^{(t)}-\mathcal{X}^*\|_{\mathrm{F}}^2+C\sum_{k=1}^{3}\|(\mathbf{U}_k^{(t)})^\top\mathbf{U}_k^{(t)}-\mathbf{U}_k^{*\top}\mathbf{U}_k^*\|_{\mathrm{F}}^2\le CE^{(t)}$$

under the regularity conditions to be specified in Lemma E.2.

Note that there is no equivalence between $E^{(t)}$ and $\|\mathcal{X}^{(t)}-\mathcal{X}^*\|_{\mathrm{F}}^2$ unless we force $\mathbf{U}_k$ and $\mathbf{U}_k^*$ have similar singular structures, and this is the reason why we introduce the regularizer term in (2.4) to keep $\mathbf{U}_k^{(t)}$ from being singular.

Based on Lemma 3.1, the proof of Theorem 3.1 reduces to establishing an error contraction inequality between $E^{(t)}$ and $E^{(t+1)}$:

$$E^{(t+1)}\le(1-\gamma)E^{(t)}+C\xi^2 \tag{3.6}$$

for constants $0<\gamma<1$ and $C>0$. Define the best rotation matrices

$$(\mathbf{R}_1^{(t)},\mathbf{R}_2^{(t)},\mathbf{R}_3^{(t)})=\operatorname*{arg\,min}_{\substack{\mathbf{R}_k\in\mathbb{O}_{p_k,r_k}\\ k=1,2,3}}\left\{\sum_{k=1}^{3}\|\mathbf{U}_k^{(t)}-\mathbf{U}_k^*\mathbf{R}_k\|_{\mathrm{F}}^2+\|\mathcal{S}^{(t)}-[\![\mathcal{S}^*;\mathbf{R}_1^\top,\mathbf{R}_2^\top,\mathbf{R}_3^\top]\!]\|_{\mathrm{F}}^2\right\}.$$

By plugging in the gradient of $L(\mathcal{X})$ and $\mathcal{X}=\mathcal{S}\times_1\mathbf{U}_1\times_2\mathbf{U}_2\times_3\mathbf{U}_3$, we can show

$$\begin{aligned}&\|\mathbf{U}_k^{(t+1)}-\mathbf{U}_k^*\mathbf{R}_k^{(t+1)}\|_{\mathrm{F}}^2-\|\mathbf{U}_k^{(t)}-\mathbf{U}_k^*\mathbf{R}_k^{(t)}\|_{\mathrm{F}}^2\\&\quad\approx-2\eta\langle\mathcal{X}^{(t)}-\mathcal{X}_k^{(t)},\nabla L(\mathcal{X}^{(t)})\rangle-\frac{a\eta}{2}\|\mathbf{U}_k^{(t)}\mathbf{U}_k^{(t)}-b^2\mathbf{I}_{r_k}\|_{\mathrm{F}}^2,\end{aligned} \tag{3.7}$$

$$\begin{aligned}&\|\mathcal{S}^{(t+1)}-[\![\mathcal{S}^*;\mathbf{R}_1^{(t+1)\top},\mathbf{R}_2^{(t+1)\top},\mathbf{R}_3^{(t+1)\top}]\!]\|_{\mathrm{F}}^2-\|\mathcal{S}^{(t)}-[\![\mathcal{S}^*;\mathbf{R}_1^{(t)\top},\mathbf{R}_2^{(t)\top},\mathbf{R}_3^{(t)\top}]\!]\|_{\mathrm{F}}^2\\&\quad\approx-2\eta\langle\mathcal{X}^{(t)}-\mathcal{X}_{\mathcal{S}}^{(t)},\nabla L(\mathcal{X}^{(t)})\rangle,\end{aligned} \tag{3.8}$$

where

$$\mathcal{X}_k^{(t)}:=\mathcal{S}^{(t)}\times_k(\mathbf{U}_k^*\mathbf{R}_k^{(t)})\times_{k+1}\mathbf{U}_{k+1}^{(t)}\times_{k+2}\mathbf{U}_{k+2}^{(t)},\quad k=1,2,3;\quad\text{and}$$

$$\mathcal{X}_{\mathcal{S}}^{(t)}=[\![\mathcal{S}^*;\mathbf{U}_1^{(t)}\mathbf{R}_1^{(t)\top},\mathbf{U}_2^{(t)}\mathbf{R}_2^{(t)\top},\mathbf{U}_3^{(t)}\mathbf{R}_3^{(t)\top}]\!].$$

Note $E^{(t+1)}-E^{(t)}$ corresponds to the summation of (3.7) and (3.8), whose right-hand sides are dominated by the inner product between $\nabla L(\mathcal{X}^{(t)})$ and $3\mathcal{X}^{(t)}-\sum_{k=1}^3\mathcal{X}_k^{(t)}-\mathcal{X}_{\mathcal{S}}^{(t)}$. We develop a new tensor perturbation lemma to characterize $3\mathcal{X}^{(t)}-\sum_{k=1}^3\mathcal{X}_k^{(t)}-\mathcal{X}_{\mathcal{S}}^{(t)}$.

LEMMA 3.2 (An informal version of Lemma E.3). *Under regularity conditions to be specified in Lemma E.*3, *we have*

$$\mathcal{X}^{(t)} - \mathcal{X}^* = (\mathcal{X}^{(t)} - \mathcal{X}_{\mathcal{S}}^{(t)}) + \sum_{k=1}^{3} (\mathcal{X}^{(t)} - \mathcal{X}_k^{(t)}) + \mathcal{H}_\varepsilon, \tag{3.9}$$

where $\mathcal{H}_\varepsilon$ *is some low-rank residual tensor with* $\|\mathcal{H}_\varepsilon\|_{\mathrm{F}}^2 = o((E^{(t)})^2)$.

Combining (3.7), (3.8) and Lemma 3.2, we can connect $E^{(t+1)}$ and $E^{(t)}$ as

$$E^{(t+1)} \approx E^{(t)} - 2\eta\langle \mathcal{X}^{(t)} - \mathcal{X}^* - \mathcal{H}_\varepsilon, \nabla L(\mathcal{X}^{(t)})\rangle - \frac{a\eta}{2}\sum_{k=1}^{3} \|\mathbf{U}_k^{(t)\top}\mathbf{U}_k^{(t)} - \mathbf{U}_k^{*\top}\mathbf{U}_k^*\|_{\mathrm{F}}^2. \tag{3.10}$$

Then, we introduce another decomposition

$$\begin{aligned} &\langle \mathcal{X}^{(t)} - \mathcal{X}^* - \mathcal{H}_\varepsilon, \nabla L(\mathcal{X}^{(t)})\rangle \\ &= \underbrace{\langle \mathcal{X}^{(t)} - \mathcal{X}^*, \nabla L(\mathcal{X}^{(t)}) - \nabla L(\mathcal{X}^*)\rangle}_{A_1} \\ &\quad - \underbrace{\langle \mathcal{H}_\varepsilon, \nabla L(\mathcal{X}^{(t)}) - \nabla L(\mathcal{X}^*)\rangle}_{A_2} + \underbrace{\langle \mathcal{X}^{(t)} - \mathcal{X}^* + \mathcal{H}_\varepsilon, \nabla L(\mathcal{X}^*)\rangle}_{A_3} \end{aligned} \tag{3.11}$$

The three terms can be bounded separately:

$$\begin{aligned} A_1 &\geq \alpha\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 + \beta\|\nabla L(\mathcal{X}^{(t)}) - \nabla L(\mathcal{X}^*)\|_{\mathrm{F}}^2, \\ |A_2| &\leq \frac{\beta}{2}\|\nabla L(\mathcal{X}^{(t)}) - \nabla L(\mathcal{X}^*)\|_{\mathrm{F}}^2 + \frac{1}{2\beta}\|\mathcal{H}_\varepsilon\|_{\mathrm{F}}^2, \\ |A_3| &\leq \sqrt{E^{(t)}}\cdot\xi \leq cE^{(t)} + C\xi^2. \end{aligned} \tag{3.12}$$

Here the first inequality comes from RCG condition; the second inequality comes from Cauchy–Schwarz inequality and the fact $ab \leq \frac{1}{2}(a^2 + b^2)$; and the last inequality utilizes the definition of ξ, as well as Lemma 3.2. Combining (3.11) and (3.12), we obtain

$$\begin{aligned} &\langle \mathcal{X}^{(t)} - \mathcal{X}^* - \mathcal{H}_\varepsilon, \nabla L(\mathcal{X}^{(t)})\rangle + \langle \mathcal{X}^{(t)} - \mathcal{X}^* + \mathcal{H}_\varepsilon, \nabla L(\mathcal{X}^*)\rangle \\ &\geq (\alpha\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 + \beta\|\nabla L(\mathcal{X}^{(t)}) - \nabla L(\mathcal{X}^*)\|_{\mathrm{F}}^2) \\ &\quad - \frac{\beta}{2}\|\nabla L(\mathcal{X}^{(t)}) - \nabla L(\mathcal{X}^*)\|_{\mathrm{F}}^2 - (cE^{(t)} + C\xi^2). \end{aligned} \tag{3.13}$$

Then by choosing a suitable step size η and applying (3.10) together with (3.13), one obtains

$$E^{(t+1)} \leq E^{(t)} + cE^{(t)} + C\xi^2 - c_2\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 - c_3\sum_{k=1}^{3}\|\mathbf{U}_k^{(t)\top}\mathbf{U}_k^{(t)} - \mathbf{U}_k^{*\top}\mathbf{U}_k^*\|_{\mathrm{F}}^2.$$

Applying the equivalence between $\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 + C\sum_k \|(\mathbf{U}_k^{(t)})^\top\mathbf{U}_k^{(t)} - \mathbf{U}_k^{*\top}\mathbf{U}_k^*\|_{\mathrm{F}}^2$ and $E^{(t)}$ (Lemma 3.1), we can obtain (3.6) and finish the proof of Theorem 3.1.

4. Applications of generalized tensor estimation. Next, we apply the deterministic result to a number of generalized tensor estimation problems, including sub-Gaussian tensor PCA, tensor regression, Poisson tensor PCA, and binomial tensor PCA to obtain the estimation error bound of (projected) gradient descent. In each case, Algorithm 1 is used with different initialization schemes specified by the problem settings. All the proofs are provided in Appendix D. In addition, the generalized tensor estimation framework covers many other problems. A nonexhaustive list is provided in the Introduction. See Section 8 for more discussions.

Algorithm 2 Initialization for Sub-Gaussian Tensor PCA

Require: $\mathcal{Y} \in \mathbb{R}^{p_1 \times p_2 \times p_3}$, Tucker rank (r_1, r_2, r_3), scaling parameter b.
$\tilde{\mathcal{X}} = \mathcal{Y}$, $\tilde{\mathbf{U}}_k = \text{HeteroPCA}_{r_k}(\mathcal{M}_k(\tilde{\mathcal{X}})\mathcal{M}_k(\tilde{\mathcal{X}})^\top)$ for $k = 1, 2, 3$
$\tilde{\mathcal{S}} = [\![\tilde{\mathcal{X}}; \tilde{\mathbf{U}}_1^\top, \tilde{\mathbf{U}}_2^\top, \tilde{\mathbf{U}}_3^\top]\!]$
$\mathcal{S}^{(0)} = \tilde{\mathcal{S}}/b^3$ $\mathbf{U}_k^{(0)} = b\tilde{\mathbf{U}}_k^{(0)}$ for $k = 1, 2, 3$
return $(\mathcal{S}^{(0)}, \mathbf{U}_1^{(0)}, \mathbf{U}_2^{(0)}, \mathbf{U}_3^{(0)})$

4.1. *Sub-Gaussian tensor PCA.* Suppose we observe $\mathcal{Y} \in \mathbb{R}^{p_1 \times p_2 \times p_3}$, where $\mathbb{E}\mathcal{Y} = \mathcal{X}^*$, $\mathcal{X}^*$ is Tucker low-rank, and $\{\mathcal{Y}_{ijk} - \mathcal{X}^*_{ijk}\}_{ijk}$ are independent and sub-Gaussian distributed. In literature, much attention has been focused on various setups related to this model, for example, $\mathcal{Y}_{ijk} - \mathcal{X}^*_{ijk}$ are i.i.d. Gaussian, $\mathcal{X}^*$ is sparse, symmetric, rank-1, or CP-low-rank, etc (Chen (2019), Lesieur, Krzakala and Zdeborová (2017), Montanari, Reichman and Zeitouni (2017), Perry, Wein and Bandeira (2020), Richard and Montanari (2014), Sun and Luo (2015), Zhang and Han (2019)). Particularly when $\{\mathcal{Y}_{ijk} - \mathcal{X}^*_{ijk}\}_{ijk}$ are i.i.d. Gaussian distributed, it has been shown that the higher-order orthogonal iteration (HOOI) (De Lathauwer, De Moor and Vandewalle (2000b)) achieves the optimal statistical performance on the estimation of $\mathcal{X}^*$ (Zhang and Xia (2018)). It is however unclear whether HOOI works in the more general heteroskedastic setting, where the entries of $\mathcal{Y}$ have different variances.

Departing from the existing methods, we consider the estimation of $\mathcal{X}^*$ via minimizing the quadratic loss function $L(\mathcal{X}) = \frac{1}{2}\|\mathcal{X} - \mathcal{Y}\|_{\text{F}}^2$ using gradient descent. It is easy to check that L satisfies $\text{RCG}(\frac{1}{2}, \frac{1}{2}, \mathbb{R}^{p_1 \times p_2 \times p_3})$, so the projection steps in Algorithm 1 can be skipped throughout the iterations. To accommodate possible heteroskedastic noise, we apply HeteroPCA (Zhang, Cai and Wu (2018)), an iterative algorithm for PCA when heteroskedastic noise appears instead of the regular PCA for initialization. (The implementation of HeteroPCA in Algorithm 2 is provided in Appendix A).

Now we can establish the theoretical guarantee of Algorithms 1 and 2 for sub-Gaussian tensor PCA based on the deterministic result in Theorem 3.1.

THEOREM 4.1. *Suppose we observe* $\mathcal{Y} \in \mathbb{R}^{p_1 \times p_2 \times p_3}$, *where* $\mathbb{E}\mathcal{Y} = \mathcal{X}^* = [\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]$ *is Tucker rank-*(r_1, r_2, r_3), $\mathbf{U}_k \in \mathbb{O}_{p_k, r_k}$ *and* $\|\mathbf{U}_k\|_{2,\infty} \le c$ *for some constant* $c > 0$. *Suppose all entries of* $\mathcal{Y} - \mathcal{X}^*$ *are independent mean-zero sub-Gaussian random variables such that*

$$\|\mathcal{Y}_{ijk} - \mathcal{X}_{ijk}\|_{\psi_2} = \sup_{q \ge 1} \mathbb{E}(|\mathcal{Z}_{ijk}|^q)^{1/q}/q^{1/2} \le \sigma.$$

Assume $\underline{\lambda}/\sigma \ge C_1 \overline{p}^{3/4}\overline{r}^{1/4}$. *Then with probability at least* $1 - \exp(c\overline{p})$, *Algorithms 1 and 2 yield*

$$\|\hat{\mathcal{X}} - \mathcal{X}^*\|_{\text{F}}^2 \le C_2\sigma^2\left(r_1r_2r_3 + \sum_{k=1}^{3} p_k r_k\right), \tag{4.1}$$

where C_1, C_2 *are constants that do not depend on* p_k *and* r_k.

REMARK 4.1. The proposed method turns out to achieve the minimax optimal rate of estimation error in a general class of sub-Gaussian tensor PCA settings since the order of upper bound (4.1) matches the lower bound in literature (Zhang and Xia (2018), Theorem 3). Moreover, the condition $\underline{\lambda}/\sigma \ge C_1\overline{p}^{3/4}\overline{r}^{1/4}$ is optimal w.r.t. $\bar{p}$ in the sense that all the polynomial-time feasible algorithms cannot achieve consistent estimation when $\underline{\lambda}/\sigma < \bar{p}^{3/4-\varepsilon}$ for any $\varepsilon > 0$ (Zhang and Xia (2018), Theorem 4).

Algorithm 3 Initialization of Low-rank Tensor Regression

Require: $\{\mathcal{A}_i, y_i\}$, $i = 1, \ldots, n$, rank (r_1, r_2, r_3), scaling parameter b.
$\tilde{\mathcal{X}} = \frac{1}{n}\sum y_i \mathcal{A}_i$
$(\tilde{\mathcal{S}}, \tilde{\mathbf{U}}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3) = \mathrm{HOSVD}(\tilde{\mathcal{X}})$ or $(\tilde{\mathcal{S}}, \tilde{\mathbf{U}}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3) = \mathrm{HOOI}(\tilde{\mathcal{X}})$
$\mathbf{U}_k^{(0)} = b\tilde{\mathbf{U}}_k$, for $k = 1, 2, 3$
$\mathcal{S}^{(0)} = \tilde{\mathcal{S}}/b^3$
return $(\mathcal{S}^{(0)}, \mathbf{U}_1^{(0)}, \mathbf{U}_2^{(0)}, \mathbf{U}_3^{(0)})$

4.2. *Low-rank tensor regression*. Motivated by applications of neuroimaging analysis (Guhaniyogi, Qamar and Dunson (2017), Li and Zhang (2017), Zhou, Li and Zhu (2013)), spatio-temporal forecasting (Bahadori, Yu and Liu (2014)), high-order interaction pursuit (Hao, Zhang and Cheng (2020)), longitudinal relational data analysis (Hoff (2015)), 3D imaging processing (Guo, Kotsia and Patras (2012)), among many others, we consider the low-rank tensor regression next. Suppose we observe a collection of data $D = \{y_i, \mathcal{A}_i\}_{i=1}^n$ that are associated through the following equation:

$$y_i = \langle \mathcal{A}_i, \mathcal{X}^* \rangle + \varepsilon_i, \qquad \varepsilon_i \overset{\text{i.i.d.}}{\sim} N(0, \sigma^2), \quad i = 1, \ldots, n. \tag{4.2}$$

By exploiting the negative log-likelihood, it is natural to set L to be the squared loss function

$$L(\mathcal{X}; D) = \sum_{i=1}^n (\langle \mathcal{A}_i, \mathcal{X} \rangle - y_i)^2.$$

To estimate $\mathcal{X}^*$, we first perform spectral method (Algorithm 3) to obtain initializer $\mathcal{X}^{(0)} = [\![\mathcal{S}^{(0)}, \mathbf{U}_1^{(0)}, \mathbf{U}_2^{(0)}, \mathbf{U}_3^{(0)}]\!]$, then perform the gradient descent (Algorithm 1) without the projection steps to obtain the final estimator $\hat{\mathcal{X}}$. A key step of Algorithm 3 is HOSVD or HOOI, which are described in detail in Appendix A.

For technical convenience, we assume the covariates $\{\mathcal{A}_i\}_{i=1}^n$ are randomly designed that all entries of $\mathcal{A}_i$ are i.i.d. drawn from sub-Gaussian distribution with mean 0 and variance 1. The following theorem gives an estimation error upper bound for Algorithms 1 and 3.

THEOREM 4.2. *Consider the low-rank tensor regression model* (4.2). *Suppose* $\sigma^2 \le C_1 \|\mathcal{X}^*\|_{\mathrm{F}}^2$, $\underline{\lambda} \ge C_2$, *and the sample size* $n \ge C_3 \max\{\overline{p}^{3/2}\overline{r}, \overline{p} \cdot \overline{r}^2, \overline{r}^4\}$ *for constants* $C_1, C_2, C_3 > 0$. *Then with probability at least* $1 - \exp\{-c(r_1r_2r_3 + \sum_{k=1}^3 p_k r_k)\}$, *the output of Algorithms 1 and 3 satisfies*

$$\|\hat{\mathcal{X}} - \mathcal{X}^*\|_{\mathrm{F}}^2 \le C\sigma^2 \left(r_1r_2r_3 + \sum_{k=1}^3 p_k r_k\right)/n,$$

where C_1, C_2, C_3, C *are constants depending only on* κ *and* $c > 0$ *is a universal constant*.

Theorem 4.2 together with the lower bound in Zhang et al. (2020), Theorem 5, shows that the proposed procedure achieves the minimax optimal rate of estimation error in the class of all p_1-by-p_2-by-p_3 tensors with rank-(r_1, r_2, r_3) for tensor regression.

REMARK 4.2. The assumption on the covariates $\{\mathcal{A}_i\}_{i=1}^n$ in Theorem 4.2 ensures that $\tilde{\mathcal{X}}$ in Algorithm 3 is an unbiased estimator of $\mathcal{X}^*$. Such a setting has been considered as a benchmark setting in the high-dimensional statistical inference literature (see, e.g., Candès and Plan (2011), Chen, Raskutti and Yuan (2019), Javanmard and Montanari (2018)). When

$\{\mathcal{A}_i\}_{i=1}^n$ are heteroskedastic, the spectral initialization may fail and an alternative idea is the following unfolded nuclear norm minimization:

$$\tilde{\mathcal{X}}' = \mathop{\arg\min}_{\mathcal{X}\in\mathbb{R}^{p_1\times p_2\times p_3}} \sum_{k=1}^{3} \|\mathcal{M}_k(\mathcal{X})\|_*, \quad \text{s.t. } \mathcal{A}(\mathcal{X}) = y. \tag{4.3}$$

(4.3) is equivalent to a semidefinite programming and can be solved by the interior-point method (Gandy, Recht and Yamada (2011)).

REMARK 4.3. There is a significant gap between the required sample size in Theorem 4.2 ($O(p^{3/2}r)$) and the possible sample size lower bound, that is, the degree of freedom of all rank-r dimension-p tensor ($O(r^3 + pr)$). The existing algorithms achieving the sample size lower bound are often NP-hard to compute and thus intractable in practice. We also note that the existence of a tractable algorithm for tensor completion that provably works with less than $p^{3/2-\varepsilon}$ measurements would disapprove an open conjecture in theoretical computer science on strongly random 3-SAT (Barak and Moitra (2016), Corollary 16). Since tensor completion can be seen as a special case of tensor recovery, this suggests that it may be impossible to substantially improve the sample complexity required in Theorem 4.2 using a polynomial-time algorithm. Therefore, our procedure can be taken as the first computationally efficient algorithm to achieve minimax optimal rate of convergence and exact recovery in the noiseless setting as illustrated in Table 1.

4.3. *Poisson tensor PCA*. Tensor data with count values commonly arise from various scientific applications, such as the photon-limited imaging (Salmon et al. (2014), Timmermann and Nowak (1999), Willett and Nowak (2007), Yankovich et al. (2016)), online click-through data analysis (Shan et al. (2016), Sun and Li (2017)), and metagenomic sequencing (Flores et al. (2014)). In this section, we consider the Poisson tensor PCA model: assume we observe $\mathcal{Y} \in \mathbb{N}^{p_1\times p_2\times p_3}$ that satisfies

$$\mathcal{Y}_{ijk} \sim \text{Poisson}\big(I\exp(\mathcal{X}^*_{ijk})\big) \quad \text{independently}, \tag{4.4}$$

where $\mathcal{X}^*$ is the low-rank tensor parameter and $I > 0$ is the intensity parameter. When $\mathcal{X}^*$ is entry-wise bounded (Assumption 4.1), one can set I as the average intensity of all entries of $\mathcal{Y}$ so that I essentially quantifies the signal-to-noise ratio. Rather than estimating $I\exp(\mathcal{X}^*)$, we focus on estimating $\mathcal{X}^*$, the key tensor that captures the salient geometry or structure of the data.

Then, the following negative log-likelihood is a natural choice of the loss function for estimating $\mathcal{X}^*$:

$$L(\mathcal{X}) = \sum_{i=1}^{p_1}\sum_{j=1}^{p_2}\sum_{k=1}^{p_3}\big(-\mathcal{Y}_{ijk}\mathcal{X}_{ijk} + I\exp(\mathcal{X}_{ijk})\big). \tag{4.5}$$

Unfortunately, $L(\mathcal{X})$ defined in (4.5) satisfies RCG$(\alpha, \beta, \mathcal{C})$ only for a bounded set $\mathcal{C}$ since the Poisson likelihood function is not strongly convex and smooth in the unbounded domain. We thus introduce the following assumption on $\mathcal{X}^*$ to ensure that $\mathcal{X}^*$ is in a bounded set $\mathcal{C}$.

ASSUMPTION 4.1. Suppose $\mathcal{X}^* = [\![\mathcal{S}^*; \mathbf{U}_1^*, \mathbf{U}_2^*, \mathbf{U}_3^*]\!]$, where $\mathbf{U}_k^* \in \mathbb{O}_{p_k,r_k}$ is a p_k-by-r_k orthogonal matrix for $k = 1, 2, 3$. There exist some constants $\{\mu_k\}_{k=1}^3, B$ such that $\|\mathbf{U}_k^*\|_{2,\infty}^2 \le \frac{\mu_k r_k}{p_k}$ for $k = 1, 2, 3$ and $\overline{\lambda} \le B\sqrt{\frac{\Pi_{k=1}^3 p_k}{\Pi_{k=1}^3 \mu_k r_k}}$ where $\overline{\lambda} := \max_k \|\mathcal{M}_k(\mathcal{S}^*)\|$. Here, $\|\mathbf{U}_k^*\|_{2,\infty} = \max_i \|(\mathbf{U}_k^*)_{i\cdot}\|_2$ is the largest row-wise ℓ_2 norm of $\mathbf{U}_k^*$.

Algorithm 4 Initialization for Poisson Tensor PCA

Require: Initialization observation tensor $\mathcal{Y} \in \mathbb{N}^{p_1 \times p_2 \times p_3}$, Tucker rank (r_1, r_2, r_3), scaling parameter b, intensity parameter I.

$\tilde{\mathcal{X}} = \log((\mathcal{Y}_{jkl} + \frac{1}{2})/I)$

$(\tilde{\mathcal{S}}, \tilde{\mathbf{U}}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3) = \mathrm{HOSVD}(\tilde{\mathcal{X}})$ or $(\tilde{\mathcal{S}}, \tilde{\mathbf{U}}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3) = \mathrm{HOOI}(\tilde{\mathcal{X}})$

$\mathbf{U}_k^{(0)} = b\tilde{\mathbf{U}}_k$, for $k = 1, 2, 3$

$\mathcal{S}^{(0)} = \tilde{\mathcal{S}}/b^3$

return $(\mathcal{S}^{(0)}, \mathbf{U}_1^{(0)}, \mathbf{U}_2^{(0)}, \mathbf{U}_3^{(0)})$

Assumption 4.1 requires that the loading $\mathbf{U}_k$ satisfies the incoherence condition, that is, the amplitude of the tensor is "balanced" in all parts. Previously, the incoherence condition and its variations were commonly used in the matrix estimation literature (Candès and Recht (2009), Ma and Ma (2017)) and Poisson-type inverse problems (e.g., Poisson sparse regression Jiang, Raskutti and Willett (2015), Assumption 2.1, Poisson matrix completion Cao and Xie (2016), Equation (10), compositional matrix estimation Cao, Zhang and Li (2020), Equation (7), Poisson auto-regressive models (Hall, Raskutti and Willett (2016))). Assumption 4.1 also requires an upper bound on the spectral norm of each matricization of the core tensor $\mathcal{S}^*$. Together with the incoherence condition on $\mathbf{U}_k^*$, this condition guarantees that $\mathcal{X}^*$ is entry-wise upper bounded by B. In fact, the entry-wise bounded assumption is also widely used in high-dimensional matrix/tensor generalized linear models since it guarantees the local strong convexity and smoothness of the negative log-likelihood function (Ma and Ma (2017), Wang and Li (2020), Xu, Hu and Wang (2019)).

Next, we set $(\{\mathcal{C}_k\}_{k=1}^3, \mathcal{C}_{\mathcal{S}})$ as follows:

$$\mathcal{C}_k = \left\{ \mathbf{U}_k \in \mathbb{R}^{p_k \times r_k} : \|\mathbf{U}_k\|_{2,\infty} \le b\sqrt{\frac{\mu_k r_k}{p_k}} \right\},$$

$$\mathcal{C}_{\mathcal{S}} = \left\{ \mathcal{S} \in \mathbb{R}^{r_1 \times r_2 \times r_3} : \max_k \|\mathcal{M}_k(\mathcal{S})\| \le b^{-3} B \sqrt{\frac{\Pi_{k=1}^3 p_k}{\Pi_{k=1}^3 \mu_k r_k}} \right\}. \tag{4.6}$$

Specifically for the Poisson tensor PCA, we can prove that if Assumption 4.1 holds, the loss function (4.5) satisfies $\mathrm{RCG}(\alpha, \beta, \{\mathcal{C}_k\}_{k=1}^3, \mathcal{C}_{\mathcal{S}})$ for constants α, β that only depend on I and B (see the proof of Theorem 4.3 for details). We can also show that the following Algorithm 4 provides a sufficiently good initialization with high probability.

Now we establish the estimation error upper bound for Algorithms 1 and 4.

THEOREM 4.3. *Suppose Assumption 4.1 holds and* $I > C_1 \max\{\bar{p}, \underline{\lambda}^{-2} \sum_{k=1}^3 (p_{-k} r_k + p_k r_k)\}$, *where* $p_{-k} := p_1 p_2 p_3 / p_k$. *Then with probability at least* $1 - c/(p_1 p_2 p_3)$, *the output of Algorithms 1 and 4 yields*

$$\|\hat{\mathcal{X}} - \mathcal{X}^*\|_{\mathrm{F}}^2 \le C_2 I^{-1} \left(r_1 r_2 r_3 + \sum_{k=1}^3 p_k r_k \right).$$

Here C_1, C_2 *are constants that do not depend on* p_k *or* r_k.

We further consider the following class of low-rank tensors $\mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}$, where the restrictions in $\mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}$ correspond to the conditions in Theorem 4.3:

$$\mathcal{F}_{\boldsymbol{p},\boldsymbol{r}} = \left\{ \mathcal{X} = [\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!] : \begin{array}{l} \mathbf{U}_k \in \mathbb{O}_{p_k, r_k}, \|\mathbf{U}_k\|_{2,\infty}^2 \le \frac{\mu_k r_k}{p_k}, \\ \max_k \|\mathcal{M}_k(\mathcal{S})\| \le B\sqrt{\frac{\Pi_{k=1}^3 p_k}{\Pi_{k=1}^3 \mu_k r_k}} \end{array} \right\}. \tag{4.7}$$

With some technical conditions on tensor rank and the intensity parameter, we can develop the following lower bound in estimation error for Poisson PCA.

THEOREM 4.4 (Lower Bound for Poisson tensor PCA). *Assume* $\bar{r} \le C_1 \underline{p}^{1/2}$, $\underline{r} > C_2$ *and* $\min_k \mu_k \ge C_3$ *for constants* $C_1, C_2, C_3 > 1$. *Suppose one observes* $\mathcal{Y} \in \mathbb{R}^{p_1 \times p_2 \times p_3}$, *where* $\mathcal{Y}_{jkl} \sim \text{Poisson}(I \exp(\mathcal{X}_{jkl}))$ *independently,* $\mathcal{X} \in \mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}$, *and* $I \ge c_0$. *There exists a uniform constant* c *that does not depend on* p_k *or* r_k, *such that*

$$\inf_{\hat{\mathcal{X}}} \sup_{\mathcal{X} \in \mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}} \mathbb{E} \|\hat{\mathcal{X}} - \mathcal{X}\|_{\mathrm{F}}^2 \ge c I^{-1} \left(r_1 r_2 r_3 + \sum_{k=1}^{3} p_k r_k \right).$$

Theorems 4.3 and 4.4 together yield the optimal rate of estimation error for Poisson tensor PCA problem over the class of $\mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}$:

$$\inf_{\hat{\mathcal{X}}} \sup_{\mathcal{X} \in \mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}} \mathbb{E} \|\hat{\mathcal{X}} - \mathcal{X}\|_{\mathrm{F}}^2 \asymp I^{-1} \left(r_1 r_2 r_3 + \sum_{k=1}^{3} p_k r_k \right).$$

4.4. *Binomial tensor PCA.* The binomial tensor data commonly arise in the analysis of proportion when raw counts are available. For example, in the Human Mortality Database (Wilmoth and Shkolnikov (2006)), the number of deaths and the total number of population are summarized into a three-way tensor, where the x-, y-, z-coordinates are counties, ages, and years, respectively. Given the sufficiently large number of population in each country, one can generally assume that each entry of this data tensor satisfies the binomial distribution independently.

Suppose we observe a count tensor $\mathcal{Y} \in \mathbb{N}^{p_1 \times p_2 \times p_3}$ and a total population tensor $\mathcal{N} \in \mathbb{N}^{p_1 \times p_2 \times p_3}$ such that $\mathcal{Y}_{jkl} \sim \text{Binomial}(\mathcal{N}_{jkl}, \mathcal{P}^*_{jkl})$ independently. Here, $\mathcal{P}^* \in [0,1]^{p_1 \times p_2 \times p_3}$ is a probability tensor linked to an underlying latent parameter $\mathcal{X}^* \in \mathbb{R}^{p_1 \times p_2 \times p_3}$ through $\mathcal{P}^*_{jkl} = s(\mathcal{X}^*_{jkl})$, where $s(x) = 1/(1+e^{-x})$ is the sigmoid function. Our goal is to estimate $\mathcal{X}^*$. To this end, we consider to minimize the following loss function:

$$L(\mathcal{X}) = -\sum_{jkl} (\hat{\mathcal{P}}_{jkl} \mathcal{X}_{jkl} + \log(1 - \sigma(\mathcal{X}_{jkl}))),$$

where $\hat{\mathcal{P}}_{jkl} := \mathcal{Y}_{jkl}/\mathcal{N}_{jkl}$.

We assume $\mathcal{X}^*$ satisfies Assumption 4.1 for the same reasons as in Poisson tensor PCA. We propose to estimate $\mathcal{P}^*$ by applying Algorithm 5 (initialization) and Algorithm 1 (projected gradient descent) with the following constraint sets:

$$\mathcal{C}_k = \left\{ \mathbf{U}_k \in \mathbb{R}^{p_k \times r_k} : \|\mathbf{U}_k\|_{2,\infty} \le b \sqrt{\frac{\mu_k r_k}{p_k}} \right\},$$

$$\mathcal{C}_{\mathcal{S}} = \left\{ \mathcal{S} \in \mathbb{R}^{r_1 \times r_2 \times r_3} : \max_k \|\mathcal{M}_k(\mathcal{S})\| \le b^{-3} B \sqrt{\frac{\Pi_{k=1}^3 p_k}{\Pi_{k=1}^3 \mu_k r_k}} \right\}.$$

We have the following theoretical guarantee for the estimator obtained by Algorithms 1 and 5 in binomial tensor PCA.

THEOREM 4.5 (Upper Bound for Binomial Tensor PCA). *Suppose Assumption 4.1 is satisfied and* $N = \min_{jkl} \mathcal{N}_{jkl}$ *satisfies* $N \ge C_1 \max\{\overline{p}, \underline{\lambda}^{-2} \sum_k (p_{-k} r_k + p_k r_k)\}$. *Then with*

Algorithm 5 Initialization for Binomial Tensor PCA

Require: $\mathcal{Y}, \mathcal{N} \in \mathbb{N}^{p_1 \times p_2 \times p_3}$, Tucker rank (r_1, r_2, r_3), scaling parameter b

$\tilde{\mathcal{X}}_{jkl} = \log(\frac{\mathcal{Y}_{jkl}+1/2}{\mathcal{N}_{jkl}-\mathcal{Y}_{jkl}+1/2}),\ \forall j,k,l$

$(\tilde{\mathcal{S}}, \tilde{\mathbf{U}}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3) = \mathrm{HOSVD}(\tilde{\mathcal{X}})$ or $(\tilde{\mathcal{S}}, \tilde{\mathbf{U}}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3) = \mathrm{HOOI}(\tilde{\mathcal{X}})$

$\mathbf{U}_k^{(0)} = b\tilde{\mathbf{U}}_k$, for $k = 1, 2, 3$

$\mathcal{S}^{(0)} = \tilde{\mathcal{S}}/b^3$

return $(\mathcal{S}^{(0)}, \mathbf{U}_1^{(0)}, \mathbf{U}_2^{(0)}, \mathbf{U}_3^{(0)})$

probability at least $1 - c/(p_1p_2p_3)$, *we have the following estimation upper bound for the output of Algorithms 1 and 5*:

$$\|\hat{\mathcal{X}} - \mathcal{X}^*\|_{\mathrm{F}}^2 \le C_2 N^{-1}\left(r_1r_2r_3 + \sum_{k=1}^{3} p_k r_k\right).$$

Here, C_1, C_2 *are some absolute constants that do not depend on* p_k *or* r_k.

REMARK 4.4. We assume $N \ge C_1 \max\{\overline{p}, \underline{\lambda}^{-2}\sum_k (p_{-k}r_k + p_k r_k)\}$ in Theorem 4.5 as a technical condition to prove the estimation error upper bound of $\hat{\mathcal{X}}$. N here essentially characterizes the signal-noise ratio of the binomial tensor PCA problem.

Let $\mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}$ be the class of low-rank tensors defined in (4.7). We can prove the following lower bound result, which establishes the minimax optimality of the proposed procedure over the class of $\mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}$ in binomial tensor PCA.

THEOREM 4.6 (Lower Bound for Binomial Tensor PCA). *Denote* $N = \min_{jkl} \mathcal{N}_{jkl}$. *Assume* $\bar{r} \le C_1 \underline{p}^{1/2}$, $\underline{r} > C_2$ *and* $\min_k \mu_k \ge C_3$ *for some constants* $C_1, C_2, C_3 > 1$. *Suppose one observes* $\mathcal{Y} \in \mathbb{N}^{p_1 \times p_2 \times p_3}$, *where* $\mathcal{Y}_{jkl} \sim \mathrm{Binomial}(\mathcal{N}_{jkl}, \sigma(\mathcal{X}_{jkl}))$ *independently,* $\mathcal{X}^* \in \mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}$, *and* $\max_{jkl} \mathcal{N}_{jkl} \le C \min_{jkl} \mathcal{N}_{jkl}$. *There exists constant* c *that does not depend on* p_k *or* r_k, *such that*

$$\inf_{\hat{\mathcal{X}}} \sup_{\mathcal{X} \in \mathcal{F}_{\boldsymbol{p},\boldsymbol{r}}} \|\hat{\mathcal{X}} - \mathcal{X}\|_{\mathrm{F}}^2 \ge cN^{-1}\left(r_1r_2r_3 + \sum_{k=1}^{3} p_k r_k\right).$$

5. Rank selection. The tensor rank (r_1, r_2, r_3) is required as an input to Algorithm 1 and plays a crucial role in the proposed nonconvex optimization framework. While various empirical methods have been proposed for rank selection in specific applications of low-rank tensor estimation (e.g., Yokota, Lee and Cichocki (2017)), there is a paucity of theoretical guarantees in the literature. In this section, we provide a rank estimation procedure with provable guarantees. Recall that in each application in Section 4, we first specify the initialization $\mathcal{X}^{(0)} = [\![\mathcal{S}^{(0)}; \mathbf{U}_1^{(0)}, \mathbf{U}_2^{(0)}, \mathbf{U}_3^{(0)}]\!]$ based on a spectral algorithm on some preliminary tensor $\tilde{\mathcal{X}}$ (see their definitions in Algorithms 2–5). Since $\tilde{\mathcal{X}}$ reflects the target tensor $\mathcal{X}^*$ and $\sigma_s(\mathcal{M}_k(\mathcal{X}^*)) = 0$ for $s \ge r_k + 1$, we consider the following rank selection method by exploiting the singular values of $\tilde{\mathcal{X}}$:

$$\hat{r}_k = \max\{r : \sigma_r(\mathcal{M}_k(\tilde{\mathcal{X}})) \ge t_k\}, \quad k = 1, 2, 3. \tag{5.1}$$

Here, $t_k > 0$ is the thresholding level whose value depends on specific problem settings. Next, we specifically consider the sub-Gaussian tensor PCA and tensor regression.

PROPOSITION 5.1. *Suppose* $\operatorname{rank}(\mathcal{X}^*) = (r_1, r_2, r_3)$, $r_k = o(p_k)$, $p_k = o(p_{-k})$. *Define* $\hat{r}_k$ *as* (5.1) *and*

$$\delta_k := \operatorname{Median}\{\sigma_1(\mathcal{M}_k(\tilde{\mathcal{X}})), \ldots, \sigma_{p_k}(\mathcal{M}_k(\tilde{\mathcal{X}}))\}. \tag{5.2}$$

(a) *In sub-Gaussian tensor PCA* (*Section* 4.1), *we set* $\tilde{\mathcal{X}} = \mathcal{Y}$, $t_k = 1.5\delta_k$. *Suppose* $\underline{\lambda} \geq C\overline{p}\sigma$. *Then, we have* $\hat{r}_k = r_k, \forall k \in [d]$ *with probability at least* $1 - Ce^{-c\bar{p}}$.

(b) *In low-rank tensor regression* (*Section* 4.2), *we set* $\tilde{\mathcal{X}} = \frac{1}{n}\sum_{i=1}^n y_i \mathcal{A}_i$, $t_k = 1.5\delta_k$. *Suppose* $n \geq C\kappa\bar{p}^2\bar{r}^{3/2}\log\bar{r}$ *and* $\underline{\lambda} \geq C\frac{\bar{p}\sigma}{\sqrt{n}}$. *Then,* $\hat{r}_k = r_k, \forall k \in [d]$ *with probability at least* $1 - Ce^{-c\bar{p}}$.

In practice, we can also apply a simple criterion of the *cumulative percentage of total variation* (Jolliffe (1986), Chapter 6.1.1) originating from principle component analysis:

$$\hat{r}_k = \arg\min\left\{r : \sum_{i=1}^r \sigma_i^2(\mathcal{M}_k(\tilde{\mathcal{X}}))/\sum_{i=1}^{p_k}\sigma_i^2(\mathcal{M}_k(\tilde{\mathcal{X}})) \geq \rho\right\} \tag{5.3}$$

Here, $\rho \in (0, 1)$ is some empirical thresholding level. We will illustrate this principle on real data analysis in Section 7.2. Under the general deterministic setting, the accurate (or optimal) estimation of tensor rank may be much more challenging and we leave it as future work.

6. Extensions to general order-*d* tensors. While our previous sections mainly focus on order-3 tensor estimation, our results can be generalized to the order-d low-rank tensor estimation with the key ideas outline in this section. First, the constraint set $\mathcal{C}$, RCG condition and noise quantity ξ can be defined similarly by replacing the order-3 tensor with the general low-rank order-d tensors; second, the local convergence analysis can be similarly conducted as Theorem 3.1. Define

$$E^{(t)} := \min_{\substack{\mathbf{R}_k \in \mathbb{O}_{p_k, r_k} \\ k=1,\ldots,d}} \left\{\sum_{k=1}^d \|\mathbf{U}_k^{(t)} - \mathbf{U}_k^*\mathbf{R}_k\|_{\mathrm{F}}^2 + \|\mathcal{S}^{(t)} - [\![\mathcal{S}^*; \mathbf{R}_1^\top, \ldots, \mathbf{R}_d^\top]\!]\|_{\mathrm{F}}^2\right\}.$$

We can build the equivalence between $E^{(t)}$ and

$$\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 + \frac{a}{2}\sum_{k=1}^d \|\mathbf{U}_k^{(t)\top}\mathbf{U}_k^{(t)} - b^2\mathbf{I}_{r_k}\|_{\mathrm{F}}^2$$

by setting $b \asymp \overline{\lambda}^{1/(d+1)}$ and $a \asymp \overline{\lambda}$. Then, we can establish the following theoretical guarantee under a good initialization.

THEOREM 6.1 (Informal). *Suppose* L *satisfies RCG*$(\alpha, \beta, \mathcal{C})$ *and assume* κ, α, β *are constants. Assume* $\mathcal{X}^* = [\![\mathcal{S}^*; \mathbf{U}_1^*, \ldots, \mathbf{U}_d^*]\!]$ *and the initialization* $\mathcal{X}^{(0)} = [\![\mathcal{S}^{(0)}; \mathbf{U}_1^{(0)}, \ldots, U_d^{(0)}]\!]$ *satisfy*

$$\mathbf{U}_k^{*\top}\mathbf{U}_k^* = \mathbf{U}_k^{(0)\top}\mathbf{U}_k^{(0)} = b^2\mathbf{I}_{r_k}, \qquad \mathbf{U}_k^*, \mathbf{U}_k^{(0)} \in \mathcal{C}_k, \quad k = 1, \ldots, d; \mathcal{S}^*, \mathcal{S}^{(k)} \in \mathcal{C}_{\mathcal{S}}. \tag{6.1}$$

Suppose the initialization error satisfies $\|\mathcal{X}^{(0)} - \mathcal{X}^*\| \leq c_d\underline{\lambda}^2$ *and the signal-noise-ratio satisfies* $\underline{\lambda}^2 \geq C_d\xi^2$. *Then, by taking step size* $\eta \leq c_d b^{-2d}$, *the output of Algorithm* 3.1 *satisfies*

$$\|\mathcal{X}^{(t)} - \mathcal{X}^*\|_{\mathrm{F}}^2 \leq C_d'(\xi^2 + (1 - c_d'\eta)^t\|\mathcal{X}^{(0)} - \mathcal{X}^*\|_{\mathrm{F}}^2).$$

Here, C_d, c_d, C_d', c_d' *are constants only depending on* d.

TABLE 2
Minimax optimal estimation error bounds for order-d low-rank tensor estimation in specific applications ($d \geq 2$). Here, for simplicity, $r_1 = r_2 = r_3 = r$, $p_1 = p_2 = p_3 = p$ and $r \leq p^{1/2}$

Application	SNR condition	Estimation error	Lower bound
sub-Gaussian tensor-PCA	$\underline{\lambda}/\sigma \gtrsim p^{d/4} r^{1/(d+1)}$	$\sigma^2(pr + r^d)$	$\sigma^2(pr + r^d)$
Tensor regression	$n \gtrsim p^{d/2} r$	$n^{-1}\sigma^2(pr + r^d)$	$n^{-1}\sigma^2(pr + r^d)$
Poisson tensor-PCA	$\sqrt{I}\underline{\lambda} \gtrsim p^{(d-1)/2} r^{1/2}$	$I^{-1}(pr + r^d)$	$I^{-1}(pr + r^d)$
Binomial tensor-PCA	$\sqrt{N}\underline{\lambda} \gtrsim p^{(d-1)/2} r^{1/2}$	$N^{-1}(pr + r^d)$	$N^{-1}(pr + r^d)$

It is worth mentioning that Theorem 6.1 applies to low-rank matrix estimation (i.e., $d = 2$): suppose $\mathbf{X}^* = \mathbf{U}_1^* \mathbf{S}^* \mathbf{U}_2^{*\top}$ and $\mathbf{X}^{(0)} = \mathbf{U}_1^{(0)} \mathbf{S}^{(0)} \mathbf{U}_2^{(0)\top}$ for $\mathbf{U}_k, \mathbf{U}_k^{(0)} \in \mathbb{R}^{p_k \times r}$, $\mathbf{S}^*, \mathbf{S}^{(0)} \in \mathbb{R}^{r \times r}$, we have

$$\|\mathbf{X}^{(t)} - \mathbf{X}^*\|_{\mathrm{F}}^2 \leq C\big(\xi^2 + (1 - c\eta)^t \|\mathbf{X}^{(0)} - \mathbf{X}^*\|_{\mathrm{F}}^2\big).$$

While this framework is more complicated than necessary since one can always decompose a low-rank matrix as the product of two factor matrices $\mathbf{X} = \mathbf{U}_1 \mathbf{U}_2^\top$ without explicitly introducing the "core matrix" $\mathbf{S} \in \mathbb{R}^{r \times r}$ (see our previous discussions in Section 3.3).

Based on Theorem 6.1, we can further extend the minimax optimal bounds for the proposed procedure in each application of Section 4, that is, Theorems 4.1–4.6 to high-order scenarios. We summarize the results to Table 2.

7. Numerical studies.

7.1. *Synthetic data analysis.* In this section, we investigate the numerical performance of the proposed methods on the problems discussed in Section 4 with simulated data. We assume the true rank (r_1, r_2, r_3) is known to us and the algorithm only involves two tuning parameters: a and b. According to Theorem 3.1, a proper choice of a and b primarily depends on the unknown value $\overline{\lambda}$. In practice, we propose to use the initial estimate $\mathcal{X}^{(0)}$ as an approximation of $\mathcal{X}^*$, use $\bar{\lambda}^{(0)} = \max_k \|\mathcal{M}_k(\mathcal{X}^{(0)})\|$ as a plug-in estimate of $\overline{\lambda}$, then choose $a = \overline{\lambda}^{(0)}$, $b = (\overline{\lambda}^{(0)})^{1/4}$. We consider the following root mean squared error (RMSE) to assess the estimation accuracy in all settings:

$$\mathrm{Loss}(\hat{\mathcal{X}}, \mathcal{X}^*) = (p_1 p_2 p_3)^{-1/2} \|\hat{\mathcal{X}} - \mathcal{X}^*\|_{\mathrm{F}}. \tag{7.1}$$

Average loss over 100 repetitions are reported in following different scenarios.

Tensor Regression. We investigate the numerical performance of the proposed procedure in low-rank tensor regression discussed in Section 4.2. For all simulation settings, we first generate an r_1-by-r_2-by-r_3 core tensor $\bar{\mathcal{S}}$ with i.i.d. standard Gaussian entries and rescale it as $\mathcal{S} = \bar{\mathcal{S}} \cdot \lambda / \min_{k=1}^3 \sigma_r(\mathcal{M}_k(\bar{\mathcal{S}}))$. Here, λ quantifies the signal level and will be specified later. Then we generate $\mathbf{U}_k$ uniformly at random from the Stiefel manifold $\mathbb{O}_{p_k, r_k}$ and calculate the true parameter as $\mathcal{X}^* = [\![\mathcal{S}; \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3]\!]$. The rescaling procedure here ensures that $\min_k \sigma_{r_k}(\mathcal{M}_k(\mathcal{X}^*)) \geq \lambda$. Now, we draw a random sample based on the regression model (4.2).

We aim to compare the proposed method (Algorithms 1 and 3) with the initialization estimator (Algorithm 3 solely), Tucker regression method,[1] and MLE. Since the MLE corresponds to the global minimum of the rank-constrained optimization (2.2) and is often computationally intractable, we instead consider a warm-start gradient descent estimator, that is,

[1] The implementation is based on (Zhou (2017), Zhou, Li and Zhu (2013)).

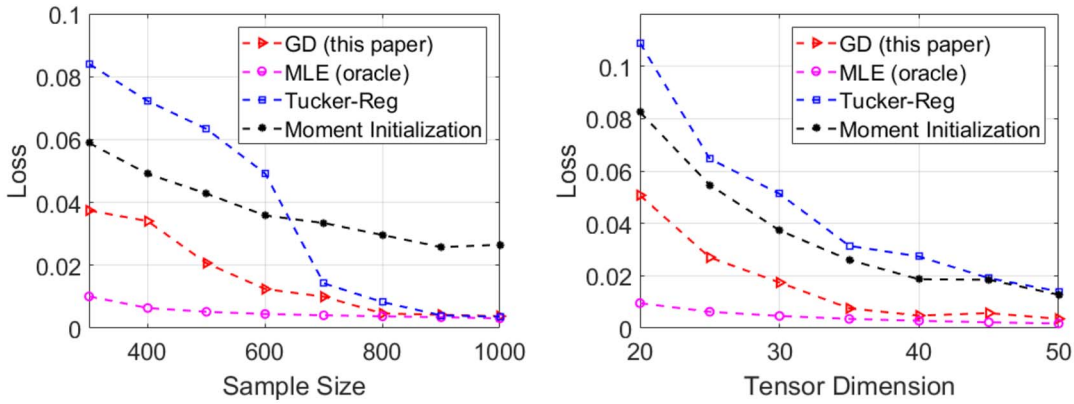

FIG. 1. *Average estimation errors in low-rank tensor regression. The MLE (oracle) is approximated by running gradient descent with the initialization chosen at $\mathcal{X}^*$, which would not be useable in practice. Here, $r = 3$, $\lambda = 2$, $\sigma = 1$. Left panel: $p = 30$, $n \in [300, 1000]$. Right panel: $p \in [25, 50]$, $n = 1.2p^{3/2}r$.*

performing Algorithm 1 starting from the true parameter $\mathcal{X}^*$. We expect that the output of this procedure can well approximate MLE. We implement all four procedures under two settings: (a) $p_1 = p_2 = p_3 = p = 30$, $r_1 = r_2 = r_3 = r = 5$, $\lambda = 2$, $\sigma = 1$, n varies from 300 to 1000; and (b) p varies from 20 to 50, $r = 3, \lambda = 2, \sigma = 1$, $n = 1.2p^{3/2}r$. The results are collected in Figure 1. We see from the left panel that for small sample size ($n \le 600$), the proposed gradient descent method significantly outperforms the Tucker regression and initialization estimator while has larger estimation errors than MLE. When the sample size increases ($n \ge 700$), the performance of the proposed gradient descent and Tucker regression algorithms tend to be as good as MLE. Compared to the initialization, gradient descent achieves a great improvement on the estimation accuracy. The right panel of Figure 1 shows that the gradient descent performs as good as the warm-start gradient descent asymptotically and is significantly better than the initialization and Tucker regression estimators.

Poisson Tensor PCA. Next, we study the numerical performance of the proposed procedure on Poisson tensor PCA. As mentioned earlier in Section 4.3, we found that the projection steps in Algorithm 1 are not essential to the numerical performance, thus we apply Algorithm 1 without the projection steps there in all numerical experiments for Poisson tensor PCA.

For each experiment, we first generate a random core tensor $\mathcal{S} \in \mathbb{R}^{r_1 \times r_2 \times r_3}$ with i.i.d. standard normal entries and random orthogonal matrices $\mathbf{U}_k$ uniformly on Stiefel manifold $\mathbb{O}^{p_k \times r_k}$. Then we calculate $\bar{\mathcal{X}} = \mathcal{S} \times_1 \mathbf{U}_1 \times_2 \mathbf{U}_2 \times_3 \mathbf{U}_3$ and rescale it as $\mathcal{X}^* = \bar{\mathcal{X}} \cdot B / \|\bar{\mathcal{X}}\|_\infty$ to ensure that each entry of $\mathcal{X}^*$ is bounded by B. Now, we generate a count tensor $\mathcal{Y} \in \mathbb{N}^{p_1 \times p_2 \times p_3}$: $\mathcal{Y}_{jkl} \sim \text{Poisson}(I \exp(\mathcal{X}^*_{jkl}))$ independently, and aim to estimate the low-rank tensor $\mathcal{X}^*$ based on $\mathcal{Y}$. In addition to the proposed method, we also consider the baseline methods of Poisson-HOSVD and Poisson-HOOI that perform HOSVD and HOOI on $\log((\mathcal{Y} + 1/2)/I)$ (i.e., Algorithm 4).

First, we fix $p = 50, r = 5$, vary the intensity value I, and study the effect of I to the numerical performance. As we can see from Figure 2, for low intensity, the gradient method is significantly better than two baselines (left panel); for high intensity, three methods are comparable while the Poisson gradient descent is the best (right panel). Next, we study their performance for different tensor dimensions and ranks. In the left panel of Figure 3, we set $r = 5$ and vary p from 30 to 100; in the right panel of Figure 3, we fix $p = 50$ and vary r from 5 to 15. As one can see, our method significantly outperforms the baselines in all settings. All these simulation results illustrate the benefits of applying gradient descent on the Poisson likelihood function.

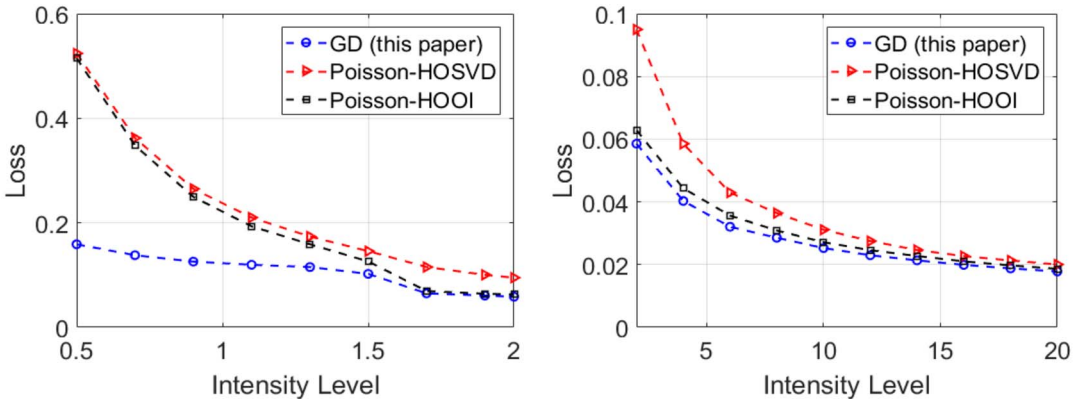

FIG. 2. *Average estimation error of Poisson tensor PCA. Here* $p = 50$, $r = 5$, $B = 2$. *Left panel*: *intensity parameter* $I \in [0.5, 2]$. *Right panel*: $I \in [2, 20]$.

Binomial Tensor PCA. We generate $\mathcal{X}^*$ in the same way as the Poisson tensor PCA settings. Suppose we observe $\mathcal{Y} \in \mathbb{N}^{p_1 \times p_2 \times p_3}$ generated from

$$\mathcal{Y}_{jkl} \sim \text{Binomial}(\mathcal{N}_{jkl}, s(\mathcal{X}^*_{jkl})), \quad \text{independently.}$$

We take all entries with the same population size (i.e., $\mathcal{N}_{ijk} = N$) for simplification. We can see from the simulation results in Figure 4 that a larger population size N yields smaller estimation error. In addition, according to Theorem 3.1, the estimation error in theory is of order $O(p^{-1/2}r^{1/2})$, which matches the trend of estimation error curves in Figure 4.

7.2. *Real data analysis.* In this section, we apply the proposed framework to real data applications in 4D-STEM image denoising. An additional real data example on click-through prediction is postponed to Appendix B in the Supplementary Material.

The 4D-scanning transmission electron microscopy (4D-STEM) is an important technique in modern material science that has been used to detect local material composition of structures such as films, defects and nanostructures (Krivanek, Dellby and Lupini (1999), Yankovich et al. (2016)). In 4D-STEM imaging technology, a focused probe is usually rastered across part of the specimen and an X-ray and/or electron energy loss spectrum is recorded at each probe position, generating a series of photon-limited images. The data generated from 4D-STEM technique are typically order-4 tensors with approximate periodic

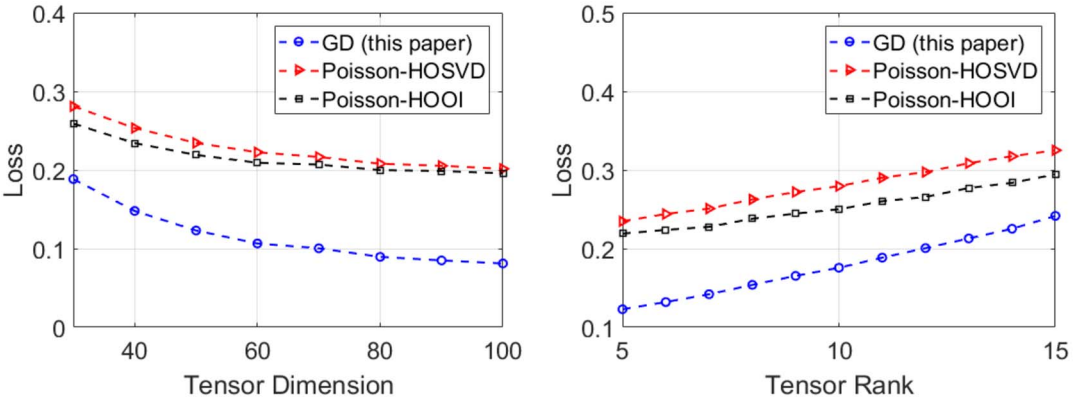

FIG. 3. *Average estimation error of Poisson tensor PCA with different dimensions and ranks. Here,* $B = 2$, $I = 1$. *Left panel*: $r = 5$, $p \in [30, 300]$. *Right panel*: $p = 50$, $r \in [5, 15]$.

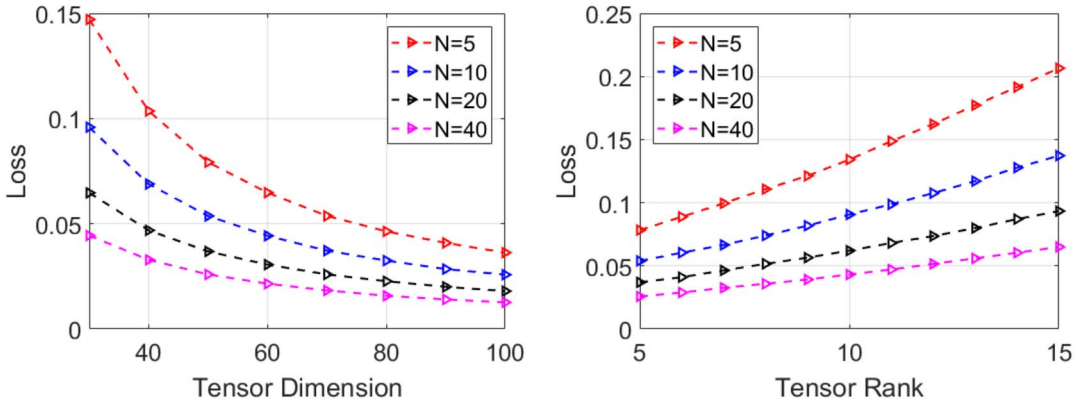

FIG. 4. *Average estimation error of binomial tensor PCA. Here,* $B = 2$. *Left panel*: $r = 5$, $p \in [30, 100]$. *Right panel*: $p = 50, r \in [5 : 15]$.

structures, as a focused probe is located on a 2-D grid and one 2-D image is generated for each probe position (see Yankovich et al. (2016) for more details). Due to the physical conditions, the observable images are often photon-limited, highly noisy, and in the form of count matrices (see the second row of Figure 5 for an example). A sufficient imaging denoising is often a crucial first step before the subsequent procedures.

We aim to illustrate the merit of the proposed method through denoising of data in 4D-STEM experiments. Specifically, we collect 160 images generated from a row of electron probe positions.[2] Since the resolution of each image is 183×183, the data images can be stacked into a nonnegative tensor of size $160 \times 183 \times 183$. We assume the observational images $\mathcal{Y}$ are generated from Poisson distribution $\mathcal{Y}_{ijk} \overset{\text{i.i.d.}}{\sim} \text{Poisson}(\exp(\mathcal{X}^*_{ijk}))$. Our goal is to recover the original images based on the photon-limited observation $\mathcal{Y}$. Since $\mathcal{Y}$ is sparse ($\approx 88\%$ pixels are zero), we take the pre-initializer $\tilde{\mathcal{X}} = \log(\mathcal{Y} + 1/30)$ and take the input rank according to the (5.3) with $\rho = 0.98$. We apply the proposed gradient descent (Algorithms 1 and 4) with the rank estimation $(\hat{r}_1, \hat{r}_2, \hat{r}_3) = (44, 36, 34)$ to obtain the estimator $\hat{\mathcal{X}}$, then calculate $\exp(\hat{\mathcal{X}})$ as the collection of denoised images. We also denoise these images one by one via the matrix Procrustes flow (Park et al. (2018)), a variant of the matrix-version gradient descent method.[3] The original, observational, and recovered images are provided in Figure 5. In addition, we calculate the recovery loss for each of the 160 images, that is, $\|\exp(\mathcal{X}^*_{i::}) - \exp(\hat{\mathcal{X}}_{i::})\|_{\text{F}} / \|\exp(\mathcal{X}^*_{i::})\|_{\text{F}}$, and the averaged recovery loss (and standard error) of matrix and tensor methods are 0.861 (0.183) and 0.303 (0.054), respectively. One can clearly see the advantage of the proposed tensor method that utilizes the tensor structure of the whole set of images.

8. Discussions. In this paper, we introduce a nonconvex optimization framework for the generalized tensor estimation. Compared to the convex relaxation methods in the literature, the proposed scheme is computationally efficient and achieves desirable statistical error rate under suitable initialization and signal-to-noise ratio conditions. We apply the proposed

[2]Simultaneously denoising the order-4 image data requires extremely large memory and computation source. Thus we focus on one row of images. It is also common to perform row-wise image denoising in 4D-STEM imaging analysis (Yankovich et al. (2016)).

[3]This algorithm also requires the specification of matrix rank. We empirically choose $\hat{r} = \arg\min\{r : \sum_{i=1}^{r} \sigma_i^2(\mathbf{X}) / \sum_{i=1}^{p_k} \sigma_i^2(\mathbf{X}) \geq 0.98\}$ for each slice of the pre-initializer $\mathbf{X} = \tilde{\mathcal{X}}_{i::}$.

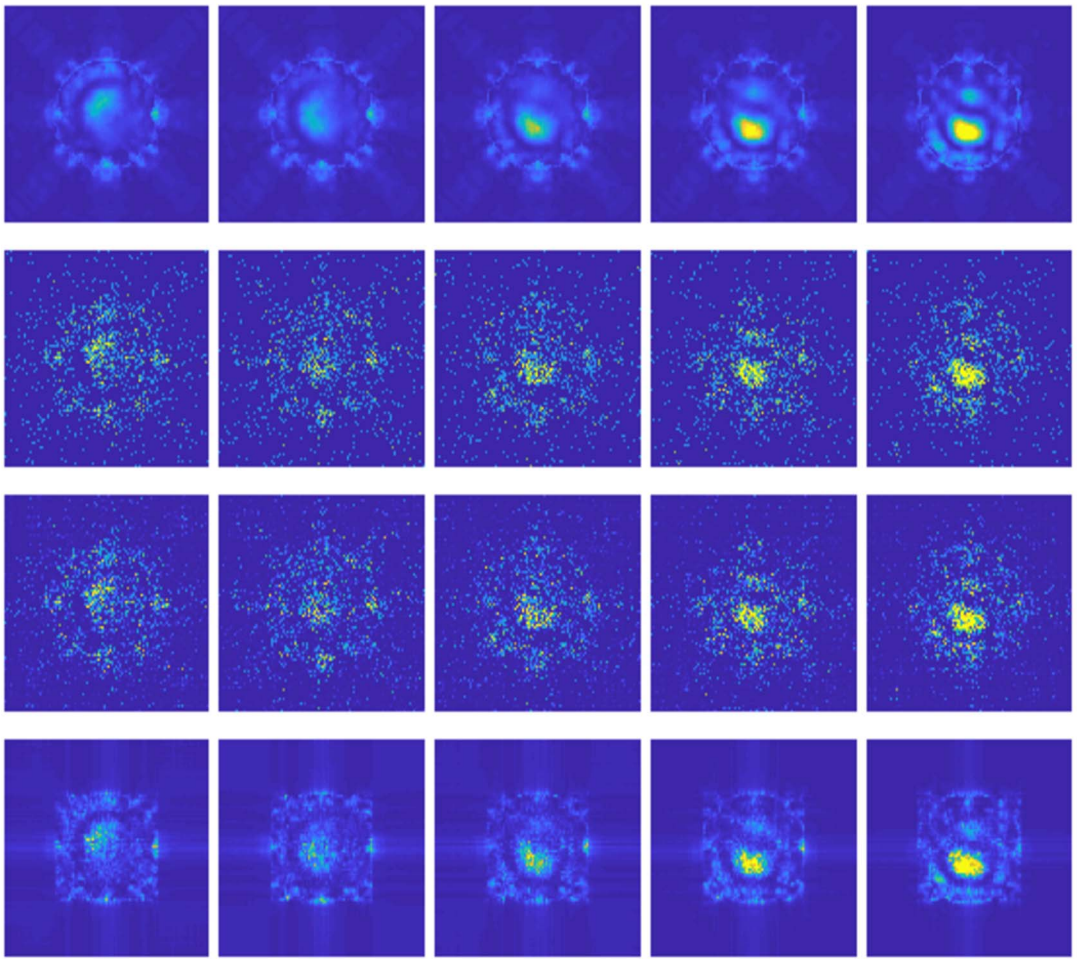

FIG. 5. *Recovery results for the first five images of* 4*D-STEM data. First row*: *original images*; *second row*: *photon-limited observations*; *third row*: *denoised images by matrix method*; *forth row*: *denoised images by the proposed tensor method*.

framework on several problems, including sub-Gaussian denoising, tensor regression, Poisson and binomial tensor PCA. We can show that the proposed gradient descent procedure achieves the minimax optimal rate of estimation error under these statistical models.

In addition to the above-mentioned problems, the proposed framework can incorporate a broader range of settings. For example, the developed result is applicable to solve the *noisy tensor completion* problem (Cai et al. (2019), Montanari and Sun (2018), Shah (2018), Xia, Yuan and Zhang (2021)), which aims to recover the low-rank tensor $\mathcal{X}^*$ based on a number of noisy observable entries, say $\{\mathcal{Y}_{ijk} = \mathcal{X}^*_{ijk} + \mathcal{Z}_{ijk}\}_{(i,j,k)\in\Omega}$, where Ω is a subset of indices.

Another example is *binary tensor PCA* (Wang and Li (2020)), where the central goal is to factorize the tensor from 0-1 valued observations. Suppose one observes $\mathcal{Y}_{ijk} \sim$ Bernoulli$(\mathcal{P}_{ijk})$ independently, where $\mathcal{P}_{ijk} = s(\mathcal{X}^*_{ijk})$, $\mathcal{X}^*$ is low-rank, and $s(\cdot)$ is some link function. Then the proposed projected gradient descent method can be applied to estimate $\mathcal{X}^*$ with provable guarantees.

Community detection in social network has attracted enormous recent attention. Although most of the existing results focused on a single-layer of network, the *multilayer network*, that is, the connections between different nodes are reflected in multiple modalities, also commonly appear in practice (Han, Xu and Airoldi (2015), Lei, Chen and Lynch (2020), Pensky and Zhang (2019)). Consider a stack of multilayer network data with shared community struc-

ture. It is reasonable to assume that the adjacency tensor $\mathcal{A}$ has a low-rank tensor structure: $\mathcal{A} \sim \text{Bernoulli}(\mathcal{X}^*)$ independently, where $\mathcal{X}^* = [\![\mathcal{S}^*, \mathbf{Z}^*, \mathbf{Z}^*, \mathbf{T}^*]\!]$, $\mathbf{Z}^*$ is the latent space of nodes features (or the indicator matrix for the community that each node belongs to), and $\mathbf{T}^*$ models the trend along the time. Then the community detection for multilayer networks essentially becomes the generalized tensor estimation problem.

In addition to the standard linear regression model discussed in Section 4.2, the proposed framework can be applied to a range of generalized tensor regression problems. Recall that the classical generalized linear model focuses on an exponential family, where the response y_i satisfies the following density or probability mass function (Wedderburn (1974)),

$$p(y_i|\theta_i, \phi) = \exp\left\{\frac{y\theta_i - b(\theta_i)}{a(\phi)} + c(y, \phi)\right\}. \tag{8.1}$$

Here, a, b, c are prespecified functions determined by the problem; θ_i and $\phi > 0$ are natural and dispersion parameters, respectively. For the generalized tensor regression, it is natural to relate the tensor covariate and response (Zhou, Li and Zhu (2013)) via

$$\mu_i = \mathbb{E}(y_i|\mathcal{X}^*), \qquad g(\mu_i) = \langle \mathcal{A}_i, \mathcal{X}^* \rangle, \tag{8.2}$$

where $g(\cdot)$ is a link function. To estimate $\mathcal{X}^*$, we can apply the proposed Algorithm 3 on the negative log-likelihood function

$$\sum_{i=1}^{n} \frac{y_i\theta_i - b(\theta_i)}{a(\phi)} + \sum_{i=1}^{n} c(y_i, \phi),$$

where θ_i is determined by (8.1) and (8.2).

Some other possible applications of the proposed framework include the *high-order interaction pursuit* (Hao, Zhang and Cheng (2020)), *generalized regression among multiple modes* (Xu, Hu and Wang (2019)), *mixed-data-type tensor data analysis* (Baker, Tang and Allen (2020)), etc. In all these problems, by exploring the log-likelihood of data and the domain $\mathcal{C}$ that satisfies RCG condition, the proposed projected gradient descent can be applied and the theoretical guarantees can be developed based on the proposed framework.

Acknowledgments. The authors thank Paul Voyles and Chenyu Zhang for providing the 4D-STEM dataset and for helpful discussions.

Funding. The research of R. H. and A. R. Z. was supported in part by NSF Grants DMS-1811868, NSF CAREER-1944904, and NIH R01-GM131399. The research of R. W. was supported in part by AFOSR Grants FA9550-18-1-0166, DOE DE-AC02-06CH11357, NSF OAC-1934637, and NSF DMS-2023109. The research of R.H. was also supported in part by a RAship from Institute for Mathematics of Data Science at UW-Madison.

SUPPLEMENTARY MATERIAL

Supplement to "An optimal statistical and computational framework for generalized tensor estimation" (DOI: 10.1214/21-AOS2061SUPP; .pdf). This supplement file includes the pseudo-codes of higher-order SVD (HOSVD), higher-order orthogonal iteration (HOOI), HeteroPCA, all proofs of main results, and technical lemmas.

REFERENCES

AHMED, A., RECHT, B. and ROMBERG, J. (2014). Blind deconvolution using convex programming. *IEEE Trans. Inf. Theory* **60** 1711–1732. MR3168432 https://doi.org/10.1109/TIT.2013.2294644

ANANDKUMAR, A., GE, R., HSU, D., KAKADE, S. M. and TELGARSKY, M. (2014). Tensor decompositions for learning latent variable models. *J. Mach. Learn. Res.* **15** 2773–2832. MR3270750

ARROYO, J., ATHREYA, A., CAPE, J., CHEN, G., PRIEBE, C. E. and VOGELSTEIN, J. T. (2021). Inference for multiple heterogeneous networks with a common invariant subspace. *J. Mach. Learn. Res.* **22** Paper No. 142. MR4318498

BAHADORI, M. T., YU, Q. R. and LIU, Y. (2014). Fast multivariate spatio-temporal analysis via low rank tensor learning. In *Advances in Neural Information Processing Systems* 3491–3499.

BAKER, Y., TANG, T. M. and ALLEN, G. I. (2020). Feature selection for data integration with mixed multiview data. *Ann. Appl. Stat.* **14** 1676–1698. MR4194243 https://doi.org/10.1214/20-AOAS1389

BARAK, B. and MOITRA, A. (2016). Noisy tensor completion via the sum-of-squares hierarchy. In *Conference on Learning Theory* 417–445.

BI, X., QU, A. and SHEN, X. (2018). Multilayer tensor factorization with applications to recommender systems. *Ann. Statist.* **46** 3308–3333. MR3852653 https://doi.org/10.1214/17-AOS1659

CAI, T. T., LI, X. and MA, Z. (2016). Optimal rates of convergence for noisy sparse phase retrieval via thresholded Wirtinger flow. *Ann. Statist.* **44** 2221–2251. MR3546449 https://doi.org/10.1214/16-AOS1443

CAI, C., LI, G., POOR, H. V. and CHEN, Y. (2019). Nonconvex low-rank tensor completion from noisy data. In *Advances in Neural Information Processing Systems* 1861–1872.

CANDÈS, E. J., LI, X. and SOLTANOLKOTABI, M. (2015). Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Trans. Inf. Theory* **61** 1985–2007. MR3332993 https://doi.org/10.1109/TIT.2015.2399924

CANDES, E. J. and PLAN, Y. (2010). Matrix completion with noise. *Proc. IEEE* **98** 925–936.

CANDÈS, E. J. and PLAN, Y. (2011). Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Trans. Inf. Theory* **57** 2342–2359. MR2809094 https://doi.org/10.1109/TIT.2011.2111771

CANDÈS, E. J. and RECHT, B. (2009). Exact matrix completion via convex optimization. *Found. Comput. Math.* **9** 717–772. MR2565240 https://doi.org/10.1007/s10208-009-9045-5

CAO, Y. and XIE, Y. (2016). Poisson matrix recovery and completion. *IEEE Trans. Signal Process.* **64** 1609–1620. MR3548877 https://doi.org/10.1109/TSP.2015.2500192

CAO, Y., ZHANG, A. and LI, H. (2020). Multisample estimation of bacterial composition matrices in metagenomics data. *Biometrika* **107** 75–92. MR4064141 https://doi.org/10.1093/biomet/asz062

CHEN, W.-K. (2019). Phase transition in the spiked random tensor with Rademacher prior. *Ann. Statist.* **47** 2734–2756. MR3988771 https://doi.org/10.1214/18-AOS1763

CHEN, Y. and CANDES, E. (2015). Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Advances in Neural Information Processing Systems* 739–747.

CHEN, Y. and CHI, Y. (2018). Harnessing structures in big data via guaranteed low-rank matrix estimation. ArXiv preprint. Available at arXiv:1802.08397.

CHEN, H., RASKUTTI, G. and YUAN, M. (2019). Non-convex projected gradient descent for generalized low-rank tensor regression. *J. Mach. Learn. Res.* **20** Paper No. 5. MR3911412

CHI, E. C. and KOLDA, T. G. (2012). On tensors, sparsity, and nonnegative factorizations. *SIAM J. Matrix Anal. Appl.* **33** 1272–1299. MR3023474 https://doi.org/10.1137/110859063

CHI, Y., LU, Y. M. and CHEN, Y. (2019). Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Trans. Signal Process.* **67** 5239–5269. MR4016283 https://doi.org/10.1109/TSP.2019.2937282

CHI, E. C., GAINES, B. R., SUN, W. W., ZHOU, H. and YANG, J. (2020). Provable convex co-clustering of tensors. *J. Mach. Learn. Res.* **21** Paper No. 214. MR4209500

DE LATHAUWER, L., DE MOOR, B. and VANDEWALLE, J. (2000a). A multilinear singular value decomposition. *SIAM J. Matrix Anal. Appl.* **21** 1253–1278. MR1780272 https://doi.org/10.1137/S0895479896305696

DE LATHAUWER, L., DE MOOR, B. and VANDEWALLE, J. (2000b). On the best rank-1 and rank-$(R_1, R_2, \ldots, R_N)$ approximation of higher-order tensors. *SIAM J. Matrix Anal. Appl.* **21** 1324–1342. MR1780276 https://doi.org/10.1137/S0895479898346995

FAN, J., GONG, W. and ZHU, Z. (2019). Generalized high-dimensional trace regression via nuclear norm regularization. *J. Econometrics* **212** 177–202. MR3994013 https://doi.org/10.1016/j.jeconom.2019.04.026

FAUST, K., SATHIRAPONGSASUTI, J. F., IZARD, J., SEGATA, N., GEVERS, D., RAES, J. and HUTTENHOWER, C. (2012). Microbial co-occurrence relationships in the human microbiome. *PLoS Comput. Biol.* **8** e1002606.

FAZEL, M. (2002). Matrix rank minimization with applications.

FLORES, G. E., CAPORASO, J. G., HENLEY, J. B., RIDEOUT, J. R., DOMOGALA, D., CHASE, J., LEFF, J. W., VÁZQUEZ-BAEZA, Y., GONZALEZ, A. et al. (2014). Temporal variability is a personalized feature of the human microbiome. *Genome Biol.* **15** 531. https://doi.org/10.1186/s13059-014-0531-y

FRIEDLAND, S. and LIM, L.-H. (2018). Nuclear norm of higher-order tensors. *Math. Comp.* **87** 1255–1281. MR3766387 https://doi.org/10.1090/mcom/3239

GANDY, S., RECHT, B. and YAMADA, I. (2011). Tensor completion and low-n-rank tensor recovery via convex optimization. *Inverse Probl.* **27** 025010. MR2765628 https://doi.org/10.1088/0266-5611/27/2/025010

GUHANIYOGI, R., QAMAR, S. and DUNSON, D. B. (2017). Bayesian tensor regression. *J. Mach. Learn. Res.* **18** Paper No. 79. MR3714242

GUO, W., KOTSIA, I. and PATRAS, I. (2012). Tensor learning for regression. *IEEE Trans. Image Process.* **21** 816–827. MR2932176 https://doi.org/10.1109/TIP.2011.2165291

HALL, E. C., RASKUTTI, G. and WILLETT, R. (2016). Inference of high-dimensional autoregressive generalized linear models. ArXiv preprint. Available at arXiv:1605.02693.

HAN, R., WILLETT, R. and ZHANG, A. R. (2022). Supplement to "An optimal statistical and computational framework for generalized tensor estimation." https://doi.org/10.1214/21-AOS2061SUPP

HAN, Q., XU, K. and AIROLDI, E. (2015). Consistent estimation of dynamic and multi-layer block models. In *International Conference on Machine Learning* 1511–1520.

HAO, B., ZHANG, A. and CHENG, G. (2020). Sparse and low-rank tensor estimation via cubic sketchings. *IEEE Trans. Inf. Theory* **66** 5927–5964. MR4158653 https://doi.org/10.1109/TIT.2020.2982499

HENRIQUES, R. and MADEIRA, S. C. (2019). Triclustering algorithms for three-dimensional data analysis: A comprehensive survey. *ACM Computing Surveys* (*CSUR*) **51** 95.

HILLAR, C. J. and LIM, L.-H. (2013). Most tensor problems are NP-hard. *J. ACM* **60** Art. 45. MR3144915 https://doi.org/10.1145/2512329

HOFF, P. D. (2015). Multilinear tensor regression for longitudinal relational data. *Ann. Appl. Stat.* **9** 1169–1193. MR3418719 https://doi.org/10.1214/15-AOAS839

HONG, D., KOLDA, T. G. and DUERSCH, J. A. (2020). Generalized canonical polyadic tensor decomposition. *SIAM Rev.* **62** 133–163. MR4064532 https://doi.org/10.1137/18M1203626

HOPKINS, S. B., SHI, J. and STEURER, D. (2015). Tensor principal component analysis via sum-of-square proofs. In *Proceedings of the* 28*th Conference on Learning Theory*, *COLT* 3–6.

JAVANMARD, A. and MONTANARI, A. (2018). Debiasing the Lasso: Optimal sample size for Gaussian designs. *Ann. Statist.* **46** 2593–2622. MR3851749 https://doi.org/10.1214/17-AOS1630

JIANG, X., RASKUTTI, G. and WILLETT, R. (2015). Minimax optimal rates for Poisson inverse problems with physical constraints. *IEEE Trans. Inf. Theory* **61** 4458–4474. MR3372365 https://doi.org/10.1109/TIT.2015.2441072

JOHNDROW, J. E., BHATTACHARYA, A. and DUNSON, D. B. (2017). Tensor decompositions and sparse log-linear models. *Ann. Statist.* **45** 1–38. MR3611485 https://doi.org/10.1214/15-AOS1414

JOLLIFFE, I. T. (1986). *Principal Component Analysis*. *Springer Series in Statistics*. Springer, New York. MR0841268 https://doi.org/10.1007/978-1-4757-1904-8

KESHAVAN, R. H., MONTANARI, A. and OH, S. (2010). Matrix completion from noisy entries. *J. Mach. Learn. Res.* **11** 2057–2078. MR2678022

KOLDA, T. G. and BADER, B. W. (2009). Tensor decompositions and applications. *SIAM Rev.* **51** 455–500. MR2535056 https://doi.org/10.1137/07070111X

KOLTCHINSKII, V., LOUNICI, K. and TSYBAKOV, A. B. (2011). Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.* **39** 2302–2329. MR2906869 https://doi.org/10.1214/11-AOS894

KRIVANEK, O., DELLBY, N. and LUPINI, A. (1999). Towards sub-A electron beams. *Ultramicroscopy* **78** 1–11.

KROONENBERG, P. M. (2008). *Applied Multiway Data Analysis*. *Wiley Series in Probability and Statistics*. Wiley Interscience, Hoboken, NJ. MR2378349 https://doi.org/10.1002/9780470238004

LEI, J., CHEN, K. and LYNCH, B. (2020). Consistent community detection in multi-layer network data. *Biometrika* **107** 61–73. MR4064140 https://doi.org/10.1093/biomet/asz068

LESIEUR, T., KRZAKALA, F. and ZDEBOROVÁ, L. (2017). Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications. *J. Stat. Mech. Theory Exp.* **7** 073403. MR3683819 https://doi.org/10.1088/1742-5468/aa7284

LI, N. and LI, B. (2010). Tensor completion for on-board compression of hyperspectral images. In 2010 *IEEE International Conference on Image Processing* 517–520. IEEE, New York.

LI, L. and ZHANG, X. (2017). Parsimonious tensor response regression. *J. Amer. Statist. Assoc.* **112** 1131–1146. MR3735365 https://doi.org/10.1080/01621459.2016.1193022

LI, X., XU, D., ZHOU, H. and LI, L. (2018). Tucker tensor regression and neuroimaging analysis. *Stat. Biosci.* **10** 520–545.

LIU, J., MUSIALSKI, P., WONKA, P. and YE, J. (2013). Tensor completion for estimating missing values in visual data. *IEEE Trans. Pattern Anal. Mach. Intell.* **35** 208–220.

LUBICH, C., ROHWEDDER, T., SCHNEIDER, R. and VANDEREYCKEN, B. (2013). Dynamical approximation by hierarchical Tucker and tensor-train tensors. *SIAM J. Matrix Anal. Appl.* **34** 470–494. MR3053575 https://doi.org/10.1137/120885723

MA, Z. and MA, Z. (2017). Exploration of large networks via fast and universal latent space model fitting. ArXiv preprint. Available at arXiv:1705.02372.

McMahan, H. B., Holt, G., Sculley, D., Young, M., Ebner, D., Grady, J., Nie, L., Phillips, T., Davydov, E. et al. (2013). Ad click prediction: A view from the trenches. In *Proceedings of the* 19*th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 1222–1230. ACM, New York.

Montanari, A., Reichman, D. and Zeitouni, O. (2017). On the limitation of spectral methods: From the Gaussian hidden clique problem to rank one perturbations of Gaussian tensors. *IEEE Trans. Inf. Theory* **63** 1572–1579. MR3625981 https://doi.org/10.1109/TIT.2016.2637959

Montanari, A. and Sun, N. (2018). Spectral algorithms for tensor completion. *Comm. Pure Appl. Math.* **71** 2381–2425. MR3862094 https://doi.org/10.1002/cpa.21748

Oymak, S., Jalali, A., Fazel, M., Eldar, Y. C. and Hassibi, B. (2015). Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Trans. Inf. Theory* **61** 2886–2908. MR3342310 https://doi.org/10.1109/TIT.2015.2401574

Park, D., Kyrillidis, A., Caramanis, C. and Sanghavi, S. (2018). Finding low-rank solutions via nonconvex matrix factorization, efficiently and provably. *SIAM J. Imaging Sci.* **11** 2165–2204. MR3860118 https://doi.org/10.1137/17M1150189

Pensky, M. and Zhang, T. (2019). Spectral clustering in the dynamic stochastic block model. *Electron. J. Stat.* **13** 678–709. MR3914178 https://doi.org/10.1214/19-ejs1533

Perry, A., Wein, A. S. and Bandeira, A. S. (2020). Statistical limits of spiked tensor models. *Ann. Inst. Henri Poincaré Probab. Stat.* **56** 230–264. MR4058987 https://doi.org/10.1214/19-AIHP960

Raskutti, G., Yuan, M. and Chen, H. (2019). Convex regularization for high-dimensional multiresponse tensor regression. *Ann. Statist.* **47** 1554–1584. MR3911122 https://doi.org/10.1214/18-AOS1725

Rauhut, H., Schneider, R. and Stojanac, Ž. (2015). Tensor completion in hierarchical tensor representations. In *Compressed Sensing and Its Applications. Appl. Numer. Harmon. Anal.* 419–450. Birkhäuser/Springer, Cham. MR3382114

Rauhut, H., Schneider, R. and Stojanac, Ž. (2017). Low rank tensor recovery via iterative hard thresholding. *Linear Algebra Appl.* **523** 220–262. MR3624675 https://doi.org/10.1016/j.laa.2017.02.028

Richard, E. and Montanari, A. (2014). A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems* 2897–2905.

Salmon, J., Harmany, Z., Deledalle, C.-A. and Willett, R. (2014). Poisson noise reduction with nonlocal PCA. *J. Math. Imaging Vision* **48** 279–294. MR3152105 https://doi.org/10.1007/s10851-013-0435-6

Sewell, D. K. and Chen, Y. (2015). Latent space models for dynamic networks. *J. Amer. Statist. Assoc.* **110** 1646–1657. MR3449061 https://doi.org/10.1080/01621459.2014.988214

Shah, D. (2018). Matrix estimation, latent variable model and collaborative filtering. In 37*th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. LIPIcs. Leibniz Int. Proc. Inform.* **93** Art. No. 4. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern. MR3774048

Shan, L., Lin, L., Sun, C. and Wang, X. (2016). Predicting ad click-through rates via feature-based fully coupled interaction tensor factorization. *Electronic Commerce Research and Applications* **16** 30–42.

Signoretto, M., Van de Plas, R., De Moor, B. and Suykens, J. A. (2011). Tensor versus matrix completion: A comparison with application to spectral data. *IEEE Signal Process. Lett.* **18** 403–406.

Sun, W. W. and Li, L. (2017). STORE: Sparse tensor response regression and neuroimaging analysis. *J. Mach. Learn. Res.* **18** Paper No. 135. MR3763769

Sun, R. and Luo, Z.-Q. (2015). Guaranteed matrix completion via nonconvex factorization. In 2015 *IEEE* 56*th Annual Symposium on Foundations of Computer Science—FOCS* 2015 270–289. IEEE Computer Soc., Los Alamitos, CA. MR3473312 https://doi.org/10.1109/FOCS.2015.25

Sun, W. W., Lu, J., Liu, H. and Cheng, G. (2017). Provable sparse tensor decomposition. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 899–916. MR3641413 https://doi.org/10.1111/rssb.12190

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. MR1379242

Timmermann, K. E. and Nowak, R. D. (1999). Multiscale modeling and estimation of Poisson processes with application to photon-limited imaging. *IEEE Trans. Inf. Theory* **45** 846–862. MR1682515 https://doi.org/10.1109/18.761328

Tomioka, R. and Suzuki, T. (2013). Convex tensor decomposition via structured Schatten norm regularization. In *Advances in Neural Information Processing Systems* 1331–1339.

Tomioka, R., Suzuki, T., Hayashi, K. and Kashima, H. (2011). Statistical performance of convex tensor decomposition. In *Advances in Neural Information Processing Systems* 972–980.

Tu, S., Boczar, R., Simchowitz, M., Soltanolkotabi, M. and Recht, B. (2016). Low-rank solutions of linear matrix equations via procrustes flow. In *International Conference on Machine Learning* 964–973.

Wang, M., Fischer, J. and Song, Y. S. (2017). Three-way clustering of multi-tissue multi-individual gene expression data using constrained tensor decomposition. *BioRxiv* 229245.

Wang, M. and Li, L. (2020). Learning from binary multiway data: Probabilistic tensor decomposition and its statistical optimality. *J. Mach. Learn. Res.* **21** Paper No. 154. MR4209440

WANG, M. and ZENG, Y. (2019). Multiway clustering via tensor block models. In *Advances in Neural Information Processing Systems* 713–723.

WEDDERBURN, R. W. M. (1974). Quasi-likelihood functions, generalized linear models, and the Gauss-Newton method. *Biometrika* **61** 439–447. MR0375592 https://doi.org/10.1093/biomet/61.3.439

WEN, Z., YIN, W. and ZHANG, Y. (2012). Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Math. Program. Comput.* **4** 333–361. MR3006618 https://doi.org/10.1007/s12532-012-0044-1

WILLETT, R. M. and NOWAK, R. D. (2007). Multiscale Poisson intensity and density estimation. *IEEE Trans. Inf. Theory* **53** 3171–3187. MR2417680 https://doi.org/10.1109/TIT.2007.903139

WILMOTH, J. R. and SHKOLNIKOV, V. (2006). Human mortality database. Available at http://www.mortality.org.

XIA, D. and YUAN, M. (2019). On polynomial time methods for exact low-rank tensor completion. *Found. Comput. Math.* **19** 1265–1313. MR4029842 https://doi.org/10.1007/s10208-018-09408-6

XIA, D., YUAN, M. and ZHANG, C.-H. (2021). Statistically optimal and computationally efficient low rank tensor completion from noisy entries. *Ann. Statist.* **49** 76–99. MR4206670 https://doi.org/10.1214/20-AOS1942

XU, Z., HU, J. and WANG, M. (2019). Generalized tensor regression with covariates on multiple modes. ArXiv preprint. Available at arXiv:1910.09499.

YANKOVICH, A. B., ZHANG, C., OH, A., SLATER, T. J., AZOUGH, F., FREER, R., HAIGH, S. J., WILLETT, R. and VOYLES, P. M. (2016). Non-rigid registration and non-local principle component analysis to improve electron microscopy spectrum images. *Nanotechnology* **27** 364001.

YOKOTA, T., LEE, N. and CICHOCKI, A. (2017). Robust multilinear tensor rank estimation using higher order singular value decomposition and information criteria. *IEEE Trans. Signal Process.* **65** 1196–1206. MR3584316 https://doi.org/10.1109/TSP.2016.2620965

YONEL, B. and YAZICI, B. (2020). A deterministic theory for exact non-convex phase retrieval. *IEEE Trans. Signal Process.* **68** 4612–4626. MR4144925 https://doi.org/10.1109/TSP.2020.3007967

YU, M., GUPTA, V. and KOLAR, M. (2020). Recovery of simultaneous low rank and two-way sparse coefficient matrices, a nonconvex approach. *Electron. J. Stat.* **14** 413–457. MR4054252 https://doi.org/10.1214/19-EJS1658

YUAN, M. and ZHANG, C.-H. (2016). On tensor completion via nuclear norm minimization. *Found. Comput. Math.* **16** 1031–1068. MR3529132 https://doi.org/10.1007/s10208-015-9269-5

ZHANG, A. (2019). Cross: Efficient low-rank tensor completion. *Ann. Statist.* **47** 936–964. MR3909956 https://doi.org/10.1214/18-AOS1694

ZHANG, A., CAI, T. T. and WU, Y. (2018). Heteroskedastic PCA: Algorithm, optimality, and applications. ArXiv preprint. Available at arXiv:1810.08316.

ZHANG, A. and HAN, R. (2019). Optimal sparse singular value decomposition for high-dimensional high-order data. *J. Amer. Statist. Assoc.* **114** 1708–1725. MR4047294 https://doi.org/10.1080/01621459.2018.1527227

ZHANG, A. and XIA, D. (2018). Tensor SVD: Statistical and computational limits. *IEEE Trans. Inf. Theory* **64** 7311–7338. MR3876445 https://doi.org/10.1109/TIT.2018.2841377

ZHANG, A. R., LUO, Y., RASKUTTI, G. and YUAN, M. (2020). ISLET: Fast and optimal low-rank tensor regression via importance sketching. *SIAM J. Math. Data Sci.* **2** 444–479. MR4106613 https://doi.org/10.1137/19M126476X

ZHAO, T., WANG, Z. and LIU, H. (2015). A nonconvex optimization framework for low rank matrix estimation. In *Advances in Neural Information Processing Systems* 559–567.

ZHOU, H. (2017). Matlab TensorReg Toolbox Version 1.0. Available at https://hua-zhou.github.io/TensorReg/.

ZHOU, H., LI, L. and ZHU, H. (2013). Tensor regression with applications in neuroimaging data analysis. *J. Amer. Statist. Assoc.* **108** 540–552. MR3174640 https://doi.org/10.1080/01621459.2013.776499

ZHU, Z., LI, Q., TANG, G. and WAKIN, M. B. (2017). The global optimization geometry of nonsymmetric matrix factorization and sensing. ArXiv preprint. Available at arXiv:1703.01256.