

Optimal Estimation of Bacterial Growth Rates Based on Permuted Monotone Matrix

BY RONG MA

*Department of Biostatistics, Epidemiology and Informatics, Perelman School of Medicine,
University of Pennsylvania, Philadelphia, Pennsylvania 19104, U.S.A.*

rongm@upenn.edu

T. TONY CAI

*Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia,
Pennsylvania 19104, U.S.A.*

tcai@wharton.upenn.edu

HONGZHE LI

*Department of Biostatistics, Epidemiology and Informatics, Perelman School of Medicine,
University of Pennsylvania, Philadelphia, Pennsylvania 19104, U.S.A.*

hongzhe@penmedicine.upenn.edu

SUMMARY

Motivated by the problem of estimating the bacterial growth rates for genome assemblies from shotgun metagenomic data, we consider the permuted monotone matrix model $Y = \Theta\Pi + Z$, where $Y \in \mathbb{R}^{n \times p}$ is observed, $\Theta \in \mathbb{R}^{n \times p}$ is an unknown approximately rank-one signal matrix with monotone rows, $\Pi \in \mathbb{R}^{p \times p}$ is an unknown permutation matrix, and $Z \in \mathbb{R}^{n \times p}$ is the noise matrix. This paper studies the estimation of the extreme values associated to the signal matrix Θ , including its first and last columns, as well as their difference (the range vector). Treating these estimation problems as compound decision problems, minimax rate-optimal estimators are constructed using the spectral column sorting method. Numerical experiments through simulated and synthetic microbiome metagenomic data are presented, showing the superiority of the proposed methods over the alternatives. The methods are illustrated by comparing the growth rates of gut bacteria between inflammatory bowel disease patients and normal controls.

Some key words: Extreme values; Metagenomics; Minimax lower bounds; Permutation; Spectral method

1. INTRODUCTION

The statistical problem considered in this paper is motivated by the problem of estimating the bacterial growth dynamics using shotgun metagenomics data. Several methods have been developed to quantify the bacterial growth dynamics based on shotgun metagenomics data by extrapolating particular patterns in the sequencing read coverages resulted from the bidirectional microbial DNA replications (Myhrvold et al., 2015; Abel et al., 2015; Korem et al., 2015; Brown et al., 2016). For bacterial species with known complete genome sequences, Korem et al. (2015) proposed to use the peak-to-trough ratio (PTR) of read coverages to quantify the bacterial growth rates after aligning the sequencing reads to the bacterial genomes. Besides quantifying the growth

rates for the bacteria with complete genome sequences, it is also of great importance to estimate the growth rates of incomplete genome assemblies, where the coverages of contigs are observed in multiple samples. However, the order the contigs is only known up to an unknown permutation.

Recently, Gao & Li (2018) developed a computational algorithm (DEMIC) that accurately estimates the growth dynamics of a given assembled species by taking advantage of highly fragmented contigs assembled from multiple samples. DEMIC is based on the following permuted monotone matrix model:

$$Y = \Theta\Pi + Z \quad (1)$$

where the observed data $Y \in \mathbb{R}^{n \times p}$ is the matrix of the preprocessed contig coverage for a given bacterial species. Specifically, the entry Y_{ij} represents the log-transformed averaged read counts of the j -th contig of the bacterial species for the i -th sample after the preprocessing steps, including genome assemblies, GC adjustment of read counts and outlier filtering. In practice, the data set is usually high-dimensional in the sense that the number of contigs p far exceeds the sample size n , so throughout we assume $p \gg n$. The signal matrix $\Theta \in \mathbb{R}^{n \times p}$ represents the true log-transformed coverage matrix of n samples and p contigs, where each row is monotone due to the bi-directional DNA replication mechanism (Brown et al., 2016; Gao & Li, 2018), $Z \in \mathbb{R}^{n \times p}$ is the noise matrix, and $\Pi \in \mathbb{R}^{p \times p}$ is a permutation matrix, corresponding to some permutation π from the symmetric group \mathcal{S}_p . Ma et al. (2020a) developed methods for optimally recovering the underlying permutation π from Y . In particular, considering the loss function being either the 0-1 loss or the normalized Kendall's τ distance, a minimax optimal permutation estimator is proposed and theoretically analyzed under various parameter spaces.

In addition to the monotonicity constraint imposed on the rows of Θ , real metagenomic data sets also suggest approximate linear relationship between the contig positions and their log-coverages for each sample, which indicates approximately rank-one structure of Θ , after certain normalization. As an example, Figure 1 shows the normalized log-contig counts of an assembled bacterial genome for three individuals along estimated contig orders, suggesting the aforementioned approximate linear or rank-one structure (see Section 3.1 for details).

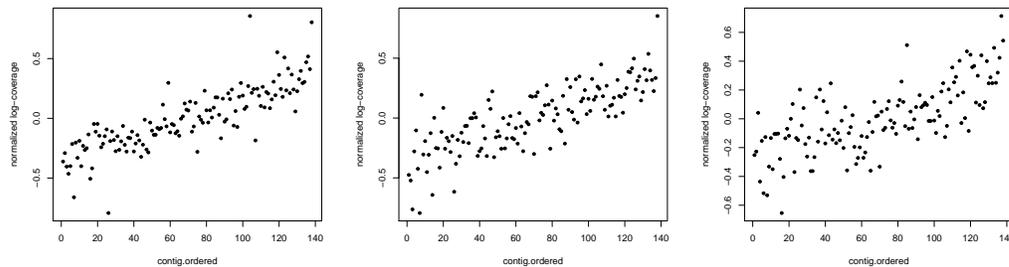


Fig. 1. The log-coverages of ordered contigs of an assembled bacteria species for from 3 individuals with inflammatory bowel disease from the iHMP study, detailed in Section 5.3.

Under the permuted monotone matrix model, one can relate the two extreme columns Θ_R and Θ_L , i.e., the first and the last columns of Θ , to the log-transformed true peak and trough coverages of a given bacterial species, and define their difference $\mathbf{R}(\Theta) = \Theta_R - \Theta_L$ as the true log-PTRs that characterize the bacterial growth rates over n samples. The goal of this paper is to provide a rigorous statistical framework for optimal estimation of the extreme values in the

approximately rank-one permuted monotone matrix model, including Θ_R and Θ_L and the range vector $\mathbf{R}(\Theta)$. Based on the idea of spectral column sorting and the theory of low-rank matrix estimation, we develop computationally efficient estimators for the extreme columns and the range vector. In particular, the minimax optimality of the proposed methods are theoretically established and empirically illustrated with numerical experiments, which also justify its applicability in analyzing real data sets such as the microbiome metagenomics data.

Throughout the paper, we define the permutation π as a bijection from the set $\{1, 2, \dots, p\}$ onto itself. For simplicity, we denote $\pi = (\pi(1), \pi(2), \dots, \pi(p))$. All permutations of the set $\{1, 2, \dots, p\}$ form a symmetric group, equipped with the function composition operation \circ , denoted as \mathcal{S}_p . For any $\pi \in \mathcal{S}_p$, we denote $\pi^{-1} \in \mathcal{S}_p$ as its group inverse, so that $\pi \circ \pi^{-1} = \pi^{-1} \circ \pi = id$. In particular, we may use π and its corresponding permutation matrix $\Pi \in \mathbb{R}^{p \times p}$ interchangeably, depending on the context. For a vector $\mathbf{a} = (a_1, \dots, a_n)^\top \in \mathbb{R}^n$, we define the ℓ_p norm $\|\mathbf{a}\|_p = (\sum_{i=1}^n |a_i|^p)^{1/p}$, and the ℓ_∞ norm $\|\mathbf{a}\|_\infty = \max_{1 \leq j \leq n} |a_j|$. For a matrix $\Theta \in \mathbb{R}^{p_1 \times p_2}$, we denote $\Theta_{\cdot i} \in \mathbb{R}^{p_1}$ as its i -th column and denote $\Theta_{i \cdot} \in \mathbb{R}^{p_2}$ as its i -th row. We write $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. Furthermore, for sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n = o(b_n)$ if $\lim_n a_n/b_n = 0$, and write $a_n = O(b_n)$, $a_n \lesssim b_n$ or $b_n \gtrsim a_n$ if there exists a constant C such that $a_n \leq Cb_n$ for all n . We write $a_n \asymp b_n$ if $a_n \lesssim b_n$ and $a_n \gtrsim b_n$. Lastly, C, C_0, C_1, \dots are constants that may vary from place to place.

2. EXTREME VALUE ESTIMATION VIA SPECTRAL SORTING

2.1. Spectral Sorting and Extreme Column Localization

A crucial step for estimating the extreme columns is to sort the permuted columns in order to identify the extreme ones. In this section, we introduce a spectral approach for localizing the permuted columns. Toward this end, for any Θ with monotone rows, we consider the row-centered matrix

$$\Theta' = \Theta \left(\mathbf{I}_p - \frac{1}{p} \mathbf{e} \mathbf{e}^\top \right) \in \mathbb{R}^{n \times p}, \quad (2)$$

where $\mathbf{e} = (1, \dots, 1)^\top \in \mathbb{R}^p$. Intuitively, Θ' is invariant to the row averages of Θ and preserves the row-monotonicity structure as well as the distances between the columns of Θ . The singular value decomposition (SVD) of Θ' can be written as

$$\Theta' = \sum_{i=1}^r \lambda_i \mathbf{u}_i \mathbf{v}_i^\top, \quad \text{for some } r \leq \min\{n, p\}, \quad (3)$$

where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_r$ are the ordered singular values of Θ' and \mathbf{u}_i and \mathbf{v}_i are the left and right singular vectors corresponding to λ_i , respectively. To overcome the identifiability issue, we assume

(A) λ_1 has multiplicity one and the first nonzero component of \mathbf{v}_1 is negative.

The following proposition provides an important insight that the row-monotonicity of a matrix actually implies the monotonicity of the components of its leading right singular vector \mathbf{v}_1 . This property plays a fundamental role in analyzing the permuted monotone matrix model.

PROPOSITION 1. *Let Θ be a row-monotone matrix, whose row-centered version Θ' defined in (2) satisfies (A). Then its first right singular vector $\mathbf{v}_1 = (v_{11}, \dots, v_{1p})^\top$ is a centered monotone vector, i.e., $\sum_{i=1}^p v_{1i} = 0$ and $v_{11} \leq v_{12} \leq \dots \leq v_{1p}$. In addition, the sign vector $\text{sgn}(\mathbf{u}_1)$ indicates the direction of monotonicity of the rows of Θ' (or Θ).*

From the above proposition, the relative orders of the columns of Θ' (and Θ) are qualitatively preserved by the leading right singular vector \mathbf{v}_1 , whereas the directions of monotonicity for different rows are coded by the leading left singular vector \mathbf{u}_1 . As a result, given a column-permuted and noisy matrix Y in (1), one could localize the extreme columns Θ_R and Θ_L in $\Theta\Pi$ by considering the row-normalized observation matrix $X = Y(\mathbf{I}_p - \frac{1}{p}\mathbf{e}\mathbf{e}^\top)$ and its first right singular vector, i.e.,

$$\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_p)^\top = \arg \max_{\mathbf{v} \in \mathbb{R}^p: \|\mathbf{v}\|_2=1} \mathbf{v}^\top X^\top X \mathbf{v}. \quad (4)$$

In accordance with Proposition 1, it was shown by Ma et al. (2020a) that the order statistics $\{\hat{v}_{(1)}, \dots, \hat{v}_{(p)}\}$ can be used to optimally recover the permutation π , or the original column orders, by tracing back the permutation map between the elements of $\hat{\mathbf{v}}$ and their order statistics. Clearly, for extreme column localization, the extreme values statistics $\hat{v}_{(1)}$ and $\hat{v}_{(p)}$ are more relevant. In fact, it is shown in the subsequent section that, minimax optimal estimators can be constructed using such spectral extreme values estimates.

2.2. Compound Decision Problem and the Proposed Estimators

The problem of estimating Θ_R , Θ_L or \mathbf{R} consists of n individual sub-problems, namely, estimating each of its n coordinates. Following the concept proposed by Robbins (1951, 1964) and further elaborated in Samuel (1967); Copas (1969); Zhang (2003) and Brown & Greenshtein (2009), among many others, we observe that the problem of finding their minimax optimal estimators is a compound statistical decision problem, as the n individual sub-problems are amalgamated into one larger problem through the combined risk (10). Moreover, although the observations over n samples are independent, it has been argued that, in general, for a compound decision problem, usually the simple estimators, where only the i -th sample is used to estimate the i -th coordinate, are suboptimal; in contrast, a minimax optimal estimator should be compound in the sense that multiple samples are used for the estimation of each coordinate.

In light of our discussion in Section 2.1 as to the fundamental role of $(\lambda_1, \mathbf{u}_1, \mathbf{v}_1)$, we introduce our proposed estimators for the extreme columns as

$$\hat{\Theta}_R^* = \hat{v}_{(p)} X \hat{\mathbf{v}} + \frac{1}{p} Y \mathbf{e} \in \mathbb{R}^n, \quad \hat{\Theta}_L^* = \hat{v}_{(1)} X \hat{\mathbf{v}} + \frac{1}{p} Y \mathbf{e} \in \mathbb{R}^n, \quad (5)$$

and our proposed range estimator as

$$\hat{\mathbf{R}}^* = \hat{\Theta}_R^* - \hat{\Theta}_L^* = (\hat{v}_{(p)} - \hat{v}_{(1)}) X \hat{\mathbf{v}}, \quad (6)$$

where we recall that $\hat{\mathbf{v}}$ is defined in (4) and $\hat{v}_{(i)}$ is the i -th smallest order statistic among $\{\hat{v}_1, \dots, \hat{v}_p\}$. By construction, the proposed extreme column estimators (5) are compound estimators, and each of them consists of two parts: the first part estimates the extreme columns of the row-centered matrix Θ' whereas the second part compensates the row-specific mean effects. In particular, in accordance with the observations made in Section 1, to construct the first parts of $\hat{\Theta}_R^*$ and $\hat{\Theta}_L^*$, the approximately rank-one structure $\Theta'_{\ell} \approx \lambda_1 v_{1\ell} \mathbf{u}$ for $\ell \in \{1, p\}$, is incorporated with $v_{1\ell}$ estimated by $\hat{v}_{(\ell)}$ and $\lambda_1 \mathbf{u}_1$ estimated by $X \hat{\mathbf{v}}$.

Ma et al. (2020a) developed an optimal estimator for the permutation π under the model (1). Specifically, let $\tau: \mathbb{R}^p \rightarrow \mathcal{S}_p$ be the ranking operator, which is defined such that for any vector $\mathbf{x} \in \mathbb{R}^p$, $\tau(\mathbf{x})$ is the vector of ranks for components of \mathbf{x} in increasing order – whenever there are ties, increasing orders are assigned from left to right. The best linear projection estimator of π was defined as

$$\hat{\pi} = [\tau(\hat{\mathbf{v}})]^{-1}. \quad (7)$$

This permutation estimator can be used to construct a natural two-step estimator of the two extreme columns. In the first step, we recover/sort the columns of Y to obtain the sorted matrix $\check{Y} = [Y_{\cdot\hat{\pi}(1)} \ Y_{\cdot\hat{\pi}(2)} \ \dots \ Y_{\cdot\hat{\pi}(p)}]$. Intuitively, the column-sorted matrix \check{Y} is expected to be close to Θ . In the second step, we fit a simple linear regression between each row of \check{Y} and the sorted projection scores $(\hat{v}_{(1)}, \hat{v}_{(2)}, \dots, \hat{v}_{(p)})$, which characterize the column relative locations. Denote the fitted intercepts as $\alpha = (\alpha_1, \dots, \alpha_n)^\top$ and the slopes as $\beta = (\beta_1, \dots, \beta_n)^\top$. We define the two-step regression estimators as

$$\hat{\Theta}_L^{Reg} = \alpha + \beta \hat{v}_{(1)}, \quad \hat{\Theta}_R^{Reg} = \alpha + \beta \hat{v}_{(p)}, \quad \hat{\mathbf{R}}^{Reg} = \beta(\hat{v}_{(p)} - \hat{v}_{(1)}). \quad (8)$$

It is easy to check (see Section 3 of Supplementary Material (Ma et al., 2020b)) that under the conditions of Proposition 1, it holds that

$$\hat{\Theta}_L^{Reg} = \hat{\Theta}_L^*, \quad \hat{\Theta}_R^{Reg} = \hat{\Theta}_R^*, \quad \hat{\mathbf{R}}^{Reg} = \hat{\mathbf{R}}^*, \quad (9)$$

This interesting observation provides another way of understanding our proposed estimators.

3. THEORETICAL PROPERTIES

3.1. Risk Upper Bounds for the Extreme Column Estimators

In what follows, we study the theoretical properties of our proposed estimator $\hat{\Theta}_R^*$, as the results for $\hat{\Theta}_L$ would hold in parallel. Towards this end, we consider the normalized ℓ_2 distances $\frac{1}{\sqrt{n}} \|\hat{\Theta}_R - \Theta_R\|_2$ and denote the corresponding estimation risk as

$$\mathcal{R}_R(\hat{\Theta}_R) = \frac{1}{\sqrt{n}} \mathbb{E} \|\hat{\Theta}_R - \Theta_R\|_2. \quad (10)$$

We first define the set of monotone matrices

$$\mathcal{D} = \left\{ \Theta = (\theta_{ij}) \in \mathbb{R}^{n \times p} : \begin{array}{l} \text{for each } 1 \leq i \leq n, \text{ either } \theta_{i,j} \leq \theta_{i,j+1} \text{ for all } j, \\ \text{or } \theta_{i,j} \geq \theta_{i,j+1} \text{ for all } j \end{array} \right\}.$$

Recall that the row-centered version of Θ , namely Θ' , has the SVD given by (3). Consequently, throughout, we consider the following parameter space for (Θ, π)

$$\mathcal{D}_R(t, \beta) = \left\{ (\Theta, \pi) \in \mathcal{D} \times \mathcal{S}_p : \begin{array}{l} \text{(A) holds, } 0 \leq v_{1p} \leq \beta, \\ \lambda_1 \in [t/8, 8t], \sum_{i=2}^r \lambda_i \leq \sigma \sqrt{\log p} \end{array} \right\}, \quad (11)$$

with $t \geq 0$ and $p^{-1/2} \leq \beta \leq 1$. Here the constraint on β is natural since \mathbf{v}_1 is a unit vector and β is no less than the order of its largest component. Intuitively, the hyper-parameters (t, β) characterize the global signal strength as well as the relative position of the extreme column Θ_R shared by the signal matrices in $\mathcal{D}_R(t, \beta)$, while the condition $\sum_{i=2}^r \lambda_i \leq \sigma \sqrt{\log p}$ imposes an approximately rank-one structure on the row-centered Θ .

To simplify notation, we define the rate function $\psi = \psi(n, p) = \sqrt{(\log p/n)}$. The following theorem provides a uniform risk upper bound of the proposed estimator $\hat{\Theta}_R^*$ over $\mathcal{D}_R(t, \beta)$.

THEOREM 1 (UNIFORM UPPER BOUND). *Suppose the pair (t, β_R) satisfies $p^{-1/2} \leq \beta_R \leq 1$, $t^2 \gtrsim \sigma^2 \left[\frac{1}{\beta_R^2} \wedge \left\{ \frac{1}{\psi^2} + \frac{1}{\psi} \sqrt{\left(\frac{p}{n \log p} \right)} \right\} \right] n \log p$ and the noise matrix Z has independent sub-Gaussian entries Z_{ij} with parameter σ^2 . Then*

$$\sup_{\mathcal{D}_R(t, \beta_R)} \mathcal{R}_R(\hat{\Theta}_R^*) \lesssim \frac{\beta_R t}{\sqrt{n}} \left[\frac{\sigma \sqrt{\{(t^2 + \sigma^2 p)n\}}}{t^2} \wedge 1 \right] + \sigma \psi. \quad (12)$$

The risk upper bound (12) consists of two components. In the first component, the factor $[\sigma\sqrt{\{(t^2 + \sigma^2 p)n\}/t^2} \wedge 1]$ is the error from estimating the leading left singular vector \mathbf{u}_1 by its sample counterpart, whereas the factor $\beta_R t/\sqrt{n}$ reflects the overall magnitude of the extreme column Θ_R of the matrices in $\mathcal{D}_R(t, \beta_R)$. As for the second component $\sigma\psi(n, p)$, it comes from using the order statistic $\hat{v}_{(p)}$ to estimate the largest component of \mathbf{v}_1 .

Interestingly, about the first component, we observe two phase transitions when t^2 passes $\sigma^2\sqrt{(np)}$ and $\sigma^2 p$, respectively. Specifically, in (12), we have

$$\frac{\beta_R t}{\sqrt{n}} \left[\frac{\sigma\sqrt{\{(t^2 + \sigma^2 p)n\}}}{t^2} \wedge 1 \right] \asymp \begin{cases} \beta_R t/\sqrt{n} & \text{if } t^2 \lesssim \sigma^2\sqrt{(np)}, \\ \frac{\beta_R \sigma\sqrt{(t^2 + \sigma^2 p)}}{t} & \text{if } \sigma^2\sqrt{(np)} \lesssim t^2 \lesssim \sigma^2 p, \\ \beta_R \sigma & \text{if } t^2 \gtrsim \sigma^2 p. \end{cases}$$

From the theory of low-rank matrix estimation (Cai & Zhang, 2018), the quantity $\sigma^2\sqrt{(np)}$ is the critical point, below which it is impossible to estimate the singular vector \mathbf{u}_1 . Hereafter we refer the collection of parameter spaces $\{\mathcal{D}_R(t, \beta_R) : t^2 \leq \sigma^2\sqrt{(np)}\}$, $\{\mathcal{D}_R(t, \beta_R) : \sigma^2\sqrt{(np)} \lesssim t^2 \lesssim \sigma^2 p\}$ and $\{\mathcal{D}_R(t, \beta_R) : t^2 \gtrsim \sigma^2 p\}$ as the ‘‘weak signal-to-noise ratio’’ regime, the ‘‘intermediate signal-to-noise ratio’’ regime, and the ‘‘strong signal-to-noise ratio’’ regime, respectively.

To see the implications of the condition

$$t^2 \gtrsim \sigma^2 \left[\frac{1}{\beta_R^2} \wedge \left\{ \frac{1}{\psi^2} + \frac{1}{\psi} \sqrt{\left(\frac{p}{n \log p} \right)} \right\} \right] n \log p \quad (13)$$

of Theorem 1 on the critical events $t^2 \asymp \sigma^2\sqrt{(np)}$ and $t^2 \asymp \sigma^2 p$, we note that, as long as $\beta_R \gg (n/p)^{1/4}$, by ignoring the logarithmic factors, the right-hand side of the condition (13) is asymptotically smaller than both critical points $\sigma^2\sqrt{(np)}$ and $\sigma^2 p$, so that both phase transitions exist under the condition of Theorem 1.

3.2. Minimax Rates and Optimality of the Extreme Column Estimators

Now we establish the minimax rate of convergence and the optimality of the proposed extreme column estimator $\hat{\Theta}_R^*$ over the parameter space $\mathcal{D}_R(t, \beta_R)$. Specifically, for some given (t, β) , we define the minimax risks over $\mathcal{D}_R(t, \beta_R)$ as $\inf_{\hat{\Theta}_R} \sup_{\mathcal{D}_R(t, \beta_R)} \mathcal{R}_R(\hat{\Theta}_R)$ where the infimum is over all the possible estimators obtained from the data. The following theorem provides the minimax lower bound of the estimation risk under the Gaussian noise.

THEOREM 2 (MINIMAX LOWER BOUND). *Suppose Z in model (1) has i.i.d. entries $Z_{ij} \sim N(0, \sigma^2)$. Then, for any $\mathcal{D}_R(t, \beta)$ such that $t^2 \geq c_0 \left(\frac{1 - \beta_R^2}{\beta_R^2} \sigma^2 \log p + \frac{\beta_R^2}{1 - \beta_R^2} \sigma^2 p \right)$ and $c_1 p^{-1/2} \sqrt{\log p} \leq \beta_R \leq c_2$, for sufficiently large (n, p) and some constants $c_0, c_1 > 0$ and $0 < c_2 < 1$, it holds that*

$$\inf_{\hat{\Theta}_R} \sup_{\mathcal{D}_R(t, \beta_R)} \mathcal{R}_R(\hat{\Theta}_R) \gtrsim \frac{\beta_R t}{\sqrt{n}} \left[\frac{\sigma\sqrt{\{(t^2 + \sigma^2 p)n\}}}{t^2} \wedge 1 \right] + \sigma\psi. \quad (14)$$

The proof of Theorem 2 is involved. The main difficulty lies in the non-linearity and multi-dimensionality of the maps from the original parameter Θ to its extreme columns of interest. As the lower bound contains several components, we essentially derived three distinct minimax lower bounds corresponding to different worst-case scenarios. In addition to adopting the existing techniques such as the sphere packing of the Grassmannian manifolds, we developed two novel lower bound techniques to facilitate the proof of the minimax lower bound. The details can be found in the Supplementary Material (Ma et al., 2020b).

Combining the upper and the lower bounds, we obtain the exact minimax rate for estimating Θ_R . Specifically, under the conditions of Theorems 1 and 2, i.e., for $Z_{ij} \sim_{i.i.d.} N(0, \sigma^2)$ and

$$t^2 \gtrsim \sigma^2 \left[\frac{1}{\beta_R^2} \wedge \left\{ \frac{1}{\psi^2} + \frac{1}{\psi} \sqrt{\left(\frac{p}{n \log p} \right)} \right\} \right] n \log p + \left(\frac{1 - \beta_R^2}{\beta_R^2} \sigma^2 \log p + \frac{\beta_R^2 \sigma^2 p}{1 - \beta_R^2} \right), \quad (15)$$

we have

$$\inf_{\hat{\Theta}_R} \sup_{\mathcal{D}_R(t, \beta_R)} \mathcal{R}_R(\hat{\Theta}_R) \asymp \frac{\beta_R t}{\sqrt{n}} \left[\frac{\sigma \sqrt{\{(t^2 + \sigma^2 p)n\}}}{t^2} \wedge 1 \right] + \sigma \psi, \quad (16)$$

where the optimal rate is attained by our proposed estimator $\hat{\Theta}_R^*$. To make better sense of condition (15), we note that, if we calibrate $\beta_R \asymp (n/p)^{1/4}$, then by ignoring the logarithmic factors, condition (15) is equivalent to $t^2 \gtrsim \sigma^2 \sqrt{np}$, which means the minimax rate can essentially be established over the intermediate to strong SNR regime, where the minimax rate is

$$\inf_{\hat{\Theta}_R} \sup_{\mathcal{D}_R(t, \beta_R)} \mathcal{R}_R(\hat{\Theta}_R) \asymp \frac{\sigma n^{1/4} \sqrt{(t^2 + \sigma^2 p)}}{p^{1/4} t} + \sigma \psi. \quad (17)$$

As a consequence of the phase transition phenomena pointed out earlier, some interesting insights about the interplay between the global signal strength t^2 , the dimensionality of the problem, the hardness of estimating Θ_R and that of estimating the leading left singular vector \mathbf{u}_1 , can be obtained. Specifically, we observe that (i) within the intermediate SNR regime ($\sigma^2 \sqrt{np} \lesssim t^2 \lesssim \sigma^2 p$), increasing the signal strength t^2 will reduce the difficulty of estimating \mathbf{u}_1 and therefore the rate for estimating Θ_R , and (ii) within the strong SNR regime ($t^2 \gtrsim \sigma^2 p$), the difficulty of estimating Θ_R no longer depends on t^2 , as in this case the improved estimation of \mathbf{u}_1 is neutralized by the increased magnitude of Θ_R . Especially, all the above rate analysis is subjected to a possible lower bound of $\psi(n, p)$.

Moreover, since the above minimax optimal rates are simultaneously attained by the proposed estimator $\hat{\Theta}_R^*$ regardless of the specific value of the underlying indices (t, β_R) , then, under the sub-Gaussian noise, $\hat{\Theta}_R^*$ is minimax rate-adaptive over the collection of parameter spaces $\mathcal{C} = \{\mathcal{D}_R(t, \beta_R) : p^{-1/2} c_1 \sqrt{\log p} \leq \beta_R \leq c_2 < 1, (15) \text{ holds}\}$. In particular, with the calibration $\beta_R \asymp (n/p)^{1/4}$, by ignoring the logarithmic factors, our proposed estimator is rate-optimally adaptive over the collection of parameter spaces lying in the intermediate to strong SNR regime, namely, $\mathcal{C}_{Adap} = \{\mathcal{D}_R(t, \beta_R) : t^2 \gtrsim \sigma^2 \sqrt{np}\}$.

3.3. Minimax Rates and Optimality of the Range Estimator

As a direct consequence of our previous results on the extreme column estimation, the theoretical properties of the range estimator $\hat{\mathbf{R}}^*$ can be obtained in the same manner. Again we consider the normalized ℓ_2 distances $\|\hat{\mathbf{R}} - \mathbf{R}(\Theta)\|_2 / \sqrt{n}$ and denote the corresponding estimation risk as $\mathcal{R}_W(\hat{\mathbf{R}}) = \frac{1}{\sqrt{n}} \mathbb{E} \|\hat{\mathbf{R}} - \mathbf{R}(\Theta)\|_2$. Define the parameter space

$$\mathcal{D}_W(t, \beta_R, \beta_L) = \left\{ (\Theta, \pi) \in \mathcal{D} \times \mathcal{S}_p : \begin{array}{l} \text{(A) holds, } \lambda_1 \in [t/8, 8t], \sum_{i=2}^r \lambda_i \leq \sigma \sqrt{\log p}, \\ -\beta_L \leq v_{11} \leq 0 \leq v_{1p} \leq \beta_R, \end{array} \right\}, \quad (18)$$

where $t \geq 0$, $p^{-1/2} \leq \beta_R, \beta_L \leq 1$, and define the function $q'(x, y, n, p) = \sigma^2 n \log p \left[\frac{1}{x^2} \wedge \left\{ \frac{1}{\psi^2} + \frac{1}{\psi} \sqrt{\left(\frac{p}{n \log p} \right)} \right\} \right] + \left(\frac{1-x^2}{x^2} \sigma^2 \log p + \frac{y^2 \sigma^2 p}{1-y^2} \right)$. The following theorem establishes the minimax rate of convergence for estimating $\mathbf{R}(\Theta)$ and the minimax optimality and adaptivity of our proposed estimator $\hat{\mathbf{R}}^*$.

THEOREM 3 (MINIMAX RATES). *Let $\beta_W = \beta_R + \beta_L$. Suppose $t^2 \geq c_0 q'(\beta_R \wedge \beta_L, \beta_R \vee \beta_L, n, p)$, $c_1 p^{-1/2} \sqrt{\log p} \leq \{\beta_R, \beta_L\} \leq c_2$ for sufficiently large (n, p) and some constants $c_0, c_1 > 0$ and $0 < c_2 < 1$, and Z has independent sub-Gaussian entries Z_{ij} with parameter σ^2 . Then*

$$\inf_{\hat{\mathbf{R}}} \sup_{\mathcal{D}_W(t, \beta_R, \beta_L)} \mathcal{R}_W(\hat{\mathbf{R}}) \asymp \frac{\beta_W t}{\sqrt{n}} \left[\frac{\sigma \sqrt{\{(t^2 + \sigma^2 p)n\}}}{t^2} \wedge 1 \right] + \sigma \psi. \quad (19)$$

In particular, the minimax rates are simultaneously attained by the estimator $\hat{\mathbf{R}}^$.*

4. A SPECIAL CASE: PERMUTED LINEAR GROWTH MODEL

In the previous sections, theoretical results are obtained for the general approximately rank-one matrices characterized by (11) (18) as well as the conditions of Theorems 1 to 3. One advantage is the rich row-monotonicity structures contained in such parameter spaces, which adapts well to real applications such as our motivating example in microbiome studies where the noisy data sets are generated from the shotgun metagenomic sequencing (See Boulund et al. (2018); Gao & Li (2018) and Figure 1). However, in many cases such as classical theories of the bacterial growth dynamics, an important subclass of the general permuted monotone matrix model has usually been considered for its heuristic simplicity and explanatory power. We refer this sub-model as the *permuted linear growth model*, where (1) holds over the restricted set

$$\mathcal{D}_0 = \left\{ (\Theta, \pi) \in \mathcal{D} \times \mathcal{S}_p : \begin{array}{l} \theta_{ij} = a_i \eta_j + b_i, \text{ where } a_i, b_i \in \mathbb{R} \text{ for } 1 \leq i \leq n, \\ \eta_j \leq \eta_{j+1} \text{ for } 1 \leq j \leq p-1 \text{ and } \sum_{j=1}^p \eta_j = 0. \end{array} \right\}.$$

In other words, each row of Θ has a linear growth pattern with possibly different intercepts and slopes. Denote $\mathbf{a} = (a_i)_{1 \leq i \leq n}$, $\boldsymbol{\eta} = (\eta_j)_{1 \leq j \leq p}$ and $\mathbf{b} = (b_i)_{1 \leq i \leq n}$. In this case, the parameters of interest have the expressions of $\Theta_R = \mathbf{a} \boldsymbol{\eta}_p$, $\Theta_L = \mathbf{a} \boldsymbol{\eta}_1$ and $\mathbf{R} = \mathbf{a}(\boldsymbol{\eta}_p - \boldsymbol{\eta}_1)$.

In the context of bacterial growth dynamics, the above model is commonly referred as the Cooper-Helmstetter model (Cooper & Helmstetter, 1968; Bremer & Churchward, 1977) that associates the copy number of genes with their relative distances to the replication origin. Specifically, a_i is the ratio of genome replication time and doubling time for the i th sample, η_j is the distance from the replication origin for the j th contig, and b_i is related to the read counts at the replication origin and the sequencing depth. Consequently, the extreme columns $\mathbf{a} \boldsymbol{\eta}_p$ and $\mathbf{a} \boldsymbol{\eta}_1$ correspond to the true log-transformed peak and trough coverages that are used to quantify the bacterial growth dynamics across the samples (see also Section 5.2 and 5.3 for more details).

In the following, we discuss the consequences for the estimation of Θ_R under this special linear growth model, and the results for estimating Θ_L and \mathbf{R} follow similarly. By definition, the SVD (3) for $\Theta \in \mathcal{D}_0$ has a reduced form. Specifically, the row-centered matrix Θ' is exactly rank-one, where the leading right singular vector \mathbf{v}_1 has components

$$v_{1j} = \frac{\eta_j}{\|\boldsymbol{\eta}\|_2}, \quad \text{for } j = 1, \dots, p, \quad (20)$$

and the largest singular value admits the expression

$$\lambda_1 = \|\mathbf{a}\|_2 \|\boldsymbol{\eta}\|_2. \quad (21)$$

Intuitively, the set $\{v_{1j}\}_{1 \leq i < j \leq p}$ characterize the exact normalized column positions of Θ' (and Θ), while λ_1 summarizes the slope magnitude of the rows and the overall separateness of the columns. Consequently, the risk upper bound obtained in Theorem 1 has a reduced form, which

has simpler and more intuitive interpretations. Specifically, for any given $\Theta \in \mathcal{D}_R(t, \beta_R)$, we consider the following pointwise risk upper bound

$$\mathcal{R}_R(\hat{\Theta}_R^*) \lesssim \frac{v_{1p}\lambda_1(\Theta)}{\sqrt{n}} \left[\frac{\sigma\sqrt{\{(\lambda_1^2(\Theta) + \sigma^2 p)n\}}}{\lambda_1^2(\Theta)} \wedge 1 \right] + \sigma\psi, \quad (22)$$

induced by (12) of Theorem 1. With the reparametrizations (20) and (21), we can rewrite (22) as

$$\mathcal{R}_R(\hat{\Theta}_R^*) \lesssim \frac{\eta_p \|\mathbf{a}\|_2}{\sqrt{n}} \left[\frac{\sigma\sqrt{\{(\|\mathbf{a}\|_2^2 \|\boldsymbol{\eta}\|_2^2 + \sigma^2 p)n\}}}{\|\mathbf{a}\|_2^2 \|\boldsymbol{\eta}\|_2^2} \wedge 1 \right] + \sigma\psi. \quad (23)$$

Some observations about this risk upper bound are in order.

(a) Over the low SNR regime where $\|\mathbf{a}\|_2^2 \|\boldsymbol{\eta}\|_2^2 \lesssim \sigma^2 \sqrt{(np)}$, (23) becomes

$$\mathcal{R}_R(\hat{\Theta}_R^*) \lesssim \frac{\|\mathbf{a}\|_2 \eta_p}{\sqrt{n}} + \sigma\psi, \quad (24)$$

where the first term is proportional to the overall slope magnitude $\|\mathbf{a}\|_2$, but does not rely on the locations of the other columns, i.e., η_j for $1 \leq j \leq p-1$. In this case, since the signal changes across different columns are so vague, $\hat{\Theta}_R^*$ fails to implement a good estimate for the slopes \mathbf{a} and the estimation error can only decrease when the extreme column $\Theta_R = \mathbf{a}\eta_p$ itself (and its norm $\|\mathbf{a}\|_2 \eta_p$) is close to zero.

(b) Over the intermediate SNR regime where $\sigma^2 \sqrt{(np)} \lesssim \|\mathbf{a}\|_2^2 \|\boldsymbol{\eta}\|_2^2 \lesssim \sigma^2 p$, (23) becomes

$$\mathcal{R}_R(\hat{\Theta}_R^*) \lesssim \frac{\sigma \eta_p}{\|\boldsymbol{\eta}\|_2} \left(1 + \frac{\sigma^2 p}{\|\mathbf{a}\|_2^2 \|\boldsymbol{\eta}\|_2^2} \right)^{1/2} + \sigma\psi. \quad (25)$$

In this case, as the signal differences between every consecutive columns are steep enough so that the slopes \mathbf{a} can be well estimated, increasing $\|\boldsymbol{\eta}\|_2$ or $\|\mathbf{a}\|_2$ would expand the advantage and therefore leads to a better estimate.

(c) Over the strong SNR regime where $\|\mathbf{a}\|_2^2 \|\boldsymbol{\eta}\|_2^2 \gtrsim \sigma^2 p$, the upper bound (23) becomes

$$\mathcal{R}_R(\hat{\Theta}_R^*) \lesssim \frac{\sigma \eta_p}{\|\boldsymbol{\eta}\|_2} + \sigma\psi. \quad (26)$$

In the case, the advantage of large $\|\mathbf{a}\|_2$ has been exploited to extremity so that increasing $\|\mathbf{a}\|_2$ will no longer improve the performance of $\hat{\Theta}_R^*$.

Comparing the rates from (24) to (26), an interesting discrepant role played by the overall slope magnitude $\|\mathbf{a}\|_2$ can be observed. In general, the theoretical performance of $\hat{\Theta}_R^*$ is clearly driven by the global SNR $\|\mathbf{a}\|_2^2 \|\boldsymbol{\eta}\|_2^2 / \sigma^2$, which measures the magnitude of the signal changes and the degree of monotonicity relative to the noise level.

Following the same argument as the proof of Theorem 2, the minimax optimality of our proposed estimator $\hat{\Theta}_R^*$ can be also established under the permuted linear growth model. Specifically, if we define the indexed parameter space $\mathcal{D}_{0,R}(t, \beta) = \{(\Theta, \pi) \in \mathcal{D}_0 : 0 \leq \eta_p / \|\boldsymbol{\eta}\|_2 \leq \beta, \|\mathbf{a}\|_2 \|\boldsymbol{\eta}\|_2 \in [t/8, 8t]\}$, then it can be shown that for any pair (t, β_R) such that (15) holds and $p^{-1/2} \sqrt{\log p} \lesssim \beta_R \leq c < 1$,

$$\inf_{\hat{\Theta}_R} \sup_{\mathcal{D}_{0,R}(t, \beta_R)} \mathcal{R}_R(\hat{\Theta}_R) \asymp \frac{\beta_R t}{\sqrt{n}} \left[\frac{\sigma\sqrt{\{(t^2 + \sigma^2 p)n\}}}{t^2} \wedge 1 \right] + \sigma\psi,$$

where the optimal rate is simultaneously attained by the proposed estimator $\hat{\Theta}_R^*$.

5. NUMERICAL STUDIES

5.1. *Simulation with Model-Generated Data*

To demonstrate our theoretical results and to compare with alternative methods, we generate data from model (1) with various configurations of the signal matrix Θ . Specifically, the signal matrix $\Theta = (\theta_{ij}) \in \mathbb{R}^{n \times p}$ is generated under the following two regimes:

- $S_1(n, p, \alpha)$: for any $1 \leq i \leq n$, $\theta_{ij} = a_i \eta_j + b_i$ for $1 \leq j \leq p$, where $a_i \sim \text{Unif}(0, \alpha)$, $b_i \sim \text{Unif}(0, 6)$ and $(\eta_1, \dots, \eta_p) = (-1, 0, 0, \dots, 0, 1)$;
- $S_2(n, p, \alpha)$: for any $1 \leq i \leq n$, $\theta_{ij} = \log(1 + a_i j + \beta_i)$ for $1 \leq j \leq p$ where $a_i \sim \text{Unif}(0, \alpha)$ and $b_i \sim \text{Unif}(0, 6)$.

By construction, $S_1(\alpha, n, p)$ belongs to the linear growth model whereas $S_2(\alpha, n, p)$ does not. The elements of Z are drawn from i.i.d. standard normal distributions, and, without loss of generality, we set $\Pi = \mathbf{I}_p$.

For the extreme column Θ_R , we compare the empirical performance of our proposed estimator $\hat{\Theta}_R^*$ with (i) the direct sorting estimator (DS) $\check{\Theta}_R$ defined as $\check{\Theta}_R = Y_{\cdot, \hat{\pi}(p)}$, where $\hat{\pi}$ is given by (7); and (ii) the order statistic estimator (OS) $\check{\Theta}_R = (Y_{i, (p)})_{1 \leq i \leq n}$, as all the rows of Θ are monotonic increasing. For the range vector $\mathbf{R}(\Theta)$, we compare our proposed estimator $\hat{\mathbf{R}}^*$ with (i) the direct sorting estimator (DS) $\check{\mathbf{R}}_{DS} = Y_{\cdot, \hat{\pi}(p)} - Y_{\cdot, \hat{\pi}(1)}$, and (ii) the order statistic estimator (OS) $\check{\mathbf{R}}_{OS} = (Y_{i, (p)} - Y_{i, (1)})_{1 \leq i \leq n}$. We use the empirical risk, or the averaged normalized ℓ_2 distance, to compare these methods. For each setting, we evaluate the empirical performance of each method over a range of n, p and α . Each setting is repeated for 200 times.

The results are summarized as boxplots in Figure 2 and Figure 3. The empirical results agree with our theory in the following perspectives: (i) our proposed estimators $\hat{\Theta}_R^*$ and $\hat{\mathbf{R}}^*$ perform the best among all the settings; (ii) in the middle two plots of Figure 2 and 3, the risks of our proposed estimators decrease as n grows, which agrees with our theorems. In addition, in the top left panel of Figure 2 and 3 we observe that the risks of the OS estimator decrease as α increases. This is because under $S_1(\alpha, p, n)$, the parameter α characterizes the separateness of the two extreme columns from the other columns. The OS estimators would apparently favour the cases where the separation is more significant. In addition, both our proposed estimators and the DS estimators outperform the OS estimators, showing the advantage of the compound estimators.

5.2. *Simulation with Synthetic Microbiome Metagenomic Data*

We then evaluate the empirical performance of our proposed method using a synthetic metagenomic sequencing data set (Gao & Li, 2018) by generating sequencing reads based on 45 closely related bacterial genomes in 50 independent samples. Particularly, Gao & Li (2018) presented a synthetic shotgun metagenomic sequencing data set of a community of 45 phylogenetically related species from 15 genera of five different phyla with known RefSeq ID, taxonomy and replication origin (Gao et al., 2013). To generate the metagenomic reads, reference genome sequences of randomly selected three species in each genus were downloaded from NCBI. Read coverages were generated along the genome based on an exponential distribution with a specified peak-to-trough ratio and a function of accumulative distribution of read coverages along the genome was calculated. Sequencing reads were then generated using the above accumulative distribution functions and a random location for each read on the genome, until the total read number achieved a randomly assigned average coverage between 0.5 and 10 folds for the species in a sample. Sequencing errors including substitution, insertion and deletion were simulated in a

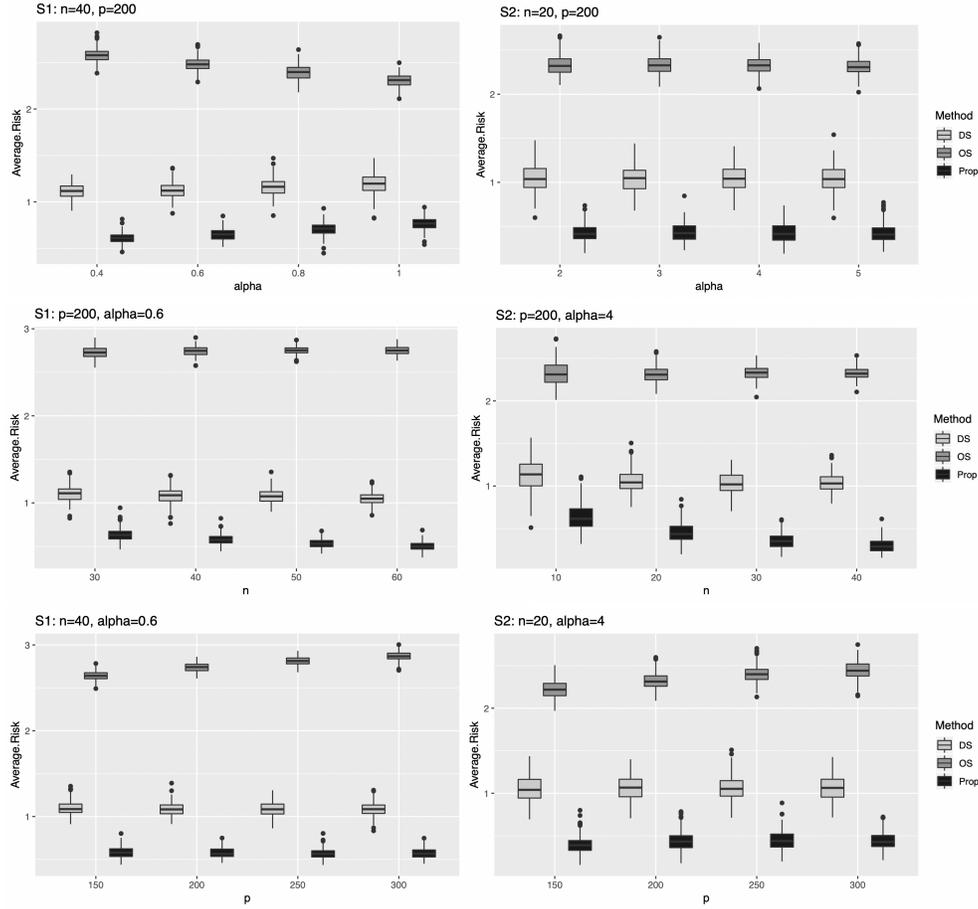


Fig. 2. Boxplots of the empirical risks for estimating Θ_R under two different models, $S_1(\alpha, p, n)$ and $S_2(\alpha, p, n)$, with DS, OS and Prop representing $\hat{\Theta}_R$, $\hat{\Theta}_R$ and $\hat{\Theta}_R^*$.

position- and nucleotide-specific pattern according to the metagenomic sequencing error profile of Illumina.

For the final data set, the average nucleotide identities (ANI) between species within each genus ranged from 66.6% to 91.2%. The probability of one species existing in each of the 50 simulated samples was set as 0.6, and a total of 1,336 average coverages and the corresponding PTRs were randomly and independently assigned. After the same processing, filtering, and CG-adjustment steps as in Gao & Li (2018), the final data set included genome assemblies of 41 species. For each species, we obtained the permuted matrix of log-contig coverage with the number of samples ranging from 29 to 46 and the number of contigs from 47 to 482.

We provide estimates of the log-PTRs of the assembled species for all the samples, or the range vector \mathbf{R} , using our previous notations. As a comparison, in addition to our proposed method $\hat{\mathbf{R}}^*$, we consider the iRep estimator proposed by Brown et al. (2016), where the contigs of a given species were ordered for each sample separately based on the observed read counts, before fitting a piece-wise linear regression function. We evaluate these methods by considering the ℓ_2 distance between the vectors of the true log-PTRs and their estimates. To generalize our evaluation to diverse metagenomic data sets, we also evaluate the effect of sample size n as well

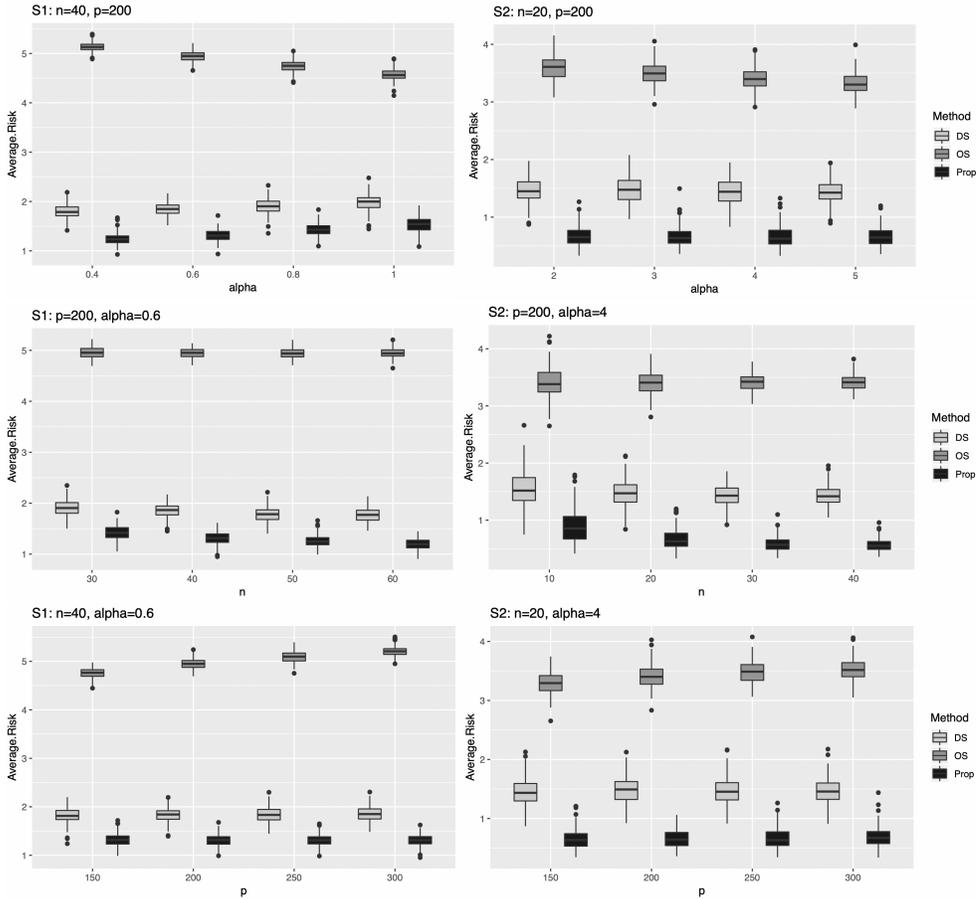


Fig. 3. Boxplots of the empirical risks for estimating \mathbf{R} under two different models, $S_1(\alpha, p, n)$ and $S_2(\alpha, p, n)$, with DS, OS and Prop representing $\tilde{\mathbf{R}}_{DS}$, $\tilde{\mathbf{R}}_{OS}$ and $\hat{\mathbf{R}}^*$.

as contig numbers p by randomly selecting subsets of samples or contigs from each data set. The selection was made with replacement.

The results are summarized in Figure 4. As n or p varies, our proposed estimator performs consistently better than iRep. Moreover, the performance of our proposed method is not sensitive to the sample size, the number of contigs from the genome assemblies or the underlying true PTRs. These results partially explain why the DEMIC algorithm has superior performance compared to the existing ones (Gao & Li, 2018).

5.3. Analysis of A Real Microbiome Metagenomic Data Set

We complete our numerical study by analyzing a real metagenomic data set from the NIH Integrative Human Microbiome Project (iHMP). As part of the iHMP, the Inflammatory Bowel Disease (IBD) Multi'omics project is a multi-institutional effort to investigate the differences in gut microbiome communities among adults and children with IBD (Lloyd-Price et al., 2019) and normal non-IBD controls. Many studies have reported strong associations between IBD, including both Crohn's disease (CD) and ulcerative colitis (UC) and gut microbiota composition. In contrast, we focus on comparing the bacterial growth rates between UC, CD and normal non-IBD individuals using the proposed methods.

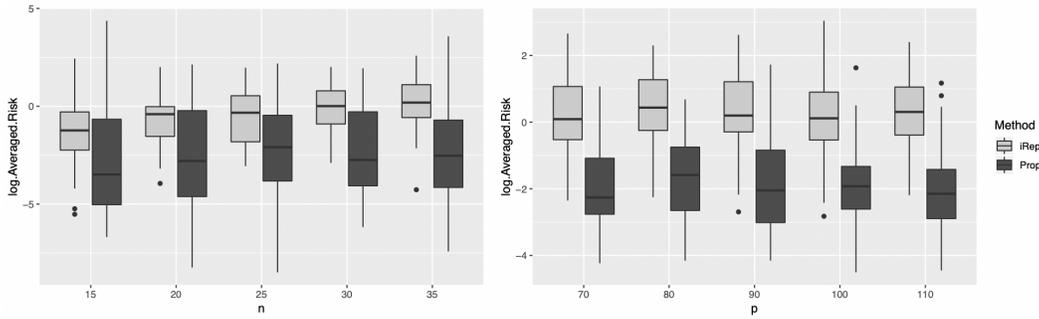


Fig. 4. Boxplots of the ℓ_2 distances between the vectors of the estimated and the true log-PTRs for different sample sizes n and different numbers of contigs p . The darker ones correspond to the proposed method and the lighter ones correspond to the iRep estimation method.

Table 1. Analysis of bacterial growth rates among CD, UC and non-IBD samples. Bins that show significantly different growth rates and their taxonomic annotations are presented. (n_1, n_2, n_3) : numbers of samples of CD, UC and non-IBD samples that carried the respective bin (assembled bacterial genome).

Bins	(n_1, n_2, n_3)	P-values	Taxonomic Annotations
bin-054	(54, 32, 54)	0.015	Roseburia (genus)
bin-090	(38, 41, 52)	0.005	Faecalibacterium (genus)
bin-091	(26, 40, 52)	0.016	Clostridiales (order)
bin-099	(30, 32, 49)	< 0.001	Subdoligranulum (genus)
bin-465	(36, 41, 33)	0.043	Dialister (genus)

The metagenomic data sets, including 300 samples of the CD, UC and non-IBD subjects, were downloaded from the IBDMDB website (<https://www.ibdmdb.org>). Specifically, we randomly select 100 samples of UC, CD and normal non-IBD samples, respectively. For each sample, the sequencing data was obtained from the stool sample using Illumina shotgun sequencing. We first apply MEGAHIT (Li et al., 2015) version 1.1.1 to perform metagenomic co-assembly. The co-assembled contigs were then clustered into metagenomic bins or genome assemblies using MaxBin (Wu et al., 2015) version 2.2.4. Finally, Bowtie 2 (Langmead & Salzberg, 2012) version 2.3.2 was used to align reads back to the assembled contigs for each of the samples, and the output alignments were then sorted by samtools (Li et al., 2009) version 0.1.19.

After these preparations, the DEMIC algorithm, incorporated with our proposed methods, was applied to obtain the estimated PTRs (ePTRs) of a given species represented by a contig cluster (bin) for each sample. As a result, ePTRs of 25 bins were obtained for subsets of the UC (n_1), CD (n_2) and non-IBD (n_3) samples with $n_1 + n_2 + n_3 \geq 100$, as some contig clusters may not be carried or abundant enough among many samples. For each bin, we compare the ePTRs among the UC, CD and non-IBD samples using an F-test. We applied the CAT/BAT algorithm (von Meijenfeldt et al., 2019) that compares the metagenomic assembled bins to a taxonomy database to obtain the taxonomic annotations of the 25 bins.

Interestingly, based on the F-test, among the 25 contig clusters, 5 of them show significant difference in ePTRs among the UC, CD and non-IBD samples (Table 1). For reasons of space,

Table 2. *p*-values from pairwise *t*-tests of differential growth rates between different groups for five genome assembly bins.

Bins	UC vs. CD	UC vs. non-IBD	CD vs. non-IBD
bin-054	0.525	0.004	0.081
bin-090	0.392	0.016	0.004
bin-091	0.012	0.054	0.335
bin-099	0.960	< 0.001	< 0.001
bin-465	0.042	0.818	0.026

Table 1 only provides the taxonomic annotation of the bins in terms of their genus – except for bin-091 which can only be determined up to orders (see our Supplementary Material (Ma et al., 2020b) for the complete annotations). We also performed pairwise comparisons using two-sample *t*-test for the 5 differential bins (Table 2). We found that the difference in the growth rates of bin-054 (Roseburia), bin-090 (Faecalibacterium) and bin-099 (Subdoligranulum) are more significant between IBD and non-IBD samples. In particular, boxplots in our Supplementary Material (Ma et al., 2020b) indicate higher growth rates of bin-054 (Roseburia) and bin-090 (Faecalibacterium), and a lower growth rates of bin-099 (Subdoligranulum) for IBD samples when compared to the non-IBD samples. Moreover, the growth rates of bin-091 (Clostridiales) is significantly higher among UC samples, whereas the growth rate of bin-465 (Dialister) is significantly higher among the CD samples, comparing to the samples of the other two categories. These results show that the gut microbiome communities in CD and UC patients or IBD and non-IBD patients differ not only in relative abundance but also in growth rates of certain bacterial species, an important insight from our data analysis.

6. DISCUSSION

The present paper focused on the permuted monotone matrix model with homoskedastic noise. If the noises are heteroskedastic, for example (i) the columns of the noise matrix are not independent, or (ii) the variances of the noise matrix entries are not identical, we argue that, as long as the marginal distributions of the noise matrix remain sub-Gaussian, the framework developed in this paper can still be applied. Specifically, in light of the recent work of Zhang et al. (2018), where heteroskedastic PCA and SVD are studied, the key analytical tools paralleling to those applied in the current work, such as concentration and perturbation inequalities associated to the heteroskedastic random matrices, can be obtained by generalizing the results of Zhang et al. (2018). Such extensions are involved and we leave them for future research.

Moreover, the current analytical framework is built upon the approximately rank-one structure observed in the real data sets from our metagenomic applications, as well as the leading singular vector property demonstrated in Section 2.1. However, there might be other interesting applications where the underlying monotone signal matrix is of general rank $r > 1$. As a result, it is also of interest to investigate whether the current method and theoretical framework can be extended to estimate the extreme values in those general rank r matrices.

SUPPLEMENTARY MATERIALS

Supplementary material available at *Biometrika* online includes the proofs of other theorems and the technical lemmas, as well as some supplementary tables and figures.

REFERENCES

- ABEL, S., ZUR WIESCH, P. A., CHANG, H.-H., DAVIS, B. M., LIPSITCH, M. & WALDOR, M. K. (2015). Sequence tag-based analysis of microbial population dynamics. *Nature Methods* **12**, 223.
- BOULUND, F., PEREIRA, M. B., JONSSON, V. & KRISTIANSSON, E. (2018). Computational and statistical considerations in the analysis of metagenomic data. In *Metagenomics*. Elsevier, pp. 81–102.
- BREMER, H. & CHURCHWARD, G. (1977). An examination of the cooper-helmstetter theory of dna replication in bacteria and its underlying assumptions. *Journal of Theoretical Biology* **69**, 645–654.
- BROWN, C. T., OLM, M. R., THOMAS, B. C. & BANFIELD, J. F. (2016). Measurement of bacterial replication rates in microbial communities. *Nature Biotechnology* **34**, 1256.
- BROWN, L. D. & GREENSHTEIN, E. (2009). Nonparametric empirical bayes and compound decision approaches to estimation of a high-dimensional vector of normal means. *The Annals of Statistics* , 1685–1704.
- CAI, T. T. & ZHANG, A. (2018). Rate-optimal perturbation bounds for singular subspaces with applications to high-dimensional statistics. *The Annals of Statistics* **46**, 60–89.
- COOPER, S. & HELMSTETTER, C. E. (1968). Chromosome replication and the division cycle of escherichia coli b/r. *Journal of Molecular Biology* **31**, 519–540.
- COPAS, J. (1969). Compound decisions and empirical bayes. *Journal of the Royal Statistical Society: Series B (Methodological)* **31**, 397–417.
- GAO, F., LUO, H. & ZHANG, C.-T. (2013). Doric 5.0: an updated database of oric regions in both bacterial and archaeal genomes. *Nucleic Acids Research* **41**, D90.
- GAO, Y. & LI, H. (2018). Quantifying and comparing bacterial growth dynamics in multiple metagenomic samples. *Nature Methods* **15**, 1041–1044.
- KOREM, T., ZEEVI, D., SUEZ, J., WEINBERGER, A., AVNIT-SAGI, T., POMPAN-LOTAN, M., MATOT, E., JONA, G., HARMELIN, A. & COHEN, N. (2015). Growth dynamics of gut microbiota in health and disease inferred from single metagenomic samples. *Science* , aac4812.
- LANGMEAD, B. & SALZBERG, S. L. (2012). Fast gapped-read alignment with bowtie 2. *Nature Methods* **9**, 357.
- LI, D., LIU, C.-M., LUO, R., SADAKANE, K. & LAM, T.-W. (2015). Megahit: an ultra-fast single-node solution for large and complex metagenomics assembly via succinct de bruijn graph. *Bioinformatics* **31**, 1674–1676.
- LI, H., HANDSAKER, B., WYSOKER, A., FENNELL, T., RUAN, J., HOMER, N., MARTH, G., ABECASIS, G. & DURBIN, R. (2009). The sequence alignment/map format and samtools. *Bioinformatics* **25**, 2078–2079.
- LLOYD-PRICE, J., ARZE, C., ANANTHAKRISHNAN, A. N., SCHIRMER, M., AVILA-PACHECO, J., POON, T. W., ANDREWS, E., AJAMI, N. J., BONHAM, K. S., BRISLAWN, C. J. et al. (2019). Multi-omics of the gut microbial ecosystem in inflammatory bowel diseases. *Nature* **569**, 655.
- MA, R., CAI, T. T. & LI, H. (2020a). Optimal permutation recovery in permuted monotone matrix model. *Journal of the American Statistical Association* , 1–15.
- MA, R., CAI, T. T. & LI, H. (2020b). Supplement to “optimal estimation of extreme values for permuted approximate rank-one monotone matrix” .
- MYHRVOLD, C., KOTULA, J. W., HICKS, W. M., CONWAY, N. J. & SILVER, P. A. (2015). A distributed cell division counter reveals growth dynamics in the gut microbiota. *Nature Communications* **6**, 10039.
- ROBBINS, H. (1951). Asymptotically subminimax solutions of compound statistical decision problems. In *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability*. The Regents of the University of California.
- ROBBINS, H. (1964). The empirical bayes approach to statistical decision problems. *The Annals of Mathematical Statistics* **35**, 1–20.
- SAMUEL, E. (1967). The compound statistical decision problem. *Sankhyā: The Indian Journal of Statistics, Series A* , 123–140.
- VON MEIJENFELDT, F. B., ARKHIPOVA, K., CAMBUY, D. D., COUTINHO, F. H. & DUTILH, B. E. (2019). Robust taxonomic classification of uncharted microbial sequences and bins with cat and bat. *bioRxiv* , 530188.
- WU, Y.-W., SIMMONS, B. A. & SINGER, S. W. (2015). Maxbin 2.0: an automated binning algorithm to recover genomes from multiple metagenomic datasets. *Bioinformatics* **32**, 605–607.
- ZHANG, A., CAI, T. T. & WU, Y. (2018). Heteroskedastic pca: Algorithm, optimality, and applications. *arXiv preprint arXiv:1810.08316* .
- ZHANG, C.-H. (2003). Compound decision theory and empirical bayes methods. *The Annals of Statistics* **31**, 379–390.