

Euclidean-Norm-Induced Schatten- p Quasi-Norm Regularization for Low-Rank Tensor Completion and Tensor Robust Principal Component Analysis

Anonymous authors

Paper under double-blind review

Abstract

The nuclear norm and Schatten- p quasi-norm are popular rank proxies in low-rank matrix recovery. However, computing the nuclear norm or Schatten- p quasi-norm of a tensor is hard in both theory and practice, preventing their application to low-rank tensor completion (LRTC) and tensor robust principal component analysis (TRPCA). In this paper, we propose a new class of tensor rank regularizers based on the Euclidean norms of the CP component vectors of a tensor and show that these regularizers are monotonic transformations of tensor Schatten- p quasi-norm. This connection enables us to minimize the Schatten- p quasi-norm in LRTC and TRPCA implicitly via the component vectors. The method does not require the singular value decomposition and hence scales to big tensors. Moreover, the method is not sensitive to the choice of initial rank empirically, and provides an arbitrarily sharper rank proxy for low-rank tensor recovery compared to nuclear norm. On the other hand, we study the generalization abilities of LRTC with Schatten- p quasi-norm regularization and LRTC with our regularizers. The theorems show that a relatively sharper regularizer leads to a tighter error bound, which is consistent with our numerical results. We also provide a recovery error bound for TRPCA. Numerical results on synthetic data and real data demonstrate the effectiveness and superiority of our methods compared to baselines.

1 Introduction

Low-rank tensor completion (LRTC) (Gandy et al., 2011; Acar et al., 2011; Liu et al., 2012; Romera-Paredes & Pontil, 2013; Kressner et al., 2014; Yuan & Zhang, 2016; Cheng et al., 2016; Zhou et al., 2017; Lacroix et al., 2018; Ghadermarzy et al., 2019; Liu & Moitra, 2020; Wimalawarne & Mamitsuka, 2021; Fan, 2022), as a high-order generalization of low-rank matrix completion (LRMC) (Candès & Recht, 2009; Hardt, 2014), aims to recover the missing entries of a low-rank tensor. One may organize LRTC methods into different categories according to the types of decomposition model, e.g., CP (CANDECOMP/PARAFAC) decomposition based LRTC (Acar et al., 2011; Jain & Oh, 2014; Zhao et al., 2015; Liu & Moitra, 2020), Tucker decomposition based LRTC (Xu et al., 2013; Xu & Yin, 2013; Kasai & Mishra, 2016; Xie et al., 2018; Kong et al., 2018), and tensor ring based LRTC (Yuan et al., 2019). This paper will focus on CP decomposition based LRTC.

It is known that tensor nuclear norm is hard to compute in practice. One has to consider other tractable solutions such as learning a CP model directly from the incomplete data (Acar et al., 2011; Jain & Oh, 2014). Jain & Oh (2014) proved that an $n \times n \times n$ tensor of rank r can be recovered from $O(n^{3/2}r^5 \log^4(n))$ randomly sampled entries, via alternating minimization. Barak & Moitra (2016), Potechin & Steurer (2017), and Foster & Risteski (2019) used the sum-of-squares hierarchy as a tractable relaxation to the tensor nuclear norm. For 3rd-order tensor completion, Bazerque et al. (2013) used the sum of squared Frobenius norms as regularizer and showed that it is related to the $\ell_{2/3}$ norm of the weights in CP decomposition. Yang et al. (2016) applied group-sparse regularization to LRTC and showed that the regularizer is related to the $\ell_{1/3}$ norm. Shi et al. (2017) proposed to directly minimize the ℓ_1 norm of the weights in CP decomposition for LRTC. In fact, the regularizers presented in the above three works are closely related to the tensor Schatten- p (quasi) norm with $p = 1, 2/3$, or $1/3$. Notice that these works are purely empirical-motivated and have

no theoretical guarantee on the tensor completion performance. Moreover, the regularizations are limited in the discrete values $\{1, 2/3, 1/3\}$ and only for 3rd-order tensors. One may expect to exploit Schatten- p quasi-norm with arbitrary p on arbitrary-order tensor and have theoretical guarantees for the recovery.

Besides LRTC, tensor robust principal component analysis (TRPCA) is another important problem of tensor recovery. TRPCA is a generalization of robust PCA (De la Torre & Black, 2001; Ding et al., 2006; Candès et al., 2011; Haeffele & Vidal, 2019) and aims to decompose a noisy tensor to the sum of a low-rank tensor and a sparse tensor. Many recent works on TRPCA can be found in (Anandkumar et al., 2016; Lu et al., 2016; Xie et al., 2018; Goldfarb & Qin, 2014; Zheng et al., 2019; Lu et al., 2019; Wang et al., 2020). For example, Lu et al. (2019) defined a new tensor nuclear norm based on the t -product (Kilmer & Martin, 2011) of tensors and provided sufficient conditions for exact recovery. Notice that these TRPCA algorithms have high computational costs on large-scale data. One approach to reducing the computational costs is using low-rank factorization, which, however, requires the estimation of rank.

In this paper, we focus on fast and accurate LRTC and TRPCA. Our contributions are two-fold.

- We propose a new class of regularizers as a tensor rank proxy for tensors of arbitrary order. These are based on the Euclidean norms of the component vectors in the form of CP decomposition. We show that the regularizers are monotonic transformations of Schatten- p quasi-norms, where p could be any positive value. We also provide asymmetric variational forms for the tensor Schatten- p quasi-norm with discrete p . When applying the regularizers to LRTC and TRPCA, empirically, the recovery performance is robust to the choice of initial rank and the recovery accuracy is high when p is much less than 1.
- We provide generalization error bounds for LRTC with Schatten- p quasi-norm regularization and LRTC with our regularizers. We show that a smaller p (but not too small) in the tensor Schatten- p quasi-norm or a smaller q in our factorization regularizers lead to a tighter generalization error bound. Note that the bounds are also applicable to order-2 tensors. In other words, our bound works for LRMC. We also provide a recovery error bound for TRPCA with Schatten- p quasi-norm regularization, which verifies the usefulness of small p in TRPCA.

The numerical results of LRTC and TRPCA on synthetic data, image inpainting, and image denoising corroborate the effectiveness and superiority of our methods over state-of-the-art baseline methods.

2 Euclidean Regularization

First of all, recall the following definitions in terms of the CP decomposition.

Definition 1. Let $\mathbf{x}_i^{(j)} \in \mathbb{R}^{n_j \times 1}$, $i \in [r]$, $j \in [d]$. The rank of a tensor $\mathcal{X} \in \mathcal{R}^{n_1 \times n_2 \times \dots \times n_d}$ is defined as the minimum number of rank-one tensors that sum to \mathcal{X} :

$$\text{rank}(\mathcal{X}) = \min \left\{ r \in \mathbb{N} : \mathcal{X} = \sum_{i=1}^r \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \dots \circ \mathbf{x}_i^{(d)} \right\}.$$

Definition 2. The nuclear norm (Friedland & Lim, 2018) of a tensor \mathcal{X} is defined as

$$\|\mathcal{X}\|_* = \inf \left\{ \sum_{i=1}^r |s_i| : \mathcal{X} = \sum_{i=1}^r s_i \mathbf{u}_i^{(1)} \circ \mathbf{u}_i^{(2)} \dots \circ \mathbf{u}_i^{(d)}, \|\mathbf{u}_i^{(j)}\| = 1, r \in \mathbb{N} \right\}.$$

Using Definition 2, we see the tensor nuclear norm is a convex relaxation of tensor rank. However, both the tensor rank and tensor nuclear norm are NP-hard to compute (Friedland & Lim, 2018). Similarly, we may define a tensor Schatten- p quasi-norm that extends matrix Schatten- p quasi-norm. The p -th power of the tensor Schatten- p quasi-norm is a nonconvex relaxation of tensor rank.

Definition 3. The Schatten- p quasi-norm ($0 < p < 1$) of a tensor \mathcal{X} is defined as

$$\|\mathcal{X}\|_{S_p} = \inf \left\{ \left(\sum_{i=1}^r |s_i|^p \right)^{1/p} : \mathcal{X} = \sum_{i=1}^r s_i \mathbf{u}_i^{(1)} \circ \mathbf{u}_i^{(2)} \dots \circ \mathbf{u}_i^{(d)}, \|\mathbf{u}_i^{(j)}\| = 1, r \in \mathbb{N}, 0 < p < 1 \right\}.$$

For convenience, we make the following definition.

Definition 4. Let $d, k \in \mathbb{N}$ and given a set of vectors $\{\mathbf{x}_i^{(j)}\}_{i \in [k], j \in [d]}$. Define $\mathcal{CP}_k^d(\mathbf{x}_i^{(j)}) := \sum_{i=1}^k \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \cdots \circ \mathbf{x}_i^{(d)}$.

Note that $\mathcal{CP}_k^d(\mathbf{x}_i^{(j)})$ is just a shorthand for the sum of the k outer products of d vectors. We provide the following variational form of the tensor Schatten- p (quasi) norm ($p > 0$):

Theorem 1 (Symmetric Regularizer¹). *The tensor Schatten- q/d (quasi) norm for any real positive q and positive integer d can be represented as a function of the Euclidean norms of the CP components:*

$$\|\mathcal{X}\|_{S_{q/d}}^{q/d} = \inf_{\mathcal{X}=\mathcal{CP}_k^d(\mathbf{x}_i^{(j)})} \frac{1}{d} \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|^q.$$

Note that in the theorem $k \geq \text{rank}(\mathcal{X})$ is assumed implicitly via the constraint $\mathcal{X} = \mathcal{CP}_k^d(\mathbf{x}_i^{(j)})$. Compared to the definition (Definition 3) of the tensor Schatten- p quasi-norm, $\|\mathcal{X}\|_{S_p}$ given by Theorem 1 internalizes the weight factors \mathbf{s} into the vectors \mathbf{u} and removes the constraints on \mathbf{u} . This reduces decision variables and the computational cost of calculating the gradient in optimization. In addition, for some choice of p , there is no nonsmooth functions on the factors in Theorem 1, while there are always nonsmooth functions in Definition 3. Hence, the formulation of Schatten- p quasi-norm in Theorem 1 is more tractable than that in Definition 3. For instance, in Theorem 1, letting $p = 1, \frac{2}{d}$, or $\frac{1}{d}$, we have $\|\mathcal{X}\|_* = \inf_{\mathcal{X}=\mathcal{CP}_k^d(\mathbf{x}_i^{(j)})} \frac{1}{d} \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|$, $\|\mathcal{X}\|_{S_{2/d}}^{2/d} = \inf_{\mathcal{X}=\mathcal{CP}_k^d(\mathbf{x}_i^{(j)})} \frac{1}{d} \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|^2$, and $\|\mathcal{X}\|_{S_{1/d}}^{1/d} = \inf_{\mathcal{X}=\mathcal{CP}_k^d(\mathbf{x}_i^{(j)})} \frac{1}{d} \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|$. These three special cases are based on convex functions of Euclidean norms of component vectors, and hence are possibly easier to handle in optimization.

According to Theorem 1, for a relatively low-order tensor, when we want to obtain a sharp enough regularizer, we have to use a sufficiently small p for all i and j . Thus, every term in $\{\|\mathbf{x}_i^{(j)}\|^{pd}\}_{i \in [k], j \in [d]}$ is nonconvex and nonsmooth, making it difficult to solve the optimization problem of low-rank tensor recovery. The following theorem provides a class of asymmetric regularizers that have fewer nonconvex and nonsmooth terms than those in Theorem 1.

Theorem 2 (Asymmetric Regularizer). *Suppose $q \in \{1, 1/2, 1/3, 1/4, \dots\}$. Let $p_1 = q/(1 + qd - q)$ and $p_2 = 2q/(2 + qd - q)$. We have*

$$\begin{aligned} (a) \quad \|\mathcal{X}\|_{S_{p_1}}^{p_1} &= \inf_{\mathcal{X}=\mathcal{CP}_k^d(\mathbf{x}_i^{(j)})} p_1 \sum_{i=1}^k \left(\frac{1}{q} \|\mathbf{x}_i^{(1)}\|^q + \sum_{j=2}^d \|\mathbf{x}_i^{(j)}\| \right); \\ (b) \quad \|\mathcal{X}\|_{S_{p_2}}^{p_2} &= \inf_{\mathcal{X}=\mathcal{CP}_k^d(\mathbf{x}_i^{(j)})} p_2 \sum_{i=1}^k \left(\frac{2}{q} \|\mathbf{x}_i^{(1)}\|^q + \sum_{j=2}^d \|\mathbf{x}_i^{(j)}\|^2 \right). \end{aligned}$$

In Theorem 2 (b), the terms of $j \geq 2$ are convex and smooth while those in Theorem 2 (a) are convex but nonsmooth. Therefore, the optimization related to Theorem 2 (b) is easier. Since 3rd-order tensors (e.g. color images and videos) are more prevalent than tensors with other orders, we list their symmetric and asymmetric regularizers with only convex terms in Table 1 for convenience. Note that the last two regularizers in the table are not the consequences of Theorem 2. The derivations are in Appendix D.3.

In sum, the regularizers we proposed in this section covers all Schatten- p (quasi) norms with any $0 < p \leq 1$ for any-order tensors. Some of the regularizers, especially those asymmetric ones, are based on convex or/and smooth functions on the tensor factors, which provide convenience for application and optimization. These regularizers can be applied to LRTC and TRPCA that enjoy a variety of real applications in machine learning and computer vision.

¹It is worth mentioning that the Theorem 2 of (Cheng et al., 2016) and the Proposition 1 of (Lacroix et al., 2018) are two special cases of our Theorem 1 (when $p = 1$ and $d = 3$ or $p = 2/3$ and $d = 3$).

Table 1: Regularizers ($\mathcal{R}(\mathcal{X})$) given by the sum of only convex terms for 3rd-order tensor ($\mathcal{X} = \mathcal{CP}_k^3(\mathbf{x})$).

	$\mathcal{R}(\mathcal{X})$	Characterization based on Euclidean norm	values of q, p_1, p_2 in Theorems 1 and 2
Symmetric	$\ \mathcal{X}\ _*$	$\frac{1}{3} \sum_{i=1}^k (\ \mathbf{x}_i^{(1)}\ ^3 + \ \mathbf{x}_i^{(2)}\ ^3 + \ \mathbf{x}_i^{(3)}\ ^3)$	$q = 3$ (Theorem 1)
	$\ \mathcal{X}\ _{S_{2/3}}^{2/3}$	$\frac{1}{3} \sum_{i=1}^k (\ \mathbf{x}_i^{(1)}\ ^2 + \ \mathbf{x}_i^{(2)}\ ^2 + \ \mathbf{x}_i^{(3)}\ ^2)$	$q = 2$ (Theorem 1)
	$\ \mathcal{X}\ _{S_{1/3}}^{1/3}$	$\frac{1}{3} \sum_{i=1}^k (\ \mathbf{x}_i^{(1)}\ + \ \mathbf{x}_i^{(2)}\ + \ \mathbf{x}_i^{(3)}\)$	$q = 1$ (Theorem 1)
Asymmetric	$\ \mathcal{X}\ _{S_{1/2}}^{1/2}$	$\frac{\sqrt{2}}{4} \sum_{i=1}^k (\ \mathbf{x}_i^{(1)}\ ^2 + \ \mathbf{x}_i^{(2)}\ ^2 + \ \mathbf{x}_i^{(3)}\ ^2)$	$q = 1, p_2 = 1/2$ (Theorem 2)
	$\ \mathcal{X}\ _{S_{2/5}}^{2/5}$	$\frac{16^{1/5}}{5} \sum_{i=1}^k (\ \mathbf{x}_i^{(1)}\ ^2 + \ \mathbf{x}_i^{(2)}\ ^2 + \ \mathbf{x}_i^{(3)}\ ^2)$	derived from Appendix D.3
	$\ \mathcal{X}\ _{S_{3/7}}^{3/7}$	$\frac{81^{1/7}}{7} \sum_{i=1}^k (\ \mathbf{x}_i^{(1)}\ ^3 + \ \mathbf{x}_i^{(2)}\ ^3 + \ \mathbf{x}_i^{(3)}\ ^3)$	derived from Appendix D.3

The tensor regularizers we presented in this paper are closely related to the variational forms of matrix nuclear norm and Schatten- p quasi-norms (Shang et al., 2016; 2017; Fan et al., 2019; Giampouras et al., 2020). For instance, the squared sum of Frobenius norms of two factors of a matrix is lower bounded by the nuclear norm of the matrix, which is a special case of our Theorem 1 with $d = 2$ and $p = 1$. In (Fan et al., 2019), the authors provided a class of SVD-free variational forms of the matrix Schatten- p quasi-norm with discrete p that can be arbitrarily small. Our tensor regularizers can be regarded as a generalization of the variational form of Schatten- p quasi-norm from matrix to tensor. For example, in Theorem 1, when $d = 2$ and $p = 1/2$, the regularizer is equivalent to the FGSR-1/2 regularizer of (Fan et al., 2019), which is related to the Schatten-1/2 quasi-norm of matrix. In Theorem 2(b), when $d = 2$ and $q = 1$, the regularizer is equivalent to the FGSR-2/3 regularizer of (Fan et al., 2019), which is related to the Schatten-2/3 quasi-norm of matrix. Nevertheless, theoretical guarantees about these tensor regularizers in low-rank tensor recovery are more difficult to derive than their matrix counterparts.

3 Low-Rank Tensor Completion

3.1 LRTC-ENR algorithm

Let $\mathcal{X}^* \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$ be a rank- r tensor. Suppose we observed a few noisy entries of \mathcal{X}^* randomly (without replacement):

$$[\mathcal{D}]_{j_1 j_2 \dots j_d} = [\mathcal{X}^*]_{j_1 j_2 \dots j_d} + [\mathcal{N}^*]_{j_1 j_2 \dots j_d}, \quad (j_1, \dots, j_d) \in \Omega$$

where Ω consists of the locations of the observed entries and each entry of the noise tensor \mathcal{N}^* is drawn from $\mathcal{N}(0, \sigma^2)$. To recover \mathcal{X}^* from \mathcal{D} , one may solve

$$\underset{\mathcal{X}}{\text{minimize}} \frac{1}{2} \|\mathcal{P}_\Omega(\mathcal{D} - \mathcal{X})\|_F^2 + \lambda \|\mathcal{X}\|_{S_p}^p, \quad (1)$$

where $[\mathcal{P}_\Omega(\mathcal{Y})]_{j_1 j_2 \dots j_d} = [\mathcal{Y}]_{j_1 j_2 \dots j_d}$ if $(j_1, \dots, j_d) \in \Omega$ and $[\mathcal{P}_\Omega(\mathcal{Y})]_{j_1 j_2 \dots j_d} = 0$ otherwise. The solution of (1) is an estimate of \mathcal{X}^* . However, Problem (1) is intractable because the computation of the Schatten- p quasi-norm is NP-hard. Instead, based on the analysis in Section 2, assuming $k \geq r$, we propose to solve

$$\underset{\{\mathbf{x}_i^{(j)}\}}{\text{minimize}} \frac{1}{2} \left\| \mathcal{P}_\Omega \left(\mathcal{D} - \mathcal{CP}_k^d(\mathbf{x}_i^{(j)}) \right) \right\|_F^2 + \lambda \mathcal{R}(\{\mathbf{x}_i^{(j)}\}), \quad (2)$$

where $\mathcal{R}(\{\mathbf{x}_i^{(j)}\})$ denotes a certain regularizer in Theorem 1, Theorem 2, or Table 1, e.g.,

$$\mathcal{R}(\{\mathbf{x}_i^{(j)}\}) = \frac{1}{d} \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|^q = \frac{1}{d} \sum_{j=1}^d \|\mathbf{X}^{(j)}\|_{2,q}^q, \quad (3)$$

where $\mathbf{X}^{(j)} = [\mathbf{x}_1^{(j)}, \dots, \mathbf{x}_k^{(j)}]$. Similarly, the regularizers in Theorem 2 (a) and (b) can be formulated as $p_1(\frac{1}{q} \|\mathbf{X}^{(1)}\|_{2,q}^q + \sum_{j=2}^d \|\mathbf{X}^{(j)}\|_{2,1})$ and $p_2(\frac{2}{q} \|\mathbf{X}^{(1)}\|_{2,q}^q + \sum_{j=2}^d \|\mathbf{X}^{(j)}\|_F^2)$ respectively. Note that even finding a stationary point of (1) is hard due to the computation of the Schatten- p quasi-norm while finding a stationary

point of (2) is computationally feasible. For convenience, we call problem (2) Low-Rank Tensor Completion with Euclidean Norm Regularization (**LRTC-ENR**).

We propose to solve (2) by Block Coordinate Descent (BCD) with Extrapolation (BCDE for short) (Xu & Yin, 2013), which is usually more efficient than BCD alone in practice. The decision variables can be organized into d blocks (d matrices, i.e., $\mathbf{X}^{(j)}$, $j = 1, \dots, d$). Then we need to iteratively solve d subproblems related to the ℓ_q norm minimization. When $q < 1$, we integrate BCDE with iteratively reweighted method (Lu, 2014). If we use the symmetric regularizers given by Theorem 1 with $q \leq 1$, we need to solve d nonsmooth subproblems (related to $\|\mathbf{X}^{(j)}\|_{2,q}$, $j = 1, \dots, d$) in every iteration and when $q < 1$ we need to perform iteratively reweighted method for $\|\mathbf{X}^{(j)}\|_{2,q}$, $j = 1, \dots, d$. In contrast, if we use the asymmetric regularizers given by Theorem 2 (b), the number of nonsmooth subproblems (related to $\|\mathbf{X}^{(1)}\|_{2,q}$) is reduced to 1 and when $q < 1$ we only need to perform iteratively reweighted method for $\|\mathbf{X}^{(1)}\|_{2,q}$. Therefore, the optimization related to the asymmetric regularizers is more efficient than that related to the symmetric regularizers.

We may also use quasi-Newton methods such as L-BFGS (Liu & Nocedal, 1989) to solve problem (2) even when the objective function is nonsmooth. The details of the optimization for (2) are in Appendix A.

It is worth mentioning that according to Definition 3, one may consider the following LRTC problem

$$\begin{aligned} & \underset{\{\mathbf{x}_i^{(j)}\}, \mathbf{s}}{\text{minimize}} \quad \frac{1}{2} \left\| \mathcal{P}_\Omega \left(\mathcal{D} - \widetilde{\mathcal{CP}}_k^d(\{\mathbf{x}_i^{(j)}\}, \mathbf{s}) \right) \right\|_F^2 + \lambda \sum_{i=1}^k |\mathbf{s}_i|^p, \\ & \text{subject to} \quad \|\mathbf{x}_i^{(j)}\| = 1, \forall i \in [k], j \in [d]. \end{aligned} \quad (\text{LRTC-Schatten-}p)$$

where $\widetilde{\mathcal{CP}}_k^d(\{\mathbf{x}_i^{(j)}\}, \mathbf{s}) = \sum_{i=1}^k s_i \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \dots \circ \mathbf{x}_i^{(d)}$. Nevertheless, it is more difficult to solve LRTC-Schatten- p than LRTC-ENR due to the following reasons. First, LRTC-Schatten- p has one more block of variables \mathbf{s} , of which the gradient computation is costly (it requires evaluating every rank-one component of \mathcal{X}). Second, constrained optimization LRTC-Schatten- p is generally harder than unconstrained optimization. Heuristically, we suggest using BCDE with iteratively reweighted minimization embedded to solve LRTC-Schatten- p . It is compared with LRTC-ENR in Figure 3 of Section 5.1.1.

3.2 Generalization Error Bound of LRTC

In this section, we study the generalization error bounds of LRTC-Schatten- p and LRTC-ENR. The generalization error of tensor completion is defined as the difference between the prediction error for all entries and the training error for the observed entries, i.e.,

$$\mathcal{GE}_{\text{TC}} := \mathcal{E}(\mathcal{D}, \mathcal{X}, [n] \times \dots \times [n]) - \mathcal{E}(\mathcal{D}, \mathcal{X}, \Omega).$$

In this study, we let $\mathcal{E}(\mathcal{D}, \mathcal{X}, [n] \times \dots \times [n]) = \frac{1}{\sqrt{n^d}} \|\mathcal{D} - \mathcal{X}\|_F$ and $\mathcal{E}(\mathcal{D}, \mathcal{X}, \Omega) = \frac{1}{\sqrt{|\Omega|}} \|\mathcal{P}_\Omega(\mathcal{D} - \mathcal{X})\|_F$, though their squares can also be used. It should be pointed out that, compared to LRMC, it is much more difficult to analyze the generalization ability of LRTC. The main reason is that we may not have orthogonal factors in a CP decomposition. Thus, we first consider the case of orthogonal CP decomposition. We will also provide generalization error bounds for the general cases (without the restriction of orthogonality).

Without loss of generality, we consider the case of hyper-cubic tensors, denoted by $\mathcal{X} \in \mathbb{R}^{n^{\otimes d}}$. For convenience, we define the following two tensor sets.

Definition 5 (Orthogonal CP tensor sets). *Let $\mathcal{S}_{d,n}^\perp$ be the set of hyper-cubic tensors with orthogonal CP factors. Denote the Schatten- p quasi-norm of $\mathcal{X} \in \mathcal{S}_{d,n}^\perp$ by $\|\mathcal{X}\|_{\mathcal{S}_p^\perp}$. Let $\mathcal{S}_{d,n,p}^\perp$ be the set of tensors in $\mathcal{S}_{d,n}^\perp$ and with bounded (by ψ_p) Schatten- p quasi norm. To be more precise,*

$$\begin{aligned} (a) \quad \mathcal{S}_{d,n}^\perp &:= \left\{ \mathcal{X} \in \mathbb{R}^{n^{\otimes d}} : \mathcal{X} = \sum_{i=1}^r s_i \mathbf{u}_i^{(1)} \circ \mathbf{u}_i^{(2)} \dots \circ \mathbf{u}_i^{(d)}; \forall j \in [d], \mathbf{u}_i^{(j)\top} \mathbf{u}_l^{(j)} = 1 \text{ if } i = l, \mathbf{u}_i^{(j)\top} \mathbf{u}_l^{(j)} = 0 \text{ if } i \neq l \right\}, \\ (b) \quad \mathcal{S}_{d,n,p}^\perp &:= \{ \mathcal{X} \in \mathcal{S}_{d,n}^\perp : \|\mathcal{X}\|_{\mathcal{S}_p^\perp} \leq \psi_p \}. \end{aligned}$$

We have the following covering number result.

Theorem 3. *The covering numbers of $\mathcal{S}_{d,n,p}^\perp$ with respect to the Frobenius norm satisfy*

$$\log \mathcal{N}(\mathcal{S}_{d,n,p}^\perp, \|\cdot\|_F, \epsilon) \leq \left(\frac{1}{2} + \frac{1}{p}\right) nd (\log(d+1)) \left(\frac{c\psi_p}{\epsilon}\right)^{\frac{2p}{2-p}},$$

where $c > 0$ is a universal constant and $0 < p < 2$.

Note that although Theorem 3 does not include the case $p = 2$, it is much easier to obtain the covering number of $\mathcal{S}_{d,n,2}^\perp$, which will not be detailed in this paper because $\|\mathcal{X}\|_{S_2^\perp}$ is useless in regularizing the tensor rank.

Based on Theorem 3, we derive the following generalization error bound for Schatten- p quasi-norm regularized orthogonal tensor completion (minimize $\mathcal{X} \frac{1}{2} \|\mathcal{P}_\Omega(\mathcal{D} - \mathcal{X})\|_F^2 + \lambda \|\mathcal{X}\|_{S_p^\perp}^p$).

Theorem 4. *Suppose $\mathcal{D} \in \mathbb{R}^{n^{\otimes d}}$, $\mathcal{X} \in \mathcal{S}_{d,n}^\perp$, $\max\{\|\mathcal{D}\|_\infty, \|\mathcal{X}\|_\infty\} \leq \epsilon$, and $0 < p < 2$. Then there exists a numerical constant c such that with probability at least $1 - 2n^{-d}$,*

$$\frac{1}{\sqrt{n^d}} \|\mathcal{D} - \mathcal{X}\|_F - \frac{1}{\sqrt{|\Omega|}} \|\mathcal{P}_\Omega(\mathcal{D} - \mathcal{X})\|_F \leq c\epsilon \left(\frac{(\frac{1}{2} + \frac{1}{p})nd \log(d+1)}{|\Omega|} \left(\frac{\|\mathcal{X}\|_{S_p^\perp}}{\epsilon \sqrt{dn}} \right)^{\frac{2p}{2-p}} \right)^{1/4}. \quad (4)$$

In the theorem, since $\|\mathcal{X}\|_{S_p^\perp}/(\epsilon \sqrt{dn}) > 1$, the bound is a U-shape function of p when $0 < p < 2$. Therefore, there exists a threshold $\bar{p} = \operatorname{argmin}_{0 < p < 2} (1/2 + 1/p) (\|\mathcal{X}\|_{S_p^\perp}/(\epsilon \sqrt{dn}))^{2p/(2-p)}$ such that when $\bar{p} \leq p < 2$, a smaller p will lead to a tighter error bound. However, it is difficult to obtain \bar{p} analytically. Here we use a toy example to illustrate the result of Theorem 4. According to Definition 3, we generate \mathbf{s} using the following procedures: 1) $\mathbf{s} = [2, 1.9, 1.8, \dots, 0.2, 0.1]$; 2) $s_i \leftarrow s_i', i = 1, 2, \dots, 20$; 3) $\mathbf{s} \leftarrow \mathbf{s} / \sum_i s_i$. We choose ν from $\{0.1, 1, 5, 10, 20, 50\}$ and generate six different \mathbf{s} . We see that a larger ν leads to a higher the decay rate of the element of \mathbf{s} . We use these six \mathbf{s} together with three random orthogonal matrices \mathbf{U} of size 100×20 to form six low-rank tensors of size $100 \times 100 \times 100$ and rank 20. These six tensors have different "low-rankness" because of the different ν . For instance, the tensor corresponding to $\nu = 50$ can be well approximated by a rank-4 tensor, where the relative error is less than 0.01%. We let $|\Omega| = 0.1n^3$ and compute $((\frac{1}{2} + \frac{1}{p})nd \log(d+1)|\Omega|^{-1} (\|\mathcal{X}\|_{S_p^\perp}/(\epsilon \sqrt{dn}))^{\frac{2p}{2-p}})^{1/4} \triangleq \Delta$ for each tensor. The six \mathbf{s} and average Δ (corresponding to each \mathbf{s}) of 100 repeated trials are shown in Figure 1. We see that in each case, a smaller p but not too small will lead to a smaller Δ . When the decay rate is higher, the approximate rank is lower, which leads to a smaller Δ . On the other hand, when the decay rate is lower, reducing p is more useful for reducing the error bound Δ .

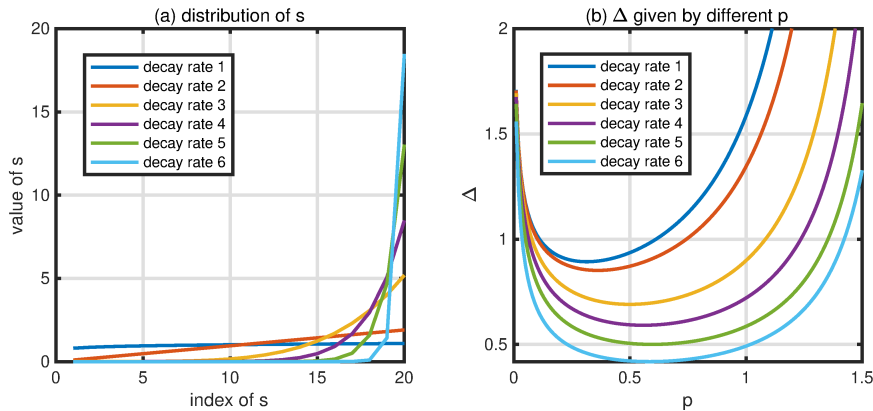


Figure 1: An intuitive example of the error bound of Theorem 4

Now we study the generalization error of LRTC-ENR (2). For convenience, let $\|\mathbf{X}^{(j)}\|_{2,q} := \left(\sum_i \|\mathbf{x}_i^{(j)}\|^q\right)^{1/q}$ and define

$$\mathcal{S}_{d,n,q}^k = \left\{ \mathbf{X} \in \mathbb{R}^{n \otimes d} : \mathbf{X} = \mathcal{CP}_k^d(\mathbf{x}_i^{(j)}); \|\mathbf{X}^{(j)}\|_{2,q} \leq \alpha_q^{(j)}, \|\mathbf{X}^{(j)}\|_{op} \leq \gamma_j, j \in [d] \right\}, \quad (5)$$

where $\mathbf{X}^{(j)} = [\mathbf{x}_1^{(j)}, \mathbf{x}_2^{(j)}, \dots, \mathbf{x}_k^{(j)}]$. We have

Theorem 5. *The covering numbers of $\mathcal{S}_{d,n,q}^k$ with respect to the Frobenius norm satisfy*

$$\log \mathcal{N}(\mathcal{S}_{d,n,q}^k, \|\cdot\|_F, \epsilon) \leq \frac{c(n + \log(ek))}{q} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j\epsilon} \right)^{\frac{2q}{2-q}},$$

for any $0 < q < 2$, and

$$\log \mathcal{N}(\mathcal{S}_{d,n,q}^k, \|\cdot\|_F, \epsilon) \leq c'nk^{2(q-1)/q} \log(2nk) \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j\epsilon} \right)^2$$

for any $q \geq 2$, where $\phi = \prod_{j=1}^d \gamma_j$ and c, c' are universal constants.

Based on Theorem 5, we have the following generalization error bound for LRTC-ENR (2) with $\mathcal{R}(\{\mathbf{x}_i^{(j)}\}) = \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|^q$.

Theorem 6. *Suppose $\mathcal{D} \in \mathbb{R}^{n \otimes d}$, $\mathbf{X} \in \mathcal{S}_{d,n,q}^k$, $\phi = \prod_{j=1}^d \gamma_j$, and $\max\{\|\mathcal{D}\|_\infty, \|\mathbf{X}\|_\infty\} \leq \epsilon$. Denote $\bar{\Omega}$ the index set of the missing entries of \mathcal{D} and suppose $|\bar{\Omega}| > |\Omega| > 50$. Then with probability at least $1 - \delta$ over the random partition Ω and $\bar{\Omega}$, we have*

$$\frac{1}{|\bar{\Omega}|} \|\mathcal{P}_{\bar{\Omega}}(\mathcal{D} - \mathbf{X})\|_F^2 \leq \frac{1}{|\Omega|} \|\mathcal{P}_{\Omega}(\mathcal{D} - \mathbf{X})\|_F^2 + \epsilon^2(B_{\mathcal{R}} + B_{\delta})$$

where $B_{\delta} = \frac{44n^d}{\sqrt{|\Omega||\bar{\Omega}|}} + 12\sqrt{\frac{n^d}{|\Omega||\bar{\Omega}|} \log \frac{1}{\delta}}$ and

$$B_{\mathcal{R}} = \begin{cases} \frac{c_1}{|\Omega|} \sqrt{nk^{2(q-1)/q} \log(2nk) \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\epsilon\gamma_j} \right)^2 \log |\Omega|} & \text{if } q \geq 1 \\ \frac{c_2}{|\Omega|} \sqrt{\frac{(n + \log(ek))(2-q)^2}{q(2-2q)^2} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\epsilon\gamma_j} \right)^{\frac{2q}{2-q}} |\Omega|^{\frac{1-q}{4-2q}}} & \text{if } 0 < q < 1 \end{cases}$$

with constants c_1 and c_2 .

We see that in Theorem 6, the bound is not continuous around $q = 1$. In the theorem when $0 < q < 1$, the bound is a U-shape function of q . It indicates that a smaller q but not too small leads to a tighter error bound and for a smaller $|\Omega|$ reducing the value of q becomes more useful. We show an intuitive example by Figure 2 to illustrate the role of q and $|\Omega|$ in the error bound of Theorem 6. In the example, we use $d = 3$, $n = 100$, and $k = 20$ to generate synthetic random tensors $\mathbf{X} = \sum_{i=1}^{20} \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \circ \mathbf{x}_i^{(3)}$, where $\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(3)}$ are drawn from Gaussian distribution and $\|\mathbf{x}_i^{(1)}\| = \dots = \|\mathbf{x}_i^{(3)}\| = \delta_i$. We consider six different decay rates for $\delta_1, \dots, \delta_{20}$ corresponding to different levels of low-rankness. In general, a larger decay rate means an easier tensor recovery problem. We let $\rho = \frac{|\Omega|}{n^d}$ and

$$\Delta = \frac{1}{|\Omega|} \sqrt{\frac{(n + \log(ek))(2-q)^2}{q(2-2q)^2} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\epsilon\gamma_j} \right)^{\frac{2q}{2-q}} |\Omega|^{\frac{1-q}{4-2q}}}, \quad (6)$$

where $0 < q < 1$. In the sub-figures (b,c,d), we see that a smaller q but not too small provides a tighter error bound. When the decay rate is smaller, reducing the value of q becomes more useful, which is consistent

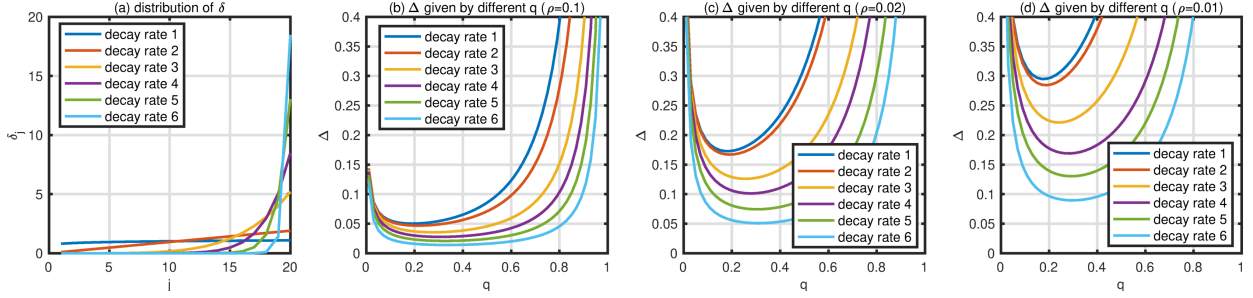


Figure 2: An intuitive exxample for the error bound of Theorem 6.

with the result of Figure 1. In addition, when ρ is smaller, namely the problem is harder, reducing the value of q becomes much more effective.

Based on Theorem 6, we can obtain the generalization error bound for (1), which is shown in the following corollary.

Corollary 1. Suppose $\mathcal{D} \in \mathbb{R}^{n \times d}$, $\mathcal{X} \in \mathbb{R}^{n \times d}$, $\max\{\|\mathcal{D}\|_\infty, \|\mathcal{X}\|_\infty\} \leq \varepsilon$, $\|\mathcal{X}\|_{S_p}$ is attained at $\mathcal{X} = \mathcal{CP}_k^d(\mathbf{x}_i^{(j)})$, and $\max_{j \in [d]} \|\mathbf{X}^{(j)}\|_{op} \leq \bar{\gamma}$. Denote $\bar{\Omega}$ the index set of the missing entries of \mathcal{D} and suppose $|\bar{\Omega}| > |\Omega| > 50$. Then with probability at least $1 - \delta$ over the random sampling of Ω , we have

$$\frac{1}{|\bar{\Omega}|} \|\mathcal{P}_{\bar{\Omega}}(\mathcal{D} - \mathcal{X})\|_F^2 \leq \frac{1}{|\Omega|} \|\mathcal{P}_{\Omega}(\mathcal{D} - \mathcal{X})\|_F^2 + \varepsilon^2(B_{\mathcal{R}} + B_{\delta})$$

where $B_{\delta} = \frac{44n^d}{\sqrt{|\Omega||\bar{\Omega}|}} + 12\sqrt{\frac{n^d}{|\Omega||\bar{\Omega}|} \log \frac{1}{\delta}}$ and

$$B_{\mathcal{R}} = \begin{cases} \frac{c_1}{|\bar{\Omega}|} \sqrt{nk^{2(pd-1)/pd} \log(2nk)d \left(\frac{d\bar{\gamma}-1}{\varepsilon}\right)^2 \|\mathcal{X}\|_{S_p}^{2/d} \log |\Omega|} & \text{if } p \geq 1/d \\ \frac{c_2}{|\bar{\Omega}|} \sqrt{\frac{(n + \log(ek))(2 - pd)^2}{p(2 - 2pd)^2} \left(\frac{d\bar{\gamma}-1}{\varepsilon}\right)^{2pd/(2-pd)} \|\mathcal{X}\|_{S_p}^{2p/(2-pd)} |\Omega|^{\frac{1-pd}{4-2pd}}} & \text{if } 0 < p < 1/d \end{cases}$$

with constants c_1 and c_2 .

In Corollary 1 as well as Theorem 6, the error bounds are for $\frac{1}{|\bar{\Omega}|} \|\mathcal{P}_{\bar{\Omega}}(\mathcal{D} - \mathcal{X})\|_F^2$ and are almost linear with $\frac{1}{|\bar{\Omega}|}$. In contrast, in Theorem 4, the error bound is for $\frac{1}{\sqrt{n^d}} \|\mathcal{D} - \mathcal{X}\|_F$ and is linear with $(\frac{1}{|\bar{\Omega}|})^{1/4}$. Therefore, the bounds of Corollary 1 and Theorem 6 are tighter than that of Theorem 4, however, at a price of breaking the continuity at $q = 1$ or $p = 1/d$.

Consider the terms related to p in $B_{\mathcal{R}}$ of Corollary 1 with $p \geq 1/d$, i.e., $k^{2(pd-1)/pd} \|\mathcal{X}\|_{S_p}^{2/d} \triangleq \pi_p$. Let $1/d \leq p_2 < p_1$. Using the power-mean inequality², we have $k^{-1/p_2} \|\mathcal{X}\|_{S_{p_2}} \leq k^{-1/p_1} \|\mathcal{X}\|_{S_{p_1}}$. It follows that $\pi_{p_2} = k^{2(p_2d-1)/p_2d} \|\mathcal{X}\|_{S_{p_2}}^{2/d} \leq k^{2(p_2d-1)/p_2d} (k^{1/p_2-1/p_1})^{2/d} \|\mathcal{X}\|_{S_{p_1}}^{2/d} = \pi_{p_1}$, where the equality holds if and only if $s_1 = s_2 = \dots = s_k$ (it will not happen in practice almost surely). Therefore, when $1/d \leq p_2 < p_1$, the error bound given by p_2 is always tighter than that given by p_1 . This provides a rule of thumb for real applications that p should be at most $1/d$. Together with the result of Corollary 1 for $0 < p < 1/d$, we suggest using $p = 1/2d$.

Comparison with previous work Barak & Moitra (2016), Foster & Risteski (2019) and Wimalawarne & Mamitsuka (2021) have also studied the generalization error of CP based tensor completion. Their bounds

²For any $z_1, z_2, \dots, z_n \geq 0$, if $\infty > \alpha > \beta > 0$, we have $(\frac{1}{n} \sum_{i=1}^n z_i^\alpha)^{\frac{1}{\alpha}} \geq (\frac{1}{n} \sum_{i=1}^n z_i^\beta)^{\frac{1}{\beta}}$ with equality if and only if $z_1 = z_2 = \dots = z_n$.

are based on the nuclear norm (or Schatten-1 norm equivalently). We compare their bounds with our bound in Corollary 1, in terms of third-order tensor with unit nuclear norm and r rank. The generalization error bound given by (Barak & Moitra, 2016) scales as $O(rn^{3/2} \log^4 n/|\Omega|)$. Foster & Risteski (2019) provided an empirical risk bound for orthogonal tensor completion that scales as $O(r^2 n^{3/2} \log^6 n/|\Omega|)$. The generalization error bound (Theorem 1(b)) given by (Wimalawarne & Mamitsuka, 2021) is $O(rn^{3/2} s_{\max}/|\Omega|)$, where $1/r < s_{\max} < 1$. In our Corollary 1, let $d = 3$, $p = 1$, $k = r$, $\bar{\gamma} = 1$, and $\|\mathcal{X}\|_{S_p} = 1$, our bound becomes $O(r^2 \sqrt{n \log(2nr)} \log |\Omega|/|\Omega|)$. We see our bound is tighter than all others if $r \sqrt{\log(2nr)} \log |\Omega| \leq n s_{\max}$.

4 Tensor Robust PCA

4.1 TRPCA-ENR algorithm

Let $\mathcal{X}^* \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$ be a rank- r tensor. Suppose \mathcal{X}^* is corrupted as

$$\mathcal{D} = \mathcal{X}^* + \mathcal{N}^* + \mathcal{E}^*, \quad (7)$$

where \mathcal{N}^* is a dense noise tensor drawn from $\mathcal{N}(0, \sigma^2)$ and \mathcal{E}^* is a sparse noise tensor with randomly distributed nonzero entries. To recover \mathcal{X}^* from \mathcal{D} , we wish to solve

$$\underset{\mathcal{X}, \mathcal{E}}{\text{minimize}} \quad \frac{1}{2} \|\mathcal{D} - \mathcal{X} - \mathcal{E}\|_F^2 + \lambda_x \|\mathcal{X}\|_{S_p}^p + \lambda_e \|\mathcal{E}\|_1, \quad (8)$$

but the problem is intractable. Then, based on the analysis in Section 2, we propose to solve

$$\underset{\{\mathbf{x}_i^{(j)}\}, \mathcal{E}}{\text{minimize}} \quad \frac{1}{2} \left\| \mathcal{D} - \sum_{i=1}^k \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \dots \circ \mathbf{x}_i^{(d)} - \mathcal{E} \right\|_F^2 + \lambda_x \mathcal{R}(\{\mathbf{x}_i^{(j)}\}) + \lambda_e \|\mathcal{E}\|_1, \quad (9)$$

where λ_e is a penalty parameter for the sparse tensor \mathcal{E} . For convenience, we call (9) **TRPCA-ENR**.

Regarding problem (9), we proposed to update the $d+1$ blocks of decision variables (i.e. \mathcal{E} and $\mathbf{X}^{(j)} = [\mathbf{x}_1^{(j)}, \dots, \mathbf{x}_k^{(j)}]$, $j = 1, \dots, d$) alternately. Namely, we need to iteratively solve $d+1$ subproblems related to the ℓ_q ($0 < q \leq 2$) norm minimization. Similar to LRTPC-ENR, we use the iteratively reweighted method (Lu, 2014) to solve the subproblems associated with $q < 1$. Particularly, if we use the symmetric regularizers given by Theorem 1, when $q = 2$, we use alternating minimization to solve (9) because every subproblem has a closed-form solution; when $q \leq 1$, there are d subproblems having no closed-form solution. If we use the asymmetric regularizers given by Theorem 2 (b), there is only one subproblem having no closed-form solution and hence we propose to solve (9) by the (nonconvex) alternating direction method of multipliers (ADMM) (Wang et al., 2015). We see that the asymmetric regularizers lead to much easier optimization problems than the symmetric counterparts do. The optimization for (9) is detailed in Appendix B.

4.2 Recovery Error Bound of TRPCA

Currently, it is difficult to provide recovery guarantees for problem (8) and problem (9) because of the non-convexity and non-smoothness. Instead, we consider the following constrained orthogonal tensor recovery problem

$$\begin{aligned} & \underset{\mathcal{X} \in \mathcal{S}_{d,n}^\perp, \mathcal{E}}{\text{minimize}} \quad \|\mathcal{D} - \mathcal{X} - \mathcal{E}\|_F^2 \\ & \text{subject to} \quad \|\mathcal{X}\|_{S_p^\perp}^p \leq R_x^p, \quad \|\mathcal{E}\|_{p'}^{p'} \leq R_e^{p'}, \end{aligned} \quad (10)$$

where $p, p' \in (0, 1]$, the orthogonal CP tensor set $\mathcal{S}_{d,n}^\perp$ was defined by Definition 5, and the observed tensor \mathcal{D} was generated by (7) with $n_1 = n_2 = \dots = n_d = n$. In terms of the regularization on \mathcal{E} , problem (10) is more general than (8) and (9) because here we consider the $\ell_{p'}$ regularizer. However, compared to (8) and (9), we have to consider the orthogonal tensor set for simplicity. We provide the following bound of recovery error.

Theorem 7. Let $p, p' \in (0, 1]$, $\|\mathcal{X}^*\|_{S_p^\perp}^p \leq R_x^p$, and $\|\mathcal{E}^*\|_{p'}^{p'} \leq R_e^{p'}$. Suppose $\hat{\mathcal{X}}, \hat{\mathcal{E}}$ is the optimal solution of (10) and $\max(\|\mathcal{X}^*\|_\infty, \|\hat{\mathcal{X}}\|_\infty) \leq \alpha$. Then with probability at least $1 - 4n^d$,

$$\begin{aligned} & \|\mathcal{X}^* - \hat{\mathcal{X}}\|_F^2 + \|\mathcal{E}^* - \hat{\mathcal{E}}\|_F^2 \\ & \leq (24 + 16\sqrt{2}) \left(\left(2\sqrt{2}\sigma\sqrt{dn\log(5d) + d\log(n)} \right)^{2-p} R_x^p + \left(2\alpha + 2\sigma\sqrt{d\log(n)} \right)^{2-p'} R_e^{p'} \right). \end{aligned}$$

In this theorem, for convenience, let $c_1 = 2\sqrt{2}\sigma\sqrt{dn\log(5d) + d\log(n)}$ and $c_2 = 2\alpha + 2\sigma\sqrt{d\log(n)}$. When p and p' decrease, c_1^{2-p} and $c_2^{2-p'}$ increase but R_x^p and $R_e^{p'}$ decrease, provided that there are a few elements in \mathbf{s} (refer to Definition 3) and elements in \mathcal{E}^* larger than 1. Therefore, when c_1 and c_2 are not too large, reducing the value of p or p' will lead a tighter recovery error bound for \mathcal{X}^* and \mathcal{E}^* . Although Theorem 7 is for (10) rather than (8) and (9), it verified the superiority of smaller p of the tensor Schatten- p quasi-norm in TRPCA.

5 Numerical Results

In this section, we conduct experiments of low-rank tensor completion and tensor robust PCA on both synthetic data and real image data. All the tensor computations are implemented in the MATLAB Tensor Tool Box of (Bader et al., 2019). We use a MacBook Pro with 2.6 GHz Intel Core (i7) and 16 GB RAM.

5.1 Experiments of LRTC

We compare the proposed LRTC methods with a few baselines on synthetic low-rank tensors and multi-spectral images. We use Relative recovery error $\|\mathcal{P}_{\hat{\Omega}}(\mathcal{T} - \hat{\mathcal{T}})\|_F / \|\mathcal{P}_{\hat{\Omega}}(\mathcal{T})\|_F$ to evaluate the tensor completion performance, where \mathcal{T} is the original complete tensor and $\hat{\mathcal{T}}$ denotes the recovered tensor from a tensor completion method. The minimum value of the metric is zero. A meaningful value of the metric is in the range of $[0, 1]$.

5.1.1 Synthetic data

With the size for each dimension $n = 50$, we generate noisy low-rank synthetic tensors $\mathcal{D} = \mathcal{T} + \mathcal{N} \in \mathbb{R}^{n \times n \times n}$, where $\mathcal{T} = \sum_{i=1}^r \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \circ \mathbf{x}_i^{(3)}$. The entries of $\mathbf{x}_i^{(j)} \in \mathbb{R}^n$ ($i \in [r]$, $j \in [3]$) are drawn from $\mathcal{N}(0, 1)$. The entries of the noise tensor \mathcal{N} are drawn from $\mathcal{N}(0, (0.1\sigma_{\mathcal{T}})^2)$, where $\sigma_{\mathcal{T}}$ denotes the standard deviation of the entries of \mathcal{T} . We randomly mask a fraction (which we call missing rate) of the entries to test the performance of tensor completion.

We compare the performance of LRTC with Schatten- p (quasi) norm regularization (LRTC-Schatten- p) and Euclidean norm regularization (solved by BCDE and LBFGS) in Figure 3, in which we set $r = 10$ and considered different noise levels, missing rates, and p values. In all methods, we set $k = 2r = 20$, $t_{\max} = 500$, and select λ from $\{0.01, 0.1, 0.5, 1, 5, 10, 50, 100, 500\}$. We can see that:

- Regularization is helpful.* In Figure 3 (a-d), the recovery error when there is no regularization (i.e., $\lambda = 0$) is higher than the others.
- Smaller p leads to lower recovery error?* In almost all cases, $p < 1$ outperforms $p = 1$, verifying the superiority of Schatten- p quasi-norm over nuclear norm in LRTC (both LRTC-ENR and LRTC-Schatten- p). In LRTC-ENR solved by BCDE and LBFGS, smaller p provides lower recovery error provided that p is not too small (e.g. $p \geq 1/6$), which matches Theorem 6 and Corollary 1.
- LRTC-ENR v.s. LRTC-Schatten- p* Compared to Euclidean norm minimization, direct Schatten- p norm minimization has higher recovery error and time cost because the optimization problem is more difficult to solve.

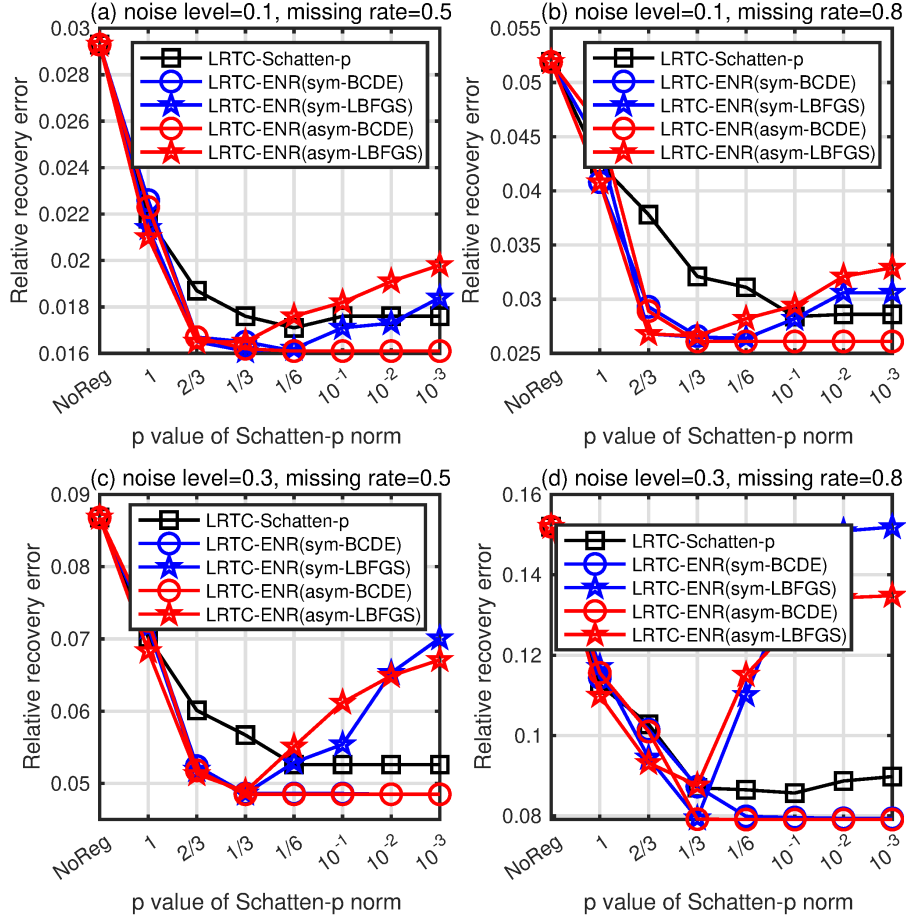


Figure 3: Performance (average of 20 repeated trials) of Schatten- p (quasi) norm with different p in the case of different noise level v for the noise distribution $\mathcal{N}(0, (v\sigma_{\mathcal{T}})^2)$; “NoReg” means LRTC without regularization (Acar et al., 2011); the time costs (500 iterations) of the five methods are 8.6s, 3.0s, 2.1s, 2.8s, and 2.3s respectively).

- d. *BCDE v.s. LBFGS* When p is too small (i.e., smaller than $1/6$), the recovery error of LRTC-ENR solved by BCDE does not change, but the recovery error of LRTC-ENR solved by LBFGS increases. One possible reason is that LBFGS is not effective in handling nonsmooth objective function and is often stuck in bad local minima or even saddle points, especially when p is very small. When $p < 1/3$, LRTC-ENR has at least one nonsmooth nonconvex subproblem, which brings difficulty to LBFGS. However, the time cost of LBFGS is less than that of BCDE.
- e. *Symmetric regularizer v.s. asymmetric regularizer* It can be found that the performance of asym-BCDE and sym-BCDE are very similar and sym-LBFGS outperforms asym-LBFGS in many cases. These empirical results are not consistent with our intuition that the asymmetric regularizers are easier to optimize than the symmetric regularizers. A possible reason is that the soft-thresholding for $q = 1$ and iteratively reweighted method for $q < 1$ of the symmetric regularizers are as effective as the gradient descent for $q = 2$ of the asymmetric regularizers. When $p < 1/3$, the q in asym-LBFGS is much less than the q in sym-LBFGS and hence leads to higher instability in computing the gradient and Hessian for LBFGS.

Based on the above results, we suggest using BCDE to solve LRTC-ENR on small datasets and using LBFGS to solve LRTC-ENR on large datasets. In Schatten- p quasi norm, we suggest using $p = 1/3$ or $p = 1/6$.

Comparison with other tensor completion methods We compare LRTC-ENR with HaLRTC (Liu et al., 2012), TenALS (Jain & Oh, 2014), TMac (Xu et al., 2013), BCPF (Zhao et al., 2015), Rprecon (Kasai & Mishra, 2016), KBR-TC (Xie et al., 2018), TRLRF (Yuan et al., 2019). In TenALS, TMac, BCPF, Rprecon, TRLRF, and LRTC-ENR, we need to determine the factorization size (or the initial rank in other word) beforehand. These methods, compared to HaLRTC and KBR-TC, may be applicable to large tensors efficiently provided that the ranks are sufficiently small. Particularly, in TMac, BCPF, and LRTC-ENR, the rank is adjusted adaptively.

Figure 4 shows the recovery of LRTC-ENR ($p = 1/3$, solved by LBFGS) and seven baselines on the synthetic data with $r = 10$. In all methods except HaLRTC and KBR-TC, one has to determine the initial rank beforehand. For these methods except Rprecon, we have set the initial rank to $2r$ because in practice it is difficult to know the true rank. The multi-rank of Rprecon is set to (r, r, r) because it performs the best. The results in the table show that Rprecon, BCPF, and LRTC-ENR specifically outperformed other methods.




Figure 4: Low-rank tensor completion on synthetic data ($r = 10$): LRTC-ENR ($p = 1/3$) v.s. seven baseline methods.




Figure 5: Low-rank tensor completion on synthetic data: LRTC-ENR ($p = 1/3$) v.s. baseline methods, $r = 50$; (a) recovery error with different missing rates; (b) average time cost in all cases of missing rates.

When the missing rate is 0.9 in Figure 4, we observed 12,500 entries, much more than the minimum number (1,500) of freedom parameters required to determine the tensor uniquely. It means the tensor completion problem is too easy. Hence we increase r to 50 and report the recovery error and computational time of KBR-TC, TenALS, BCPF, TMac, and LRTC-ENR in Figure 5. We see the recovery error of BCPF and LRTC-ENR are lower than those of other methods. When the missing rate is sufficiently high, LRTC-ENR has smaller recovery error than BCPF. Moreover, LRTC-ENR (200 iterations) is at least 15 times faster than BCPF and 5 times faster than KBR-TC.

5.1.2 Multi-spectral image inpainting

We use the Columbia multi-spectral image database (Yasuma et al., 2008) to show the performance of tensor completion methods in image inpainting problem. The dataset consists of the multi-spectral images of 32 real-world scenes of a variety of real-world materials and objects. The images have a spatial resolution of 512×512 and the number of spectral bands as 31. Thus the data of each image is a tensor of size $512 \times 512 \times 31$.

In our experiments, we pre-scale the image of every channel to $[0, 1]$. Since TRLRF (Yuan et al., 2019) and Rprecon (Kasai & Mishra, 2016) have extremely high computational costs on these tensors, we omit their implementations here. In (Xie et al., 2018), it has been shown that BCPF (Zhao et al., 2015) significantly outperformed HaLRTC (Liu et al., 2012) as well as a few other baselines on this dataset, so we will not compare HaLRTC. The initial rank of TMac (Xu et al., 2013) and LRTC-ENR ($\lambda = 0.1$) are set to 100. When the initial rank is 100, it takes TenALS (Jain & Oh, 2014) and BCPF (Zhao et al., 2015) more than 1500 seconds on each image, thus we set the initial rank to 50. Since the tensor rank is very small compared

to the size and the images are noiseless, we here consider highly incomplete images of missing rate no less than 0.95.

As an example, Figure 6 shows the recovery errors of LRTC-ENR (at different p 's) and competing methods on the first image of the dataset when the missing rate is 0.95, 0.97, or 0.99. We can see that when the missing rate is 0.95, KBR-TC outperformed LRTC-ENR. When the missing rate is 0.99, LRTC-ENR outperformed all other methods. In Figure 6 (a), for LRTC-ENR, $p = 1/6$ is the best choice, while in Figure 6 (b) and (c), $p = 1/3$ has least recovery error. These results are consistent with the theorems: a smaller p but not too small leads to a lower recovery error.




Figure 6: Recovery error on the first image of the MSI dataset.

Figure 7 and 8 visualize the recovery performance on the first image when the missing rates are 0.97 and 0.99 respectively. Visually, the recovery performance of LRTC-ENR ($p = 1/3$) is much better than those of other methods. The average results (PSNR) on the 32 images are reported in Table 2. KBR-TC and LRTC-ENR ($p = 1/3$) outperformed other methods significantly. When the missing rate is 0.97 or 0.99, LRTC-ENR outperformed KBR-TC, and the improvement is significant according to the paired t-test. In addition, the time cost of LRTC-ENR is only 15% of that of KBR-TC.

Table 2: Average recovery performance (PSNR) on the Columbia MSI dataset of 32 images (the p -value is from the paired t-test between KBR-TC and LRTC-ENR)

Missing rate	0.95	0.97	0.99	Time
TenALS	30.1±4.4	25.9±3.22	13.3±4.9	500s
BCPF	30.4±5.8	26.8±5.6	20.2±3.3	450s
TMac	33.3±5.4	30.6±4.9	17.9±3.1	190s
KBR-TC	37.1±4.9	31.4±4.9	20.5±2.8	730s
LRTC-ENR	35.4±4.4	32.4±5.1	26.1±5.4	110s
p -value (t-test)	5×10^{-11}	5×10^{-4}	2×10^{-8}	0

5.2 Experiments of Tensor Robust PCA

5.2.1 Synthetic data

We generate synthetic tensors by $\mathcal{D} = \mathcal{T} + \mathcal{N} + \mathcal{E}$. The low-rank tensor \mathcal{T} is given by $\mathcal{T} = \sum_{i=1}^r w_i \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \circ \mathbf{x}_i^{(3)}$, where $w_i = i/r$ and the entries of $\mathbf{x}_i^{(j)} \in \mathbb{R}^n$ ($i \in [r]$, $j \in [3]$) are drawn from $\mathcal{N}(0, 1)$. \mathcal{N} is a dense noise tensor drawn from $\mathcal{N}(0, (0.1\sigma_{\mathcal{T}})^2)$ and \mathcal{E} is a sparse noise tensor drawn from $\mathcal{N}(0, \sigma_{\mathcal{T}}^2)$. We set $n = 50$, $r = 25$ and evaluate the TRPCA performance in recovering \mathcal{T} in the case of different sparsity




Figure 7: Inpainting performance on the first image (band 16) of the MSI dataset when the missing rate is 0.97.




Figure 8: Inpainting performance on the first image (band 16) of the MSI dataset when the missing rate is 0.99.

(termed as noise density) of \mathcal{E} . The evaluation metric is the relative recovery error defined by

$$\text{Relative recovery error} = \|\mathcal{T} - \hat{\mathcal{T}}\|_F / \|\mathcal{T}\|_F.$$

We compare TRPCA-ENR with KBR-PCA (Xie et al., 2018). TRPCA (denoted by T-TRPCA) proposed by (Lu et al., 2019) and the robust tensor decomposition method OITNN-L proposed by (Wang et al., 2020) are under the assumption of t -product (not CP decomposition) and hence not compatible in our setting of synthetic data. We only compare with them on real data. The hyper-parameters of KBR-RPCA and TRPCA-ENR³ are sufficiently tuned to provide their best performance. Figure 9 (a) shows the recovery errors of KBR-RPCA and TRPCA-ENR (with $p = 2/3, 1/3$, and $1/6$) when the noise level increased from 0.1 to 0.7. We see that TRPCA-ENR consistently outperformed KBR-RPCA. In TRPCA-ENR, $p = 1/6$ is better than $p = 2/3$ and $1/3$ but the difference is not obvious. In Figure 9 (b), we show the performance of different p 's. It can be found that a smaller p (but not too small, e.g. larger than 10^{-2}) yields less recovery error.

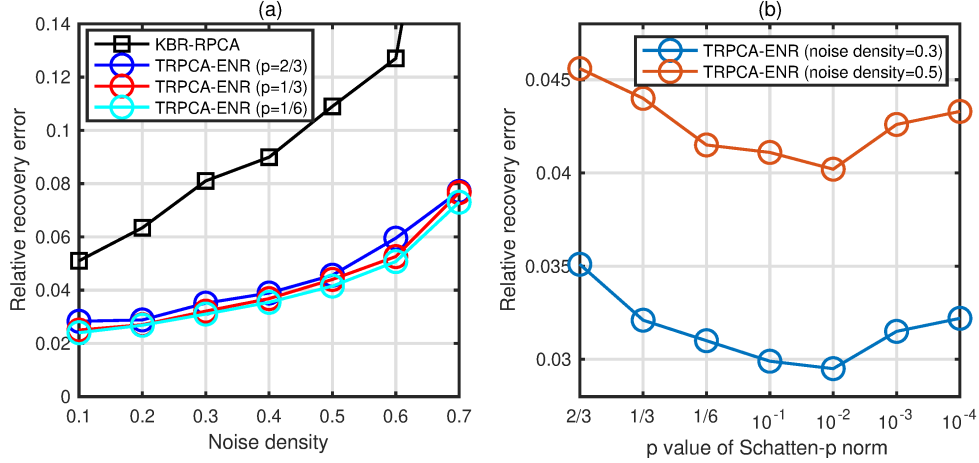


Figure 9: Performance of TRPCA on synthetic data (average of 20 repeated trials): (a) KBR-RPCA v.s. TRPCA-ENR (initial rank= $2r$); (b) relative recovery error of TRPCA-ENR with different p .

In Figure 10, we compare the performance of symmetric regularizers and asymmetric regularizers in TRPCA-ENR. We see that in all cases, the recovery error and times cost given by the asymmetric regularizers are lower than those given by the symmetric regularizers. The reason is that there are more subproblems in the optimization associated with the symmetric regularizers that have no closed-form solutions.

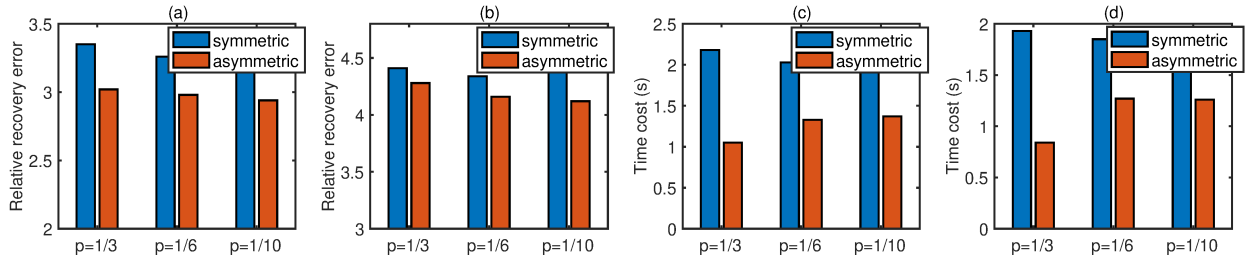


Figure 10: Comparison of symmetric and asymmetric regularizers in TRPCA-ENR on synthetic data (average of 20 repeated trials): (a) relative recovery error when the noise level is 0.3; (b) relative recovery error when the noise level is 0.5; (c) time cost when the noise density is 0.3; (d) time cost when the noise density is 0.5.

5.2.2 Color image denoising

We compare our TRPCA-ENR with KBR-PCA (Xie et al., 2018), T-TRPCA (Lu et al., 2019), and OITNN-L (Wang et al., 2020) in the task of color image denoising on nine color images (shown in (11)) of size

³Throughout this paper, we set $\mu = 10$ and $t_{\max} = 500$ for Algorithms 4 and 5.



Figure 11: Nine color images for denoising

$256 \times 256 \times 3$ used in (Wang et al., 2020). For each image, we randomly set 10% of the tensor elements to random values in $[0, 1]$. We tuned all hyper-parameters carefully within wide ranges to provided the best denoising performance of the methods. As a result, in KBR-RPCA, we set $\lambda = 2$ or 3. In OITNN-L, we set $\lambda_L = 4$ or 5 and $\lambda_S = 0.16$. In T-TRPCA, we set $\lambda = 1/\sqrt{768}$ or $2/\sqrt{768}$. In our TRPCA-ENR, we set the initial rank to a relative large value 200 because the approximate rank of the image tensors is not too small and our methods can adjust the rank adaptively; we set $\lambda_e = 0.02$ and consider the cases of $p = 2/3, 1/3, 1/6$, and $1/10$, for which $\lambda_x = 0.1, 0.1, 0.02$, and 0.015 .

The average \pm standard deviation PSNR across 20 repeated trials are reported in Tables 3. We see that the PSNRs of our TRPCA-ENR methods are much lower than those of other methods in all cases. In addition, TRPCA-ENR with smaller p has higher PSNR and TRPCA-ENR with $p = 1/10$ performs the best. Figure 12 visualizes the denoising performance on Image 1, showing that TRPCA-ENR methods are better than other methods. These results provided evidence that tensor Schatten- p quasi-norm minimization is able to provide higher recovery accuracy in TRPCA when a relatively smaller p is used.

Table 3: PSNR of color image denoising

	Image 1	Image 2	Image 3	Image 4	Image 5
KBR-RPCA (Xie et al., 2018)	30.25 \pm 0.06	33.36 \pm 0.22	31.93 \pm 0.12	31.87 \pm 0.08	35.08 \pm 0.12
T-TRPCA (Lu et al., 2019)	31.16 \pm 0.05	32.75 \pm 0.06	32.87 \pm 0.11	32.52 \pm 0.06	33.91 \pm 0.10
OITNN-O (Wang et al., 2020)	28.46 \pm 0.04	29.99 \pm 0.03	29.16 \pm 0.05	28.13 \pm 0.04	29.86 \pm 0.03
TRPCA-ENR ($p=\frac{2}{3}$, sym)	34.23 \pm 0.12	35.44 \pm 0.09	34.82 \pm 0.18	34.23 \pm 0.14	37.14 \pm 0.06
TRPCA-ENR ($p=\frac{1}{3}$, sym)	34.68 \pm 0.18	36.88 \pm 0.16	37.52 \pm 0.12	34.53 \pm 0.31	37.78 \pm 0.04
TRPCA-ENR ($p=\frac{1}{3}$, asym)	34.81 \pm 0.16	37.04 \pm 0.06	37.61 \pm 0.19	34.92 \pm 0.28	38.16 \pm 0.10
TRPCA-ENR ($p=\frac{1}{6}$, asym)	34.96 \pm 0.40	36.94 \pm 0.12	37.57 \pm 0.33	35.12 \pm 0.17	38.14 \pm 0.11
TRPCA-ENR ($p=\frac{1}{10}$, asym)	35.45\pm0.12	37.14\pm0.10	37.76\pm0.14	35.38\pm0.46	38.28\pm0.13
	Image 6	Image 7	Image 8	Image 9	
KBR-RPCA (Xie et al., 2018)	31.36 \pm 0.09	30.08 \pm 0.10	32.41 \pm 0.05	34.28 \pm 0.12	
T-TRPCA (Lu et al., 2019)	28.96 \pm 0.06	29.36 \pm 0.09	31.49 \pm 0.11	33.67 \pm 0.14	
OITNN-O (Wang et al., 2020)	27.39 \pm 0.05	27.46 \pm 0.05	28.61 \pm 0.07	30.06 \pm 0.06	
TRPCA-ENR ($p=\frac{2}{3}$, sym)	35.09 \pm 0.12	33.67 \pm 0.10	35.66 \pm 0.16	36.12 \pm 0.05	
TRPCA-ENR ($p=\frac{1}{3}$, sym)	33.21 \pm 0.54	34.10 \pm 0.23	36.09 \pm 0.14	37.95 \pm 0.16	
TRPCA-ENR ($p=\frac{1}{3}$, asym)	33.54 \pm 0.71	34.29 \pm 0.18	36.38 \pm 0.09	38.13\pm0.12	
TRPCA-ENR ($p=\frac{1}{6}$, asym)	33.62 \pm 0.76	34.44 \pm 0.12	36.49 \pm 0.08	37.96 \pm 0.10	
TRPCA-ENR ($p=\frac{1}{10}$, asym)	35.18\pm0.35	34.57\pm0.14	36.71\pm0.06	38.01 \pm 0.16	




Figure 12: Denoising performance on Image 1.

6 Conclusion

We propose a new class of nonconvex rank regularizers based on the Euclidean norms of component vectors in CP decomposition. These regularizers are monotonic transformations of tensor nuclear norm and Schatten- p quasi-norms. We prove that Schatten- p quasi-norm with a decently small p provides a tight error bound for LRTC. Numerical results on synthetic tensors and natural images demonstrate the advantages of our methods in LRTC and TRPCA against state-of-the-art methods. Nevertheless, there are many avenues for improvement. One is to provide exact recovery guarantees for LRTC and TRPCA, though it is quite difficult because of the nonconvex and nonsmooth optimizations.

References

- Evrin Acar, Daniel M Dunlavy, Tamara G Kolda, and Morten Mørup. Scalable tensor factorizations for incomplete data. *Chemometrics and Intelligent Laboratory Systems*, 106(1):41–56, 2011.
- Anima Anandkumar, Prateek Jain, Yang Shi, and Uma Naresh Niranjan. Tensor vs. matrix methods: Robust tensor decomposition under block sparse perturbations. In *Artificial Intelligence and Statistics*, pp. 268–276, 2016.
- Brett W. Bader, Tamara G. Kolda, et al. Matlab tensor toolbox version 3.1. Available online, June 2019. URL <https://www.tensortoolbox.org>.
- Boaz Barak and Ankur Moitra. Noisy tensor completion via the sum-of-squares hierarchy. In *Conference on Learning Theory*, pp. 417–445, 2016.
- Peter Bartlett, Dylan J Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1706.08498*, 2017.
- Juan Andrés Bazerque, Gonzalo Mateos, and Georgios B Giannakis. Rank regularization and bayesian inference for tensor completion and extrapolation. *IEEE transactions on signal processing*, 61(22):5689–5703, 2013.
- Rajendra Bhatia. *Matrix analysis*, volume 169. Springer Science & Business Media, 2013.
- Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009. ISSN 1615-3383. doi: 10.1007/s10208-009-9045-5.

- Emmanuel J. Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *J. ACM*, 58(3):1–37, 2011. ISSN 0004-5411. doi: 10.1145/1970392.1970395.
- Hao Cheng, Yaoliang Yu, Xinhua Zhang, Eric Xing, and Dale Schuurmans. Scalable and sound low-rank tensor learning. In Arthur Gretton and Christian C. Robert (eds.), *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pp. 1114–1123, Cadiz, Spain, 09–11 May 2016. PMLR.
- Fernando De la Torre and Michael J Black. Robust principal component analysis for computer vision. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 1, pp. 362–369. IEEE, 2001.
- Chris Ding, Ding Zhou, Xiaofeng He, and Hongyuan Zha. R 1-pca: rotational invariant l 1-norm principal component analysis for robust subspace factorization. In *Proceedings of the 23rd international conference on Machine learning*, pp. 281–288, 2006.
- Ran El-Yaniv and Dmitry Pechyony. Transductive rademacher complexity and its applications. *Journal of Artificial Intelligence Research*, 35:193–234, 2009.
- Jicong Fan. Multi-mode deep matrix and tensor factorization. In *International Conference on Learning Representations*, 2022.
- Jicong Fan, Lijun Ding, Yudong Chen, and Madeleine Udell. Factor group-sparse regularization for efficient low-rank matrix recovery. In *Advances in Neural Information Processing Systems 32*, pp. 5104–5114. Curran Associates, Inc., 2019.
- Dylan J. Foster and Andrej Risteski. Sum-of-squares meets square loss: Fast rates for agnostic tensor completion. In Alina Beygelzimer and Daniel Hsu (eds.), *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pp. 1280–1318. PMLR, 25–28 Jun 2019. URL <https://proceedings.mlr.press/v99/foster19a.html>.
- Shmuel Friedland and Lek-Heng Lim. Nuclear norm of higher-order tensors. *Mathematics of Computation*, 87(311):1255–1281, 2018.
- Silvia Gandy, Benjamin Recht, and Isao Yamada. Tensor completion and low-n-rank tensor recovery via convex optimization. *Inverse Problems*, 27(2):025010, 2011.
- Navid Ghadermarzy, Yaniv Plan, and Özgür Yilmaz. Near-optimal sample complexity for convex tensor completion. *Information and Inference: A Journal of the IMA*, 8(3):577–619, 2019.
- Paris Giampouras, René Vidal, Athanasios Rontogiannis, and Benjamin Haeffele. A novel variational form of the schatten- p quasi-norm. *arXiv preprint arXiv:2010.13927*, 2020.
- Donald Goldfarb and Zhiwei Qin. Robust low-rank tensor recovery: Models and algorithms. *SIAM Journal on Matrix Analysis and Applications*, 35(1):225–253, 2014.
- Benjamin D Haeffele and René Vidal. Structured low-rank matrix factorization: Global optimality, algorithms, and applications. *IEEE transactions on pattern analysis and machine intelligence*, 42(6):1468–1482, 2019.
- Moritz Hardt. Understanding alternating minimization for matrix completion. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 651–660. IEEE, 2014.
- Aicke Hinrichs, Joscha Prochno, and Jan Vybiral. Entropy numbers of embeddings of schatten classes. *Journal of Functional Analysis*, 273(10):3241–3261, 2017. ISSN 0022-1236. doi: <https://doi.org/10.1016/j.jfa.2017.08.008>. URL <https://www.sciencedirect.com/science/article/pii/S0022123617303221>.
- Bo Huang, Cun Mu, Donald Goldfarb, and John Wright. Provable models for robust low-rank tensor completion. *Pacific Journal of Optimization*, 11(2):339–364, 2015.

- Prateek Jain and Sewoong Oh. Provable tensor factorization with missing data. In *Advances in Neural Information Processing Systems*, pp. 1431–1439, 2014.
- Hiroyuki Kasai and Bamdev Mishra. Low-rank tensor completion: a riemannian manifold preconditioning approach. In *International Conference on Machine Learning*, pp. 1012–1021, 2016.
- Misha E Kilmer and Carla D Martin. Factorization strategies for third-order tensors. *Linear Algebra and its Applications*, 435(3):641–658, 2011.
- Hao Kong, Xingyu Xie, and Zhouchen Lin. t-schatten- p norm for low-rank tensor recovery. *IEEE Journal of Selected Topics in Signal Processing*, 12(6):1405–1419, 2018.
- Daniel Kressner, Michael Steinlechner, and Bart Vandereycken. Low-rank tensor completion by riemannian optimization. *BIT Numerical Mathematics*, 54(2):447–468, 2014.
- Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. Canonical tensor decomposition for knowledge base completion. In *International Conference on Machine Learning*, pp. 2863–2872. PMLR, 2018.
- Allen Liu and Ankur Moitra. Tensor completion made practical. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 18905–18916. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/dab1263d1e6a88c9ba5e7e294def5e8b-Paper.pdf>.
- Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- Ji Liu, Przemyslaw Musialski, Peter Wonka, and Jieping Ye. Tensor completion for estimating missing values in visual data. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):208–220, 2012.
- Canyi Lu, Jiashi Feng, Yudong Chen, Wei Liu, Zhouchen Lin, and Shuicheng Yan. Tensor robust principal component analysis: Exact recovery of corrupted low-rank tensors via convex optimization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5249–5257, 2016.
- Canyi Lu, Jiashi Feng, Yudong Chen, Wei Liu, Zhouchen Lin, and Shuicheng Yan. Tensor robust principal component analysis with a new tensor nuclear norm. *IEEE transactions on pattern analysis and machine intelligence*, 42(4):925–938, 2019.
- Zhaosong Lu. Iterative reweighted minimization methods for l_p regularized unconstrained nonlinear programming. *Mathematical Programming*, 147(1-2):277–307, 2014.
- Sebastian Mayer and Tino Ullrich. Entropy numbers of finite dimensional mixed-norm balls and function space embeddings with small mixed smoothness. *Constructive Approximation*, 53(2):249–279, 2021.
- Feiping Nie, Heng Huang, and Chris Ding. Low-rank matrix recovery via efficient Schatten p -norm minimization. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, AAAI’12, pp. 655–661. AAAI Press, 2012.
- Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3): 127–239, 2014.
- Aaron Potechin and David Steurer. Exact tensor completion with sum-of-squares. In Satyen Kale and Ohad Shamir (eds.), *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 1619–1673. PMLR, 07–10 Jul 2017.
- Bernardino Romera-Paredes and Massimiliano Pontil. A new convex relaxation for tensor completion. In *Proceedings of the 26th International Conference on Neural Information Processing Systems-Volume 2*, pp. 2967–2975, 2013.
- Nicolas Schreuder. Bounding the expectation of the supremum of empirical processes indexed by h^{α} older classes. *arXiv preprint arXiv:2003.13530*, 2020.

- Fanhua Shang, Yuanyuan Liu, and James Cheng. Tractable and scalable Schatten quasi-norm approximations for rank minimization. In *Artificial Intelligence and Statistics*, pp. 620–629, 2016.
- Fanhua Shang, James Cheng, Yuanyuan Liu, Zhi-Quan Luo, and Zhouchen Lin. Bilinear factor matrix norm minimization for robust pca: Algorithms and applications. *IEEE transactions on pattern analysis and machine intelligence*, 40(9):2066–2080, 2017.
- Qiquan Shi, Haiping Lu, and Yiu-ming Cheung. Tensor rank estimation and completion via cp-based nuclear norm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pp. 949–958, 2017.
- Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *International Conference on Computational Learning Theory*, pp. 545–560. Springer, 2005.
- Ryota Tomioka and Taiji Suzuki. Spectral norm of random tensors. *arXiv preprint arXiv:1407.1870*, 2014.
- Andong Wang, Chao Li, Zhong Jin, and Qibin Zhao. Robust tensor decomposition via orientation invariant tubal nuclear norms. In *AAAI*, pp. 6102–6109, 2020.
- Yu Wang, Wotao Yin, and Jinshan Zeng. Global convergence of admm in nonconvex nonsmooth optimization. *Journal of Scientific Computing*, pp. 1–35, 2015.
- Kishan Wimalawarne and Hiroshi Mamitsuka. Reshaped tensor nuclear norms for higher order tensor completion. *Machine Learning*, 110(3):507–531, 2021.
- Q. Xie, Q. Zhao, D. Meng, and Z. Xu. Kronecker-basis-representation based tensor sparsity and its applications to tensor recovery. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(8):1888–1902, 2018.
- Yangyang Xu and Wotao Yin. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences*, 6(3):1758–1789, 2013.
- Yangyang Xu, Ruru Hao, Wotao Yin, and Zhixun Su. Parallel matrix factorization for low-rank tensor completion. *arXiv preprint arXiv:1312.1254*, 2013.
- Bo Yang, Gang Wang, and Nicholas D Sidiropoulos. Tensor completion via group-sparse regularization. In *2016 50th Asilomar Conference on Signals, Systems and Computers*, pp. 1750–1754. IEEE, 2016.
- F. Yasuma, T. Mitsunaga, D. Iso, and S.K. Nayar. Generalized Assorted Pixel Camera: Post-Capture Control of Resolution, Dynamic Range and Spectrum. Technical report, Nov 2008.
- Longhao Yuan, Chao Li, Danilo Mandic, Jianting Cao, and Qibin Zhao. Tensor ring decomposition with rank minimization on latent space: An efficient approach for tensor completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 9151–9158, 2019.
- Ming Yuan and Cun-Hui Zhang. On tensor completion via nuclear norm minimization. *Foundations of Computational Mathematics*, 16(4):1031–1068, 2016.
- Zemin Zhang and Shuchin Aeron. Exact tensor completion using t-svd. *IEEE Transactions on Signal Processing*, 65(6):1511–1526, 2016.
- Qibin Zhao, Liqing Zhang, and Andrzej Cichocki. Bayesian cp factorization of incomplete tensors with automatic rank determination. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1751–1763, 2015.
- Yu-Bang Zheng, Ting-Zhu Huang, Xi-Le Zhao, Tai-Xiang Jiang, Tian-Hui Ma, and Teng-Yu Ji. Mixed noise removal in hyperspectral image via low-fibered-rank regularization. *IEEE Transactions on Geoscience and Remote Sensing*, 58(1):734–749, 2019.
- Pan Zhou, Canyi Lu, Zhouchen Lin, and Chao Zhang. Tensor factorization for low-rank tensor completion. *IEEE Transactions on Image Processing*, 27(3):1152–1163, 2017.

A Optimization for LRTC-ENR

A.1 Block Coordinate Descent with Extrapolation

There are d blocks of decision variables in problem (2), i.e. $\{\mathbf{X}^{(j)}\}_{j \in [d]}$, where $\mathbf{X}^{(j)} = [\mathbf{x}_1^{(j)}, \mathbf{x}_2^{(j)}, \dots, \mathbf{x}_k^{(j)}]$. We propose to find a critical point of (2) by Block Coordinate Descent (BCD) with Extrapolation (BCDE for short) (Xu & Yin, 2013), which is more efficient than BCD. For simplicity, here we only present the optimization of (2) with the regularizer shown in (3). The optimization can be easily extended to (2) with other regularizers we proposed in Section 2.

Let

$$\begin{aligned} \mathcal{L}(\{\mathbf{X}^{(j)}\}_{j \in [d]}) &:= \frac{1}{2} \left\| \mathcal{M} * \left(\mathcal{D} - \sum_{i=1}^k \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \dots \circ \mathbf{x}_i^{(d)} \right) \right\|_F^2 \\ &= \frac{1}{2} \left\| \mathbf{M}_{(j)} * \left(\mathbf{D}_{(j)} - \mathbf{X}^{(j)} [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}]^\top \right) \right\|_F^2 \end{aligned}$$

and

$$\mathcal{R}_{pd}(\mathbf{X}^{(j)}) := \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|^{pd},$$

where $[(\mathbf{X}^{(i)})^{\odot_{i \neq j}}] = \mathbf{X}^{(d)} \odot \dots \odot \mathbf{X}^{(j+1)} \odot \mathbf{X}^{(j-1)} \odot \dots \odot \mathbf{X}^{(1)}$. We initialize $\{\mathbf{X}^{(j)}\}_{j \in [d]}$ randomly and rescale \mathbf{x} to have unit Euclidean norm:

$$\mathbf{x}_i^{(j)} \sim \mathcal{N}(0, 1), \quad \mathbf{x}_i^{(j)} \leftarrow \mathbf{x}_i^{(j)} / \|\mathbf{x}_i^{(j)}\|, \quad i \in [k], j \in [d]. \quad (11)$$

Then at iteration t , for $j = 1, 2, \dots, d$, we first perform an extrapolation

$$\hat{\mathbf{X}}_{t-1}^{(j)} = \mathbf{X}_{t-1}^{(j)} + \omega_{j,t-1}(\mathbf{X}_{t-1}^{(j)} - \mathbf{X}_{t-2}^{(j)}), \quad (12)$$

where $\omega_{j,t-1} \geq 0$ controls the size of the extrapolation at iteration t . Then we perform a proximal step as

$$\begin{aligned} \mathbf{X}_t^{(j)} &= \underset{\mathbf{X}^{(j)}}{\operatorname{argmin}} \left\langle \nabla_{\hat{\mathbf{X}}_{t-1}^{(j)}} \mathcal{L}, \mathbf{X}^{(j)} - \hat{\mathbf{X}}_{t-1}^{(j)} \right\rangle + \frac{\hat{L}_{j,t}}{2} \|\mathbf{X}^{(j)} - \hat{\mathbf{X}}_{t-1}^{(j)}\|_F^2 + \lambda \mathcal{R}_{pd}(\mathbf{X}^{(j)}) \\ &\triangleq \operatorname{prox}_{pd}^\lambda \left(\hat{\mathbf{X}}_{t-1}^{(j)} - \hat{L}_{j,t}^{-1} \nabla_{\hat{\mathbf{X}}_{t-1}^{(j)}} \mathcal{L} \right). \end{aligned} \quad (13)$$

In (13), $\nabla_{\hat{\mathbf{X}}_{t-1}^{(j)}} \mathcal{L} = \left(\mathbf{M}_{(j)} * (\mathbf{D}_{(j)} - \hat{\mathbf{X}}_{t-1}^{(j)} [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}]^\top) \right) (-[(\mathbf{X}^{(i)})^{\odot_{i \neq j}}])$, $[(\mathbf{X}^{(i)})^{\odot_{i \neq j}}] = \mathbf{X}_{t-1}^{(d)} \odot \dots \odot \mathbf{X}_{t-1}^{(j+1)} \odot \mathbf{X}_{t-1}^{(j-1)} \odot \dots \odot \mathbf{X}_{t-1}^{(1)}$, and $\hat{L}_{j,t}$ is larger than the Lipschitz constant of $\nabla_{\hat{\mathbf{X}}_{t-1}^{(j)}} \mathcal{L}$. We approximate it by

$$\hat{L}_{j,t} = \varrho \sqrt{\frac{|\Omega|}{\prod_{i=1}^d n_i}} \left\| [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}] \right\|_2^2, \quad (14)$$

where $\varrho = 0.5, 1$, or 2 . The derivation of (14) is as follows. For convenience, we denote a perturbed copy of $\hat{\mathbf{X}}_{t-1}^{(j)}$ by $\hat{\mathbf{Z}}_{t-1}^{(j)}$. We have

$$\begin{aligned}
& \|\nabla_{\hat{\mathbf{X}}_{t-1}^{(j)}} \mathcal{L} - \nabla_{\hat{\mathbf{Z}}_{t-1}^{(j)}} \mathcal{L}\|_F \\
&= \left\| \left(\mathbf{M}_{(j)} * (\mathbf{D}_{(j)} - \hat{\mathbf{X}}_{t-1}^{(j)} [(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}]^\top) \right) \left(- [(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}] \right) \right. \\
&\quad \left. - \left(\mathbf{M}_{(j)} * (\mathbf{D}_{(j)} - \hat{\mathbf{Z}}_{t-1}^{(j)} [(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}]^\top) \right) \left(- [(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}] \right) \right\|_F \\
&\leq \|[(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}]\|_2 \|\mathbf{M}_{(j)} * (\hat{\mathbf{X}}_{t-1}^{(j)} [(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}]^\top - \hat{\mathbf{Z}}_{t-1}^{(j)} [(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}]^\top)\|_F \\
&\leq \|[(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}]\|_2 \varrho \sqrt{\frac{|\Omega|}{\prod_{i=1}^d n_i}} \|(\hat{\mathbf{X}}_{t-1}^{(j)} - \hat{\mathbf{Z}}_{t-1}^{(j)}) [(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}]^\top\|_F \\
&\leq \varrho \sqrt{\frac{|\Omega|}{\prod_{i=1}^d n_i}} \|[(\mathbf{X}^{(i)})_t^{\odot_{i \neq j}}]\|_2^2 \|\hat{\mathbf{X}}_{t-1}^{(j)} - \hat{\mathbf{Z}}_{t-1}^{(j)}\|_F.
\end{aligned}$$

Here $0 < \varrho \leq \sqrt{\frac{\prod_{i=1}^d n_i}{|\Omega|}}$ is some suitable constant.

In (13), when $p = 1/d$, $\text{prox}_1^\lambda(\mathbf{Y}) = \Phi_\lambda(\mathbf{Y})$, where Φ is the column-wise soft-thresholding operator (Parikh et al., 2014) defined by

$$\Phi_\lambda(\mathbf{y}) = \begin{cases} \frac{(\|\mathbf{y}\| - \lambda)\mathbf{y}}{\|\mathbf{y}\|}, & \text{if } \|\mathbf{y}\| > \lambda; \\ \mathbf{0}, & \text{otherwise.} \end{cases}$$

When $p = 2/d$, $\text{prox}_2^\lambda(\mathbf{Y}) = \frac{\hat{L}_{j,t}}{\hat{L}_{j,t} + 2\lambda} \mathbf{Y}$. When $p > 2/d$, we can estimate $\mathbf{X}_t^{(j)}$ by gradient descent. When $p < 1/d$, (13) is nonconvex and nonsmooth. Then we update $\mathbf{X}^{(j)}$ with iteratively reweighted method (Lu, 2014), which is given by Algorithm 1.

Algorithm 1 $\min_{\mathbf{Y}} \frac{1}{2} \|\mathbf{Y} - \mathbf{G}\|_F^2 + \tilde{\lambda} \sum_{i=1}^k \|\mathbf{y}_i\|^q$

Require: \mathbf{G} , q , $\tilde{\lambda}$, t_q , ϵ .

1: $\mathbf{Y} \leftarrow \mathbf{G}$.

2: **for** $t = 1, 2, \dots, t_q$ **do**

3: $\mathbf{W} = \text{diag}((\|\mathbf{y}_1\| + \epsilon)^{\frac{q-2}{2}}, \dots, (\|\mathbf{y}_k\| + \epsilon)^{\frac{q-2}{2}})$.

4: $\mathbf{Y} = \mathbf{G}(\mathbf{I} + 2\tilde{\lambda}\mathbf{W}\mathbf{W}^T)^{-1}$.

5: **end for**

Ensure: \mathbf{Y} .

In (12), the parameter ω_{jt} is determined as

$$\omega_{j,t-1} = \delta \sqrt{\hat{L}_{j,t-2} / \hat{L}_{j,t-1}}, \quad (15)$$

where $\delta < 1$. We set $\delta = 0.95$ for simplicity. The whole procedure is summarized in Algorithm 2. The convergence analysis when $p \geq 1/d$ can be found in (Xu & Yin, 2013). When $p < 1/d$, suppose Algorithm 1 returns the optimal solution, we can get similar convergence result as the case $p = 1/d$. Empirically, we found that there is no need to find the exact solution for the subproblem and instead we just perform Algorithm 1 for a few iterations, which can still provide satisfactory result. Currently, proving the convergence is out of the scope of our paper.

A.2 L-BFGS

Though faster than BCD, the computational cost per iteration of BCDE is high when d is not small because we need to compute $\mathbf{X}^{(j)} [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}]^\top$ for d times in every iteration. One may consider the Jacobi-type

Algorithm 2 solve LRTC-ENR by BCDE**Input:** $\mathcal{D}, \mathcal{M}, k, \lambda, t_{\max}$.

```

1: Initialize  $\{\mathbf{x}_i^{(j)}\}_{i \in [k]}^{j \in [d]}$  with (11), let  $t = 0$ .
2: repeat
3:    $t \leftarrow t + 1$ .
4:   for  $j = 1, 2, \dots, d$  do
5:     if  $t \leq 2$  then
6:        $\omega_{j,t-1} = 0$ .
7:     else
8:       Compute  $\omega_{j,t-1}$  by (15).
9:     end if
10:    Compute  $\hat{\mathbf{X}}_{t-1}^{(j)}$  by (12).
11:    Compute  $L_{j,t-1}$  by (14).
12:    Compute  $\mathbf{X}_t^{(j)}$  by (13) or Algorithm 1.
13:  end for
14:  Remove the zero columns of  $\mathbf{X}_t^{(j)}$ ,  $j \in [d]$ .
15: until converged or  $t = t_{\max}$ 
Output:  $\hat{\mathcal{T}} = \mathcal{I} \times_1 \mathbf{X}_t^{(1)} \times_2 \mathbf{X}_t^{(2)} \dots \times_d \mathbf{X}_t^{(d)}$ .
```

iteration (cheaper computation but slower convergence) rather than the Gauss-Seidel iteration in BCD and BCDE. In practice, we can use quasi-Newton methods such as L-BFGS (Liu & Nocedal, 1989) to solve problem (2) even though the objective function is nonsmooth. Particularly, we can drop the columns of $\mathbf{X}_t^{(j)}$ ($j \in [d]$) with nearly-zero Euclidean norms for acceleration. The corresponding algorithm⁴ is shown in Algorithm 3.

Algorithm 3 solve LRTC-ENR by LBFGS**Input:** $\mathcal{D}, \mathcal{M}, k, \lambda, t_{\max}$.

```

1: Initialize  $\{\mathbf{x}_i^{(j)}\}_{i \in [k]}^{j \in [d]}$  with (11), let  $t = 0$ .
2: repeat
3:    $t \leftarrow t + 1$ .
4:   Compute the search directions by LBFGS.
5:   Use line search to determine the step size.
6:   Update  $\mathbf{X}_t^{(j)}$ ,  $j \in [d]$ .
7:   For  $j \in [d]$ , remove the columns of  $\mathbf{X}_t^{(j)}$  with Euclidean norms less than a small threshold, e.g.  $10^{-5}$ .
8: until converged or  $t = t_{\max}$ 
Output:  $\hat{\mathcal{T}} = \mathcal{I} \times_1 \mathbf{X}_t^{(1)} \times_2 \mathbf{X}_t^{(2)} \dots \times_d \mathbf{X}_t^{(d)}$ .
```

B Optimization for TRPCA-ENR

When $p = 2/d$, we use alternating minimization to solve the optimization of TRPCA-ENR because every subproblem has a closed-form solution. When $p = 1/d$, we may, like Section A.1, use block coordinate descent with extrapolation (Xu & Yin, 2013). However, the additional variable \mathcal{E} further slows down the convergence. We hence propose to solve (9) by the (nonconvex) alternating direction method of multipliers (ADMM) (Wang et al., 2015). Specifically, by adding auxiliary variables $\{\mathbf{Y}^{(j)}\}_{j=1}^d$, we reformulate (9) as

$$\begin{aligned}
& \underset{\{\mathbf{X}^{(j)}, \mathbf{Y}^{(j)}\}_{j=1}^d, \mathcal{E}}{\text{minimize}} && \frac{1}{2} \left\| \mathbf{D}_{(j)} - \mathbf{X}^{(j)} [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}]^\top - \mathbf{E}_{(j)} \right\|_F^2 + \lambda_x \mathcal{R} \left(\{\mathbf{y}_i^{(j)}\}_{i \in [k]}^{j \in [d]} \right) + \lambda_e \|\mathcal{E}\|_1 \\
& \text{subject to} && \mathbf{Y}^{(j)} = \mathbf{X}^{(j)}, j = 1, 2, \dots, d.
\end{aligned}$$

⁴The implementation in this paper is based on the *minFunc* MATLAB toolbox of M. Schmidt: <http://www.cs.ubc.ca/~schmidtm/Software/minFunc.html>.

Let $\{\mathbf{Z}^{(j)}\}_{j=1}^d$ be Lagrange multipliers and solve

$$\begin{aligned} & \underset{\{\mathbf{X}^{(j)}, \mathbf{Y}^{(j)}\}_{j=1}^d, \mathcal{E}}{\text{minimize}} \quad \frac{1}{2} \left\| \mathbf{D}_{(j)} - \mathbf{X}^{(j)} [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}]^\top - \mathbf{E}_{(j)} \right\|_F^2 \\ & + \lambda_x \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{y}_i^{(j)}\|^{pd} + \lambda_e \|\mathcal{E}\|_1 + \sum_{j=1}^d \left\langle \mathbf{Y}^{(j)} - \mathbf{X}^{(j)}, \mathbf{Z}^{(j)} \right\rangle + \frac{\mu}{2} \|\mathbf{Y}^{(j)} - \mathbf{X}^{(j)}\|_F^2, \end{aligned} \quad (16)$$

where μ is the augmented Lagrange penalty parameter. Then update $\{\mathbf{X}^{(j)}, \mathbf{Y}^{(j)}\}_{j=1}^d$ and \mathcal{E} sequentially to minimize (16) and update $\{\mathbf{Z}^{(j)}\}_{j=1}^d$ lastly. The procedures are summarized into Algorithm 4, where

$$\mathbf{X}_t^{(j)} = \left((\mathbf{D}_{(j)} - \mathbf{E}_{(j)}) [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}] + \mu \mathbf{Y}_{t-1}^{(j)} - \mathbf{Z}^{(j)} \right) \left([(\mathbf{X}^{(i)})^{\odot_{i \neq j}}]^T [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}] + \mu \mathbf{I} \right)^{-1}, \quad (17)$$

$$\mathbf{Y}_t^{(j)} = \Phi_{\lambda_x/\mu} \left(\mathbf{X}_t^{(j)} - \mathbf{Z}^{(j)} / \mu \right), \quad (18)$$

$$\mathbf{Z}^{(j)} \leftarrow \mathbf{Z}^{(j)} + \mu (\mathbf{Y}_t^{(j)} - \mathbf{X}_t^{(j)}), \quad (19)$$

$$\mathcal{E}_t = \Psi_{\lambda_e} \left(\mathcal{D} - \mathcal{I} \times_1 \mathbf{X}_t^{(1)} \times_2 \mathbf{X}_t^{(2)} \dots \times_d \mathbf{X}_t^{(d)} \right). \quad (20)$$

Ψ is the element-wise soft-thresholding operator (Parikh et al., 2014) defined by

$$\Psi_{\lambda_e}(v) = \text{sign}(v) \max(0, |v| - \lambda_e).$$

Algorithm 4 TRPCA-ENR ($p = 1/d$) solved by ADMM

Input: \mathcal{D} , k , λ_x , λ_e , μ , t_{\max} .

- 1: Initialize $\{\mathbf{x}_i^{(j)}\}_{i \in [k]}^{j \in [d]}$ with (11); for $j = 1, \dots, d$, let $\mathbf{Y}_0^{(j)} = \mathbf{X}_0^{(j)}$ and $\mathbf{Z}^{(j)} = \mathbf{0}$; let $t = 0$ and $\mathcal{E} = \mathbf{0}$.
 - 2: **repeat**
 - 3: $t \leftarrow t + 1$.
 - 4: **for** $j = 1, 2, \dots, d$ **do**
 - 5: Compute $\mathbf{X}_t^{(j)}$ by (17).
 - 6: Compute $\mathbf{Y}_t^{(j)}$ by (18).
 - 7: Update $\mathbf{Z}^{(j)}$ by (19).
 - 8: **end for**
 - 9: Compute \mathcal{E}_t by (20).
 - 10: **until** converged or $t = t_{\max}$
- Output:** $\hat{\mathcal{T}} = \mathcal{I} \times_1 \mathbf{X}_t^{(1)} \times_2 \mathbf{X}_t^{(2)} \dots \times_d \mathbf{X}_t^{(d)}$.
-

In (9), when $p \notin \{1/d, 2/d\}$, the optimization becomes more difficult because all d groups of the regularizers are nonconvex and nonsmooth. Thanks to Theorem 2, especially its (b), we can obtain arbitrarily sharp Schatten- p quasi-norm regularization by using only one group of nonconvex and nonsmooth regularizers on the component vectors. The corresponding problem is

$$\underset{\{\mathbf{x}_i^{(j)}\}_{i \in [k]}^{j \in [d]}, \mathcal{E}}{\text{minimize}} \quad \frac{1}{2} \left\| \mathcal{D} - \sum_{i=1}^k \mathbf{x}_i^{(1)} \circ \mathbf{x}_i^{(2)} \dots \circ \mathbf{x}_i^{(d)} - \mathcal{E} \right\|_F^2 + \lambda_x \sum_{i=1}^k \left(\frac{1}{q} \|\mathbf{x}_i^{(1)}\|^q + \frac{1}{2} \sum_{j=2}^d \|\mathbf{x}_i^{(j)}\|^2 \right) + \lambda_e \|\mathcal{E}\|_1. \quad (21)$$

We propose to solve (21) by ADMM with iteratively reweighted update (Algorithm 1) embedded, which is shown in Algorithm 5. In the algorithm, for $j = 2, \dots, d$, the subproblem of $\mathbf{X}_t^{(j)}$ has a closed-form solution, which makes the algorithm more efficient than Algorithm 4, especially when d is large.

Algorithm 5 TRPCA-ENR when $p \notin \{1/d, 2/d\}$ **Input:** \mathcal{D} , k , q , λ_x , λ_e , μ , t_{\max} .

- 1: Initialize $\{\mathbf{x}_i^{(j)}\}_{i \in [k]}^{j \in [d]}$ with (11); let $t = 0$ and $\mathcal{E} = \mathbf{0}$; let $\mathbf{Y}_0^{(1)} = \mathbf{X}_0^{(1)}$ and $\mathbf{Z}^{(1)} = \mathbf{0}$.
 - 2: **repeat**
 - 3: $t \leftarrow t + 1$.
 - 4: Compute $\mathbf{X}_t^{(1)}$ by (17).
 - 5: Compute $\mathbf{Y}_t^{(1)}$ by Algorithm 1.
 - 6: Update $\mathbf{Z}^{(1)}$ by (19).
 - 7: **for** $j = 2, 3, \dots, d$ **do**
 - 8: Compute $\mathbf{X}_t^{(j)}$ by $\mathbf{X}_t^{(j)} = (\mathbf{D}_{(j)} - \mathbf{E}_{(j)})[(\mathbf{X}^{(i)})^{\odot_{i \neq j}}] \left([(\mathbf{X}^{(i)})^{\odot_{i \neq j}}]^T [(\mathbf{X}^{(i)})^{\odot_{i \neq j}}] + \lambda_x \mathbf{I} \right)^{-1}$.
 - 9: **end for**
 - 10: Compute \mathcal{E}_t by (20).
 - 11: **until** converged or $t = t_{\max}$
- Output:** $\hat{\mathcal{T}} = \mathcal{I} \times_1 \mathbf{X}_t^{(1)} \times_2 \mathbf{X}_t^{(2)} \dots \times_d \mathbf{X}_t^{(d)}$.

C More experimental results**C.1 Theoretical bound v.s. empirical error**

We then apply our method to the synthetic tensors used in Figure 1 and Figure 2, where we let the noise-signal ratio be 0.2. We let $\tilde{\Delta} = \left(\frac{1}{|\Omega|} \|\mathcal{P}_{\Omega}(\mathcal{D} - \mathcal{X})\|_F^2 - \frac{1}{|\Omega|} \|\mathcal{P}_{\Omega}(\mathcal{D} - \mathcal{X})\|_F^2 \right) / \varepsilon^2$, where $\varepsilon = \max\{\|\mathcal{D}\|_{\infty}, \|\mathcal{X}\|_{\infty}\}$. Now in Figure 13, we compare $\tilde{\Delta}$ with the theoretical error upper bound Δ defined by (6). $\tilde{\Delta}$ is in log scale for better visualization. We see the empirical reconstruction error is much less than the theoretical error (upper) bound. One reason is that the bound involves ε^2 , which is much larger (often 100 times) than the average of square entries of \mathcal{D} and \mathcal{X} .

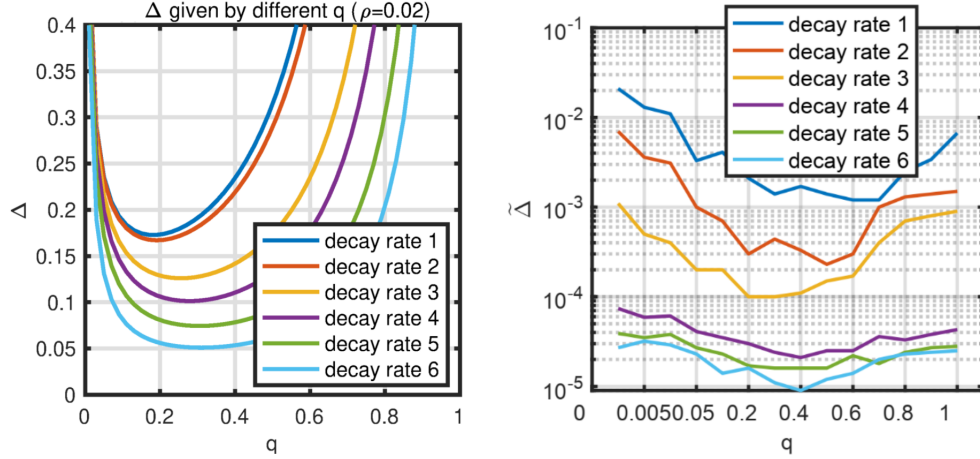


Figure 13: Left: theoretical upper bound Δ (Theorem 6). Right: empirical reconstruction error $\tilde{\Delta}$. The sampling rate ρ is 0.02.

C.2 Iterative performance of LRTC-ENR

Figure 14 shows the value of the objective function and relative recovery error of LRTC-ENR in each iteration on the synthetic data used in Section 5.1.1. We see that the optimization converged quickly and the relative recovery error decreased when the iteration number increased.




Figure 14: Iterative performance of LRTC-ENR

Relative recovery error
on 4-order tensor

	missing rate	
	0.7	0.9
HaLRTC	0.98	0.99
KBR-TC	0.34	0.67
BCPF	0.11	0.21
LRTC-ENR	0.03	0.05

Figure 15: Tensor completion on 4-order tensor of size $30 \times 30 \times 30 \times 30$. The data generating model is similar to the one used in Section 5.1.1.

C.3 LRTC-ENR on higher-order tensors

We compare our methods with HaLRTC, KBR-TC, and BCPF on a synthetic 4-order tensor of size $30 \times 30 \times 30 \times 30$ and rank 50. The relative recovery errors (average of ten trials) are reported in Figure 15. Our method outperformed the baselines.

D Proof of the theorems

D.1 Proof of Theorem 1

Proof. Let $\mathbf{x}_i^{(j)} = \alpha_{ij} \bar{\mathbf{x}}_i^{(j)}$, where $\|\bar{\mathbf{x}}_i^{(j)}\| = 1$. Then $\prod_{j=1}^d \|\mathbf{x}_i^{(j)}\| = \prod_{j=1}^d \alpha_{ij} = |\lambda_i|$. We have

$$\begin{aligned} \frac{1}{d} \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|^q &\geq \sum_{i=1}^k \left(\prod_{j=1}^d \|\mathbf{x}_i^{(j)}\|^q \right)^{1/d} \\ &= \sum_{i=1}^k \prod_{j=1}^d \|\mathbf{x}_i^{(j)}\|^{q/d} = \sum_{i=1}^k |\lambda_i|^{q/d}, \end{aligned}$$

in which the inequality holds according to the AM-GM inequality. When $\alpha_{i1} = \alpha_{i2} = \dots \alpha_{id} = |\lambda_i|^{1/d}$, the equality of AM-GM inequality is true. Replacing q with pd , we have

$$\frac{1}{d} \sum_{i=1}^k \sum_{j=1}^d \|\mathbf{x}_i^{(j)}\|^{pd} \geq \sum_{i=1}^k |\lambda_i|^p \geq \|\mathbf{x}\|_{S_p}^p.$$

□

D.2 Proof of Theorem 2

Proof. (1) We have

$$\begin{aligned}
& \sum_{i=1}^k \left(\frac{1}{q} \|\mathbf{x}_i^{(1)}\|^q + \sum_{j=2}^d \|\mathbf{x}_i^{(j)}\| \right) = \sum_{i=1}^k \left(\sum_{l=1}^{1/q} \|\mathbf{x}_i^{(1)}\|^q + \sum_{j=2}^d \|\mathbf{x}_i^{(j)}\| \right) \\
& \geq \sum_{i=1}^k \left(\frac{1}{q} + d - 1 \right) \left(\prod_{l=1}^{1/q} \|\mathbf{x}_i^{(1)}\|^q \prod_{j=2}^d \|\mathbf{x}_i^{(j)}\| \right)^{\frac{q}{1+qd-q}} \\
& = \sum_{i=1}^k \left(\frac{1}{q} + d - 1 \right) \left(\prod_{j=1}^d \|\mathbf{x}_i^{(j)}\| \right)^{\frac{q}{1+qd-q}} \\
& = \frac{1+qd-q}{q} \sum_{i=1}^k |\lambda_i|^{\frac{q}{1+qd-q}} \geq \frac{1+qd-q}{q} \|\mathbf{x}\|_{S_{q/(1+qd-q)}}^{q/(1+qd-q)}.
\end{aligned}$$

(2) We have

$$\begin{aligned}
& \sum_{i=1}^k \left(\frac{2}{q} \|\mathbf{x}_i^{(1)}\|^q + \sum_{j=2}^d \|\mathbf{x}_i^{(j)}\|^2 \right) = \sum_{i=1}^k \left(\sum_{l=1}^{2/q} \|\mathbf{x}_i^{(1)}\|^q + \sum_{j=2}^d \|\mathbf{x}_i^{(j)}\|^2 \right) \\
& \geq \sum_{i=1}^k \left(\frac{2}{q} + d - 1 \right) \left(\prod_{l=1}^{2/q} \|\mathbf{x}_i^{(1)}\|^q \prod_{j=2}^d \|\mathbf{x}_i^{(j)}\|^2 \right)^{\frac{q}{2+qd-q}} \\
& = \sum_{i=1}^k \left(\frac{2}{q} + d - 1 \right) \left(\prod_{j=1}^d \|\mathbf{x}_i^{(j)}\|^2 \right)^{\frac{q}{2+qd-q}} \\
& = \frac{2+qd-q}{q} \sum_{i=1}^k |\lambda_i|^{\frac{2q}{2+qd-q}} \geq \frac{2+qd-q}{q} \|\mathbf{x}\|_{S_{2q/(2+qd-q)}}^{2q/(2+qd-q)}.
\end{aligned}$$

The equality in the second row of formula (1) (also (2)) holds when all terms in the parentheses are equal, for each i . \square

D.3 Proof of the last two in Table 1

Proof. Recall the definition of $\lambda_1, \dots, \lambda_k$. We have

$$\begin{aligned}
& \frac{16^{1/5}}{5} \sum_{i=1}^k \left(\|\mathbf{x}_i^{(1)}\|^2 + \|\mathbf{x}_i^{(2)}\| + \|\mathbf{x}_i^{(3)}\| \right) \\
&= \frac{16^{1/5}}{5} \sum_{i=1}^k \left(\|\mathbf{x}_i^{(1)}\|^2 + \frac{1}{2} \|\mathbf{x}_i^{(2)}\| + \frac{1}{2} \|\mathbf{x}_i^{(2)}\| \right. \\
&\quad \left. + \frac{1}{2} \|\mathbf{x}_i^{(3)}\| + \frac{1}{2} \|\mathbf{x}_i^{(3)}\| \right) \\
&\geq \frac{16^{1/5}}{5} \sum_{i=1}^k 5 \left(\frac{1}{16} \|\mathbf{x}_i^{(1)}\|^2 \|\mathbf{x}_i^{(2)}\| \|\mathbf{x}_i^{(2)}\| \|\mathbf{x}_i^{(3)}\| \|\mathbf{x}_i^{(3)}\| \right)^{1/5} \\
&= \sum_{i=1}^k \left(\|\mathbf{x}_i^{(1)}\|^2 \|\mathbf{x}_i^{(2)}\|^2 \|\mathbf{x}_i^{(3)}\|^2 \right)^{1/5} \\
&= \sum_{i=1}^k |\lambda_i|^{2/5} \geq \|\mathbf{x}\|_{S_{2/5}}^{2/5}.
\end{aligned}$$

Similarly, we have

$$\begin{aligned}
& \frac{81^{1/7}}{7} \sum_{i=1}^k \left(\|\mathbf{x}_i^{(1)}\|^3 + \|\mathbf{x}_i^{(2)}\| + \|\mathbf{x}_i^{(3)}\| \right) \\
&= \frac{81^{1/7}}{7} \sum_{i=1}^k \left(\|\mathbf{x}_i^{(1)}\|^3 + \sum_{j=1}^3 \left(\frac{1}{3} \|\mathbf{x}_i^{(j)}\| + \frac{1}{3} \|\mathbf{x}_i^{(j)}\| \right) \right) \\
&\geq \frac{81^{1/7}}{7} \sum_{i=1}^k 7 \left(\frac{1}{3^6} \|\mathbf{x}_i^{(1)}\|^3 \|\mathbf{x}_i^{(2)}\|^3 \|\mathbf{x}_i^{(3)}\|^3 \right)^{1/7} \\
&= \sum_{i=1}^k \left(\|\mathbf{x}_i^{(1)}\| \|\mathbf{x}_i^{(2)}\| \|\mathbf{x}_i^{(3)}\| \right)^{2/7} \\
&= \sum_{i=1}^k |\lambda_i|^{3/7} \geq \|\mathbf{x}\|_{S_{3/7}}^{3/7}.
\end{aligned}$$

□

D.4 Proof of Theorem 3

First, we give the following lemma, which is proved in Appendix E.1.

Lemma 1. *Let $\mathcal{S}_{d,n,2}^{k,\perp} = \{\mathbf{x} \in \mathcal{S}_{d,n}^\perp : \text{rank}(\mathbf{x}) \leq k, \|\mathbf{x}\|_{S_2^\perp} \leq 1\}$. Then the covering numbers of $\mathcal{S}_{d,n,2}^{k,\perp}$ with respect to $\|\cdot\|_{S_2^\perp}$ satisfy*

$$\log \mathcal{N}(\mathcal{S}_{d,n,2}^{k,\perp}, \|\cdot\|_{S_2^\perp}, \epsilon) \leq \left(\frac{c(d+1)}{\epsilon} \right)^{dk(n-(k+1)/2)+k},$$

where $c > 0$ is a universal constant.

Proof. We follow the idea of Theorem 4.3 of (Hinrichs et al., 2017) and analyze the entropy number first. According to the monotonicity of dyadic entropy numbers, it is enough to let $n = 2^\alpha$ and $\eta = 2^\alpha \cdot 2^\beta$, where

$1 \leq \alpha < \beta$ are natural numbers. We have $n/\eta = 2^{-\beta}$. Let $\mathcal{X} \in \mathcal{S}_{d,n,p}^\perp$ with $\psi = 1$. The orthogonal CP decomposition of \mathcal{X} is denoted by $\mathcal{X} = \sum_{i=1}^n s_i \mathbf{u}_i^{(1)} \circ \mathbf{u}_i^{(2)} \cdots \circ \mathbf{u}_i^{(d)}$ and $s_1 \geq s_2 \geq \cdots \geq s_n$. We define

$$\mathcal{X}_1 := \sum_{i=1}^n \bar{s}_{1i} \mathbf{u}_i^{(1)} \circ \mathbf{u}_i^{(2)} \cdots \circ \mathbf{u}_i^{(d)}, \quad \bar{s}_1 = (s_1, 0, \dots, 0)^\top, \quad (22)$$

and, for $j = 2, 3, \dots, \beta$,

$$\mathcal{X}_j = \sum_{i=1}^n \bar{s}_{ji} \mathbf{u}_i^{(1)} \circ \mathbf{u}_i^{(2)} \cdots \circ \mathbf{u}_i^{(d)}, \quad \bar{s}_j = (0, \dots, 0, \overbrace{s_{2^{j-1}}, \dots, s_{2^j-1}}^{2^{j-1}}, 0, \dots, 0)^\top. \quad (23)$$

It follows that $\text{rank}(\mathcal{X}_j) \leq 2^{j-1}$, $j \in [\beta]$. Then we have $\|\mathcal{X}_1\|_{S_p^\perp} \leq 1$, and $j = 2, 3, \dots, \beta$,

$$\begin{aligned} \|\mathcal{X}_j\|_{S_q^\perp} &= \left(\sum_{t=2^{j-1}}^{2^j-1} s_t^q \right)^{1/q} \leq (2^{j-1} s_{2^{j-1}}^q)^{1/q} = 2^{(j-1)/q} (s_{2^{j-1}}^p)^{1/p} \\ &\leq 2^{(j-1)/q} \left(\frac{1}{2^{j-1}} \sum_{t=1}^{2^{j-1}} s_t^p \right)^{1/p} \leq 2^{(j-1)(1/q-1/p)}, \end{aligned} \quad (24)$$

where the last inequality holds because of $\sum_{t=1}^{2^{j-1}} s_t^p \leq 1$. Now we can decompose \mathcal{X} as

$$\mathcal{X} = \mathcal{X}_1 + \mathcal{X}_2 + \cdots + \mathcal{X}_\beta + \bar{\mathcal{X}}. \quad (25)$$

Suppose $q \geq p$, we have

$$\begin{aligned} \|\bar{\mathcal{X}}\|_{S_q^\perp}^q &= \sum_{t=2^\beta}^{\tilde{n}} s_t^q = \sum_{t=2^\beta}^{\tilde{n}} s_t^p s_t^{q-p} \leq s_{2^\beta}^{q-p} \sum_{t=2^\beta}^{\tilde{n}} s_t^p \\ &\leq \left(\frac{1}{2^\beta} \sum_{t=1}^{2^\beta} s_t^p \right)^{\frac{q-p}{p}} \left(\sum_{t=2^\beta}^{\tilde{n}} s_t^p \right) \\ &\leq 2^{-\beta(q-p)/p} \|\mathcal{X}\|_{S_p^\perp}^{q-p} \|\mathcal{X}\|_{S_p^\perp}^p = 2^{-\beta(q-p)/p} \|\mathcal{X}\|_{S_p^\perp}^q \leq 2^{-\beta(q-p)/p}. \end{aligned} \quad (26)$$

It follows that

$$\left\| \mathcal{X} - \sum_{j=1}^{\beta} \mathcal{X}_j \right\|_{S_q^\perp} = \|\bar{\mathcal{X}}\|_{S_q^\perp} \leq 2^{\beta(1/q-1/p)}.$$

We focus on the case $q = 2$. We see that $2^{(j-1)(1/p-1/2)} \mathcal{X}_j \in \mathcal{S}_{d,n,2}^{2^{j-1}, \perp}$. Let $\mathcal{N}_j \subseteq \mathcal{S}_{d,n,2}^{2^{j-1}, \perp}$ be an ϵ_j -net, $j \in [\beta]$, and define

$$\mathcal{N} := \left\{ \sum_{j=1}^{\beta} 2^{-(j-1)(1/p-1/2)} \mathcal{Z}_j : \mathcal{Z}_j \in \mathcal{N}_j, j \in [\beta] \right\}.$$

Then \mathcal{N} is an ϵ -net of $\mathcal{S}_{d,n,p}^\perp$ with $\psi = 1$ in $\|\cdot\|_{S_2^\perp}$, where

$$\epsilon = \sum_{j=1}^{\beta} 2^{-(j-1)(1/p-1/2)} \epsilon_j + 2^{\beta(1/2-1/p)}$$

and $|\mathcal{N}| = \prod_{j=1}^{\beta} |\mathcal{N}_j|$.

Now let $\epsilon_j = c2^{(j-\beta)(1/p-1/2+1)}$, where c is the constant in Lemma 1. We have

$$\begin{aligned}\epsilon &= \sum_{j=1}^{\beta} c2^{(1-\beta)(1/p-1/2)} 2^{(j-\beta)} + 2^{\beta(1/2-1/p)} \\ &= c2^{(1-\beta)(1/p-1/2)} \sum_{j=1}^{\beta} 2^{(j-\beta)} + 2^{\beta(1/2-1/p)} \\ &\leq c' 2^{-\beta(1/p-1/2)} \\ &= c' \left(\frac{n}{\eta}\right)^{1/p-1/2},\end{aligned}$$

where c' is a constant. Then we have

$$\eta \leq n \left(\frac{c'}{\epsilon}\right)^{2p/(2-p)}. \quad (27)$$

Using Lemma 1 and the definition of \mathcal{N}_j , we obtain

$$\begin{aligned}\log |\mathcal{N}| &= \log \prod_{j=1}^{\beta} |\mathcal{N}_j| \leq \log \prod_{j=1}^{\beta} \left(\frac{c(d+1)}{c2^{(j-\beta)(1/p-1/2+1)}} \right)^{d2^{j-1}(n-(2^{j-1}+1)/2)+2^{j-1}} \\ &= \sum_{j=1}^{\beta} \left(d2^{j-1} \left(n - \frac{2^{j-1}+1}{2} \right) + 2^{j-1} \right) \log \left((d+1)2^{(\beta-j)(1/p-1/2+1)} \right) \\ &\leq \sum_{j=1}^{\beta} dn2^{j-1} ((\beta-j)(1/p-1/2+1) + \log(d+1)) \\ &\leq (1/p-1/2+1)d \log(d+1)n \sum_{j=1}^{\beta} 2^{j-1} (\beta-j) \\ &= (1/p-1/2+1)d \log(d+1)n2^{\beta} \sum_{j=1}^{\beta} 2^{j-\beta-1} (\beta-j) \\ &\leq (1/p-1/2+1)d \log(d+1)\eta \\ &\leq (1/p-1/2+1)nd \log(d+1) \left(\frac{c'}{\epsilon}\right)^{2p/(2-p)}.\end{aligned} \quad (28)$$

Now using a general ψ_p instead of 1, we finish the proof. \square

D.5 Proof of Theorem 4

Before proving the theorem, we give the following lemma.

Lemma 2. *Let \mathcal{S} be a set of d -order hyper-cubic tensors of side length n . Suppose the ϵ -covering numbers of \mathcal{S} with respect to the Frobenius norm are upper-bounded by B . Suppose $\max\{\|\mathcal{D}\|_{\infty}, \|\mathcal{X}\|_{\infty}\} \leq \epsilon$. Then the following inequality holds with probability at least $1 - 2n^{-d}$,*

$$\sup_{\mathcal{X} \in \mathcal{S}} \left| \frac{1}{\sqrt{n^d}} \|\mathcal{D} - \mathcal{X}\|_F - \frac{1}{\sqrt{|\Omega|}} \|\mathcal{P}_{\Omega}(\mathcal{D} - \mathcal{X})\|_F \right| \leq \frac{2\epsilon}{\sqrt{|\Omega|}} + 2\epsilon \left(\frac{d \log n + \log |B|}{2|\Omega|} \right)^{1/4}.$$

The proof of the lemma can be found in Appendix E.2. Now we prove Theorem 4 as follows.

Proof. According to Theorem 3, the covering numbers of $\mathcal{S}_p^{n,\perp}$ with respect to the Frobenius norm satisfy

$$\log |\mathcal{S}_p^{n,\perp}| \leq (1/2 + 1/p)nd (\log(d+1)) \left(\frac{c\|\mathcal{X}\|_{\mathcal{S}_p^{\perp}}}{\epsilon} \right)^{2p/(2-p)} \triangleq \log B, \quad (29)$$

where $c > 0$ is a universal constant. Now substituting $\log B$ into Lemma 2 and letting $\epsilon = c\tau\epsilon\sqrt{dn}$, we get

$$\begin{aligned}
& \sup_{\mathcal{X} \in \mathcal{S}} \left| \frac{1}{\sqrt{n^d}} \|\mathcal{D} - \mathcal{X}\|_F - \frac{1}{\sqrt{|\Omega|}} \|\mathcal{P}_\Omega(\mathcal{D} - \mathcal{X})\|_F \right| \\
& \leq \frac{2\epsilon}{\sqrt{|\Omega|}} + 2\epsilon \left(\frac{d \log n + (\frac{1}{2} + \frac{1}{p})nd (\log(d+1)) \left(\frac{c\|\mathcal{X}\|_{S_p^\perp}}{\epsilon} \right)^{2p/(2-p)}}{2|\Omega|} \right)^{1/4} \\
& = \frac{2c\tau\epsilon\sqrt{dn}}{\sqrt{|\Omega|}} + 2\epsilon \left(\frac{d \log n + (\frac{1}{2} + \frac{1}{p})nd (\log(d+1)) \left(\frac{\|\mathcal{X}\|_{S_p^\perp}}{\epsilon\tau\sqrt{dn}} \right)^{2p/(2-p)}}{2|\Omega|} \right)^{1/4} \tag{30} \\
& \leq c'\epsilon \left(\frac{(\frac{1}{2} + \frac{1}{p})nd (\log(d+1)) \left(\frac{\|\mathcal{X}\|_{S_p^\perp}}{\epsilon\sqrt{dn}} \right)^{2p/(2-p)}}{|\Omega|} \right)^{1/4},
\end{aligned}$$

where we have let $\tau = 1$ and c' is a universal constant. This finished the proof. \square

D.6 Proof of Theorem 5

The following lemma (proved in Appendix E.3) will be used in the proof of the theorem.

Lemma 3. Define $B_{2,q}^{n,k} := \{\mathbf{X} \in \mathbb{R}^{n \times k} : \|\mathbf{X}\|_{2,q} \leq 1\}$. Then the covering numbers of $B_{2,q}^{n,k}$ with respect to the Frobenius norm satisfy:

(a) $\log \mathcal{N}(\mathcal{B}_{2,q}^{n,k}, \|\cdot\|_F, \epsilon) \leq c_q(n + \log(ek))\epsilon^{-2q/(2-q)}$, when $0 < q < 1$,

(b) $\log \mathcal{N}(\mathcal{B}_{2,q}^{n,k}, \|\cdot\|_F, \epsilon) \leq \lceil nk^{2(q-1)/q} \epsilon^{-2} \rceil \log(2nk)$, when $q \geq 1$,

where $c_q = O\left(\frac{1}{q}\right)$.

Proof. Let $\mathcal{X} = \mathcal{I} \times_1 \mathbf{X}^{(1)} \times_2 \cdots \times_d \mathbf{X}^{(d)}$ be the CP (or Tucker equivalently) decomposition of $\mathcal{X} \in \mathcal{S}_{k,p}^n$, where \mathcal{I} is super-diagonal. tensor of 1s.

Let $\bar{\mathcal{X}} = \mathcal{I} \times_1 \bar{\mathbf{X}}^{(1)} \times_2 \dots \times_d \bar{\mathbf{X}}^{(d)}$ and denote $\varsigma = \|\mathcal{C}\|_F$. Let $\|\mathbf{X}^{(j)} - \bar{\mathbf{X}}^{(j)}\|_F \leq \frac{\epsilon}{d} \prod_{i \neq j} \gamma_i^{-1}$, $j \in [d]$. We have

$$\begin{aligned}
& \|\mathcal{X} - \bar{\mathcal{X}}\|_F \\
&= \|\mathcal{I} \times_1 \mathbf{X}^{(1)} \times_2 \dots \times_d \mathbf{X}^{(d)} - \mathcal{I} \times_1 \bar{\mathbf{X}}^{(1)} \times_2 \dots \times_d \bar{\mathbf{X}}^{(d)}\|_F \\
&= \|\mathcal{I} \times_1 \mathbf{X}^{(1)} \times_2 \dots \times_d \mathbf{X}^{(d)} \pm \mathcal{I} \times_1 \mathbf{X}^{(1)} \times_2 \dots \times_d \bar{\mathbf{X}}^{(d)} \\
&\quad \pm \mathcal{I} \times_1 \mathbf{X}^{(1)} \times_2 \dots \times_{d-1} \bar{\mathbf{X}}^{(d-1)} \times_d \bar{\mathbf{X}}^{(d)} \\
&\quad \pm \dots \pm \mathcal{I} \times_1 \bar{\mathbf{X}}^{(1)} \times_2 \dots \times_{d-1} \bar{\mathbf{X}}^{(d-1)} \times_d \bar{\mathbf{X}}^{(d)} \\
&\quad - \mathcal{I} \times_1 \bar{\mathbf{X}}^{(1)} \times_2 \dots \times_d \bar{\mathbf{X}}^{(d)}\|_F \\
&\leq \|\mathcal{I} \times_1 \mathbf{X}^{(1)} \times_2 \dots \times_d (\mathbf{X}^{(d)} - \bar{\mathbf{X}}^{(d)})\|_F \\
&\quad + \|\mathcal{I} \times_1 \mathbf{X}^{(1)} \times_2 \dots \times_{d-1} (\mathbf{X}^{(d-1)} - \bar{\mathbf{X}}^{(d-1)}) \times_d \bar{\mathbf{X}}^{(d)}\|_F \\
&\quad + \dots + \|\mathcal{I} \times_1 (\mathbf{X}^{(1)} - \bar{\mathbf{X}}^{(1)}) \times_2 \bar{\mathbf{X}}^{(2)} \dots \times_d \bar{\mathbf{X}}^{(d)}\|_F \\
&\leq \|\mathcal{I}\|_{op} \|\mathbf{X}^{(1)}\|_{op} \dots \|\mathbf{X}^{(d-1)}\|_{op} \|\mathbf{X}^{(d)} - \bar{\mathbf{X}}^{(d)}\|_F \\
&\quad + \|\mathcal{I}\|_{op} \|\mathbf{X}^{(1)}\|_{op} \dots \|\mathbf{X}^{(d-1)} - \bar{\mathbf{X}}^{(d-1)}\|_F \|\bar{\mathbf{X}}^{(d)}\|_{op} \\
&\quad + \dots + \|\mathcal{I}\|_{op} \|\mathbf{X}^{(1)} - \bar{\mathbf{X}}^{(1)}\|_F \|\bar{\mathbf{X}}^{(2)}\|_{op} \dots \|\bar{\mathbf{X}}^{(d-1)}\|_{op} \|\bar{\mathbf{X}}^{(d)}\|_{op} \\
&\leq \frac{\epsilon}{d} + \frac{\epsilon}{d} + \dots + \frac{\epsilon}{d} \\
&= \epsilon.
\end{aligned}$$

Denote $\phi = \prod_{j=1}^d \gamma_j$. When $0 < q < 1$, the the covering number of $\mathcal{S}_{d,n,q}^k$ can be bounded as

$$\mathcal{N}(\mathcal{S}_{d,n,q}^k, \|\cdot\|_F, \epsilon) \leq \prod_{i=1}^d \exp \left(c_q (n + \log(ek)) \left(\frac{d\phi\alpha_q^{(j)}\gamma_j^{-1}}{\epsilon} \right)^{2q/(2-q)} \right),$$

where c_q is a universal constant. It follows that

$$\log \mathcal{N}(\mathcal{S}_{d,n,q}^k, \|\cdot\|_F, \epsilon) \leq \frac{c(n + \log(ek))}{q} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}\gamma_j^{-1}}{\epsilon} \right)^{2q/(2-q)},$$

where c is a constant.

When $q \geq 1$, the the covering number of $\mathcal{S}_{d,n,q}^k$ can be bounded as

$$\mathcal{N}(\mathcal{S}_{d,n,q}^k, \|\cdot\|_F, \epsilon) \leq \prod_{i=1}^d \exp \left(c' n k^{2(q-1)/q} \log(2nk) \left(\frac{d\phi\alpha_q^{(j)}\gamma_j^{-1}}{\epsilon} \right)^2 \right),$$

where we have converted the *ceil* operation into multiplying a constant c' for simplicity. It follows that

$$\log \mathcal{N}(\mathcal{S}_{d,n,q}^k, \|\cdot\|_F, \epsilon) \leq c' n k^{2(q-1)/q} \log(2nk) \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}\gamma_j^{-1}}{\epsilon} \right)^2.$$

□

D.7 Proof for Theorem 6

The following lemma provides a sample complexity bound for transductive learning, which is consistent with the objective function and evaluation metric (RMSE) widely used in collaborative filtering.

Lemma 4 (Corollary 1 of (El-Yaniv & Pechyony, 2009), reformulated). *Let \mathcal{F} be a fixed hypothesis set and suppose $\sup_{i,j|\mathbf{X} \in \mathcal{H}} |\ell(Y_{ij}, X_{ij})| \leq \tau_\ell$. Suppose a fixed set S of distinct indices is uniformly and randomly split to two subsets S_{train} and S_{test} , where⁵ $|S_{\text{test}}| > |S_{\text{train}}| > 50$. Then with probability at least $1 - \delta$ over the random split, we have*

$$\begin{aligned} \frac{1}{|S_{\text{test}}|} \sum_{(i,j) \in S_{\text{test}}} \ell(Y_{ij}, X_{ij}) &\leq \frac{1}{|S_{\text{train}}|} \sum_{(i,j) \in S_{\text{train}}} \ell(Y_{ij}, X_{ij}) + 4\mathcal{R}_S(\ell \circ \mathcal{F}) \\ &\quad + \frac{11\tau_\ell(|S_{\text{train}}| + |S_{\text{test}}|)}{\sqrt{|S_{\text{train}}||S_{\text{test}}|}} + 3\tau_\ell \sqrt{\frac{(|S_{\text{train}}| + |S_{\text{test}}|)}{|S_{\text{train}}||S_{\text{test}}|} \log \frac{1}{\delta}}. \end{aligned} \quad (31)$$

Before proof, we give the following lemma, which is a variant of the Dudley entropy integral bound on Rademacher complexity.

Lemma 5 (Theorem 3 of (Schreuder, 2020)). *Let $\mathcal{F} \subset \{f : \mathcal{X} \mapsto \mathbb{R}\}$ be any class of measurable functions containing the uniformly zero function and let $S_n(\mathcal{F}) = \sup_{f \in \mathcal{F}} \|f\|_{L_2(P_n)}$. Then*

$$\mathcal{R}_n(\mathcal{F}) \leq \inf_{\tau > 0} \left(4\tau + \frac{12}{\sqrt{n}} \int_{\tau}^{S_n(\mathcal{F})} \sqrt{\log \mathcal{N}(\mathcal{F}, L_2(P_n), \zeta)} d\zeta \right). \quad (32)$$

In the lemma, $\|f\|_{L_2(P_n)}$ is defined as $\|f\|_{L_2(P_n)}^2 = \frac{1}{n} \sum_{i=1}^n f(X_i)^2$, which means

$$\mathcal{N}(\mathcal{F}, L_2(P_n), \zeta) = \mathcal{N}(\mathcal{F}, \|\cdot\|_F, \sqrt{n}\zeta). \quad (33)$$

Then we have

$$\begin{aligned} \mathcal{R}_n(\mathcal{F}) &\leq \inf_{\tau > 0} \left(4\tau + \frac{12}{\sqrt{n}} \int_{\tau}^{S_n(\mathcal{F})} \sqrt{\log \mathcal{N}(\mathcal{F}, \|\cdot\|_F, \sqrt{n}\zeta)} d\zeta \right) \\ &= \inf_{\tau > 0} \left(\frac{4\tau}{\sqrt{n}} + \frac{12}{n} \int_{\tau}^{S_n(\mathcal{F})\sqrt{n}} \sqrt{\log \mathcal{N}(\mathcal{F}, \|\cdot\|_F, \epsilon)} d\epsilon \right). \end{aligned} \quad (34)$$

⁵We use these assumptions to simplify the theorem.

Now we use Theorem 5 and (34) to obtain the Rademacher complexity of the tensor decomposition model in LRTC-ENR. When $0 < q < 1$, we have

$$\begin{aligned}
\mathcal{R}_{|\Omega|}(\mathcal{F}) &\leq \inf_{\tau > 0} \left(\frac{4\tau}{\sqrt{|\Omega|}} + \frac{12}{|\Omega|} \int_{\tau}^{\varepsilon\sqrt{|\Omega|}} \sqrt{\frac{c(n + \log(ek))}{q} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j \epsilon} \right)^{\frac{2q}{2-q}}} d\epsilon \right) \\
&= \inf_{\tau > 0} \left(\frac{4\tau}{\sqrt{|\Omega|}} + \frac{12}{|\Omega|} \sqrt{\frac{c(n + \log(ek))}{q} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j} \right)^{\frac{2q}{2-q}}} \int_{\tau}^{\varepsilon\sqrt{|\Omega|}} \epsilon^{-\frac{q}{2-q}} d\epsilon \right) \\
&= \inf_{\tau > 0} \left(\frac{4\tau}{\sqrt{|\Omega|}} + \frac{12}{|\Omega|} \sqrt{\frac{c(n + \log(ek))}{q} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j} \right)^{\frac{2q}{2-q}}} \left(\frac{2-q}{2-2q} \right) \left((\varepsilon\sqrt{|\Omega|})^{\frac{2-2q}{2-q}} - \tau^{\frac{2-2q}{2-q}} \right) \right) \quad (35) \\
&\leq \frac{4\varepsilon}{|\Omega|} + \frac{12}{|\Omega|} \sqrt{\frac{c(n + \log(ek))}{q} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j} \right)^{\frac{2q}{2-q}}} \left(\frac{2-q}{2-2q} \right) (\varepsilon\sqrt{|\Omega|})^{\frac{2-2q}{2-q}} \\
&\leq \frac{c'}{|\Omega|} \sqrt{\frac{(n + \log(ek))(2-q)^2 (\varepsilon\sqrt{|\Omega|})^{\frac{2-2q}{2-q}}}{q(2-2q)^2} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j} \right)^{\frac{2q}{2-q}}} \\
&= \frac{c'}{|\Omega|} \sqrt{\frac{(n + \log(ek))(2-q)^2 \varepsilon^2 |\Omega|^{\frac{1-q}{2-q}}}{q(2-2q)^2} \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\varepsilon\gamma_j} \right)^{\frac{2q}{2-q}}}
\end{aligned}$$

where c' is a suitable numerical constant.

When $q \geq 1$, we have

$$\begin{aligned}
\mathcal{R}_{|\Omega|}(\mathcal{F}) &\leq \inf_{\tau > 0} \left(\frac{4\tau}{\sqrt{|\Omega|}} + \frac{12}{|\Omega|} \int_{\tau}^{\varepsilon\sqrt{|\Omega|}} \sqrt{cnk^{2(q-1)/q} \log(2nk) \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)} q^{(j)}}{\gamma_j \epsilon} \right)^2} d\epsilon \right) \\
&= \inf_{\tau > 0} \left(\frac{4\tau}{\sqrt{|\Omega|}} + \frac{12}{|\Omega|} \sqrt{cnk^{2(q-1)/q} \log(2nk) \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j} \right)^2} \int_{\tau}^{\varepsilon\sqrt{|\Omega|}} \epsilon^{-1} d\epsilon \right) \quad (36) \\
&\leq \inf_{\tau > 0} \left(\frac{4\tau}{|\Omega|} + \frac{12}{|\Omega|} \sqrt{cnk^{2(q-1)/q} \log(2nk) \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j} \right)^2} \log \frac{\varepsilon\sqrt{|\Omega|}}{\tau} \right) \\
&\leq \frac{c'}{|\Omega|} \sqrt{nk^{2(q-1)/q} \log(2nk) \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\gamma_j} \right)^2} \log |\Omega|,
\end{aligned}$$

where c' is a suitable numerical constant.

Note that

$$\mathcal{R}_{|\Omega|}(\ell \circ \mathcal{F}) \leq \eta_{\ell} \mathcal{R}_{|\Omega|}(\mathcal{F}), \quad (37)$$

where η_{ℓ} is the Lipschitz constant of function ℓ . In this work, ℓ is the square loss, which means $\eta_{\ell} = 4\varepsilon$. Finally, integrating (35), (36), and (37) with Lemma 4 and renaming the constants, we get the desired results.

D.8 Proof of Corollary 1

Proof. In Theorem 6, letting $\theta = \sum_{j=1}^d \left(\frac{d\phi\alpha_q^{(j)}}{\varepsilon\gamma_j} \right)^t$, we have

$$\theta \leq \left(\frac{d\bar{\gamma}^{d-1}}{\varepsilon} \right)^t \sum_{j=1}^d (\alpha_q^{(j)})^t, \quad (38)$$

where $t = 2$ or $t = 2q/(2-q)$. Recall that in the proof for Theorem 1, the equality holds only when $\alpha_{i1} = \alpha_{i2} = \dots = \alpha_{id} = |\lambda_i|^{1/d}$, for all $i \in [k]$. Here we have $\alpha_q^{(j)} = (\sum_{i \in [k]} \alpha_{ij}^p)^{1/p}$. That means we can get $\|\mathcal{X}\|_{S_{p/d}}$ only when $\alpha_q^{(1)} = \alpha_q^{(2)} = \dots = \alpha_q^{(d)} = \|\mathcal{X}\|_{S_{p/d}}^{1/d}$. Then we have

$$\theta \leq d \left(\frac{d\bar{\gamma}^{d-1}}{\varepsilon} \right)^t \|\mathcal{X}\|_{S_{q/d}}^{t/d}. \quad (39)$$

Now in Theorem 6, we have

$$B_{\mathcal{R}} \leq \begin{cases} \frac{c_1}{|\Omega|} \sqrt{nk^{2(q-1)/q} \log(2nk) d \left(\frac{d\bar{\gamma}^{d-1}}{\varepsilon} \right)^2 \|\mathcal{X}\|_{S_{q/d}}^{2/d} \log |\Omega|} & \text{if } q \geq 1 \\ \frac{c_2}{|\Omega|} \sqrt{\frac{(n+\log(ek))(2-q)^2 |\Omega|^{\frac{1-q}{2-q}}}{q(2-2q)^2} d \left(\frac{d\bar{\gamma}^{d-1}}{\varepsilon} \right)^{2q/(2-q)} \|\mathcal{X}\|_{S_{q/d}}^{2q/(2-q)/d}} & \text{if } 0 < q < 1 \end{cases}$$

Letting $q = pd$, we arrive at

$$B_{\mathcal{R}} \leq \begin{cases} \frac{c_1}{|\Omega|} \sqrt{nk^{2(pd-1)/pd} \log(2nk) d \left(\frac{d\bar{\gamma}^{d-1}}{\varepsilon} \right)^2 \|\mathcal{X}\|_{S_p}^{2/d} \log |\Omega|} & \text{if } p \geq 1/d \\ \frac{c_2}{|\Omega|} \sqrt{\frac{(n+\log(ek))(2-pd)^2}{pd(2-2pd)^2} d \left(\frac{d\bar{\gamma}^{d-1}}{\varepsilon} \right)^{2pd/(2-pd)} \|\mathcal{X}\|_{S_p}^{2p/(2-pd)} |\Omega|^{\frac{1-pd}{4-2pd}}} & \text{if } 0 < p < 1/d \end{cases}$$

Now rename $N_{\mathcal{R}}$ as the upper bound, we finish the proof. \square

D.9 Proof of Theorem 7

Proof. The following three lemmas will be used. Their proof are in Section E.

Lemma 6. For any $\mathbf{z} \in \mathbb{R}_+^n$ and $\tau > 0$ for, the following inequality holds

$$\|\mathbf{z}\|_1 \leq \tau^{-p/2} \|\mathbf{z}\|_2 \sqrt{\|\mathbf{z}\|_p^p + \tau^{1-p} \|\mathbf{z}\|_p^p}. \quad (40)$$

Lemma 7. Suppose the entries of $\mathcal{N} \in \mathbb{R}^{n \otimes d}$ are drawn from $\mathcal{N}(0, \sigma^2)$ independently. Then the following inequality holds with probability at least $1 - 2n^{-d}$

$$\|\mathcal{N}\|_2 \leq 2\sqrt{2}\sigma \sqrt{dn \log(5d) + d \log(n)}.$$

Lemma 8. Suppose the entries of $\mathcal{N} \in \mathbb{R}^{n \otimes d}$ are drawn from $\mathcal{N}(0, \sigma^2)$ independently. The the following inequality holds with probability at least

$$\|\mathcal{N}\|_{\infty} \leq 2\sigma \sqrt{d \log(n)}.$$

Let $\hat{\mathcal{X}}, \hat{\mathcal{E}}$ be the optimal solution of (10). We have

$$\|\mathcal{D} - \hat{\mathcal{X}} - \hat{\mathcal{E}}\|_F^2 \leq \|\mathcal{D} - \mathcal{X}^* - \mathcal{E}^*\|_F^2. \quad (41)$$

Since $\mathcal{D} = \mathcal{X}^* + \mathcal{N}^* + \mathcal{E}^*$, we obtain

$$\|\mathcal{X}^* + \mathcal{N}^* + \mathcal{E}^* - \hat{\mathcal{X}} - \hat{\mathcal{E}}\|_F^2 \leq \|\mathcal{N}^*\|_F^2. \quad (42)$$

It follows that

$$\begin{aligned} & \|\mathcal{X}^* - \hat{\mathcal{X}}\|_F^2 + \|\mathcal{E}^* - \hat{\mathcal{E}}\|_F^2 \\ & \leq -2\langle \mathcal{X}^* - \hat{\mathcal{X}}, \mathcal{N}^* \rangle - 2\langle \mathcal{X}^* - \hat{\mathcal{X}}, \mathcal{E}^* - \hat{\mathcal{E}} \rangle - 2\langle \mathcal{E}^* - \hat{\mathcal{E}}, \mathcal{N}^* \rangle \\ & \leq 2\|\mathcal{X}^* - \hat{\mathcal{X}}\|_{S_1^\perp} \|\mathcal{N}^*\|_2 + 2\|\mathcal{E}^* - \hat{\mathcal{E}}\|_1 \left(\|\mathcal{X}^* - \hat{\mathcal{X}}\|_\infty + \|\mathcal{N}^*\|_\infty \right). \end{aligned} \quad (43)$$

Using Lemma 6, we obtain

$$\begin{aligned} \|\mathcal{X}^* - \hat{\mathcal{X}}\|_{S_1^\perp} & \leq \tau_1^{-p/2} \|\mathcal{X}^* - \hat{\mathcal{X}}\|_F \sqrt{\|\mathcal{X}^* - \hat{\mathcal{X}}\|_{S_p^\perp}^p} + \tau_1^{1-p} \|\mathcal{X}^* - \hat{\mathcal{X}}\|_{S_p^\perp}^p \\ & \leq \tau_1^{-p/2} \|\mathcal{X}^* - \hat{\mathcal{X}}\|_F \sqrt{2R_x^p} + 2\tau_1^{1-p} R_x^p \end{aligned} \quad (44)$$

and

$$\begin{aligned} \|\mathcal{E}^* - \hat{\mathcal{E}}\|_1 & \leq \tau_2^{-p'/2} \|\mathcal{E}^* - \hat{\mathcal{E}}\|_F \sqrt{\|\mathcal{E}^* - \hat{\mathcal{E}}\|_{p'}^{p'}} + \tau_2^{1-p'} \|\mathcal{E}^* - \hat{\mathcal{E}}\|_{p'}^{p'} \\ & \leq \tau_2^{-p'/2} \|\mathcal{E}^* - \hat{\mathcal{E}}\|_F \sqrt{2R_e^{p'}} + 2\tau_2^{1-p'} R_e^{p'} \end{aligned} \quad (45)$$

which hold for any $\tau_1 > 0$ and $\tau_2 > 0$. Substituting (44) and (45) into (43), we have

$$\begin{aligned} & \|\mathcal{X}^* - \hat{\mathcal{X}}\|_F^2 + \|\mathcal{E}^* - \hat{\mathcal{E}}\|_F^2 \\ & \leq 2\tau_1^{-p/2} \sqrt{2R_x^p} \|\mathcal{N}^*\|_2 \|\mathcal{X}^* - \hat{\mathcal{X}}\|_F \\ & \quad + 2\tau_2^{-p'/2} \sqrt{2R_e^{p'}} \left(\|\mathcal{X}^* - \hat{\mathcal{X}}\|_\infty + \|\mathcal{N}^*\|_\infty \right) \|\mathcal{E}^* - \hat{\mathcal{E}}\|_F \\ & \quad + 4\tau_1^{1-p} R_x^p \|\mathcal{N}^*\|_2 + 4\tau_2^{1-p'} R_e^{p'} \left(\|\mathcal{X}^* - \hat{\mathcal{X}}\|_\infty + \|\mathcal{N}^*\|_\infty \right). \end{aligned} \quad (46)$$

For convenience, let

$$\begin{aligned} u &= \|\mathcal{X}^* - \hat{\mathcal{X}}\|_F \\ v &= \|\mathcal{E}^* - \hat{\mathcal{E}}\|_F \\ c_1 &= 2\tau_1^{-p/2} \sqrt{2R_x^p} \|\mathcal{N}^*\|_2 \\ c_2 &= 2\tau_2^{-p'/2} \sqrt{2R_e^{p'}} \left(\|\mathcal{X}^* - \hat{\mathcal{X}}\|_\infty + \|\mathcal{N}^*\|_\infty \right) \\ c_3 &= 4\tau_1^{1-p} R_x^p \|\mathcal{N}^*\|_2 + 4\tau_2^{1-p'} R_e^{p'} \left(\|\mathcal{X}^* - \hat{\mathcal{X}}\|_\infty + \|\mathcal{N}^*\|_\infty \right) \\ 2\alpha &\geq \|\mathcal{X}^* - \hat{\mathcal{X}}\|_\infty \\ \beta &\geq \|\mathcal{N}^*\|_\infty \\ \gamma &\geq \|\mathcal{N}^*\|_2 \end{aligned}$$

We rewrite (46) as

$$\begin{aligned} u^2 + v^2 & \leq c_1 u + c_2 v + c_3 \\ & \leq (c_1 + c_2)(u + v) + c_3 \\ & \leq (c_1 + c_2) \sqrt{2(u^2 + v^2)} + c_3. \end{aligned} \quad (47)$$

Because $c_3 > 0$, the quadratic inequality of $\sqrt{u^2 + v^2}$ has non-empty solution. We have

$$\max \left(0, \frac{\sqrt{2}(c_1 + c_2) - \sqrt{2(c_1 + c_2)^2 + 4c_3}}{2} \right) \leq \sqrt{u^2 + v^2} \leq \frac{\sqrt{2}(c_1 + c_2) + \sqrt{2(c_1 + c_2)^2 + 4c_3}}{2}. \quad (48)$$

Let $\tau_1 = \gamma$ and $\tau_2 = 2\alpha + \beta$. We have

$$\begin{aligned}
\sqrt{u^2 + v^2} &\leq 2\gamma^{1-\frac{p}{2}} \sqrt{R_x^p} + 2(2\alpha + \beta)^{1-\frac{p'}{2}} \sqrt{R_e^{p'}} \\
&\quad + 2\sqrt{\left(\gamma^{1-\frac{p}{2}} \sqrt{R_x^p} + (2\alpha + \beta)^{1-\frac{p'}{2}} \sqrt{R_e^{p'}}\right)^2 + \gamma^{2-p} R_x^p + (2\alpha + \beta)^{2-p'} R_e^{p'}} \\
&\leq 2\gamma^{1-\frac{p}{2}} \sqrt{R_x^p} + 2(2\alpha + \beta)^{1-\frac{p'}{2}} \sqrt{R_e^{p'}} + 2\sqrt{2} \left(\gamma^{1-\frac{p}{2}} \sqrt{R_x^p} + (2\alpha + \beta)^{1-\frac{p'}{2}} \sqrt{R_e^{p'}}\right) \\
&= (2 + 2\sqrt{2}) \left(\gamma^{1-\frac{p}{2}} \sqrt{R_x^p} + (2\alpha + \beta)^{1-\frac{p'}{2}} \sqrt{R_e^{p'}}\right).
\end{aligned} \tag{49}$$

It follows that

$$\begin{aligned}
u^2 + v^2 &\leq (12 + 8\sqrt{2}) \left(\gamma^{1-\frac{p}{2}} \sqrt{R_x^p} + (2\alpha + \beta)^{1-\frac{p'}{2}} \sqrt{R_e^{p'}}\right)^2 \\
&\leq (24 + 16\sqrt{2}) \left(\gamma^{2-p} R_x^p + (2\alpha + \beta)^{2-p'} R_e^{p'}\right).
\end{aligned} \tag{50}$$

□

Using Lemma 7 and Lemma 8 for γ and β respectively, we have

$$\begin{aligned}
&\|\mathbf{X}^* - \hat{\mathbf{X}}\|_F^2 + \|\mathbf{E}^* - \hat{\mathbf{E}}\|_F^2 \\
&\leq (24 + 16\sqrt{2}) \left(\left(2\sqrt{2}\sigma\sqrt{dn\log(5d) + d\log(n)}\right)^{2-p} R_x^p + \left(2\alpha + 2\sigma\sqrt{d\log(n)}\right)^{2-p'} R_e^{p'} \right)
\end{aligned} \tag{51}$$

with probability at least $1 - 4n^{-d}$, where c is an absolute constant. This finished the proof.

E Proof of the lemmas

E.1 Proof of Lemma 1

Before proof, we restate Lemma 4.1 of (Hinrichs et al., 2017) here.

Lemma 9 (Lemma 4.1 of (Hinrichs et al., 2017)). *Define $V_k^n = \{\mathbf{U} \in \mathbb{R}^{n \times k} : \mathbf{U}^\top \mathbf{U} = \mathbf{I}_k, k, n \in \mathbb{N}, k \leq n\}$. Let $0 < \epsilon < 1$. Then*

$$\mathcal{N}(V_k^n, \|\cdot\|_{op}, \epsilon) \leq \left(\frac{c}{\epsilon}\right)^{k(n-(k+1)/2)},$$

where $c > 0$ is a universal constant.

Proof. Let $\mathbf{X} = \mathbf{C} \times_1 \mathbf{X}^{(1)} \times_2 \cdots \times_d \mathbf{X}^{(d)}$ be the orthogonal CP (or Tucker equivalently) decomposition of $\mathbf{X} \in \mathcal{S}_{d,n,2}^{k,\perp}$, where \mathbf{C} is super-diagonal. We have $\|\mathbf{C}\|_{S_2^\perp} = \|\text{diag}(\mathbf{C})\|_2 \leq 1$. Then the covering number of \mathbf{C} with respect to $\|\cdot\|_{S_2^\perp}$ is equal to the covering number of $\text{diag}(\mathbf{C})$ with respect to $\|\cdot\|_2$, i.e., $|\mathcal{N}_{\mathbf{C}}| \leq (3\epsilon^{-1})^k$. According to Lemma 9, the covering number of $\mathbf{X}^{(j)}$ with respect to $\|\cdot\|_{op}$ satisfies

$|\mathcal{N}_{\mathbf{X}^{(j)}}| \leq (c\epsilon^{-1})^{k(n-(k+1)/2)}$, $j \in [d]$. We have

$$\begin{aligned}
\|\mathbf{X} - \bar{\mathbf{X}}\|_{S_2^\perp} &= \|\mathbf{C} \times_1 \mathbf{X}^{(1)} \times_2 \cdots \times_d \mathbf{X}^{(d)} - \bar{\mathbf{C}} \times_1 \bar{\mathbf{X}}^{(1)} \times_2 \cdots \times_d \bar{\mathbf{X}}^{(d)}\|_{S_2^\perp} \\
&= \|\mathbf{C} \times_1 \mathbf{X}^{(1)} \times_2 \cdots \times_d \mathbf{X}^{(d)} \pm \mathbf{C} \times_1 \mathbf{X}^{(1)} \times_2 \cdots \times_d \bar{\mathbf{X}}^{(d)} \\
&\quad \pm \mathbf{C} \times_1 \mathbf{X}^{(1)} \times_2 \cdots \times_{d-1} \bar{\mathbf{X}}^{(d-1)} \times_d \bar{\mathbf{X}}^{(d)} \\
&\quad \pm \cdots \pm \mathbf{C} \times_1 \bar{\mathbf{X}}^{(1)} \times_2 \cdots \times_{d-1} \bar{\mathbf{X}}^{(d-1)} \times_d \bar{\mathbf{X}}^{(d)} - \bar{\mathbf{C}} \times_1 \bar{\mathbf{X}}^{(1)} \times_2 \cdots \times_d \bar{\mathbf{X}}^{(d)}\|_{S_2^\perp} \\
&\stackrel{(a)}{\leq} \|\mathbf{C} \times_1 \mathbf{X}^{(1)} \times_2 \cdots \times_d (\mathbf{X}^{(d)} - \bar{\mathbf{X}}^{(d)})\|_{S_2^\perp} \\
&\quad + \|\mathbf{C} \times_1 \mathbf{X}^{(1)} \times_2 \cdots \times_{d-1} (\mathbf{X}^{(d-1)} - \bar{\mathbf{X}}^{(d-1)}) \times_d \bar{\mathbf{X}}^{(d)}\|_{S_2^\perp} \\
&\quad + \cdots + \|(\mathbf{C} - \bar{\mathbf{C}}) \times_1 \bar{\mathbf{X}}^{(1)} \times_2 \bar{\mathbf{X}}^{(2)} \cdots \times_d \bar{\mathbf{X}}^{(d)}\|_{S_2^\perp} \\
&\stackrel{(b)}{\leq} \|\mathbf{C}\|_{S_2^\perp} \|\mathbf{X}^{(1)}\|_{op} \cdots \|\mathbf{X}^{(d-1)}\|_{op} \|\mathbf{X}^{(d)} - \bar{\mathbf{X}}^{(d)}\|_{op} \\
&\quad + \|\mathbf{C}\|_{S_2^\perp} \|\mathbf{X}^{(1)}\|_{op} \cdots \|\mathbf{X}^{(d-1)} - \bar{\mathbf{X}}^{(d-1)}\|_{op} \|\bar{\mathbf{X}}^{(d)}\|_{op} \\
&\quad + \cdots + \|\mathbf{C} - \bar{\mathbf{C}}\|_{S_2^\perp} \|\bar{\mathbf{X}}^{(1)}\|_{op} \|\bar{\mathbf{X}}^{(2)}\|_{op} \cdots \|\bar{\mathbf{X}}^{(d-1)}\|_{op} \|\bar{\mathbf{X}}^{(d)}\|_{op} \\
&= \|\mathbf{C}\|_{S_2^\perp} \|\mathbf{X}^{(d)} - \bar{\mathbf{X}}^{(d)}\|_{op} + \|\mathbf{C}\|_{S_2^\perp} \|\mathbf{X}^{(d-1)} - \bar{\mathbf{X}}^{(d-1)}\|_{op} + \cdots \\
&\quad + \|\mathbf{C}\|_{S_2^\perp} \|\mathbf{X}^{(1)} - \bar{\mathbf{X}}^{(1)}\|_{op} + \|\mathbf{C} - \bar{\mathbf{C}}\|_{S_2^\perp}.
\end{aligned}$$

Note that (a) holds owing to the triangle inequality of norms and (b) holds because the inequality $\|\mathbf{AB}\|_{S_q} \leq \|\mathbf{A}\|_{op} \|\mathbf{B}\|_{S_q}$ (Bhatia, 2013) can be easily extended to orthogonally decomposable tensors. Now let $\|\mathbf{C} - \bar{\mathbf{C}}\|_{S_2^\perp} \leq \epsilon/(d+1)$, and $\|\mathbf{X}^{(j)} - \bar{\mathbf{X}}^{(j)}\|_{op} \leq \epsilon/(d+1)$, $j \in [d]$. We arrive at

$$\|\mathbf{X} - \bar{\mathbf{X}}\|_{S_2^\perp} \leq \frac{\epsilon}{d+1} + \frac{\epsilon}{d+1} + \cdots + \frac{\epsilon}{d+1} \leq \epsilon.$$

Then the covering number of $\mathcal{S}_{d,n,2}^{k,\perp}$ can be bounded as

$$\begin{aligned}
\mathcal{N}(\mathcal{S}_{d,n,2}^{k,\perp}, \|\cdot\|_{S_2^\perp}, \epsilon) &\leq \left(\frac{3(d+1)}{\epsilon}\right)^k \prod_{i=1}^d \left(\frac{c(d+1)}{\epsilon}\right)^{k(n-(k+1)/2)} \\
&\leq \left(\frac{c'(d+1)}{\epsilon}\right)^{dk(n-(k+1)/2)+k},
\end{aligned}$$

where c' is a universal constant. This finished the proof. \square

E.2 Proof of Lemma 2

Proof. For convenience, we define

$$\hat{h}(\mathbf{X}) = \frac{1}{|\Omega|} \|\mathcal{P}_\Omega(\mathcal{D} - \mathbf{X})\|_F^2, \quad h(\mathbf{X}) = \frac{1}{n^d} \|\mathcal{D} - \mathbf{X}\|_F^2.$$

According to the following lemma

Lemma 10 ((Hoeffding inequality for sampling without replacement)). *Let X_1, X_2, \dots, X_s be a set of samples taken without replacement from a distribution $\{x_1, x_2, \dots, x_N\}$ of mean u and variance σ^2 . Denote $a = \min_i x_i$ and $b = \max_i x_i$. Then*

$$P \left[\left| \frac{1}{s} \sum_{i=1}^s X_i - u \right| \geq t \right] \leq 2 \exp \left(- \frac{2st^2}{(1 - (s-1)/N)(b-a)^2} \right).$$

we have

$$P \left[|\hat{h} - h| \geq t \right] \leq 2 \exp \left(- \frac{2|\Omega|t^2}{(1 - (|\Omega|-1)/n^d)\varsigma^2} \right),$$

where $\varsigma = 4\varepsilon^2$. Using union bound for all $\bar{\mathcal{X}} \in \mathcal{S}$ yields

$$P \left[\sup_{\bar{\mathcal{X}} \in \mathcal{S}} |\hat{h}(\bar{\mathcal{X}}) - h(\bar{\mathcal{X}})| \geq t \right] \leq 2|\mathcal{S}| \exp \left(-\frac{2|\Omega|t^2}{(1 - (|\Omega| - 1)/n^d)\varsigma^2} \right).$$

Or equivalently, with probability at least $1 - 2n^{-d}$,

$$\sup_{\bar{\mathcal{X}} \in \mathcal{S}} |\hat{h}(\bar{\mathcal{X}}) - h(\bar{\mathcal{X}})| \leq \sqrt{\frac{\varsigma^2 \log(|\mathcal{S}|n^d)}{2} \left(\frac{1}{|\Omega|} - \frac{1}{n^d} + \frac{1}{n^d|\Omega|} \right)}.$$

Then we have

$$\begin{aligned} g(\Omega) &\triangleq \sup_{\bar{\mathcal{X}} \in \mathcal{S}} |\hat{h}(\bar{\mathcal{X}}) - h(\bar{\mathcal{X}})| \\ &\leq \sqrt{\frac{\varsigma^2}{2} (d \log n + \log |\mathcal{S}|) \left(\frac{1}{|\Omega|} - \frac{1}{n^d} + \frac{1}{n^d|\Omega|} \right)}. \end{aligned}$$

Since $|\sqrt{u} - \sqrt{v}| \leq \sqrt{|u - v|}$ holds for any non-negative u and v , we have

$$\sup_{\bar{\mathcal{X}} \in \mathcal{S}} \left| \sqrt{\hat{h}(\bar{\mathcal{X}})} - \sqrt{h(\bar{\mathcal{X}})} \right| \leq \sqrt{g(\Omega)}.$$

Recall that $\epsilon \geq \|\mathcal{X} - \bar{\mathcal{X}}\|_F \geq \|\mathcal{P}(\mathcal{X} - \bar{\mathcal{X}})\|_F$, we have

$$\left| \sqrt{h(\mathcal{X})} - \sqrt{h(\bar{\mathcal{X}})} \right| = \frac{1}{\sqrt{n^d}} \left| \|\mathcal{D} - \mathcal{X}\|_F - \|\mathcal{D} - \bar{\mathcal{X}}\|_F \right| \leq \frac{\epsilon}{\sqrt{n^d}}$$

and

$$\left| \sqrt{\hat{h}(\mathcal{X})} - \sqrt{\hat{h}(\bar{\mathcal{X}})} \right| = \frac{1}{\sqrt{|\Omega|}} \left| \|\mathcal{P}_\Omega(\mathcal{D} - \mathcal{X})\|_F - \|\mathcal{P}_\Omega(\mathcal{D} - \bar{\mathcal{X}})\|_F \right| \leq \frac{\epsilon}{\sqrt{|\Omega|}}.$$

It follows that

$$\begin{aligned} &\sup_{\mathcal{X} \in \mathcal{S}} \left| \sqrt{\hat{h}(\mathcal{X})} - \sqrt{h(\mathcal{X})} \right| \\ &\leq \sup_{\mathcal{X} \in \mathcal{S}} \left| \sqrt{\hat{h}(\mathcal{X})} - \sqrt{\hat{h}(\bar{\mathcal{X}})} \right| + \left| \sqrt{\hat{h}(\bar{\mathcal{X}})} - \sqrt{h(\bar{\mathcal{X}})} \right| + \left| \sqrt{h(\bar{\mathcal{X}})} - \sqrt{h(\mathcal{X})} \right| \\ &\leq \frac{\epsilon}{\sqrt{|\Omega|}} + \sqrt{g(\Omega)} + \frac{\epsilon}{\sqrt{n^d}}. \end{aligned}$$

Using the definition of $g(\Omega)$ and letting $\epsilon = 3d$, we have

$$\begin{aligned} &\sup_{\mathcal{X} \in \mathcal{S}} \left| \sqrt{\hat{h}(\mathcal{X})} - \sqrt{h(\mathcal{X})} \right| \\ &\leq \frac{2\epsilon}{\sqrt{|\Omega|}} + \left(\frac{\varsigma^2}{2} (d \log n + \log |\mathcal{S}|) \left(\frac{1}{|\Omega|} - \frac{1}{n^d} + \frac{1}{n^d|\Omega|} \right) \right)^{1/4} \\ &\leq \frac{2\epsilon}{\sqrt{|\Omega|}} + 2\varepsilon \left(\frac{d \log n + \log |B|}{2|\Omega|} \right)^{1/4}. \end{aligned}$$

□

E.3 Proof of Lemma 3

Proof. Case (a) can be easily obtained by transforming the entropy number result of the special case $p = u = r = 2$ (we have exchanged p and q) of Theorem 13 in (Mayer & Ullrich, 2021) to covering number. Specifically, let

$$e_\eta \leq c_q \left(\frac{\log(ek/\eta) + n}{\eta} \right)^{1/q-1/2} \leq c_q \left(\frac{\log(ek) + n}{\eta} \right)^{1/q-1/2},$$

where c_q is a constant depending only on q and $c_q = O(1/q)$. It follows that

$$\eta \leq (n + \log(ek)) \left(\frac{c_q}{e_\eta} \right)^{2q/(2-q)}.$$

Then the covering number is bounded as

$$\begin{aligned} \log \mathcal{N}(\mathcal{B}_{2,q}^{n,k}, \|\cdot\|_F, \epsilon) &\leq (n + \log(ek)) \left(\frac{c_q}{\epsilon} \right)^{2q/(2-q)} \log 2 \\ &= c'_q (n + \log(ek)) \epsilon^{-2q/(2-q)}, \end{aligned}$$

where $c'_q = O(1/q)$.

Case (b) is a special case of Lemma 3.2 of (Bartlett et al., 2017). Namely, in the lemma, letting \mathbf{X} be an identity matrix and $p = q = 2$ and $\frac{1}{r} + \frac{1}{s} = 1$, renaming the variables, we have $\log \mathcal{N}(\mathcal{B}_{2,q}^{n,k}, \|\cdot\|_F, \epsilon) \leq \lceil nk^{2(q-1)/q} \epsilon^{-2} \rceil \log(2nk)$.

□

E.4 Proof of Lemma 6

Proof. Let $S = \{i : z_i > \tau\}$. We have

$$\|\mathbf{z}\|_1 = \|\mathbf{z}_S\|_1 + \sum_{i \notin S} z_i \leq \sqrt{|S|} \|\mathbf{z}\|_2 + \tau \sum_{i \notin S} \frac{z_i}{\tau}. \quad (52)$$

Since $\frac{z_i}{\tau} \leq 1$ and $0 < p \leq 1$, we have

$$\|\mathbf{z}\|_1 \leq \sqrt{|S|} \|\mathbf{z}\|_2 + \tau \sum_{i \notin S} \left(\frac{z_i}{\tau} \right)^p \leq \sqrt{|S|} \|\mathbf{z}\|_2 + \tau^{1-p} \|\mathbf{z}\|_p^p. \quad (53)$$

On the other hand, we have

$$|S| \tau^p \leq \sum_{i \in S} z_i^p \leq \|\mathbf{z}\|_p^p. \quad (54)$$

Combining (53) and (54), we arrive at

$$\|\mathbf{z}\|_1 \leq \tau^{-p/2} \|\mathbf{z}\|_2 \sqrt{\|\mathbf{z}\|_p^p} + \tau^{1-p} \|\mathbf{z}\|_p^p. \quad (55)$$

□

E.5 Proof of Lemma 7

This is a special case of Corollary 2 in Tomioka & Suzuki (2014).

E.6 Proof of Lemma 8

Proof. The Chernof bound of Gaussian distribution indicates

$$\mathbb{P}[|\mathcal{N}_{j_1 j_2 \dots j_d}| \geq t] \leq 2e^{-\frac{t^2}{2\sigma^2}}, \quad (j_1, j_2, \dots, j_d) \in [n] \times [n] \cdots \times [n]. \quad (56)$$

Using Boole's inequality (union bound), we obtain

$$\mathbb{P} \left[\bigcup_{(j_1, j_2, \dots, j_d)} |\mathcal{N}_{j_1 j_2 \dots j_d}| \geq t \right] \leq \sum_{(j_1, j_2, \dots, j_d)} \mathbb{P} [|\mathcal{N}_{j_1 j_2 \dots j_d}| \geq t]. \quad (57)$$

It follows that

$$\mathbb{P} [\|\mathcal{N}\|_\infty \leq t] \geq 1 - 2n^d 2e^{-\frac{t^2}{2\sigma^2}}. \quad (58)$$

Let $t = \sqrt{2\sigma^2 \log(n^{2d})}$. Then

$$\mathbb{P} \left[\|\mathcal{N}\|_\infty \leq 2\sigma \sqrt{d \log(n)} \right] \geq 1 - 2n^{-d}. \quad (59)$$

□