Matrix Completion using Kronecker Product Approximation

Chencheng Cai[a], Rong Chen[b] and Han Xiao[b*]

[a]Temple University, [b]Rutgers University

Abstract

A matrix completion problem is to recover the missing entries in a partially observed matrix. Most of the existing matrix completion methods assume a low rank structure of the underlying complete matrix. In this paper, we introduce an alternative and more general form of the underlying complete matrix, which assumes a low Kronecker rank instead of a low regular rank, but includes the latter as a special case. The extra flexibility allows for a much more parsimonious representation of the underlying matrix, but also raises the challenge of determining the proper Kronecker product configuration to be used. We find that the configuration can be identified using the mean squared error criterion as well as a modified cross-validation criterion. We establish the consistency of this procedure under suitable conditions on the signal-to-noise ratio. A aggregation procedure is also proposed to deal with special missing patterns and complex underlying structures. Both numerical and empirical studies are carried out to demonstrate the performance of the new method.

Keywords: Matrix completion, Kronecker product approximation, aggregated estimation, model selection, information ceriterion

*Chencheng Cai is a postdoctoral fellow at Department of Statistical Science, Fox School of Business, Temple University, Philadelphia, PA 19122. E-mail: chencheng.cai@temple.edu. Rong Chen is Professor, Department of Statistics, Rutgers University, Piscataway, NJ 08854. E-mail: rongchen@stat.rutgers.edu. Han Xiao is Associate Professor, Department of Statistics, Rutgers University, Piscataway, NJ 08854. E-mail: hxiao@stat.rutgers.edu. Han Xiao is the corresponding author. Chen's research is supported in part by National Science Foundation grants DMS-1503409, DMS-1737857, IIS-1741390, CCF-1934924, DMS-2027855, and DMS-2052949. Xiao's research is supported in part by National Science Foundation grants DMS-1454817, DMS-2027855, DMS-2052949, and a research grant from NEC Labs America.

1 Introduction

Many applications involve observations in a matrix form, and most likely of large dimensions. When the observed matrix is the sum of a signal matrix and a noise matrix, or is partially observed, a common approach in machine learning and statistics is to assume that the underlying signal matrix has a rank that is much smaller than its dimension. The low rank structure represents the interaction between matrix entries with a smaller number of parameters and reveals the core factors that drive and control the high dimensional observations, resulting in significant dimension reduction. Such a low rank assumption also makes it possible to recover the missing entries in a partially observed matrix, known as the *matrix completion problem*. Matrix completion has broad and important applications, including collaborative filtering (Goldberg et al., 1992), global positioning (Biswas et al., 2006) and remote sensing (Schmidt, 1986), among many others. One of the most famous examples is the Netflix recommendation system contest (Bennett et al., 2007), in which the winning algorithm recovers the movie-rating matrix by a rank one matrix based on the observed ratings.

Two different settings of matrix completion problems have been considered in the literature. One is the exact matrix completion problem whose goal is to recover the original matrix *exactly* when a portion of the matrix entries is missing. When the original matrix rank is known, it can be recovered through the alternating minimization algorithm proposed by Jain et al. (2013) under certain conditions. When the matrix rank is unknown, it is still possible to exactly recover the matrix through nuclear norm optimization (Candès and Recht, 2009; Candès and Tao, 2010). The nuclear norm optimization approach can also be applied to tensor completion problems whose goal is to recover a tensor structure (Yuan and Zhang, 2016). The second setting considers the circumstances when the observed entries are corrupted by noises while at the same time a portion of the entries is missing. It is known as the stable matrix completion problem. Candès and Plan (2010) extends the nuclear norm optimization approach to the stable matrix completion problem by relaxing the constraint. Assuming the matrix rank is known, Keshavan et al. (2010b) approaches the problem using a combination of spectral techniques and manifold optimization. Specifically for stable rank one matrix completion problem, Cosse and Demanet (2017) proposes to solve it using two rounds of semi-definite programming relaxation. Note that the alternating minimization algorithm in Jain et al. (2013) is applicable for the stable matrix completion problem as well.

It is observed that in many applications of image analysis, signal processing and quantum computing, the high dimensional data in matrix form often has a low-rank structure in terms of

Kronecker product decomposition instead of singular value decomposition (Werner et al., 2008; Duarte and Baraniuk, 2012; Kamm and Nagy, 1998). Approximating a matrix with a sum of a small number of matrices in Kronecker product form is an extension of the low rank approximation with a sum of rank one matrices. The flexibility provides an alternative approach for matrix completion. The key challenging factor of the approach is to determine the Kronecker product's configuration, which is the dimensions of the two matrices of the product.

In this article, we consider the matrix completion problem under the setting that the signal matrix is the sum of k Kronecker products with an unknown configuration. Although it is natural to think of using an information criterion to determine the configuration, we find strikingly that the criterion based on just the mean squared error can identify the configuration, and prove its asymptotic consistency under suitable conditions on the signal-to-noise ratio. We therefore propose a two-step procedure of the stable matrix completion problem, first identify the configuration, and then complete the matrix using the chosen configuration. We also would like to emphasize that the objective of our paper is on the completion of the matrix, and the Kronecker product structure is a vehicle for achieving it.

For a given configuration of the Kronecker product, if one block of the matrix is completely missing, then it is impossible to recover the entries in it. To alleviate this issue, we also introduce an aggregated procedure which incorporates a few different configurations, and combine the recovered matrices under each of them. This is based on an intuitive observation that entries non-recoverable under one configuration may be recovered under others, perhaps imperfectly but still usefully. The aggregation over multiple configurations can also potentially provide more robust and stable estimates of the underlying signal matrix in finite samples under the bias-variance trade-off framework. We propose an empirical procedure to implement the aggregation procedure, and use a modified cross-validation criterion to select the number of configurations to be used.

Compared to the conventional matrix completion based on the low rank structure, our approach brings two new advancements. First, it allows the user to search over configurations, and thus offers a greater flexibility. More strikingly, this is at no additional cost of the observing rate (compared to the state-of-art procedures), if the candidate set of the configurations is not too large (see Section 4.4 for details). On the other hand, a higher observing rate does allow the search of the optimal configuration over a larger candidate set. Second, the aggregation over multiple configurations can lead to more robust and stable matrix completion.

The rest of the paper is organized as follows. In Section 2, we formally introduce the matrix

completion problem, assuming the underlying complete matrix has a low Kronecker rank. The estimation and configuration determination procedures are presented in Section 3. The theoretical analysis on the consistency of the configuration selection and the error bound of recovered matrix are provided in Section 4. We explore the power of aggregation by incorporating different configurations in Section 5. Section 6 employs simulation studies and a real image example to demonstrate the performances of both the configuration selection and the matrix completion procedures. Section 7 concludes.

Notation: $[K] = \{0, 1, \ldots, K\}$ denotes the set of non-negative integers less than or equal to K. For a vector u, $\|u\|$ denotes its Euclidean norm. For a matrix $\boldsymbol{M} \in \mathbb{R}^{m\times n}$, $\|\boldsymbol{M}\|_F$ represents its Frobenius norm such that $\|\boldsymbol{M}\|_F = \sqrt{\mathrm{tr}(\boldsymbol{M}\boldsymbol{M}')}$ and $\|\boldsymbol{M}\|_S$ denotes its spectral norm defined by $\|\boldsymbol{M}\|_S = \arg\max_{u\in\mathbb{R}^n, \|u\|=1} \|\boldsymbol{M}u\|$. For positive integer P, we denote the set of all divisors of P by $d(P) = \{p \in \mathbb{Z}^+ : P \mod p = 0\}$.

2 Low K-rank Matrix Completion

2.1 K-rank and the matrix completion

Let $\boldsymbol{M} \in \mathbb{R}^{P\times Q}$ be a $P \times Q$ matrix, and suppose P and Q can be factorized as $P = pp^*$ and $Q = qq^*$. It is shown (Van Loan and Pitsianis, 1993) that $\boldsymbol{M}$ has a complete Kronecker product decomposition (KPD) in the form

$$\boldsymbol{X} = \sum_{i=1}^{(pq)\wedge(p^*q^*)} \lambda_i \boldsymbol{A}_i \otimes \boldsymbol{B}_i \tag{1}$$

where $\boldsymbol{A}_i \in \mathbb{R}^{p\times q}$ and $\boldsymbol{B}_i \in \mathbb{R}^{p^*\times q^*}$, and the Kronecker product of $\boldsymbol{A} \in \mathbb{R}^{p\times q}$ and $\boldsymbol{B} \in \mathbb{R}^{p^*\times q^*}$ is given as

$$\boldsymbol{A} \otimes \boldsymbol{B} := \begin{bmatrix} a_{1,1}\boldsymbol{B} & a_{1,2}\boldsymbol{B} & \ldots & a_{1,q}\boldsymbol{B} \\ a_{2,1}\boldsymbol{B} & a_{2,2}\boldsymbol{B} & \ldots & a_{2,q}\boldsymbol{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{p,1}\boldsymbol{B} & a_{p,2}\boldsymbol{B} & \ldots & a_{p,q}\boldsymbol{B} \end{bmatrix}, \tag{2}$$

where $a_{i,j}$ is the element of $\boldsymbol{A}$ in i-th row and j-th column. In (1) we require

$$\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_{(pq)\wedge(p^*q^*)} \geqslant 0,$$

and

$$\mathrm{tr}(\boldsymbol{A}_i\boldsymbol{A}_j^T) = \mathrm{tr}(\boldsymbol{B}_i\boldsymbol{B}_j^T) = \begin{cases} 1 & \text{when } i = j \\ 0 & \text{when } i \neq j \end{cases},$$

so that the matrices $\boldsymbol{A}_i$ and $\boldsymbol{B}_i$ are identified up to a sign change when λ_i are all distinct. The (p, q, p^*, q^*) or simply (p, q) (when P and Q are given) is called the *configuration* of the Kronecker product decomposition. When P and Q have multiple integer factors, there are multiple configurations and therefore a KPD in the form (1) exists for each configuration.

We propose to consider the matrix completion problem when the matrix $\boldsymbol{X} \in \mathbb{R}^{P\times Q}$, which we try to recover, has a low rank KPD such as

$$\boldsymbol{X} = \sum_{i=1}^{r} \lambda_i \boldsymbol{A}_i \otimes \boldsymbol{B}_i \tag{3}$$

with configuration (p_0, q_0) and rank $r \leq (p_0 q_0) \wedge (p_0^* q_0^*)$. We refer to r as the Kronecker rank (K-rank) of $\boldsymbol{X}$ with respect to the configuration (p_0, q_0). The relationship between K-rank and KPD is similar to, and an extension of the relationship between rank and singular value decomposition (SVD), which renders $\boldsymbol{X}$ in the form

$$\boldsymbol{X} = \sum_{i=1}^{r} \lambda_i u_i v_i^T, \tag{4}$$

a sum of r rank one matrices. It is seen that the SVD in (4) is the KPD in (3) with the configuration $(P, 1, 1, Q)$. On the other hand, a full rank matrix can have a low Kronecker rank. For example, the (mn)-dimensional identity matrix $\boldsymbol{I}_{mn}$ can be written as $\boldsymbol{I}_{mn} = \boldsymbol{I}_m \otimes \boldsymbol{I}_n$, and is therefore of Kronecker rank 1. In fact, there is a deeper and more intrinsic connection between (3) and (4), as will be explained in Section 2.2.

We assume that $\boldsymbol{Y}$ is a corrupted version of the underlying signal matrix $\boldsymbol{X}$ such that

$$\boldsymbol{Y} = \boldsymbol{X} + \frac{\sigma}{\sqrt{PQ}} \boldsymbol{E}, \tag{5}$$

where $\boldsymbol{E}$ is a $P \times Q$ matrix with IID standard Gaussian entries and σ denotes the noise level.

In the matrix completion problem, the matrix $\boldsymbol{Y}$ is only partially observed. In this paper, we consider the missing completely at random scheme by assuming that the observed matrix $\boldsymbol{Y}^*$ is generated as

$$[\boldsymbol{Y}^*]_{ij} = [\boldsymbol{Y}]_{ij} \delta_{ij}$$

where δ_{ij} are independent and identically distributed $\sim$ Bernoulli(τ), which are also independent with $\boldsymbol{X}$ and $\boldsymbol{E}$. The rate τ is called the *observing rate*. We also introduce another notation of the observed matrix through a projection operator. Let Ω be the set of indices of the observed entries

in $\boldsymbol{Y}^*$ and define the projection $P_\Omega(\cdot)$ as

$$[P_\Omega(\boldsymbol{M})]_{ij} = \begin{cases} [\boldsymbol{M}]_{ij} & \text{if } (i,j) \in \Omega, \\ 0 & \text{otherwise,} \end{cases}$$

for any $P \times Q$ matrix M. In particular, it follows that $\boldsymbol{Y}^* = P_\Omega(\boldsymbol{Y})$, where unobserved entries are filled with zeros by default.

2.2 Connection to low rank matrix completion

The prevailing matrix completion algorithms (Candès and Recht, 2009; Candès and Plan, 2010; Jain et al., 2013) assume a low rank structure of the signal matrix $\boldsymbol{X}$ as in (4). It corresponds to a special low K-rank form of $\boldsymbol{X}$ in (3), with the configuration $p = P, q = 1$ and $p^* = 1, q^* = Q$. In this section, we demonstrate the idea of Van Loan and Pitsianis (1993) and Cai et al. (2019b) that with a *known* configuration, the matrix completion problem with KPD in (3) can be converted to a standard low rank matrix completion problem and be solved using the existing approaches.

As discussed in Van Loan and Pitsianis (1993) and Cai et al. (2019b), the Kronecker product of two matrices and the outer product of their vectorized version are linked through a rearrangement operation. It can be seen from (2) that all the entries in $\boldsymbol{A} \otimes \boldsymbol{B}$ have the form $a_{i,j}b_{k,\ell}$ and all the entries in $\text{vec}(\boldsymbol{A})\text{vec}(\boldsymbol{B})^T$ have the same form $a_{i,j}b_{k,\ell}$, where $\text{vec}(\cdot)$ is the vectorization operation that flattens a matrix to a column vector. Hence $\boldsymbol{A} \otimes \boldsymbol{B}$ and $\text{vec}(\boldsymbol{A})\text{vec}(\boldsymbol{B})^T$ have exactly the same entries in the matrices, except the arrangement – one is a $pp^* \times qq^*$ matrix and the other is a $pq \times p^*q^*$ matrix.

We define a rearrangement operator $\mathcal{R}_{p,q}[\cdot]$ so that $\mathcal{R}[\boldsymbol{A} \otimes \boldsymbol{B}] = \text{vec}(\boldsymbol{A})\text{vec}(\boldsymbol{B})^T$. Specifically, assuming $\boldsymbol{M}$ is a $pp^* \times qq^*$ matrix, a rearrangement operation $\mathcal{R}_{p,q}[\cdot]$ is defined as

$$\mathcal{R}_{p,q}[\boldsymbol{M}] = [\text{vec}(\boldsymbol{M}_{1,1}), \ldots, \text{vec}(\boldsymbol{M}_{p,q})]^T,$$

where $\boldsymbol{M}_{i,j}$ is the (i,j)-th block of $\boldsymbol{M}$ of size $p^* \times q^*$. Then for any $p \times q$ matrix $\boldsymbol{A}$ and $p^* \times q^*$ matrix $\boldsymbol{B}$, we have

$$\mathcal{R}_{p,q}[\boldsymbol{A} \otimes \boldsymbol{B}] = \text{vec}(\boldsymbol{A})[\text{vec}(\boldsymbol{B})]^T.$$

That is, if the configuration of the rearrangement operator is the same as that of the operand, the rearrangement operator $\mathcal{R}_{p,q}$ turns a Kronecker product into a rank-one matrix. Note that the rearrangement operation is linear, in that $\mathcal{R}_{p,q}(c_1\boldsymbol{M}_1 + c_2\boldsymbol{M}_2) = c_1\mathcal{R}_{p,q}(\boldsymbol{M}_1) + c_2\mathcal{R}_{p,q}(\boldsymbol{M}_2)$ and the inverse operator $\mathcal{R}_{p,q}^{-1}$ exist so that $\mathcal{R}_{p,q}^{-1}(\mathcal{R}_{p,q}(\boldsymbol{M})) = \boldsymbol{M}$. The rearrangement operator also retains

the Frobenius norm in that $\|\mathcal{R}_{p,q}(\boldsymbol{M})\|_F = \|\boldsymbol{M}\|_F$ since both matrices contain the same set of elements. As a result, the KPD (3) becomes a SVD after the rearrangement, and many problems pertaining to KPD can therefore be solved under the realm of SVD.

If we apply $\mathcal{R}_{p_0,q_0}$ to the corrupted matrix $\boldsymbol{Y}$ in (5), assuming that $\boldsymbol{X}$ has a low K-rank $r \leq (p_0 q_0) \vee (p_0^* q_0^*)$, it returns a corrupted low rank matrix

$$\mathcal{R}_{p_0,q_0}[\boldsymbol{Y}] = \sum_{i=1}^{r} \lambda_i \text{vec}(\boldsymbol{A}_i)[\text{vec}(\boldsymbol{B}_i)]^T + \frac{\sigma}{\sqrt{PQ}} \mathcal{R}_{p_0,q_0}[\boldsymbol{E}], \tag{6}$$

where $\mathcal{R}_{p_0,q_0}[\boldsymbol{E}]$, the rearranged noise matrix, is still a matrix of IID Gaussian entries. Note that $\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}] = P_{\bar{\Omega}_{p_0,q_0}} \mathcal{R}_{p_0,q_0}[\boldsymbol{Y}]$ where $\bar{\Omega}_{p_0,q_0}$ denotes the indices of the observed entries after rearrangement. The K-rank-r matrix completion problem based on $P_\Omega \boldsymbol{Y}$ is therefore equivalent to the standard rank-r matrix completion problem based on $P_{\bar{\Omega}_{p_0,q_0}} \mathcal{R}_{p_0,q_0}[\boldsymbol{Y}]$.

The preceding discussion assumes that the true configuration (p_0, q_0) is known so that the rearrangement operator in (6) is determined. However, observing only the matrix $P_\Omega \boldsymbol{Y}$ and its dimension (P, Q) does not reveal the true configuration (p_0, q_0) — any configuration (p, q) satisfying $p \in d(P)$ and $q \in d(Q)$ might be the true configuration, where $d(P)$ is the set of factors (divisors) of P. If the rearrangement operator $\mathcal{R}$ is erroneously configured, the low rank structure in (6) might be violated since $\mathcal{R}_{p,q}[\boldsymbol{A} \otimes \boldsymbol{B}]$ typically does not preserve the low rank structure if $(p, q) \neq (p_0, q_0)$.

3 Methodology

In this section, we propose a two-step procedure to solve the matrix completion problem under the assumption that the signal matrix $\boldsymbol{X}$ is of low K-rank under a proper configuration. In the first step, we determine the configuration by maximizing a criterion function over all candidate configurations. Specifically, we use the spectral norm of the rearranged version of $P_\Omega[\boldsymbol{Y}]$ as the criterion function. In the second step, we convert the observed matrix to a standard matrix completion problem as in (6) with respect to the estimated configuration from first step. The rearranged partially observed matrix is then completed by an alternating least squares (ALS) algorithm.

3.1 Configuration Determination

Recall that (P, Q) is the dimension of $\boldsymbol{Y}^* = P_\Omega \boldsymbol{Y}$. We introduce the following candidate set

$$\mathcal{C}_\delta = \{(p, q) : p \in d(P), q \in d(Q), (PQ)^{1/4+\delta} \leqslant pq \leqslant (PQ)^{3/4-\delta}\}, \tag{7}$$

where $0 < \delta < 1/4$ is a positive constant. Again, $d(P)$ is the set of all factors of P. For each configuration (p, q) in the set $\mathcal{C}_\delta$, the rearranged $\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]$ has an aspect ratio between $(PQ)^{-1/2+2\delta}$ and $(PQ)^{1/2-2\delta}$. Here, extreme aspect ratios are excluded for several reasons. First, the extreme "corner" cases where pq is extremely small or large are of less interest. For example $p = q = 1$ and $p = P, q = Q$ represent the case of a product of scalar and the full matrix. When $p = q = 2$, the signal matrix $\boldsymbol{X}$ is assumed to be a block matrix formed by two-by-two matrices (total $PQ/4$ of them) of the same proportion, while when $p = P/2, q = Q/2$, the matrix X is a two-by-two block matrix of size $(P/2) \times (Q/2)$. The model complexity is high and the dimension reduction in these cases is limited, similar to using a very small bandwidth in kernel smoothing operations. Second, from the view of model selection using information criteria, we can view δ as a model complexity penalty parameter with the penalty function $pen(p, q) = 0$ if $(p, q) \in \mathcal{C}_\delta$ and $pen(p, q) = \infty$ otherwise. Third, we expect the partially observed matrix $\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}]$ to be recoverable under the true configuration while an extreme aspect ratio usually results in a non-negligible probability of missing an entire row or column, in which case these missing items are not recoverable. The recoverability issue will be discussed in detail later. Lastly, the aspect ratio is a crucial factor in quantifying the error of standard low rank matrix completion (Keshavan et al., 2010a; Jain et al., 2013; Gunasekar et al., 2013). Theoretical performance guarantee exists only when the aspect ratio is in a reasonable range.

We propose to determine the configuration through the following maximization problem,

$$(\hat{p}, \hat{q}) = \underset{(p,q)\in\mathcal{C}_\delta}{\arg\max} \ \|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]\|_S. \tag{8}$$

Recall that in $P_\Omega \boldsymbol{Y}$, unobserved entries are filled in with zeros.

We now briefly discuss the heuristics behind the procedure (8). The rearranged partially observed matrix $\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]$ can be viewed as a partially observed version of the rearranged matrix $P_{\bar{\Omega}_{p,q}} \mathcal{R}_{p,q}[\boldsymbol{Y}]$. Its spectral norm is therefore

$$\|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]\|_S = \|P_{\bar{\Omega}_{p,q}} \mathcal{R}_{p,q}[\boldsymbol{Y}]\|_S \approx \tau \|\mathcal{R}_{p,q}[\boldsymbol{X}]\|_S. \tag{9}$$

As discussed in Cai et al. (2019b), under mild conditions, $\|\mathcal{R}_{p,q}[\boldsymbol{X}]\|_S$ attains its maximum over $\mathcal{C}_\delta$ at the true configuration. The error in the approximation (9) comes from two sources: the noise matrix $\boldsymbol{E}$ and the missing entries. Both can be controlled when we operate within the candidate set $\mathcal{C}_\delta$ for $\delta > 0$. The analysis is presented in Section 4.

3.2 Estimation

Once the configuration is determined, the Kronecker matrix completion problem under the configuration $(\hat{p}, \hat{q})$ becomes a standard low rank matrix completion problem on the rearranged matrix whose rank is the same as the K-rank of the original matrix. Existing procedures such as that proposed by Ashraphijuo et al. (2017) for identifying the rank are readily applicable. With a fixed rank r, it becomes the following optimization problem.

$$\min_{\lambda_i, \boldsymbol{A}_i, \boldsymbol{B}_i} \left\| P_\Omega \left(\sum_{i=1}^{r} \lambda_i \boldsymbol{A}_i \otimes \boldsymbol{B}_i \right) - P_\Omega \boldsymbol{Y} \right\|_F . \tag{10}$$

where $\boldsymbol{A}_i \in \mathbb{R}^{\hat{p} \times \hat{q}}$, $\boldsymbol{B}_i \in \mathbb{R}^{\hat{p}^* \times \hat{q}^*}$ and $\|\boldsymbol{A}_i\|_F = \|\boldsymbol{B}_i\|_F = 1$ with $\hat{p}^* = P/\hat{p}$ and $\hat{q}^* = Q/\hat{q}$. Since the rearrangement operation preserves the Frobenius norm, the optimization problem (10) is equivalent to the classical rank-r matrix completion problem

$$\min_{\lambda_i, u_i, v_i} \left\| P_{\bar{\Omega}_{\hat{p},\hat{q}}} \left(\sum_{i=1}^{r} \lambda_i u_i v_i^T - \mathcal{R}_{\hat{p},\hat{q}}[\boldsymbol{Y}] \right) \right\|_F . \tag{11}$$

where $u_i = \mathrm{vec}(\boldsymbol{A}_i)$, $v_i = \mathrm{vec}(\boldsymbol{B}_i)$, and $\bar{\Omega}_{\hat{p},\hat{q}}$ records the indices of observed entries after the rearrangement. To solve the optimization in (11), we adopt the alternating least squares (ALS) algorithm proposed by Jain et al. (2013), where the initial values for u_i and v_i are directly estimated from the singular value decomposition of $P_{\bar{\Omega}_{\hat{p},\hat{q}}} \mathcal{R}_{\hat{p},\hat{q}}[\boldsymbol{Y}]$ as in Keshavan et al. (2010b). The algorithm is depicted in Algorithm 1. The recovered matrix is therefore $\hat{\boldsymbol{X}} = \mathcal{R}^{-1}_{\hat{p},\hat{q}}[\sum_{i=1}^{r} \hat{\lambda}_i \hat{u}_i \hat{v}_i^T]$, where $\mathcal{R}^{-1}_{\hat{p},\hat{q}}$ is the inverse operator of the rearrangement and $\hat{\lambda}_i$, $\hat{u}_i$ and $\hat{v}_i$ are the optimal solution of (11).

We note that there are other algorithms available for solving (11) besides the ALS (Algorithm 1). For example, the OptSpace algorithm proposed by Keshavan et al. (2010b) can be adopted as well to solve (11) once the configuration is determined. In this paper we focus on and analyze the ALS algorithm. We will also compare the ALS and the OptSpace in the numerical studies.

4 Theoretical Analysis

In this section, we present the theoretical properties of our methodology when the signal $\boldsymbol{X}$ is of K-rank 1. That is,

$$\boldsymbol{Y} = \lambda \boldsymbol{A} \otimes \boldsymbol{B} + \frac{\sigma}{\sqrt{PQ}} \boldsymbol{E}, \tag{12}$$

where $\boldsymbol{A} \in \mathbb{R}^{p_0 \times, q_0}$, $\boldsymbol{B} \in \mathbb{R}^{p_0^* \times q_0^*}$, $\|\boldsymbol{A}\|_F = \|\boldsymbol{B}\|_F = 1$ and $\boldsymbol{E}$ is a Gaussian ensemble matrix with IID standard normal entries. Here we normalize the error matrix by $(PQ)^{-1/2}$ such that the

Algorithm 1 Matrix Completion with Fixed Configuration

Input: matrix $P_\Omega \boldsymbol{Y}$, configuration (p, q).

Select rank r.

Let $\tilde{\boldsymbol{Y}} = P_{\bar{\Omega}_{p,q}} \mathcal{R}_{p,q}[\boldsymbol{Y}]$.

Initialize $\boldsymbol{\Lambda}^{(0)}$, $\boldsymbol{U}^{(0)}$, $\boldsymbol{V}^{(0)}$ such that $\boldsymbol{U}^{(0)}\boldsymbol{\Lambda}^{(0)}[\boldsymbol{V}^{(0)}]^T$ is the leading rank r SVD of $\tilde{\boldsymbol{Y}}$.

repeat

  $\boldsymbol{U}^* = \arg\min_{\boldsymbol{U}} \ \|P_{\bar{\Omega}_{p,q}}[\tilde{\boldsymbol{Y}} - \boldsymbol{U}[\boldsymbol{V}^{(k)}]^T]\|_F$

  $\boldsymbol{V}^* = \arg\min_{\boldsymbol{V}} \ \|P_{\bar{\Omega}_{p,q}}[\tilde{\boldsymbol{Y}} - \boldsymbol{U}^*[\boldsymbol{V}]^T]\|_F$

  Update $\boldsymbol{\Lambda}^{(k+1)}$, $\boldsymbol{U}^{(k+1)}$, $\boldsymbol{V}^{(k+1)}$ such that $\boldsymbol{U}^*[\boldsymbol{V}^*]^T = \boldsymbol{U}^{(k+1)}\boldsymbol{\Lambda}^{(k+1)}[\boldsymbol{V}^{(k+1)}]^T$ is in standard SVD form.

until convergence

Return $\hat{\boldsymbol{X}} = \mathcal{R}_{p,q}^{-1}[\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{V}^T]$ where $\boldsymbol{\Lambda}, \boldsymbol{U}, \boldsymbol{V}$ are from the last iteration.

signal-to-noise ratio of the fully observed $\boldsymbol{Y}$ is λ^2/σ^2. The analysis in this section can be extended to $r > 1$ cases without substantial differences. The presentation for $r = 1$ case is cleaner and easier to understand.

4.1 Assumptions

In the subsequent analysis, we consider the following asymptotic scheme.

Assumption 1 (**True configuration**). *Assume the dimension of $\boldsymbol{Y}$ goes to infinity, i.e. $PQ \to \infty$, and there is an absolute constant $0 < \delta \leqslant 1/4$ such that the true configuration (p_0, q_0) satisfies*

$$(PQ)^{1/4+\delta} \leqslant p_0 q_0 \leqslant (PQ)^{3/4-\delta}.$$

The requirement on the true configuration in Assumption 1 can be equivalently stated as $(p_0, q_0) \in \mathcal{C}_\delta$. Note that we do not require both $P \to \infty$ and $Q \to \infty$. We only require that jointly $PQ \to \infty$. In addition, $\delta > 0$ implies that the dimensions $p_0 q_0 \to \infty$ and $p_0^* q_0^* \to \infty$, or for both component matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, at least one of its dimensions goes to infinity. In the meantime, as discussed in Section 3.1, the aspect ratio of $\mathcal{R}_{p_0,q_0}[\boldsymbol{Y}]$ is restricted between $(PQ)^{-1/2+2\delta}$ and $(PQ)^{1/2-2\delta}$.

Definition 1 (Incoherence Coefficient). *The incoherence parameter μ of a $m \times n$ matrix M is defined by*

$$\mu = \frac{\sqrt{mn}\|\boldsymbol{M}\|_{\max}}{\|\boldsymbol{M}\|_F},$$

where $\|\cdot\|_{\max}$ *denotes the maximum absolute entry of a matrix.*

Assumption 2 (**Incoherence condition).** *Let* μ *be the maximum of the incoherence parameters of matrices* $\boldsymbol{A}$ *and* $\boldsymbol{B}$ *as defined in Definition 1. We assume as* $PQ \to \infty$

$$\lim_{PQ\to\infty} \frac{\mu^2}{(PQ)^\delta} = 0.$$

The incoherence condition is commonly imposed in theoretical investigations of matrix completion problems (Candès and Recht, 2009; Candès and Plan, 2010; Keshavan et al., 2010b). More specifically, Assumption 2 translates into an incoherence condition on $\mathrm{vec}(\boldsymbol{A})$ and $\mathrm{vec}(\boldsymbol{B})$, the singular vectors of the rearranged matrix, which is usually assumed for the standard low rank matrix completion problem on $\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}]$.

The incoherence parameter μ is greater or equal to 1. One would expect μ to increase slowly as $PQ \to \infty$. In the case when a matrix is generated by normalizing a $p \times q$ matrix of IID standard Gaussian entries, its incoherence parameter would be $O_p(\sqrt{\log pq})$ as $pq \to \infty$. Therefore, if both $\boldsymbol{A}$ and $\boldsymbol{B}$ are generated from random Gaussian matrices (see for example, the random scheme introduced in Cai et al. (2019b)), we have $\mu = O_p(\sqrt{\log(p_0 q_0 \vee p_0^* q_0^*)}) = O_p(\sqrt{(3/4-\delta)\log PQ})$ as $PQ \to \infty$.

Assumption 3 (**Observing rate).** *Assume the observing rate* τ *satisfies*

$$\lim_{PQ\to\infty} \frac{\tau(PQ)^{2\delta}}{(\mu^4 + \log PQ)} \to \infty.$$

Assumption 3 specifies the minimum allowable observing rate τ. Under Assumption 2, the observing rate τ is allowed to go to zero, with a decay rate depends on δ and μ.

Following Cai et al. (2019b), we define

$$\phi = \max_{(p,q)\in\mathcal{W}_\delta} \|\mathcal{R}_{p,q}[\boldsymbol{A}\otimes\boldsymbol{B}]\|_S,$$

where $\mathcal{W}_\delta := \mathcal{C}_\delta \setminus \{(p_0, q_0)\}$ is set of incorrect configurations in the candidate set. Since under the true configuration, $\|\mathcal{R}_{p_0,q_0}[\boldsymbol{A}\otimes\boldsymbol{B}]\|_S = \|\mathrm{vec}(\boldsymbol{A})\mathrm{vec}(\boldsymbol{B})^T\|_S = \|\mathrm{vec}(\boldsymbol{A})\|_F \|\mathrm{vec}(\boldsymbol{B})^T\|_F = 1$, ϕ^2 can be viewed as the maximum portion (in Frobenius norm) of the matrix $\boldsymbol{A}\otimes\boldsymbol{B}$ that can be represented by a Kronecker product of a wrong configuration. On the other hand, for any given configuration $(p,q) \in \mathcal{W}_\delta$, we always have $\|\mathcal{R}_{p,q}[\boldsymbol{A}\otimes\boldsymbol{B}]\|_S \leqslant \|\mathcal{R}_{p,q}[\boldsymbol{A}\otimes\boldsymbol{B}]\|_F = \|\boldsymbol{A}\otimes\boldsymbol{B}\|_F = 1$, resulting in $\phi \leqslant 1$. The *representation gap* ψ^2 is therefore defined as

$$\psi^2 := 1 - \phi^2. \tag{13}$$

The representation gap ψ^2, taking value on $[0, 1]$, quantifies the gap between any wrong configuration and the true configuration in the sense of representing the signal part with a K-rank 1 matrix. If $\psi^2 = 0$, there exists another configuration $(p', q') \neq (p_0, q_0)$ such that $\boldsymbol{A} \otimes \boldsymbol{B} = \boldsymbol{A}' \otimes \boldsymbol{B}'$ for some $\boldsymbol{A}' \in \mathbb{R}^{p' \times q'}$ and $\boldsymbol{B}' \in \mathbb{R}^{P/p' \times Q/q'}$, hence our model cannot be uniquely determined. A larger value of the representation gap ψ^2 makes it easier to separate the true configuration from the wrong ones. As discussed in Cai et al. (2019b), this gap is similar to the smallest eigenvalue in determining matrix ranks, or the smallest standardized non-zero coefficient in variable selection in regression analysis.

Assumption 4 (SNR and representation gap)**.** *Assume ψ^2 and λ/σ satisfy*

$$\lim_{PQ\to\infty} \frac{\tau\psi^4(PQ)^{2\delta}}{\mu^4} \to \infty, \quad \textit{and} \quad \lim_{PQ\to\infty} \frac{\tau(\lambda/\sigma)^2\psi^4(PQ)^{2\delta}}{\log PQ} \to \infty.$$

Assumption 4 gives the lower bounds on the representation gap ψ^2 and the signal-to-noise ratio λ/σ, in terms of the incoherence parameter μ and the observing rate τ. Comparing with Assumption 3 on the observing rate τ, we see that constant values of λ/σ and ψ^2 suffice for Assumption 4, though they may shrink to 0 with a rate depending on δ and τ, as $PQ \to \infty$.

4.2 Properties of the estimated configuration

To ease the notation, we denote the criterion function by

$$G(p, q) = \|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]\|_S,$$

for $(p, q) \in \mathcal{C}_\delta$. Recall that we use this function as the configuration selection criterion in (8). Theorem 1 provides the gap in the expectation of $G(p, q)$ between the true configuration and any wrong configurations.

Theorem 1 (**Gap in criterion function**)**.** *Under Assumptions 1 to 4, we have*

$$\mathbb{E}[G(p_0, q_0)] - \max_{(p,q)\in\mathcal{W}_\delta} \mathbb{E}[G(p, q)] \geqslant \tau\lambda(1 - \phi)(1 + o(1)),$$

where the expectation is taken on both random observations of entries and the random noise.

Since $\mathbb{E}[G(p_0, q_0)] \approx \tau\lambda$, Theorem 1 shows that the representation gap of the signal part is preserved using the criterion function $G(p, q)$. The proof of the theorem can be found in the Appendix.

On the face of it, Theorem 1 is similar to the one for the Kronecker product model introduced in Cai et al. (2019b). However, the underlying matrix completion problem makes it substantively different, in the form and even more so in the proof. We mention two major differences here. First of all, in the matrix completion problem, the candidate set $\mathcal{C}_\delta$ has to be smaller than the one considered in Cai et al. (2019b), where in the latter case, only two extreme configurations $(0,0),(P,Q)$ are excluded from $\mathcal{C}$. The model complexity is controlled by excluding extreme configurations. From the perspective of the aspect ratio of $\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]$, the aspect ratio is often controlled as a constant as matrix size increases to infinity in existing literature on conventional matrix completion problems (Candès and Recht, 2009; Candès and Plan, 2010; Jain et al., 2013). Within the configuration set $\mathcal{C}_\delta$, the aspect ratio is allowed to shrink to 0 or increases to infinity with a controlled rate between $(PQ)^{-1/2+2\delta}$ and $(PQ)^{1/2-2\delta}$. Configurations that result in the aspect ratios of $R_{p,q}[P_\Omega \boldsymbol{Y}]$ outside this range are excluded. Second, our assumption on the signal-to-noise ratio and representation gap (Assumption 4) is stronger than its counterpart in Cai et al. (2019b). The reason is that the missingness contributes more to the fluctuation of $G(p,q)$ than the noise matrix $\boldsymbol{E}$ itself (see equation (24) in the proof).

Theorem 2 (Consistency). *Let* $(\hat{p},\hat{q}) = \arg\max_{(p,q)\in\mathcal{C}_\delta} G(p,q)$. *Under Assumptions 1 to 4, we have*

$$P[(\hat{p},\hat{q}) = (p_0,q_0)] \geqslant 1 - C\frac{|d(P)||d(Q)|}{(PQ)^{3/4+3\delta}},$$

for some constant C, *where* $|d(P)|$ *is the size of the set* $d(P)$.

Remark: Theorem 2 is established for the model with K-rank 1. If the K-rank is higher than 1, the same procedure (8) can still select the true configuration consistently, under a larger signal-to-noise, and an additional assumption on the relative strengths of the terms in (3). On the other hand, when the K-rank is greater than 1, it is also of interest to consider the joint selection of the configuration and the K-rank. Since it exceeds the scope of this paper, we leave it to the future work.

Remark: According to the number theory (Hardy et al., 1979), for any $\epsilon > 0$, $|d(n)| = o(n^\epsilon)$. Therefore the probability of consistency in Theorem 2 converges to 1 as $PQ \to \infty$. We note that the convergence rate, as presented in Theorem 2, is only related to PQ and δ but not the signal-to-noise ratio λ/σ or the representation gap ψ^2, different from that for configuration determination for matrix denoising based on a fully observed matrix (Cai et al., 2019b). Again, this is because of the smaller candidate configuration set controlled by δ. In fact, according to the proof of Theorem 2

shown in Appendix, there is a term $\exp\{-C(\lambda/\sigma)^2\psi^4 PQ\}$ in the lower bound of the probability of correct determination. However, this term is dominated by the main term.

4.3 Property of the recovered matrix

After the configuration has been determined by maximizing the criterion function $G(p, q)$, the second step is to recover the original matrix using Algorithm 1. It is known that a completely missed row or column cannot be recovered by a low rank matrix completion. There is a similar phenomenon under the model (3). For example, if the true configuration (p_0, q_0) is used in Algorithm 1, and if a block of $\boldsymbol{Y}$ is completely missing, then the corresponding element of $\boldsymbol{A}_i$ is not recoverable, and neither is the whole missing block. The top row of Figure 1 illustrates the situation: the first block (highlighted in black in the upper left panel) is completely missing, and becomes a completely missed row (in the upper right panel) after rearrangement. Therefore, the first element of $\boldsymbol{A}_i$ in (3), or the first entry of u_i in (4) is not recoverable. We provide a formal characterization of irrecoverable entries in Definition 2.

Definition 2 (**Irrecoverable entries).** *The entry $\boldsymbol{X}_{i,j}$ is irrecoverable with respect to configuration (p, q) under the observation set Ω if either the entire row or the entire column of $\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]$ that contains $\boldsymbol{X}_{i,j}$ is completely missing.*

Equivalently, using the notation of index set Ω, the entry $\boldsymbol{X}_{i,j}$ is irrecoverable if either

$$\left\{(i', j') \in \Omega : \left\lfloor \frac{i-1}{p^*} \right\rfloor = \left\lfloor \frac{i'-1}{p^*} \right\rfloor, \ \left\lfloor \frac{j-1}{q^*} \right\rfloor = \left\lfloor \frac{j'-1}{q^*} \right\rfloor \right\} = \emptyset, \tag{14}$$

or

$$\left\{(i', j') \in \Omega : i' = i + lp^*, \ j' = j + mq^* \text{ for some } l, m \in \mathbb{Z}\right\} = \emptyset, \tag{15}$$

where $\lfloor x \rfloor$ denotes the largest integer that is no greater than x.

Remark: In Definition 2, (14) means the row in $\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]$ that contains $\boldsymbol{X}_{i,j}$ is completely missing, and (15) implies that the entire column in $\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]$ that contains $\boldsymbol{X}_{i,j}$ is missing.

Lemma 1 (**Recoverability).** *Under Assumptions 1 and 3, as $PQ \to \infty$,*

$$P[\text{all entries in } P_\Omega \boldsymbol{Y} \text{ are recoverable under } (p_0, q_0)] \geqslant 1 - 2\exp\{-C(PQ)^{1/4-\delta} \log PQ\},$$

for some constant C.

Lemma 1 shows that as long as $(p_0, q_0) \in \mathcal{C}_\delta$, with high probability, there is no missing columns or missing rows in the rearranged matrix $\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}]$ such that the optimization in (11) is well defined.

Theorem 3 (**Recovery error**). *Let $\hat{\boldsymbol{X}} = \hat{\lambda}\hat{\boldsymbol{A}} \otimes \hat{\boldsymbol{B}}$ be the recovered matrix of optimization* (11) *under the true configuration (p_0, q_0) using Algorithm 1. Then under Assumption 2, with high probability we have*

$$\|\hat{\boldsymbol{X}} - \boldsymbol{X}\|_F \leqslant \lambda\sqrt{PQ}\exp\left\{-\frac{C_1\tau(p_0q_0 \wedge p_0^*q_0^*)}{\mu^4 \log PQ}\right\} + C_2\frac{\sigma\mu}{\sqrt{\tau}}\sqrt{\log PQ}\cdot(PQ)^{-1/4}(p_0q_0 \wedge p_0^*q_0^*)^{-1}$$

for some global constant $C_1, C_2 > 0$.

Theorem 3 gives the error bound for the recovered matrix under the true configurations. Since in Assumption 1 $p_0q_0 \wedge p_0^*q_0^* \geqslant (PQ)^{1/4+\delta}$, the error bound in Theorem 3 can be revised to

$$\|\hat{\boldsymbol{X}} - \boldsymbol{X}\|_F \leqslant \lambda\exp\left\{-\frac{C_3\tau(PQ)^{1/4+\delta}}{\mu^4 \log PQ}\right\} + C_2\frac{\sigma\mu}{\sqrt{\tau}}\sqrt{\log PQ}\cdot(PQ)^{-1/2-\delta}. \tag{16}$$

Under Assumption 3, one can verify both terms on the right hand side converges to 0. The second term in (16) dominates. Using $\tau \gg \mu^4(PQ)^{-2\delta}$ from Assumption 3 and $\mu = o((PQ)^{\delta/2}$ from Assumption 2, one can further have

$$\|\hat{\boldsymbol{X}} - \boldsymbol{X}\|_F = o_p\left(\sigma\sqrt{\log PQ}(PQ)^{-1/2}\right).$$

Remark. Theorem 3 is adapted from Theorem 2.5 of Jain et al. (2013) on the ALS Algorithm for solving the optimization (11). If one chooses to use the OptSpace algorithm, similar recovery error bounds can inherit from Theorem 1.2 of Keshavan et al. (2010b) without additional assumptions. Furthermore, this paper focuses on the noisy matrix completion, and so does Theorem 3. Results for the exact matrix recovery under the noiseless setting can also be obtained, and the key is to identify the configuration correctly first. We do not include these results which are outside the scope of this paper. We refer the readers to Jain et al. (2013) for more on the exact recovery based on the low rank structure.

4.4 A discussion on the observing rate

The observing rate is critical for the task of matrix completion. A natural question to ask is: since the Kronecker product decomposition based matrix completion extends the conventional one based on the low rank structure, what is the implication on the observing rate? In this section we elaborate

on the Assumption 3 and compare it with those required by the state-of-art methods: the nuclear norm minimization algorithm of Candès and Plan (2010), the spectral matrix completion algorithm of Keshavan et al. (2010a) and the OptSpace algorithm of Keshavan et al. (2010b).

Assumption 3 can be rewritten as a lower bound of the observing rate

$$\tau \gg \log^2 PQ \ (PQ)^{-2\delta}, \tag{17}$$

where we assume $\mu < C\sqrt{\log PQ}$ with high probability for some constant C, which holds if $\boldsymbol{A}$ and $\boldsymbol{B}$ are random matrices with IID standard Gaussian entries normalized to have Frobenius norm 1. In fact, (17) can be viewed as the worst case requirement in the sense that

$$\log^2 PQ \ (PQ)^{-2\delta} = \max_{(p,q)\in\mathcal{C}_\delta, pq\leqslant(PQ)^{1/2}} \log^2 PQ \ \sqrt{\frac{p^*q^*}{(pq)^3}}, \tag{18}$$

where the maximum is attained on the boundary of $\mathcal{C}_\delta$.

Recall from Section 2.2 that although the observed low K-rank matrix $\boldsymbol{Y}$ is $P \times Q$, the dimension of its corresponding rearranged low rank matrix, $\mathcal{R}_{p_0,q_0}[\boldsymbol{Y}]$, is $p_0q_0 \times p_0^*q_0^*$. When comparing our conditions to existing ones on the completion of a $m \times n$ ($m \geq n$) low rank matrix $\boldsymbol{M}$, one should correspond (m, n) to $(p_0q_0, p_0^*q_0^*)$ instead of (P, Q). Under such a correspondence, our requirement on the observing rate is actually $\tau \gg \log^2 m \ \sqrt{m/n^3}$ as shown in (18). Again, Assumption 3 and Eq. (17) guarantees that this condition is fulfilled for any configuration in the candidate set $\mathcal{C}_\delta$.

For exact matrix completion, both the nuclear norm minimization algorithm and the spectral matrix completion approach require a minimum observing rate of order $O(n^{-1}\log^2 m)$ (see Theorem 1 in Candès and Plan (2010) and Theorem 1.2 in Keshavan et al. (2010a)), assuming a logarithm bound on the incoherence parameters.

For noisy matrix completion, consider $\boldsymbol{M}$ to be a low rank matrix added with IID Gaussian noise of standard deviation σ. The nuclear norm minimization algorithm has the same observing rate requirement, $O(n^{-1}\log^2 m)$, as the exact completion problem, with a recovery error rate $\|\boldsymbol{M} - \hat{\boldsymbol{M}}\|_F = O(\sqrt{mn^2}) \gtrsim O(\sqrt{mn})$ (see Theorem 7 in Candès and Plan (2010)). On the other hand, the OptSpace algorithm requires an observing rate of $O(\log n/n + \sqrt{m/n^3})$, achieving a smaller error rate $\|\boldsymbol{M} - \hat{\boldsymbol{M}}\|_F = O(\sqrt{mn}(n/m)^{-1/4}\log^{-1} m) = o(\sqrt{mn})$ (see Theorem 1.2 in Keshavan et al. (2010b)). Comparing these two algorithms, the OptSpace algorithm leads to a smaller recovery error under a slightly stronger requirement on the observing rate. Our condition on the observing rate is the same as that of the OptSpace algorithm,up to a logarithmic factor, if δ is chosen as $\delta = \delta^*$ such that $p_0q_0 = (PQ)^{1/4+\delta^*}$, i.e. $\mathcal{C}_{\delta^*}$ is just large enough to include the true configuration.

Our condition involves the additional logarithmic factor $\log^2 m$ because the proposed method also has a step of selecting the configuration. On the other hand, if one prefers more flexibility given by a larger candidate set $\mathcal{C}_\delta$ with $0 < \delta < \delta^*$, then Assumption 3 will require a higher observing rate. We also note that a similar comparison can be made with the recent result Abbe et al. (2020).

Following the preceding discussion, let us consider the special case where the true signal matrix is indeed low rank. In such a scenario, the true configuration of the KPD is $(P, 1)$. In order for the candidate set $\mathcal{C}_\delta$ to include the true configuration, the constant δ has to be less than or equal to

$$\delta^* = \frac{\log(P \wedge Q)}{\log PQ} - \frac{1}{4},$$

when $P \wedge Q > (PQ)^{1/4}$. If the candidate set $\mathcal{C}_{\delta^*}$ is used to search the configuration, then Assumption 3 again requires the same observing rate (up to a logarithmic factor) as that of the OptSpace algorithm.

5 Aggregated Estimation

Instead of selecting the single best configuration from the candidate set, an alternative is to incorporate several configurations and combine the recovered matrices under these configurations, leading to a aggregated recovery procedure. One immediate motivation and advantage of this approach is that it alleviates the infeasibility problem when only one configuration is used. Note that if a block is completely missing under the selected configuration $(\hat{p}, \hat{q})$, then all the missing entries in this block cannot be recovered (see Definition 2).

We explain heuristically how aggregation can help to alleviate this issue. Let (p_0, q_0) be the true configuration. Suppose the first $P/p_0 \times Q/q_0$ block of $\boldsymbol{Y}$ is completely unobserved as in Figure 1a, where missing entries in the first block are highlighted in black, and other missing entries are marked in gray. After the rearrangement using the true configuration (p_0, q_0), this block becomes a completely unobserved row in Figure 1b. If we consider using another configuration (p, q) with $p \geqslant p_0$ and $q \geqslant q_0$, there will be more completely missed rows after the rearrangement. For example, Figure 1c shows $\mathcal{R}_{p,q}[\boldsymbol{Y}]$ with $(p, q) = (2p_0, q)$, which has two missing rows. On the other hand, under a configuration (p, q) such that either $p < p_0$ or $q < q_0$, there won't be any completely missed rows after rearrangement. Figure 1d depicts the missing entries of $\mathcal{R}_{p,q}[\boldsymbol{Y}]$ with $(p, q) = (p_0/2, 2q_0)$, under which those pixels in black are recoverable. There can be multiple configurations under which the completely missing block is recoverable. Aggregate the recovered matrices under such configurations helps to recover the blocks that are completely missing.

Figure 1: Locations of unobserved entries under rearrangements of different configurations, where white pixels represent observed entries, black pixels represent the unobserved entries in the first (upper-left most) block and gray pixels represent other missing entries.

In this section we aim to provide an empirical procedure of the aggregated estimation, and discuss its properties and advantages. We also propose to use cross validation to select the number of configurations to be incorporated.

5.1 Aggregated Estimation

Let $C_k = (p_k, q_k), k = 1, 2, \ldots$ be the sequence of configurations ordered according to the criterion function (8) within the candidate configuration set $\mathcal{C}_\delta$, and let $\hat{\boldsymbol{X}}_k$ be the estimated $\boldsymbol{X}$ using configuration C_k. For the (i, j)-th entry of $\boldsymbol{X}$, and the k-th configuration C_k, define ν_{ijk} be the feasibility indicator which takes value 0 if the (i, j)-th entry is infeasible under configuration C_k, and 1 otherwise. Let $d_{ij} = \min\{k : \nu_{ijk} = 1\}$ so $C_{d_{ij}}$ is the best configuration under which (i, j)-th entry is feasible. For a given integer $d > 0$, the final estimate of the (i, j)-th entry of $\boldsymbol{X}$ is obtained as a weighted average of the (i, j)-th entries of $\hat{\boldsymbol{X}}_k$,

$$\hat{\boldsymbol{X}}[i, j] = \frac{\sum_{k=1}^{d \vee d_{ij}} w_k \nu_{ijk} \hat{\boldsymbol{X}}_k[i, j]}{\sum_{k=1}^{d \vee d_{ij}} w_k \nu_{ijk}} \tag{19}$$

where w_k is the weight assigned to configuration C_k. The simplest choice would be the constant weights such that $w_k \equiv 1$. A more refined approach is to use a set of weights that reflects the accuracy of each configuration.

For most of the missing entries, the aggregated estimator take the weighted average of the

recovered entries among the best d configurations. If an entry (i, j) is infeasible under all the d best configurations $C_1, \ldots, C_d$, the above estimator uses the best configuration $d_{ij} > d$ under which (i, j)-th entry is feasible, and fill it with the recovered entry under that configuration.

Remark: The aggregated estimation (19) can also be viewed as a model averaging procedure. The benefit of the aggregation is multi-fold. First it provides an effective approach to handle the infeasibility issue. The probability that there is an entry that is infeasible under all possible configurations is extremely small, for reasonably large (M, N), and number of factors of P and Q. Hence the procedure is able to handle higher missing rates. Second, model averaging can potentially provide more robust and stable estimators, as demonstrated in many studies in statistics literature (Buckland et al., 1997; Raftery et al., 1997). Note that for each possible configuration, there is a corresponding KPD (1), which, when truncated at a given K-rank r, provides an approximation to the signal matrix $\boldsymbol{X}$ in (5), though different approximations under different configurations and K-ranks are of different qualities. Averaging over the best performing configurations can potentially improve the quality of matrix recovery. Third, aggregation provides sharper resolution in the completed matrix, particularly in image reconstruction. Kronecker products induce a block structure in the resulting matrix hence often produce 'grainy' images. Averaging over several configurations reduces such effects. Fourth, the final predictive model after averaging is equivalent to a hybrid Kronecker product model

$$\boldsymbol{X} = \sum_{k=1}^{d} w_k \boldsymbol{X}_k,$$

where each $\boldsymbol{X}_k$ assumes a KPD of form (3) under configuration C_k. The aggregation approach bypasses the difficulty of jointly estimating such a model as well as determining the configurations in such a model. The effectiveness of the approach will be demonstrated in the empirical study.

Remark: In some real applications (for example image completion), it is difficult to justify that $\boldsymbol{X}$ has Kronecker rank 1 and it is possible that $\boldsymbol{X}$ is instead a sum of several Kronecker products with different configurations (see the hybrid decomposition in Cai et al. (2019a)). The aggregation procedure is actually identifying those most plausible configurations and recovers the matrix $\boldsymbol{X}$ as a mixture of them as in (19).

Remark: The empirical study shows that the performance is less sensitive to the number of configurations used in the aggregated recovery (19) if the corner configurations (small or large $m + n$) are excluded. More detailed investigation may be needed.

Remark: In practice the matrix dimension P and Q may not have many factors, which limits

the flexibility of the KPD approach as the candidate set $\mathcal{C}$ can be small. In this case it is possible to augment the observed matrix with additional missing rows and columns so that the new dimensions P^* and Q^* have more factors. One such choice is to make $P^* = 2^M$ and $Q^* = 2^N$. One can also use different P^* and Q^* as part of the aggregating operation. With a good configuration determination procedure and effective aggregation, significant improvement in matrix completion tasks can be obtained.

5.2 Choose Number of Terms by Cross-validation

In practice, it remains to determine the number of different configurations to be incorporated in the aggregation. Here we propose an empirical procedure for determining the number of terms by the K-fold cross-validation. Recall Ω is the set of indices of the observed entries in $\boldsymbol{Y}$. Let $\Omega^{[1]}, \cdots, \Omega^{[K]}$ be a random K-fold partition of Ω such that $\Omega^{[1]} \cup \cdots \cup \Omega^{[K]} = \Omega$ and for any $1 \leqslant i < j \leqslant K$, we have $\Omega^{[i]} \cap \Omega^{[j]} = \emptyset$ and $||\Omega^{[i]}| - |\Omega^{[j]}|| \leqslant 1$. Suppose $\hat{\boldsymbol{X}}(k; \Omega \setminus \Omega^{[i]})$ is the aggregated estimate for $\boldsymbol{X}$ as constructed in (19), which is obtained by fitting Kronecker matrix completion models on $P_{\Omega \setminus \Omega^{[i]}} \boldsymbol{Y}$ and aggregating the results from the first k configurations in $\mathcal{C}_\delta$. The K-fold cross-validation mean squared error (CV-MSE) is defined by

$$\begin{aligned} \text{CV-MSE}(k) &= \frac{1}{|\Omega|} \left\| P_\Omega \boldsymbol{Y} - \sum_{i=1}^{K} P_{\Omega^{[i]}} \hat{\boldsymbol{X}}(k; \Omega \setminus \Omega^{[i]}) \right\|_F^2 \\ &= \frac{1}{|\Omega|} \left\| \sum_{i=1}^{K} \left\{ P_{\Omega^{[i]}} \boldsymbol{Y} - P_{\Omega^{[i]}} \hat{\boldsymbol{X}}(k; \Omega \setminus \Omega^{[i]}) \right\} \right\|_F^2 . \end{aligned} \tag{20}$$

The number of models can be selected by minimizing CV-MSE(k) over $k = 1, 2, \ldots$. In principle, CV-MSE(k) can also be used to estimate the performance $\|\boldsymbol{Y} - \hat{\boldsymbol{X}}(k; \Omega)\|_F^2$, including the unobserved entries as well, which cannot be calculated directly.

In i-th cross-validation fold, we use $P_{\Omega \setminus \Omega^{[i]}} \boldsymbol{Y}$ as the training set and use $P_{\Omega^{[i]}} \boldsymbol{Y}$ as the test set. The split of data decreases the observation rate in $P_{\Omega \setminus \Omega^{[i]}} \boldsymbol{Y}$ by $1/K$ and increase the unrecoverability and the number of unrecoverable entries, which may worsen the performance. To minimize the potential impact of the decrease in observing rate, a large value of K is preferred. On the other hand, larger K involves heavier computation. In the numerical analysis that follows, we choose $K = 10$ as a compromise.

6 Empirical Examples

6.1 Simulation: Configuration Selection

In this simulation experiment, we demonstrate the performance of the configuration determination procedure under the model (5) with a $\boldsymbol{X}$ of K-rank 1. Specifically, the component matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are randomly generated as

$$\boldsymbol{A} = \begin{pmatrix} \sqrt{1-\varphi^2} \\ 0 \end{pmatrix} \otimes \boldsymbol{D}_1 + \begin{pmatrix} 0 \\ \varphi \end{pmatrix} \otimes \boldsymbol{D}_2,$$
$$\boldsymbol{B} = \begin{pmatrix} \sqrt{1-\varphi^2} \\ 0 \end{pmatrix} \otimes \boldsymbol{D}_3 + \begin{pmatrix} 0 \\ \varphi \end{pmatrix} \otimes \boldsymbol{D}_4,$$

where $\|\boldsymbol{D}_i\|_F = 1$, $i = 1, 2, 3, 4$, $\boldsymbol{D}_1$ and $\boldsymbol{D}_2$ are $2^{m_0-1} \times 2^{n_0}$ matrices such that $\mathrm{tr}(\boldsymbol{D}_1\boldsymbol{D}_2^T) = 0$, and $\boldsymbol{D}_3$ and $\boldsymbol{D}_4$ are $2^{M-m_0-1} \times 2^{N-n_0}$ matrices such that $\mathrm{tr}(\boldsymbol{D}_3\boldsymbol{D}_4^T) = 0$. We design such a construction so that the representation gap can be controlled at roughly φ^2. Let $\boldsymbol{X} = \lambda \boldsymbol{A} \otimes \boldsymbol{B}$, where λ is the parameter used to control the signal-to-noise ratio. The underlying complete matrix $\boldsymbol{Y}$ is generated according to $\boldsymbol{Y} = \boldsymbol{X} + \sigma 2^{-(M+N)/2}\boldsymbol{E}$, where $\boldsymbol{E}$ contains IID standard normal entries. The observation set Ω is sampled independently such that $P[(i,j) \in \Omega] = \tau$ for all (i,j).

We consider two dimension and configuration setups: $M = N = 9$ with true configuration $(m_0, n_0) = (4, 4)$, and $M = N = 10$ with true configuration $(m_0, n_0) = (5, 4)$. All combinations of 20 different signal-to-noise ratio values $\lambda/\sigma \in \{0.1, 0.2, \ldots, 2.0\}$, two different observing rates $\tau = 0.1$ and 0.2, three different values of $\varphi^2 = 0.3, 0.4$, and 0.5 and three different candidate sets $\mathcal{C}_5, \mathcal{C}_6$, and $\mathcal{C}_7$ are considered, where

$$\mathcal{C}_s := \{(p,q) : p \in d(P), q \in d(Q), 2^s \leqslant pq \leqslant PQ2^{-s}\}.$$

here s is an integer satisfying $(M+N)/4 < s < (M+N)/2$. With slight misuse of notations, the candidate set $\mathcal{C}_s$ is equivalent to $\mathcal{C}_\delta$ in (7) with corresponding value of $\delta = s/(M+N) - 1/4$. However, it is more convenient and intuitive to use $\mathcal{C}_s$ here. Note that $\mathcal{C}_5$ is the largest set among the three.

For each combination of λ/σ, τ, φ^2 and $\mathcal{C}$, we repeat the simulation 100 times and record how many times the true configuration (m_0, n_0) is selected by the criterion in (8). These empirical frequencies are plotted as functions of the signal-to-noise ratio in Figure 2 for different combinations of (M, N), τ, φ^2 and $\mathcal{C}$. We see the general pattern that these functions are all monotone, and when the signal-to-noise ratio is large enough, chances of choosing the correct configuration approach one.

The left panel in Figure 2 reveals that a larger signal-to-noise ratio is required when the candidate set $\mathcal{C}_s$ is larger (note that $\mathcal{C}_5$ is the largest among the three of them). This is consistent with Assumption 4, which requires a larger signal-to-noise ratio for a smaller value of δ, i.e. a larger candidate set. Theorem 2 also shows that the convergence rate is a decreasing function of δ. The middle panel of Figure 2 indicates that as the representation gap increases, the probability of making correct choice also increases. It is the reason that Assumption 4 links the representation gap ψ^2 and the required signal-to-noise ratio λ/σ in a negative relationship. The right panel gives the empirical frequency curves for different dimensions and different observing rates. The pattern is very intuitive, when the dimension is larger, or the observing rate is higher, the performance is better.

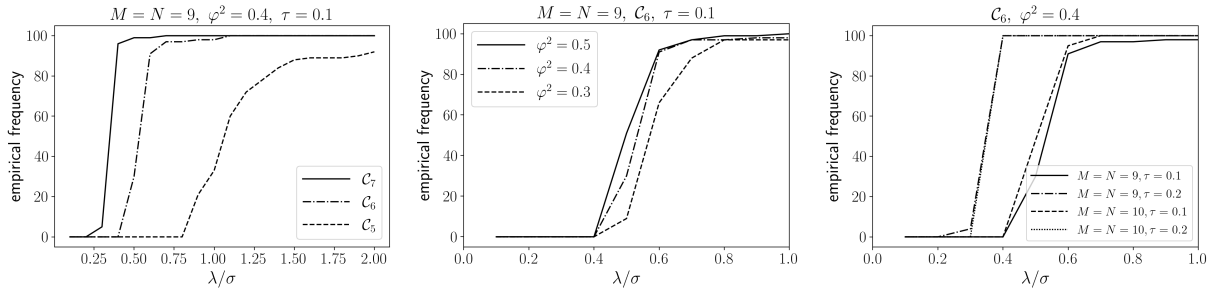

Figure 2: The empirical frequency (in %) of correct configuration selection increases monotonically as the signal-to-noise ratio for (Left) different candidate sets, (Mid) different representation gaps and (Right) different dimensions and observing rates.

To give another comparison, for any combination of (M, N), τ, φ^2 and $\mathcal{C}$, we define γ^* be the smallest signal-to-noise ratio (truncated to 0.1) such that the empirical frequency of correct configuration selection exceeds 50%. The results are reported in Table 1.

6.2 Simulation: Comparison with Conventional Matrix Completion Methods

We compare the proposed matrix completion algorithm based on the Kronecker product decomposition with some state-of-art low-rank matrix completion methods. The matrix in this simulation is generated according to

$$\boldsymbol{Y} = \boldsymbol{A} \otimes \boldsymbol{B} + 2^{-(M+N)/2}\boldsymbol{E},$$

where $\boldsymbol{A} \in \mathbb{R}^{2^{m_0} \times 2^{n_0}}$ and $\boldsymbol{B} \in \mathbb{R}^{2^{M-m_0} \times 2^{N-n_0}}$ are obtained by normalizing the IID Gaussian random matrices of corresponding dimensions so that they have Frobenius norm 1. The $\boldsymbol{E} \in \mathbb{R}^{2^M \times 2^N}$ contains

	τ	0.1			0.2		
	φ^2	0.3	0.4	0.5	0.3	0.4	0.5
$M = N = 9$	$\mathcal{C}_5$	1.6	1.1	1.0	0.6	0.5	0.5
	$\mathcal{C}_6$	0.6	0.6	0.5	0.4	0.4	0.4
	$\mathcal{C}_7$	0.4	0.4	0.4	0.3	0.3	0.3
$M = N = 10$	$\mathcal{C}_5$	1.3	1.0	0.9	0.6	0.5	0.5
	$\mathcal{C}_6$	0.6	0.6	0.5	0.4	0.4	0.4
	$\mathcal{C}_7$	0.4	0.4	0.4	0.3	0.3	0.3

Table 1: Value of γ^* for different combinations of (M, N), τ, ψ^2 and $\mathcal{C}$.

IID standard Gaussian entries. Roughly speaking, the signal-to-noise ratio in this generating model is 1.

We consider two state-of-art methods for the low rank matrix completion: OptSpace (Keshavan et al., 2010b) and ALS (Jain et al., 2013) for comparison. For the approach proposed in this paper, after the configuration has been determined according to (8), we apply both ALS and OptSpace to the rearranged matrix $\mathcal{R}_{\hat{p},\hat{q}}[P_\Omega \boldsymbol{Y}]$. These two algorithms are labelled Kron-OptSpace and Kron-ALS respectively.

Specifically, we consider four true configurations, and fix $M = 10, N = 8$ for the dimensions of $\boldsymbol{Y}$ and $\tau = 0.1$ for the observing rate. The average completion error over 100 repetitions for the four aforementioned methods are reported in Table 2. The completion error is calculated by $\|\hat{\boldsymbol{X}} - \boldsymbol{A} \otimes \boldsymbol{B}\|_F$, where $\hat{\boldsymbol{X}}$ is the completed matrix. All approaches conduct a rank-1 matrix completion, where either $\boldsymbol{Y}$ (for OptSpace and ALS) or $\mathcal{R}_{\hat{p},\hat{q}}[P_\Omega \boldsymbol{Y}]$ (for Kron-OptSpace and Kron-ALS) is assumed to have rank 1. In both OptSpace and ALS algorithms, we set the number of iterations to 100 and the relative tolerance in Frobenius norm to 1e-8.

Note that when $(m_0, n_0) = (10, 0)$, $\boldsymbol{A} \otimes \boldsymbol{B}$ is of rank 1, it turns out Kron-OptSpace and Kron-ALS have exactly the same completion errors as OptSpace and ALS respectively. For the other three choices of (m_0, n_0), Kron-OptSpace and Kron-ALS are much better than OptSpace and ALS. The inferior performance of ALS and OptSpace is due to the misspecification: the matrix $\boldsymbol{A} \otimes \boldsymbol{B}$ is not rank-1 under these configurations. In contrast, our approaches (Kron-OptSpace and Kron-ALS) successfully identifies the low Kronecker rank structure and applies the conventional low rank matrix completion algorithms on the correctly transformed matrix $\mathcal{R}_{\hat{p},\hat{q}}[P_\Omega \boldsymbol{Y}]$, whose signal parts is indeed of rank 1.

(m_0, n_0)	OptSpace	ALS	Kron-OptSpace	Kron-ALS
(5, 4)	1.1074 (0.0303)	3.2172 (2.9519)	0.2156 (0.0204)	0.2086 (0.0056)
(6, 4)	1.1159 (0.02871)	4.0533 (4.1828)	0.2508 (0.0381)	0.2370 (0.0053)
(9, 0)	0.7499 (0.0201)	0.7499 (0.0196)	0.2130 (0.0174)	0.2080 (0.0044)
(10, 0)	0.2471 (0.0357)	0.2379 (0.0055)	0.2471 (0.0357)	0.2379 (0.0055)

Table 2: Completion errors of four approaches for four different true configurations, averaged over 100 repetitions. Standard deviations among 100 repetitions are reported in parentheses. Each approach assumes the signal matrix to be either rank-1 or Kronecker rank-1.

In addition, by comparing the completion errors of OptSpace (Kron-OptSpace *resp.*) and ALS (Kron-ALS *resp.*), we observe that ALS results in a smaller error than OptSpace under the correct model specification, and OptSpace is more stable when the model is misspecified.

Since $\boldsymbol{A} \otimes \boldsymbol{B}$ is not rank-1 under the true configurations $(m_0, n_0) = (5, 4), (6, 4), (9, 0)$, we also consider the conventional OptSpace with ranks from 1 to 10, and plot the completion errors in Figure 3 (completion errors of the rank-1 Kron-OptSpace are plotted as dashed lines). When $(m_0, n_0) = (5, 4), (6, 4)$, $\boldsymbol{A} \otimes \boldsymbol{B}$ is not low rank at all. Increasing the rank also increases the completion error. When $(m_0, n_0) = (9, 0)$, $\boldsymbol{A} \otimes \boldsymbol{B}$ is exactly rank-2, a rank-2 OptSpace matrix completion results in the lowest completion error among all ranks. However, the rank-2 OptSpace is still not as good as the rank-1 Kron-OptSpace. This is because the noise matrix interferes more with $\boldsymbol{A} \otimes \boldsymbol{B}$ in OptSpace than with $\text{vec}(\boldsymbol{A})[\text{vec}(\boldsymbol{B})]^\top$ in Kron-OptSpace, because the former is rank-2 and the latter rank-1 (see Cai et al. (2019b) for more on the rank of $\boldsymbol{A} \otimes \boldsymbol{B}$ rearranged under a mis-specified configuration).

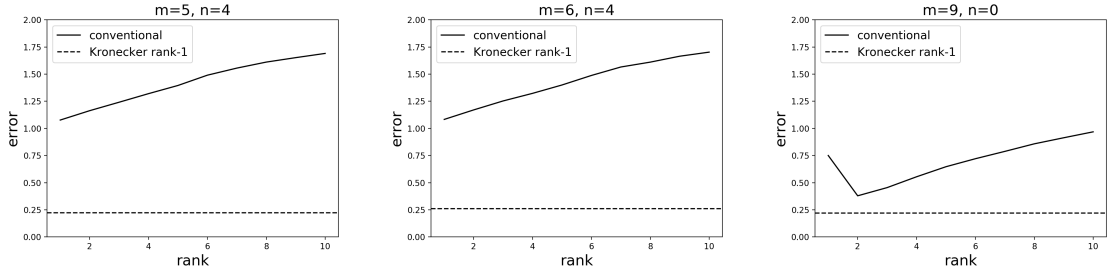

Figure 3: Completion errors versus the rank specified for conventional multi-rank OptSpace matrix completion, averaged over 100 repetitions. The completion errors from rank-1 Kronecker OptSpace algorithm are plotted as dashed lines.

To summarize, the simulation in this subsection reveals that when $\boldsymbol{X}$ has a low Kronecker rank rather than low rank, conventional matrix completion methods suffer from misspecification error while our proposed method can discover the underlying Kronecker structure and achieve a better completion. When $\boldsymbol{X}$ is indeed of low rank, our method attains a similar performance to conventional matrix approaches (with a little overhead in the computation for determining the configuration).

6.3 Simulation: Aggregated Estimation

In this simulation experiment, we look into the performance of the aggregated estimation introduced in Section 5. The signal part $\boldsymbol{X} \in \mathbb{R}^{2^M \times 2^N}$ is generated as a mixture of k_0 Kronecker products:

$$\boldsymbol{X} = \sum_{i=1}^{k_0} \boldsymbol{A}_i \otimes \boldsymbol{B}_i,$$

where $\boldsymbol{A}_i$ and $\boldsymbol{B}_i$ are $2^{m_i} \times 2^{n_i}$ and $2^{M-m_i} \times 2^{N-n_i}$ matrices, correspondingly. In other words, the i-th Kronecker product in $\boldsymbol{X}$ has the configuration (m_i, n_i). (Here for notational simplicity, we use (m, n) instead of $(2^m, 2^n)$ to indicate the configuration.) We design the structure of $\boldsymbol{A}_i$ and $\boldsymbol{B}_i$ as follows:

$$\begin{aligned}
\boldsymbol{A}_i &= \varphi \boldsymbol{D}_1^{(i)} \otimes \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} + \sqrt{1-\varphi^2} \boldsymbol{D}_2^{(i)} \otimes \begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}, \\
\boldsymbol{B}_i &= \varphi \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix} \otimes \boldsymbol{D}_3^{(i)} + \sqrt{1-\varphi^2} \begin{bmatrix} \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{bmatrix} \otimes \boldsymbol{D}_4^{(i)},
\end{aligned}$$

where $\boldsymbol{D}_1^{(i)}$ and $\boldsymbol{D}_2^{(i)}$ are random $2^{m_i} \times 2^{n_i - 1}$ matrices such that they both have Frobenius norm 1 and are orthogonal with each other, and so do $\boldsymbol{D}_3^{(i)}$ and $\boldsymbol{D}_4^{(i)}$. It follows that $\|\boldsymbol{A}_i\|_F = \|\boldsymbol{B}_i\|_F = 1$. The quantity $\varphi^2 \in (0, 1)$ controls the representation gap of the product $\boldsymbol{A}_i \otimes \boldsymbol{B}_i$. Specifically, at the configurations $(m_i, n_i - 1)$ and $(m_i + 1, n_i)$, it holds that

$$\|\mathcal{R}_{m_i, n_i - 1}[\boldsymbol{A}_i \otimes \boldsymbol{B}_i]\|_S \approx \|\mathcal{R}_{m_i + 1, n_i}[\boldsymbol{A}_i \otimes \boldsymbol{B}_i]\|_S \approx \varphi \wedge \sqrt{1 - \varphi^2}.$$

Furthermore, under the above construction, the incoherence parameters μ for $\boldsymbol{A}_i$ and $\boldsymbol{B}_i$ are roughly the same ($\approx \sqrt{(m_i + n_i) \wedge (M + N - m_i - n_i) \cdot \log 2}$) for different values of φ.

In this experiment, we set $M = N = 9$. Two specific constructions of $\boldsymbol{X}$ are considered: (i) $k_0 = 1$ (one term model) and $(m_1, n_1) = (5, 4)$, and (ii) $k_0 = 2$ (two term model) and $(m_1, n_1) = (5, 4)$, $(m_2, n_2) = (4, 5)$. Two values of $\varphi^2 = 0.5$ and 0.05 are used, corresponding to large and small representation gaps respectively. We label the four scenarios as follows for later references.

Scenario L1: One-term $\boldsymbol{X}$ ($k_0 = 1$) with a large representation gap $\varphi^2 = 0.5$.

Scenario S1: One-term $\boldsymbol{X}$ ($k_0 = 1$) with a small representation gap $\varphi^2 = 0.05$.

Scenario L2: Two-term $\boldsymbol{X}$ ($k_0 = 2$) with a large representation gap $\varphi^2 = 0.5$.

Scenario S2: Two-term $\boldsymbol{X}$ ($k_0 = 2$) with a small representation gap $\varphi^2 = 0.05$.

For each scenario, the matrix

$$\boldsymbol{Y} = \boldsymbol{X} + \frac{2}{2^{(M+N)/2}} \cdot \boldsymbol{E}$$

is partially observed with observing rate $\tau = 0.2$. We apply the aggregation approach (19) proposed in Section 5 to $P_\Omega \boldsymbol{Y}$ with equal weights using up to 10 configurations in the candidate set $\mathcal{C}_7$. For each configuration, we use either one or two Kronecker products with that configuration in the averaging, and refer to them as K-rank-1 and K-rank-2 in the sequel. The mean squared error $2^{-(M+N)}\|\hat{\boldsymbol{X}} - \boldsymbol{X}\|_F^2$ based on 100 repetitions are plotted against the number of configurations in Figure 4 for the four scenarios.

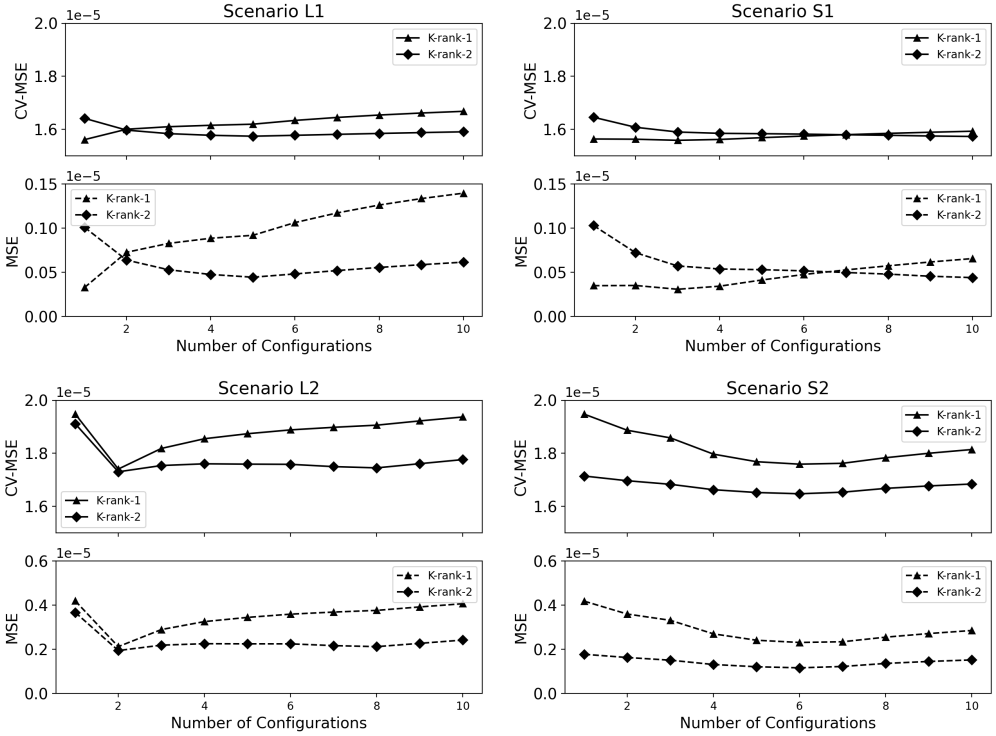

Figure 4: Cross-validated mean squared error (CV-MSE) and mean squared error (MSE) under four scenarios, using two K-ranks.

When the representation gap is large, models with wrong configurations usually possess a large bias. The benefit from averaging models under many configurations may not be enough to offset

Scenario L2

Term	Config.	Criterion Function	Error (one term)		Error (agg.)	
			Rank-1	Rank-2	Rank-1	Rank-2
1	$(5,4)$	0.2208	0.5527	0.4939	0.5527	0.4939
2	$(4,5)$	0.2124	0.5547	0.4698	0.2770	0.2520
3	$(5,5)$	0.1773	0.7811	0.6017	0.3792	0.2804
4	$(4,4)$	0.1737	0.8017	0.6226	0.4342	0.2897
5	$(6,4)$	0.1720	0.8294	0.6407	0.4670	0.2934
6	$(4,6)$	0.1703	0.8020	0.6470	0.4686	0.2732
7	$(3,5)$	0.1698	0.8133	0.6491	0.4767	0.2742
8	$(5,3)$	0.1693	0.8303	0.6366	0.4876	0.2693
9	$(6,5)$	0.1645	0.9752	0.9893	0.5140	0.2952
10	$(5,6)$	0.1606	0.9575	0.9429	0.5327	0.3172

Table 3: Configuration-wise MSE and the MSE of averaging over first k configurations are listed for both K-rank 1 and K-rank 2 matrix completion. Error (one term) refers to MSE of $\hat{\boldsymbol{X}}$ obtained in the k-th iteration. Error (agg.) refers to the MSE of the aggregated model using the first k configurations. Both are normalized by the factor $\|\boldsymbol{X}\|_F^2/2^{M+N}$.

this extra bias. Therefore, in scenarios L1 and L2 (first column of Figure 4), the MSE starts to increase when number of averaged configurations is greater than the true number of Kronecker products k_0, except when K-rank-2 models under scenario L1. However, when the representation gap is small, there exist several alternative configurations with relatively small bias. Adding those configurations into the averaging model can significantly reduce variance while adding little bias. Hence, in scenarios S1 and S2 (second column of Figure 4), adding more terms can reduce the MSE. We also observe that in general averaging using K-rank 2 models outperforms averaging K-rank 1 models in S2.

The ordered configurations, their corresponding MSE, and the MSE of the model aggregated up to it are reported in Table 3 for Scenario L2. Under L2, $\boldsymbol{X}$ is the sum of two Kronecker products of configurations $(4,5)$ and $(5,4)$, which correspond to the two largest values of the selection criterion, and hence take up the first two rows of Table 3. The gap between the second configuration $(4,5)$ and the third $(5,5)$ in terms of the criterion is much greater than the gaps between any other two

Scenario S1

Term	Config.	Criterion Function	Error (one term)		Error (agg.)	
			Rank-1	Rank-2	Rank-1	Rank-2
1	$(5, 4)$	0.2177	0.0868	0.2599	0.0868	0.2599
2	$(6, 4)$	0.2163	0.1662	0.3374	0.0931	0.1818
3	$(5, 3)$	0.2143	0.1597	0.2958	0.0792	0.1418
4	$(6, 3)$	0.2077	0.1926	0.3200	0.0888	0.1364
5	$(6, 5)$	0.1799	0.7134	0.4740	0.1088	0.1391
6	$(7, 4)$	0.1799	0.6758	0.4577	0.1247	0.1414
7	$(4, 3)$	0.1777	0.6736	0.4636	0.1342	0.1296
8	$(5, 2)$	0.1731	0.7377	0.4502	0.1466	0.1234
9	$(5, 5)$	0.1712	0.5806	0.2337	0.1593	0.1164
10	$(4, 4)$	0.1711	0.5623	0.2405	0.1666	0.1113

Table 4: Top 10 configurations according to the criterion function under Scenario S1. Error (one term) refers to MSE of $\hat{\boldsymbol{X}}$ obtained in the k-th iteration. Error (agg.) refers to the MSE of the aggregated model using the first k configurations. Both are normalized by the factor $\|\boldsymbol{X}\|_F^2/2^{M+N}$.

consecutively selected configurations. As a result, aggregation over only the first two configurations yields the best performance for both K-rank 1 and K-rank 2 matrix completion.

Similar results for Scenario S1 are reported in Table 4. Due to the small representation gap, configurations $(6, 4)$, $(5, 3)$ and $(6, 3)$ are close to the truth $(5, 4)$ in terms of the the criterion function. Therefore, aggregation over more terms can reduce the MSE.

The number of terms in the averaging can be determined by a cross-validation procedure as discussed in Section 5. We conduct a 10-fold cross-validation and report the cross-validated MSE (CV-MSE) based on 100 repetitions in Figure 4. The CV-MSE exhibits patterns similar to MSE and can approximate MSE well (up to an offset) in Figure 4, suggesting that cross-validation can be used in practice to determine the number of terms incorporated in the aggregated recovery. For Scenarios S1, L1 and L2, for all 100 repetitions, the number of configurations that minimizes the CV-MSE also minimizes the MSE, while for S2, 92 out of 100 repetitions do.

As we mentioned earlier, aggregation is an approach that tunes the bias-variance trade-off, which results in a smaller error when (i) the added terms can reduce the bias or (ii) the added terms

can reduce variance while only introducing additional bias slightly. In the preceding experiments, if $\boldsymbol{X}$ is a sum of multiple Kronecker products of different configurations, averaging over multiple configurations can reduce the bias but the bias begins to increase when number of averaged models exceeds the number of terms in $\boldsymbol{X}$ as in Scenario L2. On the other hand, if the terms in $\boldsymbol{X}$ have small representation gaps, aggregation can improve the overall performance by reducing the variance without increase the bias significantly as in Scenarios S1 and S2.

6.4 Simulation: Completely missing block

In addition to variance reduction, the aggregation over different configurations can also help recovering entries that are unrecoverable under a single configuration. In this part of the simulation study, we consider Scenario S1 in Section 5, and force that the first $2^4 \times 2^5$ block is completely unobserved. In other words, the index set of observed entries reduces from the original Ω to $\Omega^B = \Omega \setminus B$, where $B := \{(i,j) : i = 1, \cdots, 2^4; j = 1, \cdots, 2^5\}$. We repeat the whole simulation 100 times and record the mean squared error of all the entries of $\boldsymbol{X}$, and of the first block (which is completely missing) of $\boldsymbol{X}$. Specifically, the overall MSE is $2^{-(M+N)}\|\hat{\boldsymbol{X}} - \boldsymbol{X}\|_F^2$ and the MSE of the first block is $2^{-(M+N-m_1-n_1)}\|P_B[\hat{\boldsymbol{X}}] - P_B[\boldsymbol{X}]\|_F^2$, where P_B is the projection operator onto the set B. The benchmark method fills the missing entries by the average value of observed ones. If there are still irrecoverable entries with K configurations, by default the benchmark method is used to fill in the values. We report the median of the MSE over 100 repetitions as a function of the number of configurations in Figure 5. Median instead of average is reported here for the sake of robustness, since in several repetitions the MSE is tremendously large for the first block. Although the first block is irrecoverable under the true configuration $(5, 4)$ (and also irrecoverable under (m, n) with $m \geqslant 5$ and $n \geqslant 4$), Figure 5 reveals that it can be recovered partially by the aggregated estimation, where the MSE of the first block is typically greater than the overall MSE.

We pick one repetition and report the top 10 configurations according to the criterion function, and the corresponding MSE up to the given configuration in Table 5. The performance of the benchmark method is also reported. The top 10 configurations are the same as those in Table 4, but the values of the criterion function are slightly smaller due to the extra missing block. The first block is irrecoverable under the top 2 configurations and therefore a benchmark method is used to fill in the first block, resulting in a large MSE for the first two configurations. Starting from the third row, when aggregating over more than three configurations, the first block is recoverable, with a significantly smaller MSE.

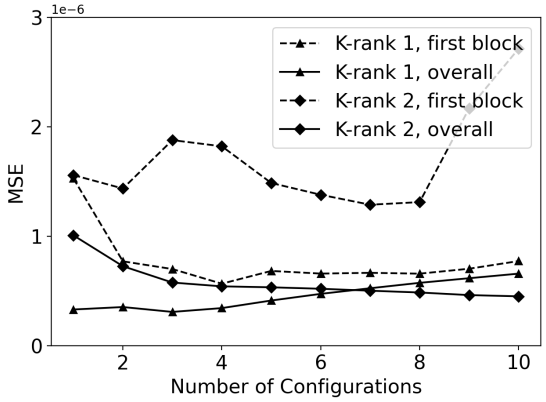

Figure 5: Median of overall MSE and first-block MSE over 100 repetitions as a function of number of configurations for Scenario S1.

Term	Config.	Criterion Function	MSE (overall)		MSE (1st block)	
			Rank-1	Rank-2	Rank-1	Rank-2
1	(5, 4)	0.2167	0.0914	0.2645	0.9986	0.9986
2	(6, 4)	0.2153	0.0977	0.1868	0.9986	0.9986
3	(5, 3)	0.2131	0.0795	0.1423	0.1191	0.1438
4	(6, 3)	0.2066	0.0889	0.1372	0.0847	0.1403
5	(7, 4)	0.1793	0.1062	0.1391	0.0847	0.1403
6	(6, 5)	0.1793	0.1248	0.1420	0.0847	0.1403
7	(4, 3)	0.1764	0.1344	0.1302	0.0710	0.1200
8	(5, 2)	0.1721	0.1472	0.1239	0.1204	0.1234
9	(5, 5)	0.1706	0.1598	0.1168	0.1204	0.1234
10	(4, 4)	0.1702	0.1685	0.1115	0.4567	0.0900
Benchmark (Fill in mean)			1.599		0.9986	

Table 5: Top 10 configurations according to the criterion function under Scenario S1 with a completely missing block under true configuration. the overall MSE and the one for the first block are normalized by $\|\boldsymbol{X}\|_F^2/2^{M+N}$ and $\|P_B\boldsymbol{X}\|_F^2/2^{(M+N-m_1-n_1)}$ accordingly.

Figure 6: (Left) Cameraman's image; (Middle) Cameraman's image with noise; (Right) Noisy image with 20% observed entries.

6.5 Real image example

In this section, we apply the proposed matrix completion approach to the cameraman's image[1], which has been widely used as a benchmark in image analysis. The original image, which is a 512×512 gray-scaled picture, is shown in Figure 6 (left panel). It is represented by a 512×512 matrix $\boldsymbol{X}$ of real numbers between 0 and 1, 0 for black and 1 for white.

We first add a noise to the original image such that $\boldsymbol{Y} = \boldsymbol{X} + 0.1\boldsymbol{E}$, where the entries of $\boldsymbol{E}$ are IID standard Gaussian noises. The corrupted image $\boldsymbol{Y}$ is shown in the middle panel of Figure 6. We set the observing rate to 20% ($\tau = 0.2$) and sample the observing set Ω as i.i.d Bernoulli(τ). The observed $\boldsymbol{Y}^* = P_\Omega \boldsymbol{Y}$ is plotted in the right panel of Figure 6, where missing entries are filled as white. We follow the configuration determination procedure proposed in Section 3.1. The maximum of the criteria function $\|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]\|_S$ is attained at the configuration $(\hat{p}, \hat{q}) = (64, 64)$, within the candidate configuration set $\mathcal{C}_6$ (defined in Section 6.1). It corresponds to decomposing $\boldsymbol{X}$ as the Kronecker product of a 64×64 and a 8×8 matrix. Recovered images under the configuration (64,64) and K-ranks 1 to 3 are shown in the upper row of Figure 7. The cameraman can be recognized from the recovered matrix of K-rank 1 and more details are added as the K-rank increases to 2 and 3.

We also compare the performance of matrix completion through KPD with the classical approach through SVD. In particular, we consider matrix completion via the alternating minimization algorithm with ranks 8, 16 and 24, matching the numbers of parameters under K-ranks 1, 2, 3 of the Kronecker matrix completion with configuration $(64, 64)$. These recovered images are shown in

[1]The image can be found in the Python package `scikit-image`, available at `https://scikit-image.org/`

KPD rank	1	2	3
Error	0.1224	0.1084	0.1136
Error (trimmed)	0.1219	0.1060	0.1071
SVD rank	8	16	24
Error	0.2167	0.4986	0.9968
Error (trimmed)	0.2060	0.3464	0.5988

Table 6: Error for Kronecker matrix completion and classical matrix completion with similar number of parameters.

the lower row of Figure 7. The superiority of the KPD approach is easily seen from the images. Besides judging the recovered images by eyesight, we quantify the quality of a recovered matrix by the *reconstruction error*

$$\|\boldsymbol{X} - \hat{\boldsymbol{X}}\|_F^2 / \|\boldsymbol{X}\|_F^2, \tag{21}$$

where $\boldsymbol{X}$ is the image without noise and $\hat{\boldsymbol{X}}$ is the recovered matrix. Table 6 reports the reconstruction errors of recovered matrices through KPD and SVD. The trimmed error is for the trimmed recovered matrix, whose pixel values are restricted to $[0, 1]$. It confirms again that Kronecker matrix completion can recover the cameraman's image more accurately compared to the SVD approach, with a similar number of parameters. The result is anticipated since KPD matrix completion has more flexibility in selecting the configurations, including SVD matrix completion as one of its special cases. The proposed configuration determination procedure is able to find a better configuration, which provides better performance than one of its special cases.

To examine the performance of the aggregation approach, we consider three restricted candidate sets $\mathcal{C}_5, \mathcal{C}_6, \mathcal{C}_7$, defined in Section 6.1, corresponding to $pq \wedge p^*q^* \geqslant 32, 64, 128$, respectively. We also consider both K-ranks 1 and 2 for each configuration. The configurations within the candidate sets are ordered based on the criterion function. Then $\hat{\boldsymbol{X}}$ is obtained by a simple average of the matrices recovered under the first d configurations, according to (19). The left sub-figure of Figure 8 shows the reconstruction error against d. The errors corresponding to K-rank 2 for the candidate set $\mathcal{C}_5$ is much worse than all other scenarios and is therefore not shown.

The right sub-figure of Figure 8 reports the cross-validation errors. We choose to use $K = 20$-fold to minimize the impact of decreased observing rate. Although CV-MSE and the reconstruction error have different scales (the former is calculated according to (20) where the error $\boldsymbol{E}$ is also involved,

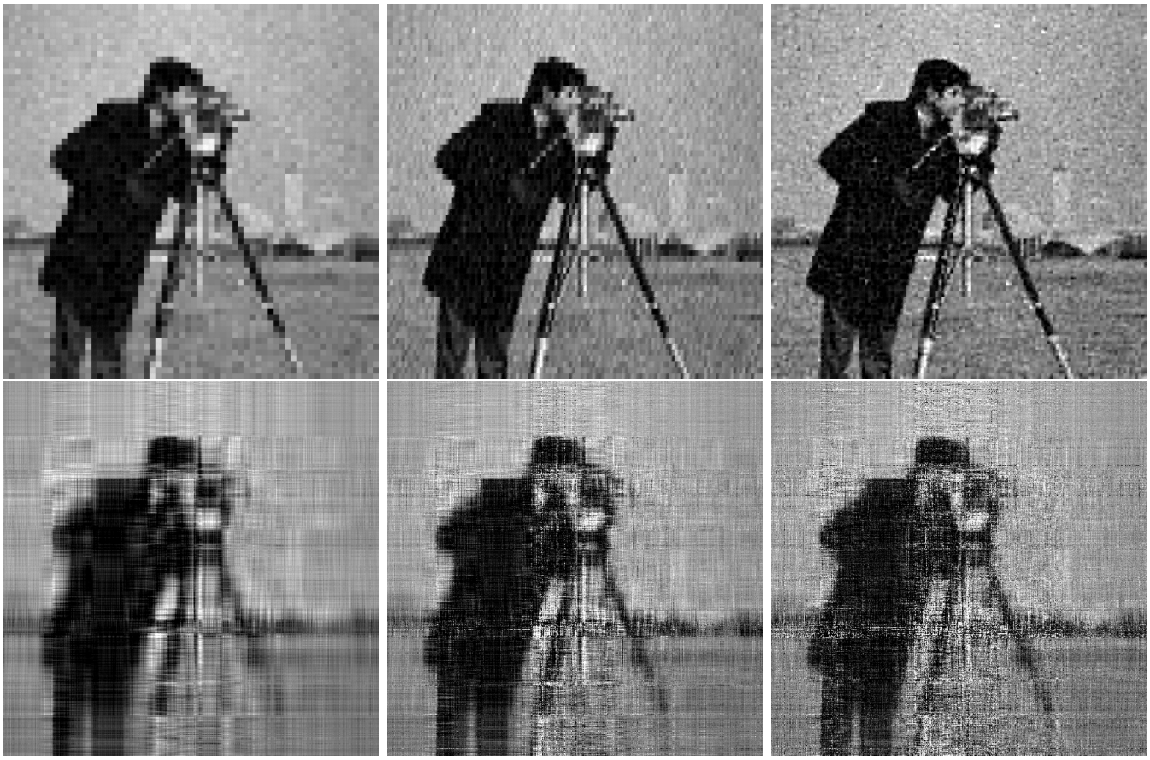

Figure 7: (Upper row) Recovered images using KPD (Lower row) Recovered images using SVD. The extreme values are trimmed such that all pixel values are between 0 and 1.

the latter does not involve $\boldsymbol{E}$ directly and is normalized by $\|\boldsymbol{X}\|_F^2$ as in (21)), CV-MSE nevertheless exhibits the same trend as the reconstruction error and the number of configurations that minimize CV-MSE coincides with the one that minimize the reconstruction error for each scenario. Similar to the simulation studies, the CV-MSE can work as a proxy of the reconstruction error, which is usually infeasible in practice. In addition, CV-MSE helps not only in determining the number of configurations with respect to a certain configuration set $\mathcal{C}$, but also in choosing the appropriate fitting K-rank and choosing the best configuration set $\mathcal{C}$.

We note that determining the best configuration set is not trivial. On the one hand, as discussed in Section 4, extreme configurations ($pq < (PQ)^{1/4}$ or $p^*q^* < (PQ)^{1/4}$) must be excluded to have a stable Kronecker product matrix completion. On the other hand, configurations close to these boundaries ($pq = (PQ)^{1/4}$ or $p^*q^* = (PQ)^{1/4}$) may have uncontrollable performance and slower convergence rates. Fitting a higher K-rank matrix completion model with those configurations can potentially suffer from sever overfitting. For example, if only K-rank 1 models are considered, $\mathcal{C}_5$ performs the best. However, fitting K-rank 2 models with respect to $\mathcal{C}_5$ results in a tremendously large error in some unobserved entries and is therefore not reported in Figure 8.

It can be seen from Figure 8 that using K-rank 2 and candidate set $\mathcal{C}_6$ performs the best. ($\mathcal{C}_5$ with K-rank 2 is severly overfitted as discussed earlier) The error of averaging the top 5 configurations is around 0.0790, which is about 20% smaller than 0.1084, obtained by the K-rank two matrix completion under a single configuration, as reported in Table 6. The error rate then stays roughly constant if more configurations are added into the averaging.

The reconstructed images averaged over 4 configurations with K-rank 1 under the candidate sets $\mathcal{C}_5$ and over 5 configurations of K-rank 2 under $\mathcal{C}_6$ are shown in Figure 9. With K-rank 1 (left panel), big pixels are observed in the reconstructed image but are less noisy than the ones in Figure 7. With K-rank 2 (right panel), more details are added, resulting in a smoother reconstructed image.

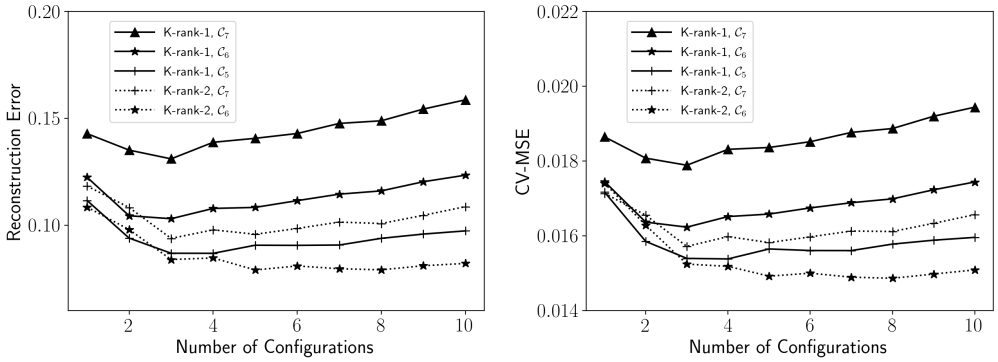

Figure 8: (Left) Reconstruction error and (Right) cross-validation error (CV-MSE) of the averaged matrix against the number of configurations.

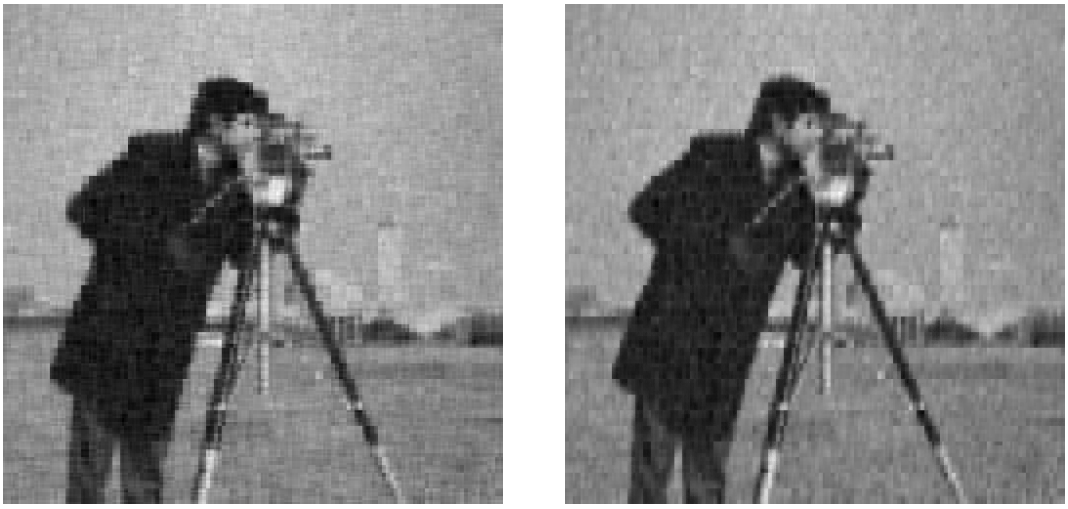

Figure 9: (Left) Average reconstructed image over 4 configurations of K-rank 1. (Right) Average reconstructed image over 5 configurations of K-rank 2.

7 Conclusion

In this article, we study matrix completion problem assuming the underlying complete signal matrix is of a low rank Kronecker product form, extending the classical assumption of low rank signal matrix. The new model includes the low rank model as a special case. Such an extension brings a significantly greater modeling flexibility which leads to more effective dimension reduction, a flexible mechanism for aggregated recovery, and a significant reduction in the number of unrecoverable entries when the observing rate is very low. It also allows a wide range of applications in signal processing, image analysis and many other fields, where data often appear to have the Kronecker product structure, which the classical low rank matrix completion methods may not be able to handle effectively. We also propose a MSE based criterion to determine the unknown configuration of the underlying Kronecker product.

There are a number of directions to further extend and explore the capability of the Kronecker product based matrix completions. First, it is of interest to devise a procedure/algorithm for the joint selection of the configuration and the K-rank. Second, to demonstrate the potential advantage of the aggregated estimation, we have used the equal weight of different configurations. A better aggregation can possibly be achieved by using weighted average, where the weights reflect the accuracy of different models, and can be determined by the cross validation. Last but not least, Similar structure can be introduced for tensors, leading to new tensor completion methods. We anticipate that the algorithm and the analysis will be substantially different, because it is known that tensor decomposition is very different from the matrix decomposition. We plan to follow out these ideas in the future work.

References

Abbe, E., Fan, J., Wang, K., and Zhong, Y. (2020). Entrywise eigenvector analysis of random matrices with low expected rank. *Annals of statistics*, 48(3):1452.

Ashraphijuo, M., Wang, X., and Aggarwal, V. (2017). Rank determination for low-rank data completion. *The Journal of Machine Learning Research*, 18(1):3422–3450.

Bennett, J., Lanning, S., et al. (2007). The netflix prize. In *Proceedings of KDD cup and workshop*, volume 2007, page 35. New York.

Biswas, P., Lian, T.-C., Wang, T.-C., and Ye, Y. (2006). Semidefinite programming based algorithms for sensor network localization. *ACM Transactions on Sensor Networks (TOSN)*, 2(2):188–220.

Buckland, S. T., Burnham, K. P., and Augustin, N. H. (1997). Model selection: An integral part of inference. *Biometrics*, 53(2):603–618.

Cai, C., Chen, R., and Xiao, H. (2019a). Hybrid Kronecker product decomposition and approximation. *arXiv preprint*. `https://arxiv.org/abs/1912.02955`.

Cai, C., Chen, R., and Xiao, H. (2019b). KoPA: Automated Kronecker product approximation. *arXiv preprint*. `https://arxiv.org/abs/1912.02392`.

Candès, E. J. and Plan, Y. (2010). Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936.

Candès, E. J. and Recht, B. (2009). Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717.

Candès, E. J. and Tao, T. (2010). The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080.

Cosse, A. and Demanet, L. (2017). Stable rank one matrix completion is solved by two rounds of semidefinite programming relaxation. *arXiv preprint arXiv:1801.00368*.

Duarte, M. F. and Baraniuk, R. G. (2012). Kronecker compressive sensing. *IEEE Transactions on Image Processing*, 21(2):494–504.

Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. (1992). Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–71.

Gunasekar, S., Acharya, A., Gaur, N., and Ghosh, J. (2013). Noisy matrix completion using alternating minimization. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 194–209. Springer.

Hardy, G. H., Wright, E. M., et al. (1979). *An introduction to the theory of numbers*. Oxford university press.

Jain, P., Netrapalli, P., and Sanghavi, S. (2013). Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 665–674. ACM.

Kamm, J. and Nagy, J. G. (1998). Kronecker product and SVD approximations in image restoration. *Linear Algebra and its Applications*, 284(1):177 – 192. International Linear Algebra Society (ILAS) Symposium on Fast Algorithms for Control, Signals and Image Processing.

Keshavan, R. H., Montanari, A., and Oh, S. (2010a). Matrix completion from a few entries. *IEEE transactions on information theory*, 56(6):2980–2998.

Keshavan, R. H., Montanari, A., and Oh, S. (2010b). Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11(Jul):2057–2078.

Raftery, A. E., Madigan, D., and Hoeting, J. A. (1997). Bayesian model averaging for linear regression models. *Journal of the American Statistical Association*, 92(437):179–191.

Schmidt, R. (1986). Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation*, 34(3):276–280.

Van Loan, C. F. and Pitsianis, N. (1993). Approximation with Kronecker products. In *Linear algebra for large scale and real-time applications*, pages 293–314. Springer.

Werner, K., Jansson, M., and Stoica, P. (2008). On estimation of covariance matrices with Kronecker product structure. *IEEE Transactions on Signal Processing*, 56(2):478–491.

Yuan, M. and Zhang, C.-H. (2016). On tensor completion via nuclear norm minimization. *Foundations of Computational Mathematics*, 16(4):1031–1068.

Appendix

We make the convention that $C, C_1, C_2, \ldots$ denote absolute constants, whose values may change from place to place.

A Proof of Theorem 1

We first prove several technical lemmas.

Lemma 2 (Over-representation)**.** *Let $\boldsymbol{Q}$ be a $m \times n$ matrix with IID Bernoulli random variables with success rate τ. Define*

$$\mathcal{A} = \left\{ i \in [m] : \sum_{j=1}^{n} Q_{ij} < 2\tau n \right\}.$$

Then we have

$$P[|\mathcal{A}| = m] \geqslant 1 - \exp\left\{ \log m - \frac{3n\tau}{8(1-\tau)} \right\}.$$

Proof. If $\tau > 1/2$, then $|\mathcal{A}| = m$ has probability one. Therefore, we only consider $\tau \leqslant 1/2$. Let $Z_{ij} = Q_{ij} - \tau$. We have $\mathbb{E}[Z_{ij}] = 0$, $\mathrm{Var}[Z_{ij}] = \tau(1-\tau)$ and $|Z_{ij}| \leqslant 1 - \tau$. By Bernstein's inequality, we have

$$P\left[\sum_{j=1}^{n} Z_{ij} \geqslant t \right] \leqslant \exp\left\{ -\frac{t^2/2}{n\tau(1-\tau) + (1-\tau)t/3} \right\}.$$

Therefore

$$P\left[\sum_{j=1}^{n} Q_{ij} \geqslant 2\tau n \right] \leqslant \exp\left\{ -\frac{3n\tau}{8(1-\tau)} \right\}.$$

Using union bound, we have

$$P\left[\max_{i \in [m]} \sum_{j=1}^{n} Q_{ij} \geqslant 2\tau n \right] \leqslant \exp\left\{ \ln m - \frac{3n\tau}{8(1-\tau)} \right\}.$$

□

The matrix $\boldsymbol{Q}$ is *row over-represented* if $|\mathcal{A}| < m$. Similarly, a matrix $P_\Omega \boldsymbol{M}$ is over-represented if either $P_\Omega \boldsymbol{M}$ is row over-represented or $[P_\Omega \boldsymbol{M}]^T$ is row over-represented. We restate the theorem on $\|P_\Omega \boldsymbol{M}\|_S$ from Keshavan et al. (2010a) in Lemma 3. Recall that $\|\cdot\|_{\max}$ denotes the maximum absolute entry of a matrix.

Lemma 3 (Keshavan et al. (2010a)). *Suppose $\boldsymbol{M}$ is a $m \times n$ $(m \geqslant n)$ matrix with elements observed IID with probability τ. Denote the partially observed matrix as $P_\Omega \boldsymbol{M}$. If $P_\Omega \boldsymbol{M}$ is not over-represented, then there exists a constant $C > 0$ such that, with probability larger than $1 - 1/n^3$*

$$\frac{|\|P_\Omega \boldsymbol{M}\|_S - \tau \|\boldsymbol{M}\|_S|}{\sqrt{\tau mn}\|\boldsymbol{M}\|_{\max}} \leqslant C m^{1/4} n^{-3/4}.$$

Lemma 4 (Max bound on Gaussian matrix). *If $\boldsymbol{E}$ is a $m \times n$ matrix with IID standard Gaussian random variables, then there exists a constant $C > 0$ such that, with probability larger than $1 - 1/(mn)^3$,*

$$\|\boldsymbol{E}\|_{\max} \leqslant C\sqrt{\log mn}.$$

Proof. Since E_{ij} is standard Gaussian, we have $P[|E_{ij}| > t] \leqslant 2\exp\{-t^2/2\}$. Using union bound, we have

$$P[\|\boldsymbol{E}\|_{\max} > t] \leqslant \exp\{\log(2mn) - t^2/2\}.$$

The lemma follows immediately by choosing $C > 2\sqrt{2}$. □

From Lemma 2, we have that for any configuration $(p, q) \in \mathcal{C}_\delta$,

$$\begin{aligned} P[\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}] \text{ is over-represented}] &\leqslant 2\exp\left\{\left(\frac{3}{4} - \delta\right)\log PQ - \frac{3\tau}{8(1-\tau)}(PQ)^{1/4+\delta}\right\} \\ &\leqslant \exp\left\{-C_1\tau(PQ)^{1/4+\delta}\right\} \end{aligned}$$

for some constant C_1. Therefore, from Lemma 3, for $(p, q) \in \mathcal{C}_\delta$, with probability larger than $1 - (PQ)^{-3/4-3\delta} - \exp\{-C_1\tau(PQ)^{1/4+\delta}\}$ we have

$$|\|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{X}]\|_S - \tau\|\mathcal{R}_{p,q}[\boldsymbol{X}]\|_S| \leqslant C_2\sqrt{\tau} \cdot \sqrt{PQ}\|\boldsymbol{X}\|_{\max} \cdot PQ^{1/4} \cdot (pq \wedge p^*q^*)^{-1}, \tag{22}$$

for some constant C_2, where $\boldsymbol{X} = \lambda \boldsymbol{A} \otimes \boldsymbol{B}$ is the signal part. Noticing that $\|\boldsymbol{X}\|_{\max} = \lambda\|\boldsymbol{A}\|_{\max}\|\boldsymbol{B}\|_{\max}$ and $pq \wedge p^*q^* \geqslant (PQ)^{1/4+\delta}$, (22) can be further revised to

$$|\|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{X}]\|_S - \tau\|\mathcal{R}_{p,q}[\boldsymbol{X}]\|_S| \leqslant C_2\lambda\sqrt{\tau}\mu^2 \cdot (PQ)^{-\delta}. \tag{23}$$

Similarly, for the noise part, according to Lemma 3 and Lemma 4, for $(p, q) \in \mathcal{C}_\delta$ with probability larger than $1 - (PQ)^{-3/4-3\delta} - (PQ)^{-3}$

$$\|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{E}]\|_S \leqslant \tau\|\mathcal{R}_{p,q}[\boldsymbol{E}]\|_S + C_3\sqrt{\tau}\sqrt{\log PQ} \cdot (PQ)^{1/2-\delta}, \tag{24}$$

for some constant C_3.

For the true configuration (p_0, q_0), using (23) and (24), we have

$$\begin{aligned}
&\mathbb{E}[\|\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}]\|_S] \\
&\geqslant \mathbb{E}[\|\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{X}]\|_S] - \sigma(PQ)^{-1/2}\mathbb{E}[\|\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{E}]\|_S] \\
&\geqslant \left(1 - (PQ)^{-3/4-3\delta} - \exp\{-C_1\tau(PQ)^{1/4+\delta}\}\right)\left(\tau\lambda - C_2\lambda\sqrt{\tau}\mu^2(PQ)^{-\delta}\right) \\
&+ \left((PQ)^{-3/4-3\delta} + \exp\{-C_1\tau(PQ)^{1/4+\delta}\right) \cdot 0 \\
&- \sigma(PQ)^{-1/2}\left(1 - (PQ)^{-3/4-3\delta} - (PQ)^{-3}\right)\left(\tau(\sqrt{p_0q_0} + \sqrt{p_0^*q_0^*}) + C_3\sqrt{\tau}\sqrt{\log PQ}(PQ)^{1/2-\delta}\right) \\
&- \sigma(PQ)^{-1/2}\left((PQ)^{-3/4-3\delta} + (PQ)^{-3}\right)\sqrt{\tau PQ} \\
&\geqslant \tau\lambda - O\left(\left(\lambda\mu^2 + \sigma\sqrt{\log PQ}\right)\sqrt{\tau}(PQ)^{-\delta}\right)
\end{aligned}$$

Similarly, for any $(p, q) \in \mathcal{W}_\delta$, we have

$$\mathbb{E}[\|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]\|_S] \leqslant \tau\lambda\phi + O\left(\left(\lambda\mu^2 + \sigma\sqrt{\log PQ}\right)\sqrt{\tau}(PQ)^{-\delta}\right).$$

Therefore,

$$\begin{aligned}
&\mathbb{E}[\|\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}]\|_S] - \max_{(p,q)\in\mathcal{W}_\delta} \mathbb{E}[\|\mathcal{R}_{p,q}[P_\Omega \boldsymbol{Y}]\|_S] \\
&\qquad\qquad \geqslant \tau\lambda(1-\phi)\cdot\left(1 + O\left(\psi^{-2}\left(\mu^2 + \frac{\sigma}{\lambda}\sqrt{\log PQ}\right)\tau^{-1/2}(PQ)^{-\delta}\right)\right).
\end{aligned}$$

Here we use $\psi^2 = 1 - \phi^2 \asymp 1 - \phi$. Assumption 4 ensures the term in big O notation is minor.

B Proof of Theorem 2

For the true configuration (p_0, q_0), using (23) and (24), we have

$$\begin{aligned}
&P[G(p_0, q_0) \leqslant \tau\lambda(1+\phi)/2] \\
\leqslant &P[\|\mathcal{R}_{p_0,q_0}[\boldsymbol{E}]\|_S \geqslant \sqrt{p_0q_0} + \sqrt{p_0^*q_0^*} + R_1] + (PQ)^{-3/4-3\delta} + \exp\{-C_1\tau(PQ)^{1/4+\delta}\} + (PQ)^{-3} \\
\leqslant &\exp\{-R_1^2/2\} + 2(PQ)^{-3/4-3\delta} + \exp\{-C_1\tau(PQ)^{1/4+\delta}\} + (PQ)^{-3} \\
\leqslant &C_4(PQ)^{-3/4-3\delta}
\end{aligned}$$

for some constant C_4, where

$$\begin{aligned}
R_1 &= (PQ)^{1/2}\left[\frac{\lambda(1-\phi)}{2\sigma} - (PQ)^{-\delta}\tau^{-1/2}\left(C_2\frac{\lambda}{\sigma}\mu^2 + C_3\sqrt{\log PQ}\right)\right] - \sqrt{p_0q_0} - \sqrt{p_0^*q_0^*} \\
&= O\left((PQ)^{1/2}\frac{\lambda(1-\phi)}{2\sigma}\right).
\end{aligned}$$

Similarly, for a wrong configuration $(p,q) \in \mathcal{W}_\delta$, using (23) and (24), we have

$$\begin{aligned}
&P[G(p,q) \geqslant \tau\lambda(1+\phi)/2] \\
\leqslant &P[\|\mathcal{R}_{p,q}[\boldsymbol{E}]\|_S \geqslant \sqrt{pq} + \sqrt{p^*q^*} + R_1] + (PQ)^{-3/4-3\delta} + \exp\{-C_1\tau(PQ)^{1/4+\delta}\} + (PQ)^{-3} \\
\leqslant &\exp\{-R_1^2/2\} + 2(PQ)^{-3/4-3\delta} + \exp\{-C_1\tau(PQ)^{1/4+\delta}\} + (PQ)^{-3} \\
\leqslant &C_4(PQ)^{-3/4-3\delta}.
\end{aligned}$$

Therefore, for any $(p,q) \in \mathcal{W}_\delta$,

$$P[G(p_0,q_0) \leqslant G(p,q)] \leqslant P[G(p_0,q_0) \leqslant \tau\lambda(1+\phi)/2] + P[G(p,q) \geqslant \tau\lambda(1+\phi)/2] \leqslant 2C_4(PQ)^{-3/4-3\delta}.$$

Using the union bound we have

$$P[G(p_0,q_0) \leqslant \max_{(p,q)\in\mathcal{W}_\delta} G(p,q)] \leqslant 2|\mathcal{W}_\delta|C_4(PQ)^{-3/4} - 3\delta,$$

where $|\mathcal{W}_\delta|$ is the number of wrong configurations. Applying $|\mathcal{W}_\delta| < |d(P)||d(Q)|$ yields the theorem.

C Proof of Lemma 1

All entries of $P_\Omega \boldsymbol{Y}$ is equivalent to $\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}]$ has no missing column or missing row. Noticing that $\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}]$ is a $p_0q_0 \times p_0^*q_0^*$ matrix with entries observed IID with probability τ, we have

$$P[\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{Y}] \text{ has no missing column}] \geqslant [1-(1-\tau)^{p_0q_0}]^{p_0^*q_0^*} \geqslant 1 - PQ\frac{(1-\tau)^{p_0q_0}}{p_0q_0}.$$

The right hand side is a increasing function of p_0q_0 and τ. From Assumption 1 we know $p_0q_0 \geqslant (PQ)^{1/4+\delta}$ and Assumption 3 gives $\tau > \log PQ \cdot (PQ)^{-2\delta}$. Therefore,

$$\frac{(1-\tau)^{p_0q_0}}{p_0q_0} = \exp\{p_0q_0\log(1-\tau) - \log p_0q_0\} \leqslant \exp\left\{-(PQ)^{1/4-\delta}\log PQ - \left(\frac{1}{4}+\delta\right)\log PQ\right\}.$$

Using a similar argument for missing row, the Lemma follows immediately.

D Proof of Theorem 3

We first restate the major result of Gunasekar et al. (2013) for rank-1 case in the following lemma.

Lemma 5. *Let $\boldsymbol{M} = \lambda uv^T$ be a rank-1 $m \times n$ $(m \leqslant n)$, incoherent matrix with both u and v being μ incoherent. Furthermore, it is assumed that the noise matrix $\boldsymbol{N}$ satisfies $\|P_\Omega \boldsymbol{N}\|_S/\tau \leqslant C\lambda$ for some constant C. Additionally, let each entry of $\tilde{\boldsymbol{M}} = \boldsymbol{M} + \boldsymbol{N}$ be observed IID with probability*

$$\tau > C\frac{\mu^4 \log n \log\frac{\lambda}{\epsilon}}{m},$$

where $C > 0$ is a global constant. Then with high probability, the output matrix $\hat{\boldsymbol{M}}$ from Algorithm 1 satisfies

$$\frac{1}{\sqrt{mn}}\|\boldsymbol{M} - \hat{\boldsymbol{M}}\|_F \leqslant \epsilon + 20\mu\frac{\|P_\Omega \boldsymbol{N}\|_S}{|\Omega|}.$$

Apply Lemma 5 directly and we have with high probability

$$\|\boldsymbol{X} - \hat{\boldsymbol{X}}\|_F \leqslant \lambda\sqrt{PQ}\exp\left\{-C\frac{\tau(p_0q_0 \wedge p_0^*q_0^*)}{\mu^4 \log PQ}\right\} + 20\mu\sigma\frac{\|\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{E}]\|_S}{|\Omega|}.$$

From Lemma 3, we have with high probability,

$$\|\mathcal{R}_{p_0,q_0}[P_\Omega \boldsymbol{E}]\|_S \leqslant C_1\sqrt{\tau}\sqrt{\log PQ}\cdot(PQ)^{3/4}(p_0q_0 \wedge p_0^*q_0^*)^{-1},$$

for some constant C_1, and with high probability

$$|\Omega| \geqslant C_2\tau PQ$$

for some constant $0 < C_2 < 1$. Therefore, with high probability,

$$\|\boldsymbol{X} - \hat{\boldsymbol{X}}\|_F \leqslant \lambda\sqrt{PQ}\exp\left\{-C\frac{\tau(p_0q_0 \wedge p_0^*q_0^*)}{\mu^4 \log PQ}\right\} + \frac{20C_1}{C_2}\mu\sigma\tau^{-1/2}\sqrt{\log PQ}\cdot(PQ)^{-1/4}(p_0q_0 \wedge p_0^*q_0^*)^{-1}.$$