

On the Optimality of Nuclear-norm-based Matrix Completion for Problems with Smooth Non-linear Structure

Yunhua Xiang*

*Department of Biostatistics
University of Washington
Seattle, WA 98195-4322, USA*

XIANGYH@UW.EDU

Tianyu Zhang*

*Department of Biostatistics
University of Washington
Seattle, WA 98195-4322, USA*

ZTY@UW.EDU

Xu Wang

*Department of Biostatistics
University of Washington
Seattle, WA 98195-4322, USA*

WANGXU85@UW.EDU

Ali Shojaie

*Department of Biostatistics
University of Washington
Seattle, WA 98195-4322, USA*

ASHOJAIE@UW.EDU

Noah Simon

*Department of Biostatistics
University of Washington
Seattle, WA 98195-4322, USA*

NRSIMON@UW.EDU

Editor: Jie Peng

Abstract

Nuclear-norm-based matrix completion was originally developed for imputing missing entries in low rank, or approximately low rank matrices. However, it has proven widely effective in many problems where there is no reason to assume low-dimensional linear structure in the underlying matrix, as would be imposed by rank constraints. In this manuscript we show that nuclear-norm-based matrix completion attains within a log factor of the minimax rate for estimating the mean structure of matrices that are not necessarily low-rank, but lie in a low-dimensional non-linear manifold, when observations are missing completely at random. In particular, we give upper bounds on the rate of convergence as a function of the number of rows, columns, and observed entries in the matrix, as well as the smoothness and dimension of the non-linear embedding. We additionally give a minimax lower bound: This lower bound agrees with our upper bound (up to a logarithmic factor), which shows that nuclear-norm penalization is (up to log terms) minimax rate optimal for these problems.

Keywords: matrix completion, nonlinear low-rank structure, nuclear-norm penalization, minimax lower bound, function class embedding

*These authors contributed equally.

1. Introduction

Matrix completion is a framework that has gained popularity in a wide range of machine learning applications, including recommender systems (Koren et al., 2009), system identification (Liu and Vandenberghe, 2010), global positioning (Singer and Cucuringu, 2010) and natural language processing (Wijaya et al., 2017). It is a useful framework for complex prediction problems, where each observation comes with a heterogeneous collection of observed features. In particular, matrix completion is applied to problems where the object of inference or prediction is a matrix whose rows correspond to observation and columns to variables/features. In many applications, only a subset of entries in this matrix are observed (often with noise), and the goal is to “complete” the matrix, filling in estimates of the unobserved entries. This “completion” is done by leveraging the known structure in the matrix. The most famous example, which brought matrix completion to prominence, is the Netflix Challenge (Koren et al., 2009), where a small sample of observed ratings for each customer was used to successfully predict future/unobserved movie ratings for Netflix customers.

More formally, suppose we have an underlying unobserved matrix $M \in \mathbb{R}^{n \times p}$: We then observe a subset of the entries from the noise-contaminated matrix $Y = M + E$, where E is a matrix of i.i.d. mean zero, finite variance noise variables. Our goal is to recover matrix M from this partially observed, noisy Y . This is known as matrix completion. Without any structure on the matrix M , recovering the values of M corresponding to unobserved entries is impossible (Laurent, 2001). Matrix completion becomes possible if one imposes some constraints on the structure of the underlying matrix: It is most common to assume that M is low rank. Directly employing this assumption by e.g., finding the minimum rank completion of Y (or corresponding rank-constrained regression) is unfortunately NP-hard and becomes computationally infeasible for problems involving large matrices (Candes and Tao, 2010; Chistov and Grigoriev, 1984). Over the last decades, computationally efficient methods using convex optimization have been developed for recovering a low rank matrix from a small number of observations with near-optimal statistical guarantees in primarily noiseless problems (Srebro et al., 2004; Recht, 2011; Candes and Tao, 2010; Recht et al., 2010), and when the observed entries are contaminated with noise (Candes and Plan, 2010a; Koltchinskii et al., 2011). These methods rely on using the nuclear norm of the matrix (Fazel, 2002; Jaggi and Sulovskỳ, 2010), that is, sum of its singular values, as a convex surrogate for the matrix rank. The low-rank structure leveraged in matrix completion can be thought of as learning a linear embedding of the data in a low-dimensional space.

In practice, the underlying matrix M may not be low rank. However, we often believe it may still have useful low-dimensional structure. It has thus become popular to learn a low-dimensional non-linear embedding of the data. This idea is used both in matrix completion and more generally for low-dimensional summaries of data. It has been applied in motion recovery (Xia et al., 2018), epigenomics (Schreiber et al., 2018), and health data analytics (Wang et al., 2015) among other areas. To recover these embeddings, Reproducing Kernel Hilbert Space (RKHS) methods (Fan and Chow, 2018), nearest neighbor methods (Li et al., 2019), and deep learning methods like autoencoders and neural-network-based variational frameworks (Fan and Cheng, 2018; Yu et al., 2013; Jiang et al., 2016) have been used.

A related literature has also discussed so-called nonlinear latent variable models. These models are often considered as the estimation targets of neural network autoencoders (Fan and Chow, 2017), and are an alternative way to frame non-linear low dimensional embedding structure. They are becoming particularly popular in biological/biomedical applications, as they allow for non-linearity and interactions between dimensions of the latent features: This is known to be a common feature of biological systems (Srinivasan, 2022; McKinney et al., 2006). Autoencoder neural networks are usually associated with two mappings (the encoder and decoder) between the higher dimension observable space (the $\mu_i \in \mathbb{R}^p$) and the much lower dimension latent feature space (the $\theta_i \in \mathbb{R}^K$). In practice, the networks are structured to estimate highly non-linear mappings. In Fan et al. (2021), the authors construct autoencoders to impute missing histological data in the context of breast cancer prognosis; in Wang et al. (2021), autoencoders are applied to various tasks in single-cell RNA sequencing analysis, including imputation of missing gene expression values. A recent work, Linderman et al. (2022), proposes a low-rank matrix factorization method for single cell RNA sequencing data imputation. The authors empirically show that their method can better improve the cluster separation of different cell types, compared with other seemingly more intricate methods (which is in-line with our theory).

Such latent variable models—large p , small K —have also gained substantial attention in biomedical research beyond missing data imputation / matrix completion. In Wang et al. (2016), non-linear latent models are applied to DNA methylation level prediction. In Greener et al. (2018) they are used for protein structure design (Curiously, the dimension of latent variable layer are set to be as low as 16 or 2). In Zhou et al. (2022), the authors consider the influence of CRISPR perturbation on gene-expression profile. They make the assumption that the perturbation would first influence a few hidden factors, which in turn spread the impact over the much larger collection of genes of interest. In fact, the increasingly popular t-distributed stochastic neighbor embedding (t-SNE) method (Van der Maaten and Hinton, 2008) for visualization directly assumes this non-linear latent variable model. In that original manuscript, the authors show that $K = 2$ or 3 latent dimensions can be sufficient for summarizing the most impactful structure in high dimensional data. For a more extensive review, see the Table 1 of Kopf and Claassen (2021) for over 60 applications of (non-linear) latent models that have mostly been published within the past 5 years.

In parallel to the realization that many important problems have low dimensional non-linear structure, there has been growing empirical evidence that matrix completion methods based on nuclear norm penalization perform well even in such scenarios where any low dimensional structure is likely non-linear (Linderman et al., 2022; Jin et al., 2022). In Fan et al. (2021), the authors applied various methods to impute missing data for breast cancer prediction. They showed neural network based latent models can achieve better prediction accuracy (AUC = 0.793), but methods applying matrix factorization can still explain a significant proportion (AUC = 0.756). And in certain settings matrix factorization methods can actually have better performance (as reported in their Section 3.1). As these methods were developed for linear low rank structure, this is, at first glance, a bit surprising. There has been some work giving theoretical justification for these empirical results (Chatterjee et al., 2015; Udell and Townsend, 2019). In particular, they note that in the presence of some types of non-linear low-dimensional structure in M , nuclear norm-based matrix completion methods can still consistently estimate M . These work additionally gives some

non-stochastic approximation error results. However, optimality of the statistical perform of nuclear-norm-based matrix completion is not considered to the best of our knowledge.

In this manuscript, we delve further into the performance of matrix completion for M with low-dimensional, non-linear structure. In particular, we consider M with rows that can be embedded in a low-dimensional smooth manifold (which directly engages with non-linear latent variable models). We then (i) show that nuclear norm-based matrix completion can consistently estimate M ; (ii) characterize the rate at which the reconstruction error converges to 0 as a function of the size of the matrix, number of observed entries, and smoothness and dimension of the underlying manifold; and (iii) prove that, up to a log term, this rate cannot be improved upon by any method; that is, our upper bound is actually the minimax rate [up to a log factor] for reconstruction error in this problem. Furthermore, our error bounds (and our techniques) also relate the matrix completion problem to more classical non-parametric estimation: Our reconstruction error bounds parallel the minimax rate for mean squared error (MSE) in the nonparametric regression setting. Results (ii) and (iii), we believe, are novel.

Our experiments on synthetic data corroborate our theoretical findings. In particular, they suggest that the finite sample empirical performance of matrix completion in non-linear low rank embeddings is consistent with the asymptotic theoretical error bounds. These empirical results also corroborate the claim that better performance is achieved when the embedding of the underlying matrix M lies in a smoother manifold.

2. Methods

2.1 Problem setup

We start by giving some notation. We use upper case letters to represent matrices and lower case letters to represent scalars. The trace inner product of any two matrices, $M, B \in \mathbb{R}^{n \times p}$, $n, p \in \mathbb{Z}^+$, is $\langle M, B \rangle = \text{tr}(M^T B)$. The element-wise infinity norm of $M \in \mathbb{R}^{n \times p}$ is defined by $\|M\|_\infty = \max_{1 \leq i \leq n, 1 \leq j \leq p} |m_{ij}|$ where m_{ij} denotes the (i, j) -th entry of M . We also denote the Frobenius norm of matrix M as $\|M\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^p m_{ij}^2}$.

In the general matrix completion problem, we randomly observe some of the entries from a matrix $M \in \mathbb{R}^{n \times p}$; the observed entries may also be contaminated with error. To support our later theoretical derivations, we will describe this process in terms of a set of mask matrices $X_t \in \mathbb{R}^{n \times p}$ and observed values $y_t \in \mathbb{R}$. Each X_t is a matrix with a single 1 whose position is indexed by t and all other entries are equal to 0 as follows:

$$X_t = \begin{pmatrix} 0 & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & & & \vdots & & \vdots \\ 0 & 0 & \cdots & 1 & \cdots & 0 \\ \vdots & & & \vdots & & \vdots \\ 0 & 0 & \cdots & 0 & \cdots & 0 \end{pmatrix}_{n \times p} . \tag{1}$$

The collection of matrices X_t fall in the set $\mathcal{X} = \{e_n(i)e_p(j)^T, \text{ for all } i = 1, \dots, n \text{ and } j = 1, \dots, p\}$, where $e_n(i) \in \mathbb{R}^n$ is the basis vector consisting of all zeros except for a single 1

at i th entry. In this formulation, X_t indicates the location in M where y_t is drawn from. That is, for $X_t = e_n(i)e_p(j)^T \in \mathcal{X}$, $\langle X_t, M \rangle = m_{ij}$.

Now, we can frame the matrix completion problem as follows: Suppose we have N pairs of observations (X_t, y_t) , $t = 1, \dots, N$, that satisfy

$$y_t = \langle X_t, M \rangle + \xi_t, \tag{2}$$

where ξ_t are i.i.d random errors distributed $N(0, \sigma^2)$, $M \in \mathbb{R}^{n \times p}$ is the underlying true matrix to be recovered, and $y_t \in \mathbb{R}$ are observed values. The observed matrix can be written as $Y = \sum_{t=1}^N y_t X_t$ where N is the number of observed entries. We assume that X_t is uniformly sampled at random from \mathcal{X} (Koltchinskii et al., 2011), that is, $X_t \sim \Pi$, and the probability that the (i, j) th entry of X_t equals to 1 is $\pi_{ij} = P(X_t = e_i(n)e_j(p)^T) = \frac{1}{np}$ for $1 \leq i \leq n, 1 \leq j \leq p$. This is essentially a missing completely at random (MCAR) assumption.

The goal is to recover M given pairs (X_t, y_t) , $t = 1, 2, \dots, N$, and we are generally interested in the setting where $N \ll np$. To solve this problem, existing methods often assume that M has low rank (or approximately low rank), that is, $M \simeq UV^T$ with $U \in \mathbb{R}^{n \times r}$ and $V \in \mathbb{R}^{p \times r}$ for some integer $r \ll \min(n, p)$. In contrast to this low rank assumption, this paper studies the problem where M is not necessarily low-rank but generated from a low-dimensional non-linear manifold. This notion is formalized in the next section.

2.2 Non-linearly Embeddable Matrices

In this section, we will formalize the previously mentioned “low-dimensional non-linear structure”. We will briefly start with the classical linear low-rank matrix completion setting and motivate ours by switching from a class of linear functions to nonlinear functions. Let $M \in \mathbb{R}^{n \times p}$ denote the true matrix. In the low-rank matrix completion setting we may assume the rank of it $r = \text{rank}(M)$ is much smaller than the size $\min(n, p)$. In this case, we can factor the true matrix M into a product of two matrices (of smaller sizes). Formally, with some $\Theta \in \mathbb{R}^{n \times r}$, $B \in \mathbb{R}^{r \times p}$, we have $M = \Theta B$. Equivalently, for each entry of M we have:

$$m_{ij} = \theta_{i,\cdot} \beta_{\cdot,j} \quad i = 1, \dots, n, j = 1, \dots, p \tag{3}$$

where m_{ij} is the (i, j) entry of M , $\theta_{i,\cdot}$ is the i -th row of Θ and $\beta_{\cdot,j}$ is the j -th column of B . This means each entry in the i -th row of M can be written as a linear function of the i -th row of Θ . In addition, there are only p linear functions involved, each one of them is related to one column of B . So we say Θ is a linear embedding of M into a lower-dimensional space.

If we replace the linear functions in (3) by non-linear functions, we will formally arrive at the concept of non-linear embeddable matrix. This generalization strongly resembles the transition from linear regression to non-linear (non-parametric) regression. Consider a matrix M , a positive integer K , and a function class $\mathcal{F} \subset \mathcal{L}^2(\mathbb{R}^K)$. We say M is \mathcal{F} -embeddable if there exist functions $f_j \in \mathcal{F} : \mathbb{R}^K \rightarrow \mathbb{R}$, $j = 1, \dots, p$, and a matrix $\Theta \in \mathbb{R}^{n \times K}$ such that

$$m_{ij} = f_j(\theta_{i,\cdot}), i = 1, \dots, n, j = 1, \dots, p, \tag{4}$$

where m_{ij} is the (i, j) entry of M and $\Theta \in \mathbb{R}^{n \times K}$ is a matrix (with θ_i indicating its i th row vector). Here, Θ gives a non-linear embedding of our observations from its original p -dimensional space into a K -dimensional space ($K \leq p$). The set of functions $\{f_j\}_{j=1}^p \subset \mathcal{F}$ identifies how to map our embedding in \mathbb{R}^K back to \mathbb{R}^p . From the perspective of latent variable models, the column dimension of Θ can be interpreted as the number of latent variables, and that of M is the number of manifest variables. The non-linearly embeddable assumption is saying that the (systematic) variability of the much higher-dimension manifest variables can be described well by the much lower-dimension latent variables. As we will see later (Theorem 5), the number of latent variables and proper smoothness of the f_j mappings ensures the possibility of estimating the matrix M under such a setting. This is curiously related to the low-dimensional bottleneck layer of neural network autoencoders (Fan and Chow, 2017; Jin et al., 2022), which is also often interpreted as the limited dimension of the latent variables.

2.3 Nuclear-norm Penalized Estimators

In classical matrix completion setting, where we assume M is low-rank, nuclear norm penalized empirical risk minimization is often used to estimate M (Argyriou et al., 2008; Candès and Plan, 2010b; Negahban and Wainwright, 2011); more specifically, the estimator is obtained by,

$$\arg \min_M \left\{ N^{-1} \sum_{t=1}^N (y_t - \langle X_t, M \rangle)^2 + \lambda \|M\|_* \right\}, \quad (5)$$

where λ is a regularization parameter which is used to balance the trade-off between fitting the unknown matrix using least squares and minimizing the nuclear norm $\|M\|_*$. This “matrix lasso” is known to have strong theoretical properties when M is low rank (Argyriou et al., 2008; Candès and Plan, 2010b; Negahban and Wainwright, 2011). However, in our scenario, M likely does not have low rank and previous work does not fully explain the effectiveness of the estimate from (5) in this setting.

While the estimator in (5) is simple and quite well known, it fails to exploit knowledge of the sampling scheme (which is often known or at least assumed to be known). To use the assumption that the mask matrices $\{X_t\}_{t=1}^N$ are i.i.d. uniformly sampled from \mathcal{X} , we study a slight modification to (5) described in Koltchinskii et al. (2011):

$$\widehat{M} \leftarrow \arg \min_M \left\{ \frac{1}{np} \|M\|_F^2 - \left\langle \frac{2}{N} \sum_{t=1}^N y_t X_t, M \right\rangle + \lambda \|M\|_* \right\} \quad (6)$$

After some simple manipulation, (6) can be further reduced to minimizing

$$\frac{1}{np} \|M - R\|_F^2 + \lambda \|M\|_*$$

where $R = \frac{np}{N} \sum_{t=1}^N y_t X_t = \frac{np}{N} Y$. Thus, \widehat{M} , the solution to (6), is merely a singular-value soft-thresholding estimator:

$$\widehat{M} = \sum_{j=1}^{\text{rank}(R)} (\Lambda_j(R) - \lambda np/2)_+ u_j(R) v_j(R)^T, \quad (7)$$

where $\Lambda_j(R)$ are the singular values and $u_j(R), v_j(R)$ are the left and right singular vectors of R such that $R = \sum_{j=1}^{\text{rank}(R)} \Lambda_j(R) u_j(R) v_j(R)^T$. Koltchinskii et al. (2011) established the rate optimality of this estimator with respect to Frobenius-norm loss when M is low rank. In this paper, we aim to ultimately claim that \widehat{M} in (6) is a rate optimal estimator of M in the case that M is non-linearly embeddable, as long as K is fixed and the function class \mathcal{F} is sufficiently smooth.

2.4 Approximation of Embeddable Matrices

Our goal is to characterize the convergence rate of estimator obtained by (6) for the estimating true underlying matrix M with respect to Frobenius-norm loss, when M is non-linearly embeddable. To this end, we first show that M can be well approximated by a series of matrices with low (and only slowly growing) rank as long as the function class \mathcal{F} is sufficiently smooth. More specifically, we will need the following condition for the function class \mathcal{F} .

Condition 1 *Given a function class \mathcal{F} , let C_0 denote a fixed positive number. Suppose that for any $\epsilon > 0$, there exists a finite set of measurable functions $\mathcal{F}_\epsilon = \{\psi_1, \psi_2, \dots, \psi_{J(\epsilon)}\}$, such that*

$$\|\psi\|_\infty \leq C_0, \quad \text{for all } \psi \in \mathcal{F}_\epsilon, \quad (8)$$

and

$$\max_{f \in \mathcal{F}} \min_{\|\beta\|_\infty \leq C_0} \left\| f - \sum_{l=1}^{J(\epsilon)} \beta_l \psi_l \right\|_\infty \leq \epsilon. \quad (9)$$

For each ϵ , we denote by \mathcal{F}_ϵ^* a set of minimal cardinality such that (8) and (9) hold. We let $J^*(\epsilon)$ denote the cardinality of \mathcal{F}_ϵ^* .

For a function class \mathcal{F} , Condition 1 characterizes the minimal number of basis functions needed to uniformly approximate functions in \mathcal{F} up to precision ϵ . In Section 3, we shall apply this condition to K -dimensional, L -th order differentiable functions, and show how this number scales as a function of ϵ .

Based on the above condition, we can establish the existence of an approximation matrix which is sufficiently close to the true matrix M and has a bounded nuclear norm.

Lemma 1 *Suppose matrix $M \in \mathbb{R}^{n \times p}$ is \mathcal{F} -embeddable, and \mathcal{F} satisfies Condition 1. Then, for any $\epsilon > 0$, there exists a matrix M^ϵ satisfying $\text{rank}(M^\epsilon) = J^*(\epsilon) \leq \min(n, p)$ such that*

$$\|M^\epsilon - M\|_\infty \leq \epsilon. \quad (10)$$

Furthermore, the nuclear norm of M^ϵ is bounded: There exists $C_1 > 0$ (independent of ϵ) such that

$$\frac{1}{\sqrt{np}} \|M^\epsilon\|_* \leq C_1 J^*(\epsilon). \quad (11)$$

The proof is given in Appendix A. Note, for the \mathcal{F} we consider later (restricted to smooth functions) we will show that $J^*(\epsilon) \ll \min(n, p)$. This parallels results in classical non-parametric regression where many function-spaces considered can be approximated uniformly with small error by linear combinations of relatively few basis functions (Tsybakov, 2009).

3. Convergence Rate

Using Lemma 1, it is relatively straightforward to evaluate the performance of our estimator \widehat{M} in (6). The performance metric simplest to theoretically analyze is $N^{-1} \sum_{i=1}^N \langle X_i, \widehat{M} - M \rangle^2$. However, this criterion only evaluates the prediction error on the *observed* entries. This is unsatisfying as our ultimate goal is to recover the entire matrix. Thus, we instead aim to evaluate the performance of \widehat{M} based on the metric $\frac{1}{np} \|\widehat{M} - M\|_F^2$. The following result gives an upper bound for the performance of our estimator \widehat{M} in this metric.

Theorem 2 *Suppose we observe N pairs $\{(y_t, X_t)\}_{t=1}^N$ satisfying data generating model (2) where X_t are i.i.d. uniformly sampled from \mathcal{X} . Assume the true matrix $M \in \mathbb{R}^{n \times p}$ is \mathcal{F} -embeddable where \mathcal{F} satisfies Condition 1. Further suppose that $N \geq (n \wedge p) \log^2(n + p)$. Then there exists a constant $C_2 > 0$ (that only depends on σ and $\|M\|_\infty$) such that if we define the regularization parameter λ by*

$$\lambda = C_2 \sqrt{\frac{\log(n + p)}{N(n \wedge p)}},$$

then, with probability at least $1 - 2(n + p)^{-1}$, the completion error of \widehat{M} in (7) is bounded by

$$\frac{1}{np} \|\widehat{M} - M\|_F^2 \leq C_2^2 \left(\frac{1 + \sqrt{2}}{2} \right)^2 \frac{(n \vee p) \log(n + p)}{N} J^*(\epsilon) + \epsilon^2, \quad (12)$$

for any $\epsilon > 0$. Here, $J^(\epsilon)$ is the rank of the approximation matrix M^ϵ with $\|M - M^\epsilon\|_\infty \leq \epsilon$, which corresponds to the minimal cardinality of \mathcal{F}^* satisfying Condition 1.*

After establishing Lemma 1, the upper bounds in Theorem 2 would follow directly using existing arguments (Koltchinskii et al., 2011). For completeness we give full details of the proof in the Appendix B. The two terms on the right-hand-side of (12) clarify the trade-off between the approximation error, ϵ , and the cardinality of the minimal linear approximation set \mathcal{F}^* , $J^*(\epsilon)$. As in the results of Koltchinskii et al. (2011), the error in our bound is decomposed into a misspecification error (ϵ^2) and a prediction error. Usually, when there is no misspecification, that is, the true matrix M is low rank, the prediction error is linearly related to the rank of M (Candes and Plan, 2011; Klopp et al., 2014). In our scenario, where the low-rank assumption is violated, the prediction error in (12) is linearly related to the rank of the approximation matrix.

Ideas similar to this occur in more traditional non-parametric estimation problems. For example, when using projection estimators in Hölder and Sobolev spaces, one of the main rate-optimal estimation approaches requires a truncated basis to be selected for projection that will grow with the sample size N (Tsybakov, 2008). However, in those examples, the number of basis vectors is a tuning parameter in the algorithm, and the set of basis functions must be selected in advance. Here, both the set of basis functions and the truncation level are rather just theoretical tools for analyzing the algorithm performance. In employing matrix completion, the analyst only needs to select λ .

We note that $N \geq (n \wedge p) \log^2(n + p)$ in the above Theorem 2 is a quite weak condition on the number of observations: N could satisfy this and still be far less than np . In

comparison, Chatterjee et al. (2015) considers a related latent space model. Their method needs to observe at least $\Theta(n^{2(K+1)/(K+2)})$ entries to guarantee the consistent recovery of an $n \times n$ matrix. This implies that one needs to observe $\Theta(n^{K/(K+2)})$ entries out of n in each row, as compared to our much weaker requirement of $\Theta(\log^2(n))$ per row.

We now specialize our results to matrices that are \mathcal{F} -embeddable for \mathcal{F} containing functions with bounded derivatives. This is a natural class of functions to work with (though one could alternatively work in a multivariate Sobolev or Hölder space).

Condition 2 M is \mathcal{F} -embeddable, where \mathcal{F} contains functions with uniformly bounded L -th order mixed partials (for some fixed $L > 0$). More formally, define $\mathcal{F}(L, \gamma, K)$, for $L, K \geq 1$ as the set of L -th order differentiable functions from $\mathbb{R}_{[0,1]}^K$ to \mathbb{R} satisfying

$$f(\mathbf{x}^0) + \left| \frac{\partial^L}{\partial x_1^{L_1} \cdots \partial x_K^{L_K}} f(\mathbf{x}) \right|_{\mathbf{x}=\mathbf{x}^0} \leq \gamma, \quad (13)$$

for all $\mathbf{x}^0 = (x_1^0, \dots, x_K^0) \in \mathbb{R}_{[0,1]}^K \subset \mathbb{R}^K$ and all integers L_1, \dots, L_K satisfying $L_1 + \dots + L_K = L$. Now, additionally define the set

$$\begin{aligned} \mathcal{M}(L, \gamma, K) = \{M \in \mathbb{R}^{n \times p} \mid m_{ij} = f_j(\boldsymbol{\theta}_i, \cdot), \\ \text{with } f_j \in \mathcal{F}(L, \gamma, K), j \leq p, \text{ and } \boldsymbol{\theta}_i, \in \mathbb{R}_{[0,1]}^K, i \leq n\} \end{aligned} \quad (14)$$

This is the set of $\mathcal{F}(L, \gamma, K)$ embeddable matrices, where the embedding lives in a compact space (for convenience we use the ℓ_∞ ball). Our formal condition here is that $M \in \mathcal{M}(L, \gamma, K)$.

Remark. In the above condition, we will often suppress the dependence on γ , and write $\mathcal{M}(L, K)$ and $\mathcal{F}(L, K)$. This is because γ does not affect the convergence rate of our estimator. Additionally, here we specify the domain of the embeddings to be $[0, 1]^K$ for ease of exposition. This is actually general as we could rescale any compactly supported embedding to live in this interval.

Condition 2 imposes an additional constraint on our embedding: The underlying manifold on which our matrix lives should be smooth. Here smoothness is characterized by a number of bounded derivatives. As we will see, this function class engages well with Condition 1 in the sense that we are able to characterize $J^*(\epsilon)$ for the function class $\mathcal{F}(L, K)$. This is essentially a multivariate Hölder class, which has been widely used in the area of non-parametric estimation (Tsybakov, 2008). One could alternatively look at this as a multivariate Sobolev class under the sup-norm, $W^{L, \infty}(\mathbb{R}^K)$.

The following lemma gives the number of basis elements needed to linearly approximate a matrix satisfying the above condition, with bounded approximation error ϵ .

Lemma 3 For the function class $\mathcal{F}(L, K)$ described in Condition 2, we have that Condition 1 is satisfied with

$$J^*(\epsilon) \leq \frac{(K+L)!}{K!L!} \left(\frac{L!}{\gamma K^L} \right)^{-K/L} \epsilon^{-K/L} = O\left(\epsilon^{-K/L}\right).$$

The proof of this lemma is given in Appendix C. Now, we can establish the final convergence result for smoothly embeddable matrices.

Theorem 4 *Under the same scenario and assumptions as in Theorem 2, assume further the $\mathcal{F}(K, L)$ -embeddable matrix M satisfies Condition 2 for given L and K . Then, the upper bound (12) is optimized at $\epsilon = (N^{-1}(n \vee p) \log(n + p))^{L/(2L+K)}$, resulting in*

$$\frac{1}{np} \left\| \widehat{M} - M \right\|_F^2 = O_P \left(\left[\frac{(n \vee p) \log(n + p)}{N} \right]^{\frac{2L}{2L+K}} \right). \quad (15)$$

In the asymptotic statement (15), we assume that L and K are fixed (there is a constant prefactor in $O_p(\cdot)$ that depends only on L and K , which we omit).

The proof is given in Appendix D (as well as a precise, non-asymptotic result). This upper bound of the convergence rate of the MSE of \widehat{M} is only based on the dimensions n and p of matrix M , the total number of observations N , as well as the degree of smoothness L and dimension of the embedding K . Previous work that assumed M was low-rank generally gave a rate of the form $N^{-1}(n \vee p) \text{rank}(M) \log(n + p)$ (Bach, 2008; Klopp et al., 2014; van de Geer, 2016). In contrast, our upper bound does not rely on the rank of M . Instead, the role of $\text{rank}(M)$ is replaced by L , and K . At the same time, the dependence on N shifts from the finite rank N^{-1} rate to our nonparametric rate $N^{-2L/(2L+K)}$. This result reaffirms that the standard matrix completion estimator based on nuclear norm minimization is consistent for matrices with low-dimensional non-linear structure. Perhaps more importantly, it also shows how the convergence rate depends on the degree of smoothness, and dimension of the manifold. This can be seen in the exponent on the RHS of (15): $2L/(2L + K)$. Increasing the degree of smoothness moves this exponent towards 1; increasing the dimension moves the exponent towards 0. This is analogous to more standard non-parametric regression problems in smooth hypothesis spaces where the minimax convergence rate for MSE looks analogous (Tsybakov, 2008).

4. Minimax Lower Bound

In this section, we use information-theoretical methods to establish a lower bound on the estimation error for completing *non-linearly embeddable* matrices with uniformly sampled at random entries when the latent embedding Θ is K -dimensional and satisfies Condition 2. The rate we find in the lower bound matches the rate obtained by nuclear norm penalization in Theorem 4 up to a log-term. Thus our upper bound is sharp (up to a logarithmic factor), and, the nuclear-norm penalization based estimator given in (6) is rate-optimal (up to polylog) for this problem.

To derive the lower bound, we consider the underlying matrices $M \in \mathcal{M}(L, \gamma, K)$ as defined in (14), that is, matrices that live in L -th order smooth, K dimensional manifolds. Let \mathbb{P}_M denote the probability distribution of the observations $\{(y_t, X_t)\}_{t=1}^N$ generated by model (2) with $E(y_t|X_t) = \langle X_t, M \rangle$. We give a minimax lower bound of the $\|\cdot\|_F^2$ -risk for estimating M in the following result.

Theorem 5 *For any given $L \geq 1$, $\gamma > 0$ and $K \geq 1$, let $\kappa := n/p$. Then, for some constant $A > 0$ that depends on K, L, γ, σ^2 and κ , the minimax risk for estimating M satisfies*

$$\inf_{\hat{M}} \sup_{M \in \mathcal{M}(L, \gamma, K)} \mathbb{P}_M \left(\frac{1}{np} \left\| \hat{M} - M \right\|_F^2 > A \left(\frac{n \vee p}{N} \right)^{\frac{2L}{2L+K}} \right) \geq 1/2, \quad (16)$$

when $c_0^{-\frac{2L+K}{K}} (n \vee p) \leq N \leq c_0^{-\frac{2L+K}{K}} 0.48^{2L+K} (n \vee p) n^{\frac{2L+K}{K}}$ for some constant c_0 which depends on K, L, γ, σ^2 and κ .

The proof is given in the Appendix E—it extends the lower bound argument of Koltchinskii et al. (2011) from linear low rank matrices to non-linearly embeddable matrices. Comparing Theorem 5 to Theorem 4, we see that the lower bound matches the upper bound (15) up to a logarithmic factor. This shows that the estimator given by (6) is actually an optimal estimator (up to a log term) for this non-linear low-dimensional matrix completion regime.

We note that the requirement $N = O((n \vee p)n^{(2L+K)/K})$ in Theorem 5 is a bit unusual. It comes from a technical constraint in our proof, required to construct a suitably large packing set. This may just be an artifact of our proof technique, and not innate to the problem. Recall that the upper bound holds as long as $N \geq (n \vee p) \log^2(n+p)$, so there is a large regime where the assumption required for our upper and lower bounds overlap.

Comparison with existing results

Matrix completion with a true underlying mean that is low rank has been widely studied (Bach, 2008; Candes and Plan, 2010b; Negahban and Wainwright, 2011; Koltchinskii et al., 2011; Klopp et al., 2014; van de Geer, 2016) in the literature and nuclear-norm penalized estimators are known to have clear theoretical guarantees in this linear setting. The rescaled Frobenius-norm of the error decreases linearly with respect to the number of observed entries N and is proportional to the rank of true matrix M .

Matrix completion without direct rank assumptions has also studied in the literature (though to a much lesser degree). In Negahban and Wainwright (2011), the authors consider both low-rank truth and near low-rank truth (formalized using the Schatten q -norms). The convergence rate under this type assumption depends on the parameter q and is no faster than the low-rank model. In Cai et al. (2016), the authors engage with some “max-norm” based estimator and allow a highly non-uniform sampling scheme, which can also be treated as some form of near low-rank structure; they showed their proposed estimator is minimax optimal with a square-root dependence $N^{-1/2}$. See Section 3.5 of Cai et al. (2016) for a discussion about the relation between max-norm and “effective rank” of Negahban and Wainwright (2011). In Chatterjee et al. (2015), the author proposed a singular-value thresholding estimator and applied to a variety of statistically interpretable settings (e.g. low rank matrix completion, block models, latent space models). Their framework is general and the author presented some theoretical guarantees for each setting—but their results do not fully recover the optimal N^{-1} rate under the finite rank assumption. The latent space model discussed in Chatterjee et al. (2015) is of a similar nature to our non-linearly embeddable matrix assumption, but they show consistency of the proposed estimator. In this paper, we detail the finite-sample convergence rate with an emphasize on how it varies with the assumed level of smoothness L and latent factor dimension K . Moreover, we showed that our theoretical rate guarantee $N^{-2L/(2L+K)}$ essentially matches the minimax lower bound.

5. Simulation Study

In this section, we empirically evaluate the effectiveness of matrix completion using the soft-thresholding estimator \widehat{M} in (7) for noisy incomplete matrices which are generated from low-dimensional non-linear embeddings. (These matrices are full rank, even though they are generated from low-dimensional non-linear embeddings). Here, we only show the case of univariate embedding ($K = 1$) and aim to empirically evaluate how the Frobenius error $\frac{1}{np} \left\| \widehat{M} - M \right\|_F^2$ changes with the dimension (n) when $n = p$. We examine scenarios where the non-linear embeddings are of different orders of smoothness.

The underlying matrices are generated as described in (4): $m_{ij} = f_j(\boldsymbol{\theta}_{i,\cdot})$ for $i = 1, \dots, n$ and $j = 1, \dots, p$. In particular, to make sure that Conditions 1 and 2 are satisfied, we generate f_j as

$$f_j(x) = \sum_{b=1}^{\infty} \beta_b \psi_b(x),$$

where $\psi_b(x)$ are orthonormal bases in $L_2[0, 1]$ defined by:

$$\begin{aligned} \psi_1(x) &= 1, \\ \psi_{2b}(x) &= \sqrt{2} \cos(2\pi bx), \\ \psi_{2b+1}(x) &= \sqrt{2} \sin(2\pi bx). \end{aligned}$$

Meanwhile, to set up the order of smoothness L and make sure that $\beta_b \psi_b(x)$ vanishes with b , we sample the coefficients β_b from a uniform distribution:

$$\beta_b \sim_{i.i.d} U \left[-b^{-(L+1)}, b^{-(L+1)} \right], \quad b = 1, 2, \dots$$

In this way, we can guarantee that $\sum_{b=1}^{\infty} b^{2L} \beta_b^2 < \infty$. Thus, f_j is a function whose L th order derivative is $O_p(1)$.

In this simulation, for computational reasons, we actually use only the first 100 basis vectors $f_j(x) = \sum_{b=1}^{100} \beta_b \psi_b(x)$. The underlying embeddings $\boldsymbol{\theta}_{i,\cdot} \in \mathbb{R}$ are also i.i.d. sampled from a uniform distribution $U(0, 1)$ for $i = 1, \dots, n$. We set the missingness rate to $\nu = 0.3$: The total number of observed entries is $N = (1 - \nu)np$. The observed entries are $y_t = \langle X_t, M \rangle + \xi_t$, where X_t are uniformly sampled from \mathcal{X} and the error terms are independently Gaussian distributed $\xi_t \sim_{i.i.d.} N(0, 1)$. We generate random data sets $\{(y_t, X_t)\}_{t=1}^N$ of size $n \in \{500, 1000, 2000, 3000, 5000\}$ and estimate M . We run 100 simulations for each size. To select λ , instead of using cross-validation, here we consider an oracle procedure: For each simulation, we estimate the MSE for a set of λ values and select the λ that minimizes the MSE. We report this MSE of the estimated matrix \widehat{M} and the corresponding λ .

Figure 1 shows the results of estimating M generated by non-linear embeddings with different orders of smoothness, L . Since $N = (1 - \nu)np$, the convergence rate in (15) reduces to $O_P \left([\log(2n)/n]^{\frac{2L}{2L+K}} \right)$. The log term inside is negligible as n increases. Hence, if we regress $\log(\text{MSE})$ on $\log(n)$, the absolute value of slope should be roughly about $2L/(2L + 1)$ ($K = 1$ in this simulation). We increase the order of smoothness of f from $L = 1$ to $L = 5$. For these values of L , the expected absolute value of the slope should be 0.67, 0.80, 0.86, 0.89, and 0.91. The rates from our simulations are respectively 0.67, 0.78,




Figure 1: Theoretical rate vs. empirical rate (in log scale) of the mean squared errors as a function of sample size. The underlying matrices M are generated by f with different orders (L) of smoothness. The low-rank embedding is one-dimensional ($K = 1$). We regress $\log(\text{MSE})$ on $\log(n)$, and compare the theoretical slopes (left) with the empirical slopes (right). For each smoothness level, L , we also obtain the 95% confidence regions using bootstrap (dash lines).

0.80, 0.88, and 0.91. There is generally strong agreement between theoretical and empirical results except for the setting of $L = 3$. We hypothesize that this is due to finite sample issues.

6. Discussion

Nuclear-norm based matrix completion methods were originally developed for scenarios where the underlying mean matrix has low rank. In this manuscript, we present theoretical results to explain the effectiveness of matrix completion in applications where the underlying mean matrix is not low rank, but instead lives in a low-dimensional smooth manifold.

Our results show that, in such scenarios, nuclear-norm regularization can still result in a procedure that is minimax rate optimal (up to a log factor) for recovering the underlying mean matrix. In particular, we give upper bounds on the rate of convergence as a function of the number of rows, columns, and observed entries in the matrix, as well as the smoothness, and dimension of the embeddings. We additionally give matching minimax lower bounds (up to a logarithmic factor) for this problem. These bounds appear analogous to the minimax rate in the case of standard non-parametric regression.

Our theoretical results relate the error bounds to the smoothness and dimension of the non-linear embedding; however, the technical proof does not provide a way to figure out the explicit form of the hidden embeddings, which may be interesting in practice, e.g., for dimension reduction. Modifying the original matrix completion method in order to estimate the hidden embeddings may be an important direction of future research.

Acknowledgments

NS and AS were supported by NSF award DMS-1915855.

Appendix A. Proof of Lemma 1

We begin with a proof of Lemma 1:

Proof

Recall that M is \mathcal{F} -embeddable and \mathcal{F} satisfies Condition 1. Thus, the entries of M are generated by $m_{ij} = f_j(\boldsymbol{\theta}_{i,\cdot})$. Consider arbitrary $\epsilon > 0$. Then there is some fixed $C_0 > 0$, and a collection of functions $\mathcal{F}_\epsilon^* = \{\tilde{\psi}_1, \tilde{\psi}_2, \dots, \tilde{\psi}_{J^*(\epsilon)}\}$ that give the finite set of minimal cardinality $J^*(\epsilon)$, with the property that

$$\max_{f \in \mathcal{F}} \min_{\|\beta\|_\infty \leq C_0} \left\| f - \sum_{l=1}^{J^*(\epsilon)} \beta_l \tilde{\psi}_l \right\|_\infty \leq \epsilon. \quad (17)$$

and $\|\tilde{\psi}_l\|_\infty \leq C_0$. For any given $f \in \mathcal{F}$, let

$$\beta^\epsilon(f) = \operatorname{argmin}_\beta \left\| f - \sum_{\tilde{\psi}_l \in \mathcal{F}_\epsilon^*} \beta_l \tilde{\psi}_l \right\|_\infty. \quad (18)$$

This implies that we can approximate M with a low rank matrix M^ϵ , with entries given by

$$m_{ij}^\epsilon \leftarrow \sum_{\tilde{\psi}_l \in \mathcal{F}_\epsilon^*} \tilde{\psi}_l(\boldsymbol{\theta}_{i,\cdot}) \cdot \beta_l^\epsilon(f_j), \quad (19)$$

such that $|m_{ij} - m_{ij}^\epsilon| \leq \epsilon$ for $i = 1, \dots, n$, $j = 1, \dots, p$. Now, let Ψ denote the matrix with $\Psi_{il} = \tilde{\psi}_l(\boldsymbol{\theta}_{i,\cdot})$ and B denote the matrix with entries $B_{lj} = \beta_l^\epsilon(f_j)$. Then, the approximation matrix can be compactly written as

$$M^\epsilon = \Psi B,$$

with $\Psi \in \mathbb{R}^{n \times J^*(\epsilon)}$ and $B \in \mathbb{R}^{J^*(\epsilon) \times p}$. Thus, $\operatorname{rank}(M^\epsilon) = J^*(\epsilon) \leq \min(n, p)$ and

$$\|M^\epsilon - M\|_\infty \leq \epsilon.$$

Finally, using a variational form of the nuclear norm (Srebro and Shraibman, 2005), we have

$$\frac{1}{\sqrt{np}} \|M^\epsilon\|_* = \frac{1}{2} \min_{UV^\top = M^\epsilon} \left(\frac{1}{n} \|U\|_F^2 + \frac{1}{p} \|V\|_F^2 \right).$$

From the statement above (18), we know that $\|\Psi\|_\infty$ and $\|B\|_\infty$ are both bounded by C_0 . Thus we have that

$$\frac{1}{\sqrt{np}} \|M^\epsilon\|_* \leq \frac{1}{2} \left(\frac{1}{n} \|\Psi\|_F^2 + \frac{1}{p} \|B\|_F^2 \right) \leq C_0^2 J^*(\epsilon).$$

Noting that C_0^2 is a constant independent of ϵ gives us our result. ■

Appendix B. Upper Bound

In this section, we derive the upper bound on the convergence rate of our estimator \widehat{M} . Recall that $\{(y_t, X_t)\}_{t=1}^N$ are generated by

$$y_t = \langle X_t, M \rangle + \xi_t, \quad (20)$$

where ξ_t are i.i.d random errors distributed $N(0, \sigma^2)$, and M is a $n \times p$ matrix. The estimator we consider is defined by

$$\begin{aligned} \widehat{M} &\leftarrow \operatorname{argmin}_{M \in \mathbb{R}^{n \times p}} \left\{ \frac{1}{np} \|M\|_F^2 - \left\langle \frac{2}{N} \sum_{t=1}^N y_t X_t, M \right\rangle + \lambda \|M\|_* \right\} \\ &\equiv \operatorname{argmin}_{M \in \mathbb{R}^{n \times p}} L_N(M) \end{aligned} \quad (21)$$

Before we begin, we note that much of the proof of Theorem 2 follows very directly from the work of Koltchinskii et al. (2011). In particular, we use Lemmas A1, and A2 from that work. We include the full proof here both for completeness, and because we need to be careful with constants—in Koltchinskii et al. (2011) certain variables were assumed constant (eg. the rank and norm of the oracle matrix) that we would like to allow to grow. The novelty in this manuscript (regarding the upper bound) comes from combining those arguments with a sharp characterization of misspecification error (to get our sharp upper-bound)

Now, to prove Theorem 2, we first introduce two technical lemmas, which will play a key role. Proving these lemmas will entail most of the work required for proving this theorem. In Lemma A1, we derive a deterministic upper bound for the estimation error (under a stochastic condition) as a function of the regularization parameter λ , when λ is sufficiently large (in this Lemma, “sufficiently large” is left as a stochastic constraint). In particular, we show that the risk can be decomposed into a misspecification error and a prediction error. Then, in Lemma A2, we identify a deterministic value for λ such that, with high probability, the condition in Lemma A1 will hold. More specifically we give probabilistic bounds for the operator norm of the stochastic error term in our generative model. We can then combine these to obtain the general oracle inequality in Theorem 2.

Before continuing, we give some additional notation: For any matrix Z , we denote $\|Z\|_{op} = \Lambda_{\max}(Z)$, where $\Lambda_{\max}^2(Z) = \Lambda_{\max}(Z^T Z)$ is the largest singular value of $Z^T Z$, also known as the operator-norm.

Lemma A1 *Suppose we observe $\{(y_t, X_t)\}_{t=1}^N$ generated by (20), where X_t are i.i.d uniformly sampled from \mathcal{X} . Further, assume the underlying true matrix $M \in \mathbb{R}^{n \times p}$ is \mathcal{F} -embeddable with Condition 1 satisfied. Let $\Delta = N^{-1} \sum_{t=1}^N [y_t X_t - \mathbb{E}(y_t X_t)]$. If $\lambda \geq 2\|\Delta\|_{op}$, then*

$$\frac{1}{np} \|\widehat{M} - M\|_F^2 \leq \epsilon^2 + \left(\frac{1 + \sqrt{2}}{2} \right)^2 J^*(\epsilon) \lambda^2 np \quad (22)$$

holds for any $\epsilon > 0$. Recall that $J^*(\epsilon)$ is the minimal rank of an approximation matrix M^ϵ with $\|M^\epsilon - M\|_\infty < \epsilon$.

Proof

The proof of this lemma is based on the strong convexity of the loss function $L_N(M)$.

Consider the the subdifferential of $L_N(M)$, which is the set of matrices of the following form:

$$\partial L_N(M) = \left\{ \frac{2}{np}M - \frac{2}{N} \sum_{t=1}^N y_t X_t + \lambda B, \quad B \in \partial \|M\|_* \right\}. \quad (23)$$

Thus, the following representation holds for $\widehat{A} \in \partial L_N(\widehat{M})$

$$\widehat{A} = \frac{2}{np} \widehat{M} - \frac{2}{N} \sum_{t=1}^N y_t X_t + \lambda \widehat{B},$$

for some $\widehat{B} \in \partial \|\widehat{M}\|_*$. Since $M \mapsto L_N(M)$ is strictly convex, \widehat{M} defined in (21) is the unique minimizer of $L_N(M)$. This implies, $\mathbf{0} \in \partial L_N(\widehat{M})$. Hence, there exists $\widehat{B} \in \partial \|\widehat{M}\|_*$ such that $\widehat{A} = \mathbf{0}$, and thus

$$\langle \widehat{A}, \widehat{M} - M^\epsilon \rangle = \langle \mathbf{0}, \widehat{M} - M^\epsilon \rangle = 0. \quad (24)$$

It further follows that

$$\begin{aligned} & \langle \widehat{A}, \widehat{M} - M^\epsilon \rangle \\ &= \frac{2}{np} \langle \widehat{M}, \widehat{M} - M^\epsilon \rangle - \frac{2}{N} \sum_{t=1}^N \langle y_t X_t, \widehat{M} - M^\epsilon \rangle + \lambda \langle \widehat{B}, \widehat{M} - M^\epsilon \rangle = 0. \end{aligned} \quad (25)$$

$M^\epsilon \in \mathbb{R}^{n \times p}$ is the approximation matrix with $\text{rank}(M^\epsilon) = J^*(\epsilon)$. So, it has spectral representation $M^\epsilon = \sum_{j=1}^{J^*(\epsilon)} \sigma_j u_j v_j^T$ where $u_j \in \mathbb{R}^n$ and $v_j \in \mathbb{R}^p$, $j = 1, \dots, J^*(\epsilon)$, are orthonormal vectors, and σ_j are the singular values of M^ϵ . Let U and V denote the linear span of $\{u_1, \dots, u_{J^*(\epsilon)}\}$ and $\{v_1, \dots, v_{J^*(\epsilon)}\}$ respectively. Then, the subdifferential of $\|M^\epsilon\|_*$ can be represented by the following set of matrices (Watson, 1992):

$$\partial \|M^\epsilon\|_* = \left\{ \sum_{j=1}^{J^*(\epsilon)} u_j v_j^T + P_{U^\perp} W P_{V^\perp} : \|W\|_{op} \leq 1 \right\},$$

where U^\perp denotes the orthogonal complements of U and P_{U^\perp} denotes the projection on the linear vector subspace U^\perp . The same argument applies to V and P_{V^\perp} . Thus, $B^\epsilon \in \partial \|M^\epsilon\|_*$ can be represented as

$$B^\epsilon = \sum_{j=1}^{J^*(\epsilon)} u_j v_j^T + P_{U^\perp} W P_{V^\perp} \quad (26)$$

for arbitrary matrix W having $\|W\|_{op} \leq 1$. Due to the trace duality, there exists W with $\|W\|_{op} \leq 1$ such that

$$\langle P_{U^\perp} W P_{V^\perp}, \widehat{M} - M^\epsilon \rangle = \langle P_{U^\perp} W P_{V^\perp}, \widehat{M} \rangle = \langle W, P_{U^\perp} \widehat{M} P_{V^\perp} \rangle = \|P_{U^\perp} \widehat{M} P_{V^\perp}\|_* \quad (27)$$

So, it follows from (25) that

$$\begin{aligned}
 & \frac{2}{np} \langle \widehat{M} - M, \widehat{M} - M^\epsilon \rangle + \frac{2}{np} \langle M, \widehat{M} - M^\epsilon \rangle + \lambda \langle \widehat{B} - B^\epsilon, \widehat{M} - M^\epsilon \rangle \\
 &= \frac{2}{N} \sum_{t=1}^N \langle \mathbb{E}(y_t X_t), \widehat{M} - M^\epsilon \rangle - \lambda \langle B^\epsilon, \widehat{M} - M^\epsilon \rangle + \frac{2}{N} \sum_{t=1}^N \langle y_t X_t - \mathbb{E}(y_t X_t), \widehat{M} - M^\epsilon \rangle
 \end{aligned} \tag{28}$$

Due to the monotonicity of subdifferentials of convex functions $M \mapsto \|M\|_*$, $\langle \widehat{B} - B^\epsilon, \widehat{M} - M^\epsilon \rangle \geq 0$. So, (28) can be further simplified:

$$\begin{aligned}
 \frac{2}{np} \langle \widehat{M} - M, \widehat{M} - M^\epsilon \rangle &\leq -\lambda \langle B^\epsilon, \widehat{M} - M^\epsilon \rangle + 2 \langle \Delta, \widehat{M} - M^\epsilon \rangle \\
 \xrightarrow{(26)} &= -\lambda \left\langle \sum_{j=1}^{J^*(\epsilon)} u_j v_j^T + P_{U^\perp} W P_{V^\perp}, \widehat{M} - M^\epsilon \right\rangle + 2 \langle \Delta, \widehat{M} - M^\epsilon \rangle \\
 \xrightarrow{(27)} &= -\lambda \left\langle \sum_{j=1}^{J^*(\epsilon)} u_j v_j^T, \widehat{M} - M^\epsilon \right\rangle + 2 \langle \Delta, \widehat{M} - M^\epsilon \rangle - \lambda \|P_{U^\perp} \widehat{M} P_{V^\perp}\|_*
 \end{aligned} \tag{29}$$

where $\Delta = N^{-1} \sum_{t=1}^N [y_t X_t - \mathbb{E}(y_t X_t)]$.

By arithmetic, we see that the left-hand side of (29) is equal to:

$$\begin{aligned}
 2 \langle \widehat{M} - M, \widehat{M} - M^\epsilon \rangle &= \langle \widehat{M} - M, \widehat{M} - M + M - M^\epsilon \rangle + \langle \widehat{M} - M^\epsilon + M^\epsilon - M, \widehat{M} - M^\epsilon \rangle \\
 &= \|\widehat{M} - M\|_F^2 - \|M^\epsilon - M\|_F^2 + \|\widehat{M} - M^\epsilon\|_F^2.
 \end{aligned} \tag{30}$$

As for the right side of (29), we use the following facts:

$$\left\| \sum_{j=1}^{J^*(\epsilon)} u_j v_j^T \right\|_{op} = 1 \quad \text{and} \quad \left\langle \sum_{j=1}^{J^*(\epsilon)} u_j v_j^T, \widehat{M} - M^\epsilon \right\rangle = \left\langle \sum_{j=1}^{J^*(\epsilon)} u_j v_j^T, P_U (\widehat{M} - M^\epsilon) P_V \right\rangle. \tag{31}$$

Given (30)-(31), (29) becomes

$$\begin{aligned}
 & \frac{1}{np} \|\widehat{M} - M\|_F^2 + \frac{1}{np} \|\widehat{M} - M^\epsilon\|_F^2 + \lambda \|P_{U^\perp} \widehat{M} P_{V^\perp}\|_* \\
 & \leq -\lambda \left\langle \sum_{j=1}^{J^*(\epsilon)} u_j v_j^T, \widehat{M} - M^\epsilon \right\rangle + \frac{1}{np} \|M^\epsilon - M\|_F^2 + 2 \langle \Delta, \widehat{M} - M^\epsilon \rangle \\
 & \leq \lambda \|P_U (M^\epsilon - \widehat{M}) P_V\|_* + \frac{1}{np} \|M^\epsilon - M\|_F^2 + 2 \langle \Delta, \widehat{M} - M^\epsilon \rangle,
 \end{aligned} \tag{32}$$

where the last inequality is due to $|\langle M_1, M_2 \rangle| \leq \|M_1\|_{op} \times \|M_2\|_*$.

In (32), the stochastic error term $\langle \Delta, \widehat{M} - M^\epsilon \rangle$ can be decomposed:

$$\begin{aligned}
 \langle \Delta, \widehat{M} - M^\epsilon \rangle &= \langle \mathcal{P}_{M^\epsilon}(\Delta), \widehat{M} - M^\epsilon \rangle + \langle P_{U^\perp} \Delta P_{V^\perp}, \widehat{M} - M^\epsilon \rangle \\
 &= \langle \mathcal{P}_{M^\epsilon}(\Delta), \mathcal{P}_{M^\epsilon}(\widehat{M} - M^\epsilon) \rangle + \langle \mathcal{P}_{M^\epsilon}(\Delta), P_{U^\perp}(\widehat{M} - M^\epsilon) P_{V^\perp} \rangle \\
 &\quad + \langle P_{U^\perp} \Delta P_{V^\perp}, \mathcal{P}_{M^\epsilon}(\widehat{M}) \rangle + \langle P_{U^\perp} \Delta P_{V^\perp}, P_{U^\perp} \widehat{M} P_{V^\perp} \rangle - \langle P_{U^\perp} \Delta P_{V^\perp}, M^\epsilon \rangle \\
 &= \langle \mathcal{P}_{M^\epsilon}(\Delta), \mathcal{P}_{M^\epsilon}(\widehat{M} - M^\epsilon) \rangle + \langle P_{U^\perp} \Delta P_{V^\perp}, P_{U^\perp} \widehat{M} P_{V^\perp} \rangle
 \end{aligned} \tag{33}$$

where $\mathcal{P}_{M^\epsilon}(\Delta) = \Delta - P_{U^\perp} \Delta P_{V^\perp}$. So it can be upper bounded by:

$$\begin{aligned}
 |\langle \Delta, \widehat{M} - M^\epsilon \rangle| &\leq \|\mathcal{P}_{M^\epsilon}(\Delta)\|_F \|\mathcal{P}_{M^\epsilon}(\widehat{M} - M^\epsilon)\|_F + \|P_{U^\perp} \Delta P_{V^\perp}\|_{op} \|P_{U^\perp} \widehat{M} P_{V^\perp}\|_* \\
 &\leq \|\mathcal{P}_{M^\epsilon}(\Delta)\|_F \|\widehat{M} - M^\epsilon\|_F + \|P_{U^\perp} \Delta P_{V^\perp}\|_{op} \|P_{U^\perp} \widehat{M} P_{V^\perp}\|_* \\
 &\leq \sqrt{2J^*(\epsilon)} \|\Delta\|_{op} \|\widehat{M} - M^\epsilon\|_F + \|\Delta\|_{op} \|P_{U^\perp} \widehat{M} P_{V^\perp}\|_*.
 \end{aligned} \tag{34}$$

The last inequality is due to the facts that

$$\begin{aligned}
 \|\mathcal{P}_{M^\epsilon}(\Delta)\|_F &\leq \sqrt{\text{rank}(\mathcal{P}_{M^\epsilon}(\Delta))} \|\Delta\|_{op} = \sqrt{\text{rank}(P_{U^\perp} \Delta P_{V^\perp} + P_U \Delta)} \|\Delta\|_{op} \\
 &\leq \sqrt{2 \text{rank}(M^\epsilon)} \|\Delta\|_{op} = \sqrt{2J^*(\epsilon)} \|\Delta\|_{op}
 \end{aligned}$$

and $\|P_{U^\perp} \Delta P_{V^\perp}\|_{op} \leq \|\Delta\|_{op}$.

Meanwhile, the first term in the right-hand side of (32) can also be bounded:

$$\|P_U(M^\epsilon - \widehat{M})P_V\|_* \leq \sqrt{\text{rank}(M^\epsilon)} \|P_U(M^\epsilon - \widehat{M})P_V\|_F \leq \sqrt{J^*(\epsilon)} \|M^\epsilon - \widehat{M}\|_F. \tag{35}$$

Combining (34) - (35), (32) becomes

$$\begin{aligned}
 \frac{1}{np} \|\widehat{M} - M\|_F^2 &+ \frac{1}{np} \|\widehat{M} - M^\epsilon\|_F^2 + (\lambda - 2\|\Delta\|_{op}) \|P_{U^\perp} \widehat{M} P_{V^\perp}\|_* \\
 &\leq \lambda \sqrt{J^*(\epsilon)} \|M^\epsilon - \widehat{M}\|_F + \epsilon^2 + 2\sqrt{2J^*(\epsilon)} \|\Delta\|_{op} \|\widehat{M} - M^\epsilon\|_F.
 \end{aligned} \tag{36}$$

If $\lambda \geq 2\|\Delta\|_{op}$, then

$$\frac{1}{np} \|\widehat{M} - M\|_F^2 + \frac{1}{np} \|\widehat{M} - M^\epsilon\|_F^2 \leq \epsilon^2 + (1 + \sqrt{2}) \lambda \sqrt{J^*(\epsilon)} \|\widehat{M} - M^\epsilon\|_F \tag{37}$$

which implies

$$\begin{aligned}
 \frac{1}{np} \|\widehat{M} - M\|_F^2 &\leq \epsilon^2 + (1 + \sqrt{2}) \lambda \sqrt{J^*(\epsilon)} \|\widehat{M} - M^\epsilon\|_F - \frac{1}{np} \|\widehat{M} - M^\epsilon\|_F^2 \\
 &\leq \epsilon^2 + \left(\frac{1 + \sqrt{2}}{2} \right)^2 J^*(\epsilon) \lambda^2 np
 \end{aligned} \tag{38}$$

as claimed. ■

The result in Lemma A1 still contains regularization parameter λ . When λ is selected too large, then entries of \widehat{M} will be overly shrunk toward zero and give poor reconstruction error. If λ is too small, then our constraint, $\lambda \geq 2\|\Delta\|_{op}$, will not be satisfied. Thus, it is important to identify a minimal value for λ such that $\lambda \geq 2\|\Delta\|_{op}$ with high probability. Here, we introduce the second lemma, which gives an upper bound for $\|\Delta\|_{op}$.

Lemma A2 Consider the same data generating mechanism as in Lemma A1, with X_t are i.i.d uniformly sampled from \mathcal{X} . Then, there exists constant c_1 (dependent on σ and $\|M\|_\infty$) such that

$$\|\Delta\|_{op} \leq c_1 \left[\sqrt{\frac{\log(n+p)}{N(n \wedge p)}} + \sqrt{\log\left(\frac{8(n \wedge p)}{3\sigma^2}\right) \frac{\log(n+p)}{N}} \right] \quad (39)$$

with probability at least $1 - 2(n+p)^{-1}$.

Furthermore, when $N \geq (n \wedge p) \log^2(n+p)$, we have $\|\Delta\|_{op} \leq 2c_1 \sqrt{\frac{\log(n+p)}{N(n \wedge p)}}$ with probability at least $1 - 2(n+p)^{-1}$.

To derive the bound of the stochastic error Δ , we shall use the matrix version of Bernstein's inequality. We now use 2 propositions from van de Geer (2016). For completeness we include statements of the propositions here below.

Proposition A3 Let $\{Z_t\}_{t=1}^N$ be i.i.d $n \times p$ matrices that satisfy for some $\alpha \geq 1$ and all t

$$\mathbb{E} Z_t = \mathbf{0}, \quad K := \|\|Z_t\|_{op}\|_{\Psi(\alpha)} < \infty,$$

where $\|\cdot\|_{\Psi(\alpha)}$ is the $\Psi(\alpha)$ -Orlicz norm defined as $\|z\|_{\Psi(\alpha)} := \inf \left\{ c > 0 : \mathbb{E} \exp\left(\frac{|z|^\alpha}{c^\alpha}\right) \leq 2 \right\}$ for a random variable $z \in \mathbb{R}$. Define

$$R^2 := \max \left\{ \left\| \frac{1}{N} \sum_{t=1}^N \mathbb{E} Z_t Z_t^T \right\|_{op}, \left\| \frac{1}{N} \sum_{t=1}^N \mathbb{E} Z_t^T Z_t \right\|_{op} \right\}.$$

Then for a constant \tilde{c} and for all $h > 0$,

$$\mathbb{P} \left(\left\| \frac{1}{N} \sum_{t=1}^N Z_t \right\|_{op} \geq \tilde{c} R \sqrt{\frac{h + \log(n+p)}{N}} + \tilde{c} \log^{1/\alpha} \left(\frac{K}{R} \right) \left(\frac{h + \log(n+p)}{N} \right) \right) \leq \exp(-h).$$

Proposition A4 Let $\{Z_t\}_{t=1}^N$ be $n \times p$ matrices that satisfy for a constant K_1

$$\mathbb{E} Z_t = \mathbf{0}, \quad \max_{1 \leq t \leq N} \|Z_t\|_{op} \leq K_1.$$

With the same definition for R as in Proposition A3 Then for all $h > 0$,

$$\mathbb{P} \left(\left\| \frac{1}{N} \sum_{t=1}^N Z_t \right\|_{op} \geq \sqrt{2} R \sqrt{\frac{h + \log(n+p)}{N}} + \frac{K_1 [h + \log(n+p)]}{3N} \right) \leq \exp(-h).$$

Given the above results, we now prove Lemma A2.

Proof [Proof of Lemma A2] $\|\Delta\|_{op}$ can be decomposed into two parts as below and we shall bound each part respectively.

$$\begin{aligned}
 \|\Delta\|_{op} &= \left\| \frac{1}{N} \sum_{t=1}^N [y_t X_t - \mathbb{E}(y_t X_t)] \right\|_{op} \\
 &= \left\| \frac{1}{N} \sum_{t=1}^N [\xi_t X_t - \mathbb{E}(\xi_t X_t) + \text{tr}[M^T X_t] X_t - \mathbb{E}(\text{tr}(M^T X_t) X_t)] \right\|_{op} \\
 &\leq \left\| \frac{1}{N} \sum_{t=1}^N \xi_t X_t \right\|_{op} + \left\| \frac{1}{N} \sum_{t=1}^N (\text{tr}(M^T X_t) X_t - \mathbb{E}(\text{tr}(M^T X_t) X_t)) \right\|_{op} \\
 &= I_1 + I_2.
 \end{aligned} \tag{40}$$

We use Proposition A3 to bound I_1 . Let $Z_{1,t} = \xi_t X_t$. Since $\xi_t \sim_{i.i.d} N(0, \sigma^2)$ and X_t are i.i.d uniformly sampled from \mathcal{X} with $\xi_t \perp X_t$, $\{Z_{1,t}\}_{t=1}^N$ are i.i.d $n \times p$ matrices having

$$\mathbb{E} Z_{1,t} = \mathbf{0}, \quad K := \|\|Z_{1,t}\|_{op}\|_{\Psi(\alpha)} = \|\xi_t\|_{\Psi(\alpha)}.$$

For a normal variable $z \sim N(0, 1)$, we have $\mathbb{E} \exp(z^2/c^2) = c/\sqrt{c^2 - 2}$ when $c > \sqrt{2}$. Thus, $\mathbb{E} \exp(z^2/c^2) \leq 2 \Rightarrow c \geq \sqrt{8/3}$. So, $K = \|\xi_t\|_{\Psi(2)} = \sqrt{8/3}$. Let

$$\begin{aligned}
 R^2 &:= \max \left\{ \left\| \frac{1}{N} \sum_{t=1}^N \mathbb{E} Z_{1,t} Z_{1,t}^T \right\|_{op}, \left\| \frac{1}{N} \sum_{t=1}^N \mathbb{E} Z_{1,t}^T Z_{1,t} \right\|_{op} \right\} \\
 &= \sigma^2 \max \left\{ \left\| \frac{1}{N} \sum_{t=1}^N \mathbb{E}(X_t X_t^T) \right\|_{op}, \left\| \frac{1}{N} \sum_{t=1}^N \mathbb{E}(X_t^T X_t) \right\|_{op} \right\} \\
 &= \frac{\sigma^2}{n \wedge p}.
 \end{aligned}$$

Due to Proposition A3, for some \tilde{c} and for all $h > 0$, we have

$$\mathbb{P} \left(I_1 \geq \tilde{c} \sigma \sqrt{\frac{h + \log(n+p)}{N(n \wedge p)}} + \tilde{c} \sqrt{\frac{1}{2} \log \left(\frac{8(n \wedge p)}{3\sigma^2} \right)} \left(\frac{h + \log(n+p)}{N} \right) \right) \leq \exp(-h). \tag{41}$$

Similarly, we use Proposition A4 to bound I_2 . Let $Z_{2,t} = \text{tr}(M^T X_t) X_t - \mathbb{E}(\text{tr}(M^T X_t) X_t)$, where $\mathbb{E}(\text{tr}(M^T X_t) X_t) = \frac{1}{np} M$. So, $\mathbb{E}(Z_{2,t}) = \mathbf{0}$, and

$$\|Z_{2,t}\|_{op} \leq \|\text{tr}(M^T X_t) X_t\|_{op} + \|\mathbb{E}(\text{tr}(M^T X_t) X_t)\|_{op} \leq 2\|M\|_{\infty}.$$

Let $K_1 = 2\|M\|_{\infty}$. Then $\max_{1 \leq t \leq N} \|Z_{2,t}\|_{op} \leq K_1$. Consider,

$$\mathbb{E}(Z_{2,t} Z_{2,t}^T) = \mathbb{E}[\text{tr}(M^T X)^2 X X^T] - \left(\frac{1}{np} \right)^2 M M^T,$$

$$\mathbb{E}(Z_{2,t}^T Z_{2,t}) = \mathbb{E}[\text{tr}(M^T X)^2 X^T X] - \left(\frac{1}{np} \right)^2 M^T M.$$

Then,

$$\begin{aligned} \|\mathbb{E}(Z_{2,t}Z_{2,t}^T)\|_{op} &\leq \|\mathbb{E}[\text{tr}(M^T X)^2 X X^T]\|_{op} + \left\| \left(\frac{1}{np}\right)^2 M M^T \right\|_{op} \\ &\leq \|M\|_\infty^2/n + \frac{\|M\|_\infty^2}{np} \leq 2\|M\|_\infty^2/n, \end{aligned}$$

and similarly $\|\mathbb{E}(Z_{2,t}^T Z_{2,t})\|_{op} \leq 2\|M\|_\infty^2/p$. Let

$$R_1^2 := \max \left\{ \left\| \frac{1}{N} \sum_{t=1}^N \mathbb{E}(Z_{2,t}Z_{2,t}^T) \right\|_{op}, \left\| \frac{1}{N} \sum_{t=1}^N \mathbb{E}(Z_{2,t}^T Z_{2,t}) \right\|_{op} \right\} \leq \frac{2\|M\|_\infty^2}{n \wedge p}.$$

Then, applying Proposition A4, we have

$$\mathbb{P} \left(I_2 \geq 2\|M\|_\infty \sqrt{\frac{h + \log(n+p)}{N(n \wedge p)}} + \frac{2\|M\|_\infty [h + \log(n+p)]}{3N} \right) \leq \exp(-h). \quad (42)$$

Combining the results of (41) and (42), for all $h > 0$

$$\begin{aligned} \mathbb{P} \left(\|\Delta\|_{op} \geq (\tilde{c}\sigma + 2\|M\|_\infty) \left[\sqrt{\frac{h + \log(n+p)}{N(n \wedge p)}} + \sqrt{\frac{1}{2} \log \left(\frac{8(n \wedge p)}{3\sigma^2} \right)} \left(\frac{h + \log(n+p)}{N} \right) \right] \right) \\ \leq 2 \exp(-h). \end{aligned} \quad (43)$$

Select $h = \log(n+p)$ and let $c_1 = \sqrt{2}(\tilde{c}\sigma + 2\|M\|_\infty)$, then

$$\mathbb{P} \left(\|\Delta\|_{op} \geq c_1 \left[\sqrt{\frac{\log(n+p)}{N(n \wedge p)}} + \sqrt{\log \left(\frac{8(n \wedge p)}{3\sigma^2} \right)} \frac{\log(n+p)}{N} \right] \right) \leq 2(n+p)^{-1}. \quad (44)$$

In particular, if $N \geq (n \wedge p) \log^2(n+p)$, we have

$$\mathbb{P} \left(\|\Delta\|_{op} \geq 2c_1 \sqrt{\frac{\log(n+p)}{N(n \wedge p)}} \right) \leq 2(n+p)^{-1},$$

as desired. ■

Based on Lemma A1 and A2, it is straightforward to prove Theorem 2.

Proof [Proof of Theorem 2] When $N \geq (n \wedge p) \log^2(n+p)$, we choose λ of the following form

$$\lambda = C_2 \sqrt{\frac{\log(n+p)}{N(n \wedge p)}} \quad (45)$$

where $C_2 > 0$ is a constant with $C_2 \geq 4c_1$, where c_1 is defined in Lemma A2 that only depends on σ and $\|M\|_\infty$. Following from (22), then

$$\frac{1}{np} \|\widehat{M} - M\|_F^2 \leq C_2^2 \left(\frac{1 + \sqrt{2}}{2} \right)^2 \frac{(n \vee p) \log(n+p)}{N} J^*(\epsilon) + \epsilon^2. \quad (46)$$

holds with probability $1 - 2(n+p)^{-1}$.

This completes the proof. \blacksquare

Appendix C. Proof of Lemma 3

We begin with an outline of the proof. To form our set of basis functions, we will tessellate our domain \mathbb{X}^K with ∞ -norm balls, and use a Taylor series centered at an arbitrary point within each ball to get a uniform approximation for functions in that ball. For a fixed center point, the Taylor series is a linear combination of fixed basis functions. To obtain our full set of basis functions, we will collect all of the terms in all of those Taylor series. We now formalize this:

Proof For functions satisfying Condition (2), we consider a Taylor series approximation to $f \in \mathcal{F}(L, \gamma, K)$ of order L at a point $\mathbf{x}^0 \in \mathbb{R}_{[0,1]}^K$. Let $l = (l_1, \dots, l_k)$ be a k -dimensional multi-index, we have the Taylor expansion:

$$T_{\mathbf{x}^0} f(\mathbf{x}) = f(\mathbf{x}^0) + \sum_{|l| \leq L-1} \frac{1}{l!} \nabla^l f(\mathbf{x}^0) (\mathbf{x} - \mathbf{x}^0)^l, \quad (47)$$

where $|l| = l_1 + l_2 + \dots + l_k$, $l! = l_1! \dots l_k!$, $\nabla^l f(\mathbf{x}) = \frac{\partial^l f}{\partial x_1^{l_1} \dots \partial x_k^{l_k}}$ and $\mathbf{x}^l = x_1^{l_1} \dots x_k^{l_k}$. There exists $\mathbf{x}' = (x'_1, \dots, x'_K)^T \in \mathbb{R}_{[0,1]}^K$ in a neighborhood of radius $\|\mathbf{x} - \mathbf{x}^0\|_2$ centered at \mathbf{x}^0 such that the approximation error obeys

$$\begin{aligned} |f(\mathbf{x}) - T_{\mathbf{x}^0} f(\mathbf{x})| &\leq \left| \sum_{L_1 + \dots + L_K = L} \frac{1}{L_1! \dots L_K!} \times \frac{\partial^L f(\mathbf{x}')}{\partial x_1^{L_1} \dots \partial x_K^{L_K}} |x_1 - x_1^0|^{L_1} \dots |x_K - x_K^0|^{L_K} \right| \\ &\stackrel{\text{Condition 2}}{\leq} \gamma \left| \sum_{L_1 + \dots + L_K = L} \frac{|x_1 - x_1^0|^{L_1} \dots |x_K - x_K^0|^{L_K}}{L_1! \dots L_K!} \right| \\ &\stackrel{\text{Multinomial Theorem}}{=} \frac{\gamma}{L!} (|x_1 - x_1^0| + \dots + |x_K - x_K^0|)^L \end{aligned} \quad (48)$$

If we consider the approximation error within an ∞ -norm ball of radius d (and choose any point in that ball as \mathbf{x}^0), then $|x_k - x_k^0| \leq d$ for $k = 1, \dots, K$. (48) has

$$|f(\mathbf{x}) - T_{\mathbf{x}^0} f(\mathbf{x})| \leq \frac{\gamma}{L!} K^L d^L. \quad (49)$$

Thus, to get an approximation error of ϵ , let $\frac{\gamma}{L!} K^L d^L = \epsilon$, we need to divide the space into balls of radius

$$d = \sqrt[L]{\frac{L!}{\gamma K^L}} \times \epsilon^{1/L}. \quad (50)$$

As the support $\mathbb{R}_{[0,1]}^K$ is bounded by 1, we need $(1/d)^K$ balls with radius d (in ∞ -norm) to cover the entirety of \mathbb{X}^K , resulting in $\binom{K+L}{L}(1/d)^K$ total terms to get an approximation error ϵ (the above Taylor series approximation contains $\binom{K+L}{L}$ terms). If we select balls of radius d in (50), this gives us a total number of terms in our linear expansion

$$J^*(\epsilon) \leq \binom{K+L}{L} \left(\frac{L!}{\gamma K^L} \right)^{-K/L} \epsilon^{-K/L}.$$

That is, $J^*(\epsilon) = O(\epsilon^{-K/L})$.

To complete the proof, we just need to confirm that the basis functions and their coefficients satisfy the boundedness conditions (8) and (9). Formally, the basis functions we are using are the (local) monomials in the Taylor expansion

$$\psi_l(x) = \prod_{j=1}^K (x_j - x_j^0)^{l_j} \cdot I_{\text{cell containing } x^0}(x), \quad (51)$$

where I is the set indicator function. It is direct to see that they satisfy (8).

Now we need to show that the coefficients $\beta_l = \nabla^l f(\mathbf{x}^0)/l!$ satisfy the condition $\|\beta\|_\infty \leq C_0$. It is known that for univariate functions, the intermediate derivative can be bounded by a higher order derivative. A general statement can be found in (DeVore and Lorentz, 1993), Theorem 5.6. For our purpose we just need the following special case:

$$\left\| f^{(k)} \right\|_{L_\infty([0,1])} \leq C(L) \left(\|f\|_{L_\infty([0,1])} + \left\| f^{(L)} \right\|_{L_\infty([0,1])} \right) \text{ for } 0 \leq k \leq L, \quad (52)$$

where $C(L)$ is a constant independent of f .

We first apply the above inequality to bound all the derivatives only involving one variable, that is, $\frac{\partial^m}{\partial x_j^m} f$, $j = 1, 2, \dots, K$, $m \leq L - 1$. And then use them as intermediate quantities to bound derivatives involving two variables. We iterate this process until we have a uniform bound for every derivative that shows up in (47). ■

Appendix D. Proof of Theorem 4

The proof of this theorem is quite straightforward by connecting a few pieces we have already built.

Proof Given Condition 2 and Lemma 3, we have $J^*(\epsilon) = C_3 \epsilon^{-\frac{K}{L}}$ for some constant C_3 relying on γ, K , and L . Plugging in this to the upper bound in Theorem 2, the upper bound (46) then becomes

$$C_2^2 C_3 \left(\frac{1 + \sqrt{2}}{2} \right)^2 \frac{(n \vee p) \log(n + p)}{N} \epsilon^{-\frac{K}{L}} + \epsilon^2, \quad (53)$$

which is optimized at

$$\begin{aligned} \frac{(n \vee p) \log(n+p)}{N} \epsilon^{-\frac{K}{L}} &= \epsilon^2 \\ \Rightarrow \epsilon &= \left(\frac{(n \vee p) \log(n+p)}{N} \right)^{\frac{L}{2L+K}}. \end{aligned} \quad (54)$$

So, we have

$$\frac{1}{np} \|\widehat{M} - M\|_F^2 \leq C^* \left(\frac{(n \vee p) \log(n+p)}{N} \right)^{\frac{2L}{2L+K}} \quad (55)$$

with probability at least $1 - 2(n+p)^{-1}$ with $C^* = C_2^2 C_3 \left(\frac{1+\sqrt{2}}{2} \right)^2 + 1$. Equivalently, we can say

$$\frac{1}{np} \|\widehat{M} - M\|_F^2 = O_P \left(\left[\frac{(n \vee p) \log(n+p)}{N} \right]^{\frac{2L}{2L+K}} \right),$$

as claimed. ■

Appendix E. Deriving the Minimax Lower Bound

In this section, we derive the minimax lower bound for estimation within $M(L, \gamma, K)$: We show that the convergence rate in Theorem 4 is optimal (up to log terms).

Recall that we assume the true M belongs to the following class of matrices:

$$\mathcal{M}(L, \gamma, K) := \{M \in \mathbb{R}^{n \times p} : m_{ij} = f_j(\boldsymbol{\theta}_{i,\cdot}), \boldsymbol{\theta}_{i,\cdot} \in \mathbb{R}_{[0,1]}^K, f_j \in \mathcal{F}(L, \gamma, K), \forall j \leq p\}, \quad (56)$$

where $\mathcal{F}(L, \gamma, K)$ is a class of functions with bounded derivatives:

$$\mathcal{F}(L, \gamma, K) := \left\{ f : \left| \frac{\partial^L}{\partial x_1^{L_1} \cdots \partial x_K^{L_K}} f(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{x}^0} \right| \leq \gamma, \forall \mathbf{x}^0 \in \mathbb{R}_{[0,1]}^K, \sum_{k=1}^K L_k = L \right\}. \quad (57)$$

For simplicity of notation, let $\boldsymbol{\theta}_i := \boldsymbol{\theta}_{i,\cdot} \in \mathbb{R}_{[0,1]}^K$ denote the i -th row vector of the embeddings $\Theta \in \mathbb{R}^{n \times K}$ in this section.

We shall obtain the lower bound based on information theory. The bound is with respect to $\|\cdot\|_F^2$ -risk. We pose things in terms of the error in a multi-way hypothesis testing problem, where the set of testing hypotheses should be a suitably large packing set for $\mathcal{M}(L, \gamma, K)$. In this section, we first show the existence of such a suitably large packing set. Then, we apply Yang's method (Yang and Barron, 1999) to prove the main results in Theorem 5.

E.1 Constructing the $2\delta_{N,n,p}$ -packing Set

For $M \in \mathcal{M}(L, \gamma, K)$, the risk of the estimator can be written as

$$\frac{1}{np} \|\widehat{M} - M\|_F^2 = \frac{1}{np} \sum_{i=1}^n \sum_{j=1}^p [\hat{m}_{ij} - f_j(\boldsymbol{\theta}_i)]^2.$$

This is to say, bounding $\frac{1}{np}\|\widehat{M} - M\|_F^2$ can be viewed as a classical nonparametric regression problem. So, we follow the construction of many hypotheses as in Section 2.6 of Tsybakov (2009). However, here we are working in a multi-dimensional setting, that is, $\boldsymbol{\theta}_i \in \mathbb{R}^K$, $K \geq 1$.

In giving our packing set, we will work with combinations of ‘‘bump functions’’. To define these, we need an archetypal ingredient—the bump functions that we will use:

$$\varphi(u) = c_L e \times \exp\left(-\frac{1}{1-4u^2}\right), \quad u \in (-1/2, 1/2), \quad (58)$$

which is infinitely differentiable and vanishes outside of $(-1/2, 1/2)$; $c_L > 0$ is a tiny constant that only depends on L such that $|\partial^l \varphi(u)/\partial u^l| \leq 1$, $\forall l = 0, 1, \dots, L$. Meanwhile, since $\int_{-1/2}^{1/2} e^2 \exp^2\left(-\frac{1}{1-4u^2}\right) du > 0.49$, (it is actually very close to 0.5), we have $\|\varphi\|_2^2 := \int_{-1/2}^{1/2} \varphi^2(u) du > 0.49c_L^2$. In addition, the maximum value of this function is $\sup_u |\varphi(u)| = \varphi(0) = c_L$.

Now, we shall work under the multidimensional setting. We use bold letters to refer to multivariate indices and regular letters to refer to the indices of each coordinate. Let $\mathbf{i} = (i_1, \dots, i_K) \in \{1, 2, \dots, \sqrt[K]{n}\}^K$ having $\sum_{i=(i_1, \dots, i_K)} 1 = \sum_{i_1=1}^{\sqrt[K]{n}} \dots \sum_{i_K=1}^{\sqrt[K]{n}} 1 = n$, where $\sqrt[K]{n}$ is assumed to be an integer. Suppose that the observed embeddings follows a fixed equispaced design, that is, $\boldsymbol{\theta}_i = \boldsymbol{\theta}_{(i_1, \dots, i_K)} = (\theta_{i_1}, \dots, \theta_{i_K})^T = (\frac{i_1}{\sqrt[K]{n}}, \dots, \frac{i_K}{\sqrt[K]{n}})^T$. Consider a multivariate function $\Phi_{\mathbf{d}}: \mathbb{R}^K \rightarrow \mathbb{R}$,

$$\begin{aligned} \Phi_{\mathbf{d}}(\boldsymbol{\theta}_i) &= \gamma b^{-L/K} \prod_{k=1}^K \varphi_{d_k}(\theta_{i_k}) \\ &:= \gamma b^{-L/K} \prod_{k=1}^K \varphi(\sqrt[K]{b}\theta_{i_k} - d_k + 1/2), \end{aligned} \quad (59)$$

where $\mathbf{d} = (d_1, \dots, d_K) \in \{1, 2, \dots, \sqrt[K]{b}\}^K$. Here $b \geq 1$ is an integer that depends on N, n, p and some constant c_0 , and will be specified later. $\varphi(u)$ is defined in (58). Then, we have the following technical lemma for $\Phi_{\mathbf{d}}$, which will later be used for constructing the packing set.

Lemma A5 *Suppose $\varphi(\cdot)$ are given by (58). Then, $\Phi_{\mathbf{d}}$ has the following properties:*

- (i) $\Phi_{\mathbf{d}}(\mathbf{x}) \in \mathcal{F}(L, \gamma, K)$.
- (ii) $\Phi_{\mathbf{d}}$ have disjoint support for different \mathbf{d} .
- (iii) There exist $C_{1,L,K} > 0$ and $C_{2,L,K} > 0$ only dependent on L and K , for any given \mathbf{d} , $\Phi_{\mathbf{d}}$ has

$$\gamma^2 C_{2,L,K} b^{-\frac{2L+K}{K}} \leq \frac{1}{n} \sum_{\mathbf{i}=(i_1, \dots, i_K)} \Phi_{\mathbf{d}}^2(\boldsymbol{\theta}_i) \leq \gamma^2 C_{1,L,K} b^{-\frac{2L+K}{K}}$$

when integer b satisfies $1 \leq b \leq 0.48^K n$.

Proof For $\varphi(\cdot)$ in (58), we have $|\frac{\partial^l}{\partial u^l} \varphi(u)| \leq 1, \forall l = 0, 1, \dots, L$ such that $\left| \frac{\partial^L}{\partial x_1^{L_1} \dots \partial x_K^{L_K}} \Phi_{\mathbf{d}}(\mathbf{x}) \right| \leq \gamma$ holds for any $L_1 + \dots + L_K = L$ for $\mathbf{x} \in \mathbb{R}_{[0,1]}^K$. Thus, $\Phi_{\mathbf{d}}(\mathbf{x}) \in \mathcal{F}(L, \gamma, K)$.

Given that $\varphi(u) > 0$ if and only if $u \in (-1/2, 1/2)$, we have $\varphi_{d_k}(x) \equiv \varphi(\sqrt[k]{b}x - d_k + 1/2) > 0$ if and only if $x \in \left(\frac{d_k-1}{\sqrt[k]{b}}, \frac{d_k}{\sqrt[k]{b}}\right)$ for $d_k \in \{1, \dots, \sqrt[k]{b}\}$. So, for each, we can divide the space $[0, 1]$ into $\sqrt[k]{b}$ intervals, that is,

$$\Delta_1 = \left[0, \frac{1}{\sqrt[k]{b}}\right], \quad \Delta_{d_k} = \left(\frac{d_k-1}{\sqrt[k]{b}}, \frac{d_k}{\sqrt[k]{b}}\right], \quad d_k = 2, \dots, \sqrt[k]{b},$$

such that $\Delta_{d_k} \cap \Delta_{d'_k} = \emptyset$ for $d_k \neq d'_k$ and $\cup_{d_k} \Delta_{d_k} = [0, 1]$. Thus, $\varphi_{d_k}(x)$ have disjoint support and their support union is the unit interval.

Because $\Phi_{\mathbf{d}}$ is the product of φ_{d_k} , they also have disjoint supports. That is, for each \mathbf{d} , $\Phi_{\mathbf{d}}(\mathbf{x}) > 0$ only when $\mathbf{x} \in \Delta_{\mathbf{d}}$ where

$$\begin{aligned} \Delta_{\mathbf{d}=(1,1,\dots,1)} &= \left[0, \frac{1}{\sqrt[k]{b}}\right] \times \dots \times \left[0, \frac{1}{\sqrt[k]{b}}\right], \\ \Delta_{\mathbf{d}=(d_1,\dots,d_K)} &= \left(\frac{d_1-1}{\sqrt[k]{b}}, \frac{d_1}{\sqrt[k]{b}}\right] \times \dots \times \left(\frac{d_K-1}{\sqrt[k]{b}}, \frac{d_K}{\sqrt[k]{b}}\right], \end{aligned}$$

$d_k = 2, \dots, \sqrt[k]{b}$ for $k = 1, \dots, K$, such that $\Delta_{\mathbf{d}} \cap \Delta_{\mathbf{d}'} = \emptyset$ if $\mathbf{d} \neq \mathbf{d}'$ and $\cup_{\mathbf{d}} \Delta_{\mathbf{d}} = [0, 1]^K$. So, the space $[0, 1]^K$ is divided into b disjoint cubes.

As for (iii), we know there exists a constant c_L that only depends on L such that $\sup_u |\varphi(u)| = \varphi(0) = c_L$, and $\|\varphi\|_2^2 > 0.49c_L^2$. Then

$$\begin{aligned}
 \frac{1}{n} \sum_{\mathbf{i}=(i_1, \dots, i_K)} \Phi_{\mathbf{d}}^2(\boldsymbol{\theta}_{\mathbf{i}}) &= \frac{1}{n} \gamma^2 b^{-2L/K} \sum_{i_1=1}^{\sqrt[K]{b}} \dots \sum_{i_K=1}^{\sqrt[K]{b}} \prod_{k=1}^K \varphi^2 \left(\sqrt[K]{\frac{b}{n}} i_k - d_k + 1/2 \right) \\
 &= \frac{1}{n} \gamma^2 b^{-2L/K} \sum_{i_2=1}^{\sqrt[K]{b}} \dots \sum_{i_K=1}^{\sqrt[K]{b}} \prod_{k=2}^K \varphi^2 \left(\sqrt[K]{\frac{b}{n}} i_k - d_k + 1/2 \right) \\
 &\quad \times \left[\sum_{i_1=1}^{\sqrt[K]{b}} \varphi^2 \left(\sqrt[K]{\frac{b}{n}} i_1 - d_1 + 1/2 \right) \right] \\
 &= \frac{1}{n} \gamma^2 b^{-2L/K} \sum_{i_3=1}^{\sqrt[K]{b}} \dots \sum_{i_K=1}^{\sqrt[K]{b}} \prod_{k=3}^K \varphi^2 \left(\sqrt[K]{\frac{b}{n}} i_k - d_k + 1/2 \right) \\
 &\quad \times \left[\sum_{i_1=1}^{\sqrt[K]{b}} \varphi^2 \left(\sqrt[K]{\frac{b}{n}} i_1 - d_1 + 1/2 \right) \right] \times \left[\sum_{i_2=1}^{\sqrt[K]{b}} \varphi^2 \left(\sqrt[K]{\frac{b}{n}} i_2 - d_2 + 1/2 \right) \right] \\
 &\quad \dots \\
 &= \frac{1}{n} \gamma^2 b^{-2L/K} \prod_{k=1}^K \left\{ \sum_{i_k=1}^{\sqrt[K]{n}} \varphi^2 \left(\sqrt[K]{\frac{b}{n}} i_k - d_k + 1/2 \right) \right\} \\
 &= \frac{1}{n} \gamma^2 b^{-2L/K} \prod_{k=1}^K \left\{ \sum_{\sqrt[K]{\frac{n}{b}}(d_k-1) < i_k \leq \sqrt[K]{\frac{n}{b}} d_k} \varphi^2 \left(\sqrt[K]{\frac{b}{n}} i_k - d_k + 1/2 \right) \right\} \\
 &\leq \frac{1}{n} \gamma^2 b^{-2L/K} \prod_{k=1}^K \left\{ \sqrt[K]{\frac{n}{b}} \times \varphi^2(0) \right\} \\
 &= \gamma^2 b^{-\frac{2L+K}{K}} c_L^{2K}.
 \end{aligned} \tag{60}$$

Therefore, $\frac{1}{n} \sum_{\mathbf{i}=(i_1, \dots, i_K)} \Phi_{\mathbf{d}}^2(\boldsymbol{\theta}_{\mathbf{i}}) \leq c_L^{2K} \gamma^2 b^{-\frac{2L+K}{K}}$ and c_L^{2K} is the constant we find for $C_{1,L,K}$.

On the other hand, we use the fact that the upper Riemann sum is greater than the integral of the function. Thus, for each coordinate,

$$\begin{aligned}
 \sqrt[\kappa]{\frac{b}{n}} \sum_{i_k=1}^{\sqrt[\kappa]{n}} \varphi^2 \left(\sqrt[\kappa]{\frac{b}{n}} i_k - d_k + 1/2 \right) &= \sqrt[\kappa]{\frac{b}{n}} \sum_{\frac{\sqrt[\kappa]{n}}{b}(d_k-1) < i_k \leq \frac{\sqrt[\kappa]{n}}{b} d_k} \varphi^2 \left[\sqrt[\kappa]{\frac{b}{n}} \left(i_k - \frac{\sqrt[\kappa]{n}(d_k - 1/2)}{\sqrt[\kappa]{b}} \right) \right] \\
 &= \sqrt[\kappa]{\frac{b}{n}} \sum_{\frac{-\sqrt[\kappa]{n}}{2\sqrt[\kappa]{b}} < t \leq \frac{\sqrt[\kappa]{n}}{2\sqrt[\kappa]{b}}} \varphi^2 \left(\sqrt[\kappa]{b/n} \times t \right) \\
 \underbrace{[\varphi(-u) = \varphi(u)]}_{\rightarrow} &= \sqrt[\kappa]{\frac{b}{n}} \left[2 \sum_{0 \leq t \leq \frac{\sqrt[\kappa]{n}}{2\sqrt[\kappa]{b}}} \varphi^2(\sqrt[\kappa]{b/n} \times t) - \varphi^2(0) \right] \\
 \underbrace{[\varphi(u) \text{ decreases for } u \geq 0]}_{\rightarrow} &\geq 2 \sqrt[\kappa]{\frac{b}{n}} \int_0^{\frac{\sqrt[\kappa]{n}}{2\sqrt[\kappa]{b}}} \varphi^2(\sqrt[\kappa]{b/n} \times t) dt - \sqrt[\kappa]{\frac{b}{n}} \varphi^2(0) \\
 &= 2 \int_0^{1/2} \varphi^2(u) du - \sqrt[\kappa]{\frac{b}{n}} c_L^2 \\
 &= \|\varphi\|_2^2 - \sqrt[\kappa]{\frac{b}{n}} c_L^2 \\
 &> c_L^2 \left(0.49 - \sqrt[\kappa]{b/n} \right) \\
 \underbrace{[1 \leq b \leq 0.49^K n]}_{\rightarrow} &\geq 0,
 \end{aligned}$$

Thus, the empirical sum can also be lower bounded by

$$\begin{aligned}
 \frac{1}{n} \sum_{\mathbf{i}=(i_1, \dots, i_K)} \Phi_{\mathbf{d}}^2(\boldsymbol{\theta}_{\mathbf{i}}) &= \gamma^2 b^{-\frac{2L+K}{K}} \prod_{k=1}^K \left\{ \sqrt[\kappa]{\frac{b}{n}} \sum_{i_k=1}^{\sqrt[\kappa]{n}} \varphi^2 \left(\sqrt[\kappa]{\frac{b}{n}} i_k - d_k + 1/2 \right) \right\} \\
 &\geq \gamma^2 b^{-\frac{2L+K}{K}} c_L^{2K} \left(0.49 - \sqrt[\kappa]{b/n} \right)^K
 \end{aligned} \tag{61}$$

When $1 \leq b \leq 0.48^K n$,

$$\frac{1}{n} \sum_{\mathbf{i}=(i_1, \dots, i_K)} \Phi_{\mathbf{d}}^2(\boldsymbol{\theta}_{\mathbf{i}}) \geq \gamma^2 (0.1c_L)^{2K} b^{-\frac{2L+K}{K}}, \tag{62}$$

and thus $(0.1c_L)^{2K}$ is the constant we find for $C_{2,L,K}$. Combining (60) and (62), we have

$$\gamma^2 C_{2,L,K} b^{-\frac{2L+K}{K}} \leq \frac{1}{n} \sum_{\mathbf{i}=(i_1, \dots, i_K)} \Phi_{\mathbf{d}}^2(\boldsymbol{\theta}_{\mathbf{i}}) \leq \gamma^2 C_{1,L,K} b^{-\frac{2L+K}{K}}$$

as claimed. ■

In proving the lower bound, we shall use Fano's method (see Section 15.3.2 in Wainwright (2019)). To do so, we first establish the connection between minimax risks and error probabilities in testing problems (for completeness), and then apply Fano's inequality to lower bound the error probabilities. To this end, we first provide the following lemma, which shows that there exists a packing set of hypotheses with suitably large cardinality, for which the mutual information (stated in terms of Kullback-Leibler divergence) can be upper bounded. We can then use Fano's inequality with this set.

Lemma A6 *Consider an arbitrary fixed L, γ and K . For some constant $C_{1,L,K}$ and $C_{2,L,K}$ that only depends on L and K , and for some other constant $c_0 > 0$, there exists a subset $\mathcal{B}^0 \subseteq \mathcal{M}(L, \gamma, K)$ with cardinality*

$$|\mathcal{B}^0| \geq 2^{\lceil c_0 \left(\frac{n \vee p}{N}\right)^{\frac{-K}{2L+K}} \rceil \times p/8} + 1,$$

when $p \geq 8$, that has the following properties:

(i) \mathcal{B}^0 is a $2\delta_{N,n,p}$ -packing set, that is, for any $M_s \neq M_{s'} \in \mathcal{B}^0$,

$$\frac{1}{np} \|M_s - M_{s'}\|_F^2 \geq 2\delta_{N,n,p} = \frac{C_{2,L,K}\gamma^2}{8} (2c_0)^{-2L/K} \left(\frac{n \vee p}{N}\right)^{\frac{2L}{2L+K}}$$

$$\text{when } c_0^{-\frac{2L+K}{K}} (n \vee p) \leq N \leq c_0^{-\frac{2L+K}{K}} 0.48^{2L+K} (n \vee p) n^{\frac{2L+K}{K}}.$$

(ii) For any $M_s, M_{s'} \in \mathcal{B}^0$,

$$K(\mathbb{P}_s || \mathbb{P}_{s'}) \leq \frac{C_{1,L,K}\gamma^2}{2\sigma^2} c_0^{-\frac{2L}{K}} N \left(\frac{n \vee p}{N}\right)^{\frac{2L}{2L+K}}$$

where $K(\mathbb{P}_s || \mathbb{P}_{s'})$ denotes the Kullback-Leibler divergence between probability distributions of observations $\{(y_t, X_t)\}_{t=1}^N$ satisfying model (20), given M_s and $M_{s'}$ respectively.

Proof

We will consider a positive integer b which depends on N, n, p and a constant c_0 . The precise specification of b will come later. Consider the multivariate function $\Phi_{\mathbf{d}}(\boldsymbol{\theta})$ in (59).

We will define a set Ω that is used to construct packing matrices where each element ω in Ω is a sequence (of length b) of diagonal matrices. We index the set in a somewhat curious way: We use a multi-index of dimension K where each index has elements in $\{1, \dots, \sqrt[b]{b}\}$. This will ease exposition later.

$$\Omega = \left\{ \boldsymbol{\omega} = (\boldsymbol{w}_{\mathbf{d}})_{\mathbf{d} \in \{1, \dots, \sqrt[b]{b}\}^K} : \text{for each } \mathbf{d}, \boldsymbol{w}_{\mathbf{d}} = \text{diag}(w_{\mathbf{d},1}, \dots, w_{\mathbf{d},p}), w_{\mathbf{d},j} \in \{0, 1\} \right\}. \quad (63)$$

From this we define the following collection of matrices,

$$\mathcal{B} = \left\{ M_{\mathbf{w}} = \sum_{d_1=1}^{\sqrt[b]{b}} \cdots \sum_{d_K=1}^{\sqrt[b]{b}} \begin{pmatrix} \Phi_{\mathbf{d}}(\boldsymbol{\theta}_1)w_{d,1} & \Phi_{\mathbf{d}}(\boldsymbol{\theta}_1)w_{d,2} & \cdots & \Phi_{\mathbf{d}}(\boldsymbol{\theta}_1)w_{d,p} \\ \Phi_{\mathbf{d}}(\boldsymbol{\theta}_2)w_{d,1} & \Phi_{\mathbf{d}}(\boldsymbol{\theta}_2)w_{d,2} & \cdots & \Phi_{\mathbf{d}}(\boldsymbol{\theta}_2)w_{d,p} \\ \vdots & \vdots & \ddots & \vdots \\ \Phi_{\mathbf{d}}(\boldsymbol{\theta}_n)w_{d,1} & \Phi_{\mathbf{d}}(\boldsymbol{\theta}_n)w_{d,2} & \cdots & \Phi_{\mathbf{d}}(\boldsymbol{\theta}_n)w_{d,p} \end{pmatrix}_{n \times p}, w_{d,j} \in \{0, 1\} \right\}$$

$$=: \left\{ M_{\mathbf{w}} = \sum_{d_1, \dots, d_K} \Phi_{\mathbf{d}}(\boldsymbol{\Theta})\mathbf{w}_{\mathbf{d}}, \text{ for } \mathbf{w} = (\mathbf{w}_{\mathbf{d}}) \in \Omega \right\}. \quad (64)$$

We see that we can compactly write each matrix in our set as the product of $\Phi_{\mathbf{d}}(\boldsymbol{\Theta})$ and $\mathbf{w}_{\mathbf{d}}$, where $\Phi_{\mathbf{d}}(\boldsymbol{\Theta})$ is a $n \times p$ matrix whose elements in the i -th row are all $\Phi_{\mathbf{d}}(\boldsymbol{\theta}_i)$. It is direct to check that the cardinality of Ω is given by $|\Omega| = |\mathcal{B}| = 2^{bp}$.

Thus, entries of $M_{\mathbf{w}} \in \mathcal{B}$ can be written as $m_{ij} = \sum_{d_1, \dots, d_K} \Phi_{\mathbf{d}}(\boldsymbol{\theta}_i)w_{d,j} = g_j(\boldsymbol{\theta}_i)$, where g_j has bounded derivatives,

$$\left| \frac{\partial^L g_j(\mathbf{x})}{\partial x_1^{L_1} \cdots \partial x_K^{L_K}} \right| \leq \sum_{d_1, \dots, d_K} \left| \frac{\partial^L \Phi_{\mathbf{d}}(\mathbf{x})}{\partial x_1^{L_1} \cdots \partial x_K^{L_K}} \right| = \left| \frac{\partial^L \Phi_{\mathbf{d}}(\mathbf{x})}{\partial x_1^{L_1} \cdots \partial x_K^{L_K}} \right| \mathbb{1}_{\{\mathbf{x} \in \Delta_{\mathbf{d}}\}} \leq \gamma$$

for $\forall \mathbf{x} \in \mathbb{R}_{[0,1]}^K$. Hence, $\mathcal{B} \subseteq \mathcal{M}(L, \gamma, K)$.

Consider a set of testing hypotheses from \mathcal{B} ,

$$\mathcal{B}^0 = \{M_{\mathbf{w}^{(0)}}, \dots, M_{\mathbf{w}^{(S)}}\} \subseteq \mathcal{B}, \quad \mathbf{w}^{(s)} \in \Omega, \quad s = 0, 1, \dots, S, \quad (65)$$

where $\mathbf{w}^{(s)} \neq \mathbf{w}^{(s')}$ for $0 \leq s \neq s' \leq S$.

For any $0 \leq s \neq s' \leq S$, and constant $C_{2,L,K}$ only dependent on L ,

$$\frac{1}{np} \|M_{\mathbf{w}^{(s)}} - M_{\mathbf{w}^{(s')}}\|_F^2 = \frac{1}{np} \sum_{\mathbf{i}=(i_1, \dots, i_K)} \sum_{j=1}^p \left[\sum_{d_1=1}^{\sqrt[b]{b}} \cdots \sum_{d_K=1}^{\sqrt[b]{b}} (w_{\mathbf{d},j}^{(s)} - w_{\mathbf{d},j}^{(s')}) \Phi_{\mathbf{d}}(\boldsymbol{\theta}_{\mathbf{i}}) \right]^2$$

the support of $\Phi_{\mathbf{d}}$'s are disjoint \rightarrow

$$= \frac{1}{p} \sum_{j=1}^p \sum_{d_1=1}^{\sqrt[b]{b}} \cdots \sum_{d_K=1}^{\sqrt[b]{b}} (w_{\mathbf{d},j}^{(s)} - w_{\mathbf{d},j}^{(s')})^2 \left(\frac{1}{n} \sum_{\mathbf{i}=(i_1, \dots, i_K)} \Phi_{\mathbf{d}}^2(\boldsymbol{\theta}_{\mathbf{i}}) \right)$$

Lemma A5-(i) \rightarrow

$$\geq \gamma^2 C_{2,L,K} b^{-\frac{2L+K}{K}} p^{-1} \rho(\mathbf{w}^{(s)}, \mathbf{w}^{(s')}) \quad (66)$$

where $\rho(\mathbf{w}^{(s)}, \mathbf{w}^{(s')}) = \sum_{j=1}^p \sum_{d_1=1}^{\sqrt[b]{b}} \cdots \sum_{d_K=1}^{\sqrt[b]{b}} (w_{\mathbf{d},j}^{(s)} - w_{\mathbf{d},j}^{(s')})^2$ is the hamming distance between $\mathbf{w}^{(s)}$ and $\mathbf{w}^{(s')}$.

Due to the Varshamov–Gilbert bound (Lemma 2.9 in Tsybakov (2009)), when $bp \geq 8$, there exists a subset $\Omega^0 = (\mathbf{w}^{(0)}, \dots, \mathbf{w}^{(S)}) \subseteq \Omega$ such that $S \geq 2^{bp/8}$ and $\rho(\mathbf{w}^{(s)}, \mathbf{w}^{(s')}) \geq bp/8$ for $0 \leq s \neq s' \leq S$. Since $b \geq 1$, $p \geq 8$ is a sufficient condition to guarantee $bp \geq 8$.

Now, in particular, we choose our testing set based on Ω^0 : That is, we place $M_{\mathbf{w}^{(s)}} \in \mathcal{B}^0$ if and only if $\mathbf{w}^{(s)} \in \Omega^0$. In particular this gives us that $\rho(\mathbf{w}^{(s)}, \mathbf{w}^{(s')}) \geq bp/8$. for all $w^{(s)}, w^{(s')} \in \mathcal{B}^0$ with $s \neq s'$. Then, following (66), we have that

$$\frac{1}{np} \|M_{\mathbf{w}^{(s)}} - M_{\mathbf{w}^{(s')}}\|_F^2 \geq \frac{\gamma^2 C_{2,L,K}}{8} b^{-\frac{2L}{K}}. \quad (67)$$

Now, we finally give the value that we use for b : Select $b = \left\lceil c_0 \left(\frac{n \vee p}{N} \right)^{\frac{-K}{2L+K}} \right\rceil$ for some constant $c_0 > 0$. We note that (66)-(67) hold only when $b \leq 0.48^K n$ as stated in Lemma A5. So, we need

$$N \leq c_0^{-\frac{2L+K}{K}} 0.48^{2L+K} (n \vee p) n^{\frac{2L+K}{K}}. \quad (68)$$

Furthermore, we also need

$$N \geq c_0^{-\frac{2L+K}{K}} (n \vee p) \quad (69)$$

such that

$$b = \left\lceil c_0 \left(\frac{n \vee p}{N} \right)^{\frac{-K}{2L+K}} \right\rceil \leq 2c_0 \left(\frac{n \vee p}{N} \right)^{\frac{-K}{2L+K}}. \quad (70)$$

This, finally gives us

$$\frac{1}{np} \|M_{\mathbf{w}^{(s)}} - M_{\mathbf{w}^{(s')}}\|_F^2 \geq \frac{C_{2,L,K} \gamma^2}{8} (2c_0)^{-2L/K} \left(\frac{n \vee p}{N} \right)^{\frac{2L}{2L+K}} =: 2\delta_{N,n,p}. \quad (71)$$

Then \mathcal{B}^0 is a $2\delta_{N,n,p}$ -packing set of $\mathcal{M}(L, \gamma, K)$ and the cardinality $|\mathcal{B}^0| = S + 1 \geq 2^{bp/8} + 1 = 2^{\lceil c_0 \left(\frac{n \vee p}{N} \right)^{\frac{-K}{2L+K}} \rceil \times p/8} + 1$ when $p \geq 8$.

We now show the second property (related to the KL distance) of \mathcal{B}^0 . For any matrices $M_{\mathbf{w}^{(s)}}, M_{\mathbf{w}^{(s')}} \in \mathcal{B}^0$, with the selected $b = \lceil c_0 \left(\frac{n \vee p}{N} \right)^{\frac{-K}{2L+K}} \rceil$, we have

$$\begin{aligned} K(\mathbb{P}_s \| \mathbb{P}_{s'}) &= \int \log \frac{d\mathbb{P}_s}{d\mathbb{P}_{s'}} d\mathbb{P}_s \\ &= \int \int \log \frac{\prod_{t=1}^N p(y_t, X_t | M_{\mathbf{w}^{(s)}})}{\prod_{t=1}^N p(y_t, X_t | M_{\mathbf{w}^{(s')}})} \left[\prod_{t=1}^N p(y_t, X_t | M_{\mathbf{w}^{(s)}}) dy_t dX_t \right] \\ &\rightarrow = \mathbb{E}_{X \sim \Pi} \sum_{t=1}^N \int [\log p(y_t | X_t, M_{\mathbf{w}^{(s)}}) - \log p(y_t | X_t, M_{\mathbf{w}^{(s')}})] p(y_t | X_t, M_{\mathbf{w}^{(s)}}) dy_t \\ \underbrace{(y_t | X_t, M) \sim_{i.i.d} N(\langle X_t, M \rangle, \sigma^2)}_{\rightarrow} &= \mathbb{E}_{X \sim \Pi} \sum_{t=1}^N \frac{\langle X_t, M_{\mathbf{w}^{(s)}} - M_{\mathbf{w}^{(s')}} \rangle^2}{2\sigma^2} \\ \underbrace{\left[\mathbb{E}_{X \sim \Pi} \langle X_t, M \rangle^2 = \frac{1}{np} \|M\|_F^2 \right]}_{\rightarrow} &= \frac{N}{2\sigma^2 np} \|M_{\mathbf{w}^{(s)}} - M_{\mathbf{w}^{(s')}}\|_F^2 \\ &\leq \frac{N}{2\sigma^2} \sum_{d_1=1}^{K\sqrt{b}} \cdots \sum_{d_K=1}^{K\sqrt{b}} \left(\frac{1}{n} \sum_{i=(i_1, \dots, i_K)} \Phi_d^2(\boldsymbol{\theta}_i) \right) \\ &\leq \frac{N\gamma^2}{2\sigma^2} b^{-\frac{2L}{K}} C_{1,L,K} \\ &\leq \frac{C_{1,L,K} \gamma^2}{2\sigma^2} c_0^{-\frac{2L}{K}} N \left(\frac{n \vee p}{N} \right)^{\frac{2L}{2L+K}}. \end{aligned} \quad (72)$$

Thus, Lemma A6 is proved. \blacksquare

E.2 Information-theoretic lower bounds

Given Lemma A6, we now apply the argument in Yang and Barron (1999) to yield a lower bound for error in our estimation problem with respect to Frobenius norm.

Proof [Proof of Theorem 5]

For a given $\delta_{N,n,p}$, let \mathcal{B}^0 be the $2\delta_{N,n,p}$ -packing set of $\mathcal{M}(L, \gamma, K)$ indicated by Lemma A6. We know that for any $M_s \neq M_{s'} \in \mathcal{B}^0$,

$$\frac{1}{np} \|M_s - M_{s'}\|_F^2 \geq 2\delta_{N,n,p}$$

with $\delta_{N,n,p} = \frac{C_{2,L,K}\gamma^2}{16} (2c_0)^{-2L/K} \left(\frac{n\vee p}{N}\right)^{\frac{2L}{2L+K}}$, when $c_0^{-\frac{2L+K}{K}} (n\vee p) \leq N \leq c_0^{-\frac{2L+K}{K}} 0.48^{2L+K} (n\vee p)n^{\frac{2L+K}{K}}$ for some constant $c_0 > 0$.

Let $d(M_1, M_2) = \frac{1}{np} \|M_1 - M_2\|_F^2$ and define

$$\tilde{M} = \arg \min_{M' \in \mathcal{B}^0} d(M', \widehat{M}) \in \mathcal{B}^0.$$

Let M be any matrix in the packing set \mathcal{B}^0 . If $d(M, \widehat{M}) < \delta_{N,n,p}$, then $\max \left\{ d(M, \widehat{M}), d(\tilde{M}, \widehat{M}) \right\} = d(M, \widehat{M}) < \delta_{N,n,p} \leq \delta_0 \equiv C_{2,L,K}\gamma^2 4^{-(L+2K)/K}$. Then, by the triangle inequality, we have $d(M, \widehat{M}) + d(\tilde{M}, \widehat{M}) \geq d(M, \tilde{M}) \geq 2\delta_{N,n,p}$ when $M \neq \tilde{M}$. This implies that $d(M, \widehat{M}) \geq \delta_{N,n,p}$, which contradicts $d(M, \widehat{M}) < \delta_{N,n,p}$. Therefore, if $M \neq \tilde{M}$, we must have $d(M, \widehat{M}) \geq \delta_{N,n,p}$. So, it follows that

$$\begin{aligned} \inf_{\widehat{M}} \sup_{M \in \mathcal{M}(L,\gamma,K)} \mathbb{P} \left\{ d(M, \widehat{M}) \geq \delta_{N,n,p} \right\} &\geq \inf_{\widehat{M}} \sup_{M \in \mathcal{B}^0} \mathbb{P} \left\{ d(M, \widehat{M}) \geq \delta_{N,n,p} \right\} \\ &= \inf_{\widehat{M}} \sup_{M \in \mathcal{B}^0} \mathbb{P} \left\{ M \neq \tilde{M} \right\} \\ &\geq \inf_{\widehat{M}} \mathbb{P}(M \neq \tilde{M}) \end{aligned} \quad (73)$$

where M is uniformly distributed over the $2\delta_{N,n,p}$ -packing set \mathcal{B}^0 with $|\mathcal{B}^0| \geq 2^{\lceil c_0 \left(\frac{n\vee p}{N}\right)^{\frac{-K}{2L+K}} \rceil \times p/8} + 1$ as in Lemma A6. this has reduced our problem essentially to a testing problem.

We now use this to obtain a lower bound, by considering KL-divergence here. By Lemma A6(iii), Fano's inequality (Cover and Thomas, 2012), or (Wainwright, 2019, Proposition 15.12), and the convexity of the Kullback–Leibler divergence (Wainwright, 2019, (15.34)),

$$\begin{aligned} \mathbb{P}(M \neq \tilde{M}) &\geq 1 - \frac{\frac{1}{|\mathcal{B}^0|^2} \sum_{M_s, M_{s'} \in \mathcal{B}^0} K(\mathbb{P}_s \| \mathbb{P}_{s'}) + \log 2}{\log |\mathcal{B}^0|} \\ \xrightarrow{\text{[Lemma A6]}} &\geq 1 - \frac{\frac{C_{1,L,K}\gamma^2}{2\sigma^2} c_0^{-\frac{2L}{K}} N \left(\frac{n\vee p}{N}\right)^{\frac{2L}{2L+K}} + \log 2}{\lceil c_0 \left(\frac{n\vee p}{N}\right)^{\frac{-K}{2L+K}} \rceil p \log 2} \\ \xrightarrow{\text{[}bp \geq 8\text{]}} &\geq \frac{7}{8} - \frac{C_{1,L,K}\gamma^2 c_0^{-\frac{2L+K}{K}} (n \vee p)}{2(\log 2)\sigma^2 p}. \end{aligned} \quad (74)$$

Consider $n = \kappa p$ for some $\kappa > 0$. Let

$$c_0 = \left(\frac{4 \max(\kappa, 1) \gamma^2 C_{1,L,K}}{3 \log 2 \sigma^2} \right)^{\frac{K}{2L+K}}. \quad (75)$$

Then,

$$\mathbb{P}(M \neq \tilde{M}) \geq 7/8 - \frac{\gamma^2 C_{1,L,K} \max(\kappa, 1) c_0^{-\frac{2L+K}{K}}}{2(\log 2) \sigma^2} = 7/8 - 3/8 = 1/2 \quad (76)$$

Thus, it follows from (73) and (76) that

$$\inf_{\widehat{M}} \sup_{M \in \mathcal{M}(L, \gamma, K)} \mathbb{P} \left\{ \frac{1}{np} \|\widehat{M} - M\|_F^2 \geq A \left(\frac{n \vee p}{N} \right)^{\frac{2L}{2L+K}} \right\} \geq 1/2. \quad (77)$$

where $A = \frac{C_{2,L,K} \gamma^2}{16} (2c_0)^{-2L/K}$. With the selection of c_0 in (75), A depends on $L, K, \gamma, \kappa, \sigma^2$. Thus, Theorem 5 is proved. ■

References

- Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.
- Francis R Bach. Consistency of trace norm minimization. *Journal of Machine Learning Research*, 9(Jun):1019–1048, 2008.
- T Tony Cai, Wen-Xin Zhou, et al. Matrix completion via max-norm constrained optimization. *Electronic Journal of Statistics*, 10(1):1493–1525, 2016.
- E. J. Candes and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, June 2010a. ISSN 0018-9219. doi: 10.1109/JPROC.2009.2035722.
- E. J. Candes and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, May 2010. ISSN 0018-9448. doi: 10.1109/TIT.2010.2044061.
- Emmanuel J Candes and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010b.
- Emmanuel J Candes and Yaniv Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359, 2011.
- Sourav Chatterjee et al. Matrix estimation by universal singular value thresholding. *The Annals of Statistics*, 43(1):177–214, 2015.

- Alexander L. Chistov and Dima Grigoriev. Complexity of quantifier elimination in the theory of algebraically closed fields. In *MFCS*, 1984.
- Thomas M Cover and Joy A Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.
- Ronald A DeVore and George G Lorentz. *Constructive approximation*, volume 303. Springer Science & Business Media, 1993.
- Jicong Fan and Jieyu Cheng. Matrix completion by deep matrix factorization. *Neural Networks*, 98:34 – 41, 2018. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2017.10.007>. URL <http://www.sciencedirect.com/science/article/pii/S0893608017302502>.
- Jicong Fan and Tommy Chow. Deep learning based matrix completion. *Neurocomputing*, 266:540–549, 2017.
- Jicong Fan and Tommy W.S. Chow. Non-linear matrix completion. *Pattern Recognition*, 77:378–394, 5 2018. ISSN 0031-3203. doi: 10.1016/j.patcog.2017.10.014.
- Ming Fan, You Zhang, Zhenyu Fu, Maosheng Xu, Shiwei Wang, Sangma Xie, Xin Gao, Yue Wang, and Lihua Li. A deep matrix completion method for imputing missing histological data in breast cancer by integrating dce-mri radiomics. *Medical Physics*, 48(12):7685–7697, 2021.
- Maryam Fazel. *Matrix rank minimization with applications*. PhD thesis, PhD thesis, Stanford University, 2002.
- Joe G Greener, Lewis Moffat, and David T Jones. Design of metalloproteins and novel protein folds using variational autoencoders. *Scientific reports*, 8(1):16189, 2018.
- Martin Jaggi and Marek Sulovský. A simple algorithm for nuclear norm regularized problems. In *ICML*, 2010.
- Zhuxi Jiang, Yin Zheng, Huachun Tan, Bangsheng Tang, and Hanning Zhou. Variational deep embedding: An unsupervised and generative approach to clustering. *arXiv preprint arXiv:1611.05148*, 2016.
- Ke Jin, Bo Li, Hong Yan, and Xiao-Fei Zhang. Imputing dropouts for single-cell rna sequencing based on multi-objective optimization. *Bioinformatics*, 38(12):3222–3230, 2022.
- Olga Klopp et al. Noisy low-rank matrix completion with general sampling distribution. *Bernoulli*, 20(1):282–303, 2014.
- Vladimir Koltchinskii, Karim Lounici, and Alexandre B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.*, 39(5):2302–2329, 10 2011. doi: 10.1214/11-AOS894. URL <https://doi.org/10.1214/11-AOS894>.
- Andreas Kopf and Manfred Claassen. Latent representation learning in biology and translational medicine. *Patterns*, 2(3):100198, 2021.

- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, August 2009. ISSN 0018-9162. doi: 10.1109/MC.2009.263. URL <http://dx.doi.org/10.1109/MC.2009.263>.
- M. Laurent. *Matrix completion problems*, pages 221–229. Kluwer Academic Publishers, Netherlands, 2001. ISBN 0792370287. Pagination: 9.
- Yihua Li, Devavrat Shah, Dogyoon Song, and Christina Lee Yu. Nearest neighbors for matrix estimation interpreted as blind regression for latent variable model. *IEEE Transactions on Information Theory*, 66(3):1760–1784, 2019.
- George C Linderman, Jun Zhao, Manolis Roulis, Piotr Bielecki, Richard A Flavell, Boaz Nadler, and Yuval Kluger. Zero-preserving imputation of single-cell rna-seq data. *Nature communications*, 13(1):192, 2022.
- Z. Liu and L. Vandenberghe. Interior-point method for nuclear norm approximation with application to system identification. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1235–1256, 2010. doi: 10.1137/090755436. URL <https://doi.org/10.1137/090755436>.
- Brett A McKinney, David M Reif, Marylyn D Ritchie, and Jason H Moore. Machine learning for detecting gene-gene interactions: a review. *Applied bioinformatics*, 5:77–88, 2006.
- Sahand Negahban and Martin J Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, pages 1069–1097, 2011.
- Benjamin Recht. A simpler approach to matrix completion. *J. Mach. Learn. Res.*, 12: 3413–3430, December 2011. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1953048.2185803>.
- Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501, August 2010. ISSN 0036-1445. doi: 10.1137/070697835. URL <http://dx.doi.org/10.1137/070697835>.
- Jacob Schreiber, Timothy J Durham, Jeffrey Bilmes, and William Stafford Noble. Multi-scale deep tensor factorization learns a latent representation of the human epigenome. *bioRxiv*, 2018. doi: 10.1101/364976. URL <https://www.biorxiv.org/content/early/2018/07/08/364976>.
- Amit Singer and Mihai Cucuringu. Uniqueness of low-rank matrix completion by rigidity theory. *SIAM Journal on Matrix Analysis and Applications*, 31(4):1621–1641, 2010.
- Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *International Conference on Computational Learning Theory*, pages 545–560. Springer, 2005.
- Nathan Srebro, Jason D. M. Rennie, and Tommi S. Jaakkola. Maximum-margin matrix factorization. In *Proceedings of the 17th International Conference on Neural Information Processing Systems*, NIPS’04, pages 1329–1336, Cambridge, MA, USA, 2004. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2976040.2976207>.

- Bharath Srinivasan. A guide to the michaelis–menten equation: steady state and beyond. *The FEBS journal*, 289(20):6086–6098, 2022.
- Alexandre B Tsybakov. *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- Alexandre B Tsybakov. Introduction to nonparametric estimation. revised and extended from the 2004 french original. translated by vladimir zaiats. 2009.
- Madeleine Udell and Alex Townsend. Why are big data matrices approximately low rank? *SIAM Journal on Mathematics of Data Science*, 1(1):144–160, 2019.
- Sara van de Geer. Estimation and testing under sparsity. *Lecture Notes in Mathematics*, 2159, 2016.
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Juexin Wang, Anjun Ma, Yuzhou Chang, Jianting Gong, Yuexu Jiang, Ren Qi, Cankun Wang, Hongjun Fu, Qin Ma, and Dong Xu. scgnn is a novel graph neural network framework for single-cell rna-seq analyses. *Nature communications*, 12(1):1882, 2021.
- Yichen Wang, Robert Chen, Joydeep Ghosh, Joshua C Denny, Abel Kho, You Chen, Bradley A Malin, and Jimeng Sun. Rubik: Knowledge guided tensor factorization and completion for health data analytics. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1265–1274. ACM, 2015.
- Yiheng Wang, Tong Liu, Dong Xu, Huidong Shi, Chaoyang Zhang, Yin-Yuan Mo, and Zheng Wang. Predicting dna methylation state of cpg dinucleotide using genome topological features and deep networks. *Scientific reports*, 6(1):1–15, 2016.
- G Alistair Watson. Characterization of the subdifferential of some matrix norms. *Linear algebra and its applications*, 170:33–45, 1992.
- Derry Tanti Wijaya, Brendan Callahan, John Hewitt, Jie Gao, Xiao Ling, Marianna Apidianaki, and Chris Callison-Burch. Learning translations via matrix completion. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1452–1463. Association for Computational Linguistics, 2017. doi: 10.18653/v1/D17-1152. URL <http://aclweb.org/anthology/D17-1152>.
- G. Xia, H. Sun, B. Chen, Q. Liu, L. Feng, G. Zhang, and R. Hang. Nonlinear low-rank matrix completion for human motion recovery. *IEEE Transactions on Image Processing*, 27(6):3011–3024, June 2018. ISSN 1057-7149. doi: 10.1109/TIP.2018.2812100.
- Yuhong Yang and Andrew Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, pages 1564–1599, 1999.

Wenchao Yu, Guangxiang Zeng, Ping Luo, Fuzhen Zhuang, Qing He, and Zhongzhi Shi. Embedding with autoencoder regularization. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 208–223. Springer, 2013.

Yifan Zhou, Kaixuan Luo, Mengjie Chen, and Xin He. A novel bayesian factor analysis method improves detection of genes and biological processes affected by perturbations in single-cell crispr screening. *bioRxiv*, pages 2022–02, 2022.