

Optimal Permutation Recovery in Permuted Monotone Matrix Model

Rong Ma T. Tony Cai and Hongzhe Li

Abstract

Motivated by recent research on quantifying bacterial growth dynamics based on genome assemblies, we consider a permuted monotone matrix model $Y = \Theta\Pi + Z$, where the rows represent different samples, the columns represent contigs in genome assemblies and the elements represent log-read counts after preprocessing steps and Guanine-Cytosine (GC) adjustment. In this model, Θ is an unknown mean matrix with monotone entries for each row, Π is a permutation matrix that permutes the columns of Θ , and Z is a noise matrix. This paper studies the problem of estimation/recovery of Π given the observed noisy matrix Y . We propose an estimator based on the best linear projection, which is shown to be minimax rate-optimal for both exact recovery, as measured by the 0-1 loss, and partial recovery, as quantified by the normalized Kendall's tau distance. Simulation studies demonstrate the superior empirical performance of the proposed estimator over alternative methods. We demonstrate the methods using a synthetic metagenomics dataset of 45 closely related bacterial species and a real metagenomic dataset to compare the bacterial growth dynamics between the responders and the non-responders of the IBD patients after 8 weeks of treatment.

KEYWORDS: Kendall's tau; Microbiome growth dynamics; Minimax lower bound; Sorting.

⁰Rong Ma is PhD Candidate, Department of Biostatistics, Epidemiology and Informatics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104 (E-mail: rongm@pennmedicine.upenn.edu). T. Tony Cai is Daniel H. Silberberg Professor of Statistics, Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104 (E-mail: tcai@wharton.upenn.edu). Hongzhe Li is Professor of Biostatistics and Statistics, Department of Biostatistics, Epidemiology and Informatics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104 (E-mail: hongzhe@upenn.edu).

1 INTRODUCTION

1.1 A Motivation Example from Microbiome Studies

The statistical problem considered in this paper is motivated by the problem of estimating the bacterial growth dynamics based on shotgun metagenomics data (Myhrvold et al. 2015; Abel et al. 2015; Korem et al. 2015; Brown et al. 2016). The growth dynamics of microbial populations reflects their physiological states and drives variation of microbial compositions, which provide important feature summary of the microbes in a given community. One way of studying such communities is through shotgun metagenomic sequencing, which involve direct DNA sequencing of all the microbiome genomes in a given microbial community. Korem et al. (2015) presented the first paper on quantifying the bacterial growth dynamics based on shotgun metagenomics data, where the uneven sequencing read coverage resulting from the bidirectional DNA replications provides information on the rates of microbial DNA replications. 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 dynamics after aligning the sequencing reads to the complete genome sequences.

However, in many applications, it is of importance to quantify the bacterial growth dynamics based on genome assemblies for the bacterial species with unknown genomes. These genome assemblies may represent new bacterial species that we have seen or sequenced before. The genome assembly of a bacterium species consists of a collection of contigs (called bin) constructed based on the overlapping of the sequencing reads (Li et al. 2015; Wu et al. 2014). Compared to the complete genome, the genome assembled bins are more fragmented and often contained errors or contaminations. The noisy read coverage data due to intraspecific variations, interspecific/intraspecific repeated sequences, limited sequencing depths and the inability of binning algorithms to correctly cluster all the contigs further complicate the estimation of growth dynamics based on read coverages of the contigs. Besides these noisy count data, one key difficulty in estimating the growth dynamic based on contig counts is that the accurate locations of the contigs on the original genome

are unknown. It is therefore not feasible to measure the microbial growth rate directly using peak-to-trough coverage ratio for the assembled genomes (Brown et al. 2016; Gao and Li 2018).

Brown et al. (2016) presented the first method (called iRep) of estimating the bacterial growth dynamics based on genome assemblies, where the contigs are ordered based on the GC-adjusted counts for each sample separately. However, due to noise in the count data, such an ordering method often leads to wrong ordering of the contigs and therefore inaccurate estimates of the growth dynamics. Gao and Li (2018) developed a computational algorithm, DEMIC, to accurately compare growth dynamics of a given assembled species existing in multiple samples by taking advantage of highly fragmented contigs assembled in typical metagenomics studies. One key step of DEMIC is to apply a principal components analysis (PCA)-based method to recover the true ordering of the contigs along the underlying unknown bacterial complete genomes. Gao and Li (2018) reported excellent empirical performance of DEMIC over existing methods. The goal of this paper is to provide a rigorous statistical framework to study the problem of optimal permutation recovery in a permuted monotone matrix model.

1.2 A Permuted Monotone Matrix Model

For a given genome assembly with p contigs, DEMIC first obtains the read coverage for each of the sliding window of size 5000 bps, denoted by X_{ijl} for the i th sample, j th contig and k th window. In order to account for the GC-content of the k th window, Gao and Li (2018) considered the following mixed-effects model,

$$\log_2 X_{ijk} = \alpha + GC_{jk}\beta + W_{ij} + e_{ijk},$$

where GC_{jk} is the centred GC count of the k th window of the j th contig, W_{ij} is the sample- and contig- specific random intercept, α is the intercept, β is the regression coefficient, and e_{ijk} is the random error. This model is fitted for each contig to obtain the best linear unbiased predictor of W_{ij} , which is used as the GC-adjusted log-read count Y_{ij} for the i th sample and j th contig. Here Y_{ij} can be regarded as average read coverage over non-overlapping windows of a contig and is approximately normally distributed.

Let Y be the GC-adjusted log-contig count matrix of n samples and p contigs of a genome assembly with Y_{ij} as its entries. Given this, we consider the following permuted monotone matrix model:

$$Y = \Theta\Pi + Z, \tag{1}$$

where $\Theta \in \mathbb{R}^{n \times p}$ is an unknown nonnegative signal matrix with nondecreasing rows, $Z \in \mathbb{R}^{n \times p}$ is a zero-mean 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 . That is, after a suitable permutation of the columns of Y , all the rows of the mean matrix are nondecreasing sequences. In microbiome applications, Θ is the matrix of true log-coverage of n samples over p contigs along the circular genome of the bacterium, which is generally hypothesized to have non-decreasing rows. Π represents a permutation due to unknown locations of the contigs relative to the replication origin. Throughout this paper, we denote the parameter space

$$(\Theta, \pi) \in \mathcal{D} = \left\{ \Theta = (\theta_{ij}) \in \mathbb{R}^{n \times p}, \pi \in \mathcal{S}_p : 0 \leq \theta_{i,j-1} \leq \theta_{i,j} < \infty \text{ for all } 1 \leq i \leq n, 2 \leq j \leq p \right\}.$$

The focus of this paper is to optimally estimate the permutation π from the noisy observation Y .

1.3 Related Problems and Other Applications

The permutation recovery problem under permuted monotone matrix model bears some similarity to other problems studied in machine learning literature, including the feature matching between two sets of observations (Collier and Dalalyan 2016) and linear regression model with permuted data, where the correspondences between the response and the predictors are unknown (Pananjady, Wainwright, and Courtade 2016; Slawski and Ben-David 2017; Pananjady, Wainwright, and Courtade 2017). More recently, Flammarion, Mao, and Rigollet (2019) considered the problem of statistical seriation, which has a close affinity to our model (1). However, the focus of Flammarion, Mao, and Rigollet (2019) is to optimally estimate the signal matrix Θ rather than the underlying permutation.

Model (1) can be thought as a natural extension of the shape constrained matrix denoising model studied in the isotonic regression literature. Specifically, under Model (1) with known $\Pi = \mathbf{I}_p$, risk bounds and the minimax rate-optimal estimator for Θ under the Frobenius norm was obtained in Chatterjee, Guntuboyina, and Sen (2015) for $n = 1$ and later in Chatterjee, Guntuboyina, and Sen (2018) for general $n > 1$. Using the idea of optimal transport, a minimax optimal estimator of the underlying signals was obtained by Rigollet and Weed (2018). However, their goal is not to recover the underlying permutation.

Besides the microbiome applications, the permuted monotone matrix model is generic and has other applications. For instance, the problem of permutation recovery is usually equivalent to statistical ranking/sorting from noisy observations, which arises commonly in finance (Currie and Pandher 2011), sport analytics (Deshpande and Jensen 2016), and recommendation systems (Rendle et al. 2009). Specifically, in the latter case, the task of tag recommendation is to provide a user with a personalized ranked list of tags for a specific item. Under the permuted monotone matrix model, we can treat the entries of Y , say Y_{ij} , as an indicator of the j th tag being related to the i th item by a given customer, and Θ as a probability matrix characterizing the customer’s tagging preferences across multiple items. As a result, recovering the underlying permutation provides a solution of a tag recommender.

1.4 Main Contributions and Organization

In this paper, we investigate the problem of permutation recovery in the permuted monotone matrix model (1), which relies on certain invariance property of the singular subspace of the monotone matrices. The properties of the proposed method in terms of both the exact and partial recovery are studied in detail. In particular, we obtained regions of the signal-to-noise ratio (defined later as Γ/σ) that are subject to exact/partial recovery (Figure 1). For both exact and partial permutation recovery, we obtained the matching minimax lower bounds and established the minimax rate-optimality of the proposed method over a wide range of parameter space (Figure 1). For partial

recovery, the proof of the lower bound relies on a version of Fano’s lemma and the sphere packing of the symmetric group equipped with the Kendall’s tau metric.

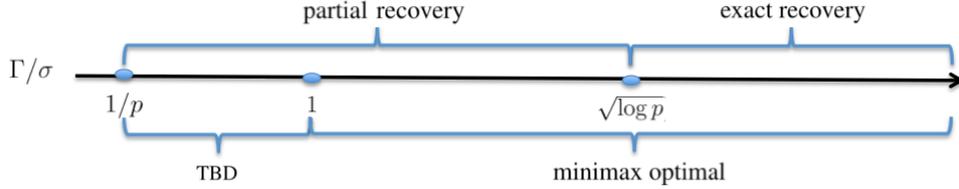


Figure 1: A graphical illustration of the main result obtained in this paper about the regions of the signal-to-noise ratio Γ/σ that correspond to exact/partial recovery, and the region with proved minimax optimality.

The rest of this paper is organized as follows. After a brief introduction of notation and definitions, we present in Section 2 the proposed permutation estimator. The theoretical properties of the proposed method are studied, first under a more illustrative linear growth model in Section 3 and then under a general growth model in Section 4. Section 5 provides results on minimax lower bounds and the optimality of the proposed estimator. We evaluate the methods using both simulated data, synthetic and real microbiome datasets and compare with other methods in Section 6. In Section 7, we discuss some implications and extensions of the methods. Finally, the proofs of our main results are given in Section 8.

1.5 Notation and Definitions

Throughout, 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 $\pi \circ \pi^{-1} = \pi^{-1} \circ \pi = id$, and denote $\text{rev}(\pi) = (\pi(p), \pi(p-1), \dots, \pi(1))$. 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 $\|\mathbf{a}\|_\infty = \max_{1 \leq i \leq n} |a_i|$. 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, $\Theta_i \in \mathbb{R}^{p_2}$ as its i -th row, and denote its (ordered) singular values as $\lambda_1(\Theta) \geq \lambda_2(\Theta) \geq \dots \geq \lambda_{\min\{p_1, p_2\}}(\Theta)$. 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 $a_n \gtrsim b_n$. For a finite set A , we denote $|A|$ as its cardinality. We use the logical symbols \wedge and \vee to represent “and” and “or,” respectively. Lastly, C, C_0, C_1, \dots are constants that may vary from place to place.

2 PERMUTATION RECOVERY VIA BEST LINEAR PROJECTION

In the following, we first make some key observations about the connection between the underlying permutation π and the column linear projections of the observed matrix Y , which motivate our construction of the proposed estimator.

2.1 Linear Projection

Given the observed noisy matrix Y , we consider the class of the linear projection statistics of the form $w^\top Y \in \mathbb{R}^p$ where $w \in \mathbb{R}^n$ and $\|w\|_2 = 1$. Intuitively, by projecting each column of Y onto the subspace generated by w , the components of $w^\top Y$ (hereafter referred as “projection scores”) would quantify the relative position of the columns of Y , so that their order statistics can be used to recover the original orders of the columns of Θ . To fix ideas, we define the following ranking operator.

Definition 1 (Ranking Operator). The ranking operator $\tau : \mathbb{R}^p \rightarrow \mathcal{S}_p$ 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.

For example, given a vector $\mathbf{x} = (2, 5, 1, 6, 2)^\top$, we have $\tau(\mathbf{x}) = (2, 4, 1, 5, 3)$. The following proposition concerning the invariance property of the column spacing of Θ is the key to our construction of the minimax optimal estimator.

Proposition 1. *Suppose $(\Theta, \pi) \in \mathcal{D}$. For any nonnegative unit vector $w \in \mathbb{R}^n$, we have*

$$\tau(w^\top \Theta \Pi) = \pi^{-1}. \quad (2)$$

Apparently, under the noiseless setting, any nonnegative unit vector would lead to the exact recovery of the underlying permutation as in this case the relative orders of the columns are exactly

coded by the relative magnitudes of the projection scores $w^\top Y = w^\top \Theta \Pi$. However, with the noisy observations, $w^\top Y = w^\top \Theta \Pi + w^\top Z$ so that the relative orders of the columns are only partially preserved by the noisy projection scores $w^\top Y$, up to some random perturbations.

Consequently, the best linear projection vector w_0 would correspond to the case where $w_0^\top \Theta \Pi$ has the most separated components such that their relative orders are most immune to the random noises. Specifically, since for any given $w \in \mathbb{R}^n$, the i -th component of $w^\top \Theta \Pi$ has the expression $w^\top \Theta \Pi e_i$ where $\{e_i\}_{i=1}^p$ is the canonical basis of the Euclidean space \mathbb{R}^p , we define

$$w_0 = \arg \max_{\substack{w \in \mathbb{R}^n \\ \|w\|_2=1}} \sum_{\substack{1 \leq i, j \leq p \\ i \neq j}} (w^\top \Theta \Pi e_i - w^\top \Theta \Pi e_j)^2 = \arg \max_{\substack{w \in \mathbb{R}^n \\ \|w\|_2=1}} \sum_{i=1}^p \left(w^\top \Theta \Pi e_i - \frac{1}{p} \sum_{j=1}^p w^\top \Theta \Pi e_j \right)^2,$$

which maximizes the pairwise distances of the components under the squared distance. Now since w_0 relies on the unknown $\Theta \Pi$ and is not computable from the data, we substitute $\Theta \Pi$ by its sample/noisy counterpart Y and define our data-driven best linear projection vector as

$$\hat{w} = \arg \max_{\substack{w \in \mathbb{R}^n \\ \|w\|_2=1}} \sum_{i=1}^p \left(w^\top Y e_i - \frac{1}{p} \sum_{i=1}^p w^\top Y e_i \right)^2, \quad (3)$$

which is actually the first eigenvector of the symmetric matrix

$$A = Y \sum_{i=1}^p \left(e_i - \frac{1}{p} \sum_{i=1}^p e_i \right) \left(e_i - \frac{1}{p} \sum_{i=1}^p e_i \right)^\top Y^\top, \quad (4)$$

and can be immediately solved by performing an eigen-decomposition on A . Once \hat{w} is obtained, we define our proposed permutation estimator as

$$\hat{\pi} = (\mathbf{r}(\hat{w}^\top Y))^{-1}. \quad (5)$$

Intuitively, the projection vector \hat{w} assigns different weights to the rows of Y so that more weight is given to the rows whose elements are better separated and therefore more informative in distinguishing the columns of Y or Θ .

2.2 Evaluation Criteria

The main focus of this paper is to investigate the theoretical properties of our proposed estimator (5) under various loss measures and parameter spaces. For any given estimator $\hat{\pi}$, we first consider

the 0-1 loss

$$\ell(\tilde{\pi}, \pi) = 1\{\tilde{\pi} \neq \pi\},$$

with the corresponding risk $\mathbb{E}\ell(\tilde{\pi}, \pi) = P(\tilde{\pi} \neq \pi)$. The 0-1 loss is used to evaluate the exact recovery, which can be a strong requirement for practical applications. As an alternative, we also consider the more flexible partial recovery, where the loss function is given by the normalized Kendall's tau distance (Kendall 1938) defined as

$$\tau_K(\pi_1, \pi_2) = \frac{\{\#\text{ of discordant pairs between } \pi_1 \text{ and } \pi_2\}}{\binom{n}{2}}. \quad (6)$$

Technically, for two permutations π_1 and π_2 , the set of discordant pairs is defined as

$$\mathcal{G}(\pi_1, \pi_2) = \{(i, j) : i < j, [\pi_1(i) < \pi_1(j) \wedge \pi_2(i) > \pi_2(j)] \vee [\pi_1(i) > \pi_1(j) \wedge \pi_2(i) < \pi_2(j)]\}$$

so that the numerator in (6) is equal to the cardinality $|\mathcal{G}(\pi_1, \pi_2)|$, which, in fact, is also the minimum number of pairwise adjacent transpositions converting π_1^{-1} into π_2^{-1} (Diaconis 1988). The denominator $\binom{n}{2}$ ensures that $\tau_K(\pi_1, \pi_2) \in [0, 1]$ where $\tau_K(\pi_1, \pi_2) = 0$ corresponds to $\pi_1 = \pi_2$.

3 A LINEAR GROWTH MODEL

We start with a simpler case where the pair (Θ, π) is from the subspace

$$\mathcal{D}_L = \left\{ (\Theta, \pi) \in \mathcal{D} : \begin{array}{l} \theta_{ij} = a_i \eta_j + b_i, \text{ where } a_i, b_i \geq 0 \text{ for } 1 \leq i \leq n, \\ 0 \leq \eta_j \leq \eta_{j+1} \text{ for } 1 \leq j \leq p - 1 \end{array} \right\}. \quad (7)$$

In other words, each row of Θ has a linear growth pattern with possibly different intercepts and different slopes. In the context of bacterial growth dynamics, this model is sometimes referred as the Cooper-Helmstetter model (Cooper and Helmstetter 1968; Bremer and 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, which can be used to quantify the bacterial growth dynamics for the i th sample, η_j is related to 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. If the bacterium is non-dividing in sample i , a_i is zero.

For the linear growth model (7), there are two key quantities that are relevant to permutation recovery.

Definition 2. For any $\Theta \in \mathcal{D}_L$, we define

$$\Gamma = \left(\sum_{i=1}^n a_i^2 \right)^{1/2} \cdot \min_{1 \leq i < j \leq p} |\eta_i - \eta_j| \quad (8)$$

as the *local minimal signal gap* of Θ , and define

$$\Lambda = \left(\sum_{i=1}^n a_i^2 \right) \cdot \frac{1}{p} \sum_{1 \leq i < j \leq p} (\eta_i - \eta_j)^2 = \left(\sum_{i=1}^n a_i^2 \right) \cdot \sum_{j=1}^p (\eta_j - \bar{\eta})^2 \quad (9)$$

as the *global signal strength* of Θ , where $\bar{\eta} = \sum_{j=1}^p \eta_j / p$.

Intuitively, both quantities involve the set $\{|\eta_j - \eta_i|\}_{1 \leq i < j \leq p}$ and the ℓ_2 norm of the vector $a = (a_1, \dots, a_n)^\top$, which characterize the column spacings and the growth rates (slopes) of Θ , respectively. Throughout this paper, we assume

(A1) the additive noise matrix $Z \in \mathbb{R}^{n \times p}$ has i.i.d. entries $z_{ij} \sim N(0, \sigma^2)$.

The Gaussian assumption simplifies our theoretical analysis. But this is not essential because all the theoretical results remain true if Z has independent sub-Gaussian entries with parameters bounded by σ^2 . The following theorem provides conditions on Γ and Λ such that exact recovery of π can be obtained by $\hat{\pi}$ in (5).

Theorem 1 (Exact Recovery, Linear). *Suppose (A1) hold, $(\Theta, \pi) \in \mathcal{D}_L$ and Θ satisfies*

$$\Gamma > C_0 \sigma \sqrt{\log p}, \quad \Lambda > C_1 \sigma^2 (n \max\{\sigma^2 n / \Gamma^2, 1\} + \sqrt{np \max\{\sigma^2 n / \Gamma^2, 1\}}) \quad (10)$$

for some $C_0, C_1 > 0$. Then with probability at least $1 - O(p^{-c})$ for some constant $c > 0$, up to a permutation reversion, we have $\hat{\pi} = \pi$.

Remark 1. Due to non-identifiability between \hat{w} and $-\hat{w}$ defined in (3), in Theorem 1, as well as all the other theoretical results concerning $\hat{\pi}$, the statement is up to a possible reversion of $\hat{\pi}$. For example, for permutation $\pi = (2, 4, 1, 5, 3)$, its reversion would be $\text{rev}(\pi) = (4, 2, 5, 1, 3)$. In fact, such indeterminacy can be avoided by noting that $a_i \geq 0$ for all i 's, but we will not pursue

such a direction in this study as the practical interest only concerns relative orders of the permuted elements.

Since Γ depends on a only through its ℓ_2 norm $\|a\|_2$, the local minimal signal gap (MSG) condition $\Gamma \geq C\sigma\sqrt{\log p}$ allows for the presence of non-informative signals in the sense that some components of a can be 0. In contrast, the condition on Λ (GSS) depends on a trade-off between Γ and $\sigma\sqrt{n}$. On the one hand, when $\Gamma > \sigma\sqrt{n}$, the condition on Λ becomes $\Lambda \geq \sigma^2(C_0n + C_1\sqrt{np})$, which is independent of Γ , and is minimax optimal for left singular subspace estimation (Cai and Zhang 2018). On the other hand, when $\Gamma < \sigma\sqrt{n}$, a stronger condition on Λ is posed, as a compensation for small Γ .

In some cases, the GSS condition in (10) can be implied by the MSG condition. We summarize our results in the following proposition.

Proposition 2. *Suppose $\Gamma/\sigma > 1/p$ and the MSG condition hold. Then the GSS condition can be implied by either one of the following conditions*

(i) $\Gamma \gtrsim \sigma\sqrt{n}$;

(ii) $\Gamma \lesssim \sigma\sqrt{n}$, and either $(\sigma^4n^2/\Gamma^4)^{1/3} \lesssim p \lesssim \sigma^2n^2/\Gamma^2$ or $p \gtrsim \sigma^2n^2/\Gamma^2 + (\sigma^3n/\Gamma^3)^{2/5}$.

We next turn to the partial recovery and study the rate of convergence of $\hat{\pi}$ measured by the normalized Kendall's tau distance under the linear growth model. In particular, we will assume an approximate uniform assignment of $\{\eta_j\}_{j=1}^p$ over some subinterval of $[0, \infty)$. In other words, the minimal element and maximal element of the set $\{|\eta_j - \eta_{j+1}|\}_{j=1}^{p-1}$ should have roughly the same magnitude, so that $\Gamma = \|a\|_2 \cdot \min_{1 \leq j \leq p-1} |\eta_j - \eta_{j+1}| \asymp \|a\|_2 \cdot \max_{1 \leq j \leq p-1} |\eta_j - \eta_{j+1}|$. This is equivalent to assuming that the contigs in genome assemblies are approximately uniformly spaced along the circular genome.

Theorem 2 (Partial Recovery, Linear). *Suppose (A1) hold, $(\Theta, \pi) \in \mathcal{D}_L$, and Θ satisfies*

(i) *there exist some $C_0 > 0$ such that $\max_{1 \leq j \leq p-1} |\eta_j - \eta_{j+1}| < C_0 \min_{1 \leq j \leq p-1} |\eta_j - \eta_{j+1}|$ for all $p > 0$, and*

(ii) $\Lambda > C_1 \sigma^2 \left(\max \left\{ \frac{\sigma^2(n+\log p)^2}{\Gamma^2}, n \right\} + \sqrt{p} \max \left\{ \frac{\sigma(n+\log p)}{\Gamma}, \sqrt{n} \right\} \right)$ for some $C_1 > 0$.

Then, up to a permutation reversion,

$$\mathbb{E}[\tau_K(\hat{\pi}, \pi)] \leq 1 \wedge \left(\frac{c_0 \sigma}{p\Gamma} \min \left\{ 1, e^{-\Gamma^2/2\sigma^2} \log \left(1 + \frac{2\sigma^2}{\Gamma^2} \right) \right\} + \frac{c_1 e^{-\Gamma^2/2\sigma^2}}{p(\Gamma/\sigma + \sqrt{8/\pi})} + \frac{c_2}{p^{c+2}} \right)$$

for some $c, c_0, c_1, c_2 > 0$.

Remark 2. The risk upper bound derived in the above theorem can be simplified as

$$\mathbb{E}[\tau_K(\hat{\pi}, \pi)] \lesssim \begin{cases} \frac{\sigma}{p\Gamma} \wedge 1 & \text{if } \Gamma/\sigma \rightarrow 0 \\ \frac{\sigma}{p\Gamma} e^{-\Gamma^2/2\sigma^2} + 1/p^{c+2} & \text{otherwise} \end{cases}$$

for some $c > 0$. In the case of $\Gamma/\sigma \rightarrow \infty$, simple calculation yields $e^{-\Gamma^2/2\sigma^2} \sigma/(p\Gamma) + 1/p^{c+2} \asymp e^{-\Gamma^2/2\sigma^2} \sigma/\Gamma$ when $\Gamma < \sigma\sqrt{2(c+1)\log p}$, whereas $e^{-\Gamma^2/2\sigma^2} \sigma/(p\Gamma) + 1/p^{c+2} \asymp 1/p^{c+2}$ when $\Gamma \geq \sigma\sqrt{2(c+1)\log p}$. As a result, we also have

$$\mathbb{E}[\tau_K(\hat{\pi}, \pi)] \lesssim \begin{cases} 1/p^{c+2} & \text{if } \Gamma/\sigma \geq \sqrt{2(c+1)\log p} \\ \frac{\sigma}{p\Gamma} e^{-\Gamma^2/2\sigma^2} & \text{if } 1 \lesssim \Gamma/\sigma < \sqrt{2(c+1)\log p} \\ \frac{\sigma}{p\Gamma} \wedge 1 & \text{if } \Gamma/\sigma \lesssim 1 \end{cases} . \quad (11)$$

See Figure 2 for an illustration.

In general, Theorem 2 shows that, even with a weaker condition on Γ that is below the requirement for the exact recovery, our proposed estimator $\hat{\pi}$ is still able to obtain a partial recovery of π with an exponential rate of convergence if $\Gamma/\sigma \gtrsim 1$ and a polynomial rate of convergence if $1/p < \Gamma/\sigma \lesssim 1$. As for Λ , the requirement is essentially the same as the exact recovery, except for an additional $\log p$ term, which is negligible in the exact recovery scenario.

Some implications about the practically preferable settings of n and p should be clarified. Firstly, although Theorem 1 implies that the difficulty for exact recovery increases as p grows (see also Table 1 from our simulations), our theory suggests a wide range of feasible choices for p . For example, if the underlying signals θ_{ij} and the noise level σ^2 are of constant order, then we have $\Gamma \asymp \sqrt{n}$ and $\Lambda \asymp np^3$, so the conditions of Theorem 1 imply that the exact recovery can be guaranteed as long as $\log p \lesssim n$. In other words, p is allowed to grow exponentially with n ,

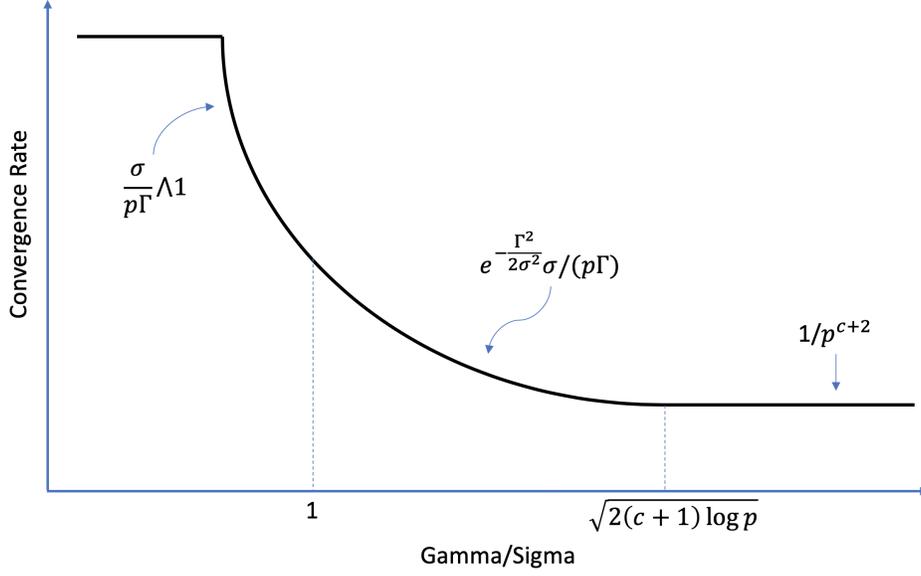


Figure 2: A graphical illustration of the risk upper bound for $\mathbb{E}[\tau_K(\hat{\pi}, \pi)]$, as a function of signal-to-noise ratio Γ/σ .

which is in line with the modern high-dimensional setting. Secondly, our Theorem 2 implies that, even if some conditions (such as MSG) for the exact recovery are not satisfied, one can still hope to partially recover the underlying permutation. In accordance to our theoretical result (11), our numerical results (Figure 4) show that, for the partial recovery, increasing p indeed reduces the overall risk of the proposed estimator. Finally, as to the sample size n , we argue that, without assuming additional structural assumptions such as row-sparsity, it is very unlikely that including more samples will result in a worse estimate (see Table 1 and Figure 4 for numerical evidences).

4 A GENERAL GROWTH MODEL

In this section we study the permutation recovery over the general parameter space \mathcal{D} where the growth pattern is not necessarily linear and therefore is more realistic inasmuch as the noisy nature of the shotgun metagenomic datasets (Boulund et al. 2018; Gao and Li 2018). The analysis relies on a deeper understanding of the relationship between the row-monotonic matrices and its leading singular vectors.

Specifically, for any $\Theta \in \mathcal{D}$, we define the row-centered matrix

$$\Theta' = \Theta(I - p^{-1}ee^\top) \in \mathbb{R}^{n \times p} \quad (12)$$

whose singular value decomposition (SVD) is given by $\Theta' = \sum_{i=1}^r \lambda_i(\Theta') u'_i v'^\top_i$, with $r \leq \min\{n, p\}$.

The following proposition is essential to our analysis of the general growth model.

Proposition 3. *Let Θ' be defined as above, then its first right singular vector v'_1 is a monotone vector, i.e., either $v'_{11} \leq v'_{12} \leq \dots \leq v'_{1p}$ or $v'_{11} \geq v'_{12} \geq \dots \geq v'_{1p}$.*

Together with Proposition 1, the above proposition justifies our construction of the permutation estimator $\hat{\pi}$ using a PCA based approach. To overcome the identifiability issue, we further assume $\lambda_1(\Theta')$ has multiplicity one. We first introduce the several quantities that play the key roles in permutation recovery over \mathcal{D} .

Definition 3. For any $\Theta \in \mathcal{D}$ and the corresponding Θ' defined as above, we define

$$\Gamma = \min_{1 \leq i < j \leq p} |u'^\top_1(\Theta'_{.i} - \Theta'_{.j})| = \lambda_1(\Theta') \min_{1 \leq i < j \leq p} |v'_{1i} - v'_{1j}|,$$

as the *local minimal signal gap*, define

$$\Xi = \max_{1 \leq i \leq p-1} \|\Theta'_{.i} - \Theta'_{.i+1}\|_2 = \max_{1 \leq i \leq p-1} \left(\sum_{j=1}^r \lambda_j^2(\Theta') |v'_{ji} - v'_{j,i+1}|^2 \right)^{1/2},$$

as the *local maximal signal gap*, and define

$$\Lambda = \lambda_1^2(\Theta') - \lambda_2^2(\Theta')$$

as the *global signal strength* of Θ .

In particular, the above definitions of Γ and Λ generalize the ones given earlier in the linear growth model as these quantities coincide for $\Theta \in \mathcal{D}_L$. The following theorem concerns the exact permutation recovery with $\hat{\pi}$ over \mathcal{D} .

Theorem 3 (Exact Recovery, General). *Suppose (A1) hold, $n \lesssim p$, $(\Theta, \pi) \in \mathcal{D}$, and Θ satisfies $\Gamma > C_0 \sigma \sqrt{\log p}$ and*

$$\Lambda > C_1 \sigma^2 \left[\left(n + \frac{\Xi^2}{\sigma^2} \right) \max \left\{ \frac{(n + \log p) \sigma^2}{\Gamma^2}, 1 \right\} + \sqrt{p} \left(\sqrt{n} + \frac{\Xi}{\sigma} \right) \max \left\{ \frac{\sigma \sqrt{n + \log p}}{\Gamma}, 1 \right\} \right]$$

for some $C_0, C_1 > 0$. Then with probability at least $1 - O(p^{-c})$ for some constant $c > 0$, up to a permutation reversion, we have $\hat{\pi} = \pi$.

As in the case of linear growth model (Theorem 1), in Theorem 3, to guarantee exact recovery, we need the MSG condition $\Gamma > C_0\sigma\sqrt{\log p}$. Unlike the linear growth model, here Γ only implicitly depends on the elements of Θ through its spectral quantities, which makes its interpretation less clear. To address this issue, we make the following observation that links the minimal singular vector gap $\min_{1 \leq i < j \leq p} |v'_{1i} - v'_{1j}|$ in the definition of Γ to the elements of Θ .

Proposition 4. *Let Θ' in (12) be such that there exists a $\delta > 0$ being the lower bound of the normalized minimum gap between any two entries in the same row, i.e.*

$$\min_{1 \leq k \leq n} \frac{|\theta'_{k,i} - \theta'_{k,j}|}{\|\Theta'_k\|_2} \geq \delta \quad \text{for some } i \neq j.$$

Then the first singular vector $v'_1 \in \mathbb{R}^p$ of Θ' satisfies $|v'_{1i} - v'_{1j}| \geq \delta$.

Consequently, the implicit requirement that $\min_{1 \leq i < j \leq p} |v'_{1i} - v'_{1j}|$ is large can be guaranteed when the normalized minimum distance $\min_{1 \leq i < j \leq p} \min_{1 \leq k \leq n} |\theta'_{k,i} - \theta'_{k,j}| / \|\Theta'_k\|_2$ is large. Our next theorem concerns the partial recovery over the general parameter space \mathcal{D} .

Theorem 4 (Partial Recovery, General). *Suppose (A1) hold, $n \lesssim p$, $(\Theta, \pi) \in \mathcal{D}$, and Θ satisfies*

(i) *there exists some $C_0 > 0$ such that $\max_{1 \leq j \leq p-1} |v'_{1j} - v'_{1,j+1}| < C_0 \min_{1 \leq j \leq p-1} |v'_{1j} - v'_{1,j+1}|$ for all $p > 0$, and*

(ii) *$\Lambda > C_1\sigma^2 \left[\max \left\{ \frac{(n+\log p)^2\sigma^2}{\Gamma^2}, n + \frac{\Xi^2}{\sigma^2} \right\} + \sqrt{p} \max \left\{ \frac{\sigma(n+\log p)}{\Gamma}, \sqrt{n} + \frac{\Xi}{\sigma} \right\} \right]$ for some $C_1 > 0$.*

Then, up to a permutation reversion,

$$\mathbb{E}[\tau_K(\hat{\pi}, \pi)] \leq 1 \wedge \left(\frac{c_0\sigma}{p\Gamma} \min \left\{ 1, e^{-\Gamma^2/2\sigma^2} \log \left(1 + \frac{2\sigma^2}{\Gamma^2} \right) \right\} + \frac{c_1 e^{-\Gamma^2/2\sigma^2}}{p(\Gamma/\sigma + \sqrt{8/\pi})} + \frac{c_2}{p^{c+2}} \right)$$

for some $c, c_0, c_1, c_2 > 0$.

Condition (i) of Theorem 4 parallels the one given in Theorem 2. It essentially requires an even distancing of the elements (the projected columns of Θ) whose ordering is to be tracked by

$\hat{\pi}$. In contrast, in both Theorem 3 and 4, the conditions on Λ are slightly more complicated than those in Theorem 1 and 2, as it further depends on the relative magnitude between Ξ/σ and \sqrt{n} . In particular, if $\Xi/\sigma \lesssim \sqrt{n}$, the conditions reduce to the ones required in the linear growth models. Interestingly, the risk upper bound obtained in Theorem 4 remains the same as in the linear growth model, which only depends on p and the signal-to-noise ratio Γ/σ .

5 MINIMAX LOWER BOUNDS AND OPTIMALITY

In this section, we establish the minimax lower bounds for both exact and partial recovery considered in previous sections, in relation to different levels of the signal-to-noise ratio Γ/σ . In the following theorem, we show the MSG condition for exact recovery is asymptotically sharp.

Theorem 5. *Suppose (A1) hold. Let $\mathcal{D}_1 = \mathcal{D}_L \cap \{(\Theta, \pi) : \Gamma \leq \frac{\sigma}{4}\sqrt{\log p}\}$ and $\mathcal{D}'_1 = \mathcal{D} \cap \{(\Theta, \pi) : \Gamma \leq \frac{\sigma}{4}\sqrt{\log p}\}$. Then for any $p \geq 10$, we have*

$$\inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}'_1} P(\hat{\pi} \neq \pi) \geq \inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}_1} P(\hat{\pi} \neq \pi) \geq 0.3,$$

where the infimum is over all the permutation estimators $\hat{\pi}$.

This theorem along with Theorem 1 and Theorem 3 indicates that our proposed estimator is minimax rate-optimal over \mathcal{D}_L and \mathcal{D} in terms of the MSG condition on Γ . In light of Proposition 2, in some situations the MSG condition can be both necessary and sufficient for the exact recovery, which includes practically important cases such as $n \asymp p$, $n < \log p$, etc. Using the information-theoretic language, we have therefore obtained both the achievability result, i.e., the existence of an algorithm or estimator that exactly recovers signal with high probability, and the converse result, namely, an upper bound on the probability of exact recovery that applies to any estimators (Cullina and Kiyavash 2016). See Figure 3 for an illustration.

Our next theorem establishes a minimax lower bound for the expected rate of convergence for the partial recovery.

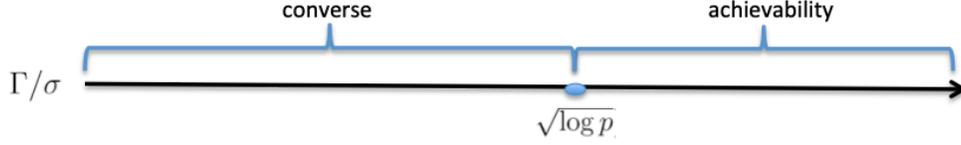


Figure 3: A graphical illustration of the achievability/converse result for exact recovery.

Theorem 6. *Suppose (A1) hold, $\mathcal{D}_2(t) = \mathcal{D}_L \cap \{(\Theta, \pi) : ct \leq \Gamma \leq Ct\}$, $\mathcal{D}'_2(t) = \mathcal{D} \cap \{(\Theta, \pi) : ct \leq \Gamma \leq Ct\}$ for some $C, c > 0$, and $t/\sigma \geq 2$. Then there exist constants $C_1, C_2 > 0$ such that*

$$\inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}'_2(t)} \mathbb{E}[\tau_K(\hat{\pi}, \pi)] \geq \inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}_2(t)} \mathbb{E}[\tau_K(\hat{\pi}, \pi)] \geq \frac{C_1 \sigma}{pt} e^{-t^2/2\sigma^2} + \frac{C_2}{p^2}.$$

Comparing the above minimax lower bound to the risk upper bounds obtained in Theorem 2 and 4, we conclude that our proposed estimator $\hat{\pi}$ is minimax rate-optimal in terms of the partial recovery for both the linear growth model and the general growth model over the range whenever Γ/σ does not diminish (Figure 1). In particular, in Theorem 5 and 6, since the minimax lower bounds only concern the worst-case scenarios, the same lower bounds should hold for any parameter spaces whenever the same worst cases are included. Similarly, the assumption (A1) does not pose a restriction to the general applicability of such results.

6 NUMERICAL STUDIES

6.1 Simulation with Model-Generated Data

To demonstrate our theoretical results and compare with alternative methods, we generate data from model (1) with various configurations of the signal matrix Θ . We compare the empirical performance of our proposed estimator $\hat{\pi}$ with the following alternatives:

- π_{mean} : Order the columns of Y by the magnitude of its column means;
- π_{max} : Order the columns of Y by the magnitude of its column maximums.

We use both the 0-1 loss and the normalized Kendall's tau distance in comparing these methods. Due to the identifiability issue, the performance of each estimator is evaluated up to a complete reversion of the permutation. For example, we use $\min\{\tau_K(\hat{\pi}, \pi), \tau_K(\hat{\pi}, \text{rev}(\pi))\}$ as the empirical

Kendall’s tau distance. By symmetry, we set the underlying permutation $\pi = id$. The signal matrix $\Theta = (\theta_{ij}) \in \mathbb{R}^{n \times p}$ is generated under the following four regimes:

- (i) $S_1(\alpha, n, p)$: For any $1 \leq j \leq p$, $\theta_{ij} = \log(1 + j\alpha_i + \beta_i)$ where $\alpha_i \sim \text{Unif}(\alpha/2, \alpha)$ for $1 \leq i \leq n/2$, $\alpha_i \sim \text{Unif}(0, 0.01)$ for $n/2 < i \leq n$, and $\beta_i \sim \text{Unif}(1, 3)$ for all $1 \leq i \leq n$;
- (ii) $S_2(\alpha, n, p)$: For any $1 \leq j \leq p$, $\theta_{ij} = j\alpha_i + \beta_i$ where $\alpha_i \sim \text{Unif}(\alpha/2, \alpha)$ for $1 \leq i \leq n/2$, $\alpha_i \sim \text{Unif}(0, \alpha/10)$ for $n/2 < i \leq n$, and $\beta_i \sim \text{Unif}(1, 3)$ for all $1 \leq i \leq n$;
- (iii) $S_3(\alpha, n, p)$: For any $1 \leq j \leq p$, $\theta_{ij} = \log(1 + j\alpha_i + \beta_i)$ where $\alpha_i \sim \text{Unif}(\alpha/2, \alpha)$ for $1 \leq i \leq 3$, $\alpha_i \sim \text{Unif}(0, 0.01)$ for $4 < i \leq n$, and $\beta_i \sim \text{Unif}(1, 3)$ for all $1 \leq i \leq n$;
- (iv) $S_4(\alpha, n, p)$: For any $1 \leq j \leq p$, $\theta_{ij} = j\alpha_i + \beta_i$ where $\alpha_i \sim \text{Unif}(\alpha/2, \alpha)$ for $1 \leq i \leq 3$, $\alpha_i \sim \text{Unif}(0, \alpha/10)$ for $4 < i \leq n$, and $\beta_i \sim \text{Unif}(1, 3)$ for all $1 \leq i \leq n$.

Specifically, under each regime, the sample-specific “growth rate” parameter α_i is randomly and uniformly generated either from the interval $[\alpha/2, \alpha]$ or an interval with much smaller values, namely, $[0, \alpha/10]$ in S_2 and S_4 and $[0, 0.01]$ in S_1 and S_3 . By construction, the four regimes consist of the nonlinear growth model where the signals spread out over many samples (S_1) or concentrate at a few rows (S_3) and the linear growth model where the signals spread out over many samples (S_2) or concentrate at a few rows (S_4). In particular, in accordance to our theory, for the supposedly “non-informative” samples, we allow the corresponding growth rates to be small but non-zero, which shows the flexibility of our proposed method. The entries of Z are drawn from i.i.d. centred normal distributions whose variance σ^2 will be given explicitly. In each setting, we evaluate the empirical performance of each method over a range of n , p or α . Each setting is repeated for 200 times.

For the exact recovery, in Table 1, we reported the empirical risks of the estimators under the 0-1 loss for various regimes and parameter combinations. The noise level σ^2 is chosen for each regime to better illustrate the differences in the empirical risks among the estimators. From our simulation results, in consistent to our theory, our proposed estimator has the smallest empirical risk over all the settings, and the estimation risk decreases as we increase α , n or decrease p .

Table 1: The empirical risks of the estimators under the 0-1 loss based on 200 simulations for various combinations of the parameters (p, n, α) . $\hat{\pi}$: proposed method; π_{mean} : mean-based method; π_{max} : max-based method.

$p = 75$	$S_1(\sigma^2 = 0.025)$	$S_2(\sigma^2 = 0.1)$		$S_3(\sigma^2 = 0.0075)$		$S_4(\sigma^2 = 0.025)$		
$n = 40$	$\alpha = 0.1$	0.2	0.1	0.2	0.1	0.2	0.1	0.2
$\hat{\pi}$	0.775	0.575	0.415	0.000	0.025	0.020	0.025	0.000
π_{mean}	0.925	0.815	0.955	0.015	0.155	0.135	0.880	0.005
π_{max}	1.000	1.000	1.000	0.995	0.995	0.970	0.840	0.430
$n = 40$	$S_1(\sigma^2 = 0.025)$	$S_2(\sigma^2 = 0.1)$		$S_3(\sigma^2 = 0.0075)$		$S_4(\sigma^2 = 0.025)$		
$\alpha = 0.1$	$p = 60$	90	60	90	60	90	60	90
$\hat{\pi}$	0.410	0.930	0.340	0.470	0.010	0.115	0.000	0.010
π_{mean}	0.720	0.985	0.910	0.980	0.070	0.245	0.775	0.900
π_{max}	1.000	1.000	1.000	1.000	0.975	1.000	0.815	0.875
$p = 75$	$S_1(\sigma^2 = 0.025)$	$S_2(\sigma^2 = 0.1)$		$S_3(\sigma^2 = 0.0075)$		$S_4(\sigma^2 = 0.025)$		
$\alpha = 0.1$	$n = 40$	60	40	60	40	60	40	60
$\hat{\pi}$	0.765	0.440	0.475	0.095	0.050	0.020	0.010	0.005
π_{mean}	0.920	0.645	0.940	0.700	0.175	0.045	0.900	0.905
π_{max}	1.000	1.000	1.000	1.000	0.995	0.995	0.855	0.820

For partial recovery, in Figure 4, we show boxplots of the empirical normalized Kendall’s tau between each estimator and the true permutation π . Again, our proposed method outperforms the alternatives in all the cases. As expected from our theory, under all the four regimes, increasing p while keeping other parameters fixed results to smaller estimation risk. As for the dependence on n , under S_1 and S_2 , increasing n leads to smaller risk as it is equivalent to increasing Γ , whereas under S_3 and S_4 , the risk roughly remains the same across different n ’s as in these case Γ doesn’t change much.

To offer more intuitive interpretation of why $\hat{\pi}$ performs better than the alternative methods, we assessed the weight vectors \hat{w} of our proposed estimator $\hat{\pi}$ under each regime after 200 rounds of simulations (Figure 3 in Supplemented Material). In comparison, the weight vector for π_{mean} is simply $(1/\sqrt{n}, \dots, 1/\sqrt{n})$, which assigns equal weight to all the samples. On the other hand, since π_{max} cannot be written in the form of $(\mathbf{r}(w^\top Y))^{-1}$ for some weight vector w and therefore does not belong to the class of linear projection estimators, we reported instead the pseudo-weight vector $\tilde{w} \in \mathbb{R}^n$ where the i -th component is the proportion that the i -th sample is used among the p coordinates. In general, we found that $\tilde{w} \in \mathbb{R}^n$ assigns larger weights to only a few samples

among those with higher signal strength, and the weight vector for π_{mean} fails to distinguish the informative samples from the non-informative ones. In contrast, the weight vectors \hat{w} for our proposed estimator $\hat{\pi}$ would automatically adapt to the varying signal strengths across the samples and assign larger weights to the samples with more significant signal changes. This also explains the interesting phenomenon in Figure 4 that, under the regime S_1 and S_2 , $\hat{\pi}$ and π_{mean} perform better than π_{max} , whereas under S_3 and S_4 , $\hat{\pi}$ and π_{max} perform better. In summary, methods that are able to detect and assign larger weight to the more informative samples would perform better than methods that are not. Observably, $\hat{\pi}$ combines the advantages of π_{mean} and π_{max} in that it finds the best weights (projection scores) in a data-driven manner.

6.2 Evaluation Using Synthetic Metagenomic Data

We evaluate the empirical performance of our proposed method using a synthetic metagenomic sequencing dataset used in Gao and Li (2018) by generating sequencing reads based on 45 bacterial genomes. Instead of estimating the PTRs, which was the focus of Gao and Li (2018), our goal is to recover the unknown relative orders of the contigs assembled in typical metagenomics studies. In addition to assisting the estimation of PTRs, such ordering of the contigs could be of independent interest for other applications, including genome assemblies based on shotgun metagenomics data.

Gao and Li (2018) presented a synthetic shotgun metagenomic sequencing dataset of a community of 45 phylogenetically related species from 15 genera of five different phyla with known RefSeq ID, taxonomy and replication origin (Gao, Luo, and Zhang 2013) (see Figure 2 in our Supplementary Material). To generate metagenomics 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 next generated using the above accumulative distribution function and a random location of 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 position- and nucleotide-specific pattern according to a recent study on metagenomic sequencing error profiles of Illumina.

For the final dataset, 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 and filtering steps and CG-adjustment step as in Gao and Li (2018), the final dataset included genome assemblies of 41 species. For each species, we obtained the permuted matrix of log-contig counts, with the number of samples ranging from 29 to 46, and the number of contigs ranging from 47 to 482.

Our proposed method ($\hat{\pi}$) was used to estimate the unknown orders of the contigs for each species and each sample. As a comparison, we also considered the iRep estimator proposed in Brown et al. (2016), where the contigs of a given species were ordered for each sample separately based on the read counts observed. We evaluate these methods by comparing the estimated contig orders to their true orders as measured by the normalized Kendall's tau distance. To generalize our evaluation to diverse metagenomic datasets, we also evaluate the effect of sample size as well as contig numbers by randomly selecting subsets of samples or contigs from each dataset. The selection was made with replacement.

The results are summarized in Figure 5 by comparing the normalized Kendall's tau distances. As n or p varies, our proposed estimator performs consistently better than iRep in recovering the true contig orders, which explains partially why the DEMIC algorithm worked better in estimating the bacterial growth dynamics. The results of our methods are not sensitive to the sample size and the number of contigs from the genome assemblies. Our estimator also shows smaller variability.

6.3 Analysis of a Real Microbiome Metagenomic Data Set

Finally, we complete our numerical studies by analyzing a real metagenomic dataset from the Pediatric Longitudinal Study of Elemental Diet and Stool Microbiome Composition (PLEASE) study, a prospective cohort study to investigate the treatment effects on the gut microbiome and

reduction of inflammation in pediatric Crohn’s disease patients (Lewis et al. 2015). In particular, sequencing data from the fecal samples of 86 Crohn’s disease children were obtained at baseline, 1 week and 8 weeks after antiTNF or enteral diet treatment. In our analysis, the sequencing data at the 8th week after treatment was used to compare the bacterial growth dynamics for non-responders ($n = 34$) and responders ($n = 47$). The reads were downloaded from NCBI short read archive (SRP057027) with the corresponding metadata. After the same coassembly, alignment and binning steps as in Gao and Li (2018), the DEMIC algorithm was applied to estimate the bacterial growth rate of a given species represented by a contig cluster (bin) for each sample. In particular, DEMIC applied our proposed method to the GC-adjusted contig coverage data to recover the original order of the contigs. After obtaining the ordered contigs, a simple linear regression was fitted to obtain estimates of the PTRs (ePTRs).

In order to compare the bacterial growth rates between responders and non-responders, our analysis focused on ePTRs of 8 contig clusters over subsets of the non-responders (n_1) and the responders (n_2) after 8 weeks of treatment with $\min\{n_1, n_2\} > 5$. Other contig clusters were rare and only appeared in a few samples. For each contig cluster, we compare the ePTRs of the responders and non-responders Wilcoxon rank sum test (Table S.1 in Supplementary Material). The taxonomic annotations of these eight contig clusters were obtained by applying the BAT algorithm (von Meijenfildt et al. 2019) that compares the metagenomic assembled bins to a taxonomy database. In Table S.1, we show the final taxonomic annotations for each bin to the finest possible resolution, with the lineage scores indicating the quality of each taxonomic classification.

Among the 8 contig clusters, bin.026 showed a significant difference in ePTRs between responders and non-responders after either antiTNF or enteral diet treatment for 8 weeks ($p=0.0418$), where the growth rate was higher in Crohn’s disease patients who did not respond to the treatment. The taxonomic classification (Table S.1) shows that this contig cluster belongs to the phylum *Firmicutes* and the order *Clostridiales*. Since BAT algorithm was not able to classify the order *Clostridiales* to finer taxonomic level of known species, this contig cluster may represent a new species that is important to the treatment outcome of Crohn’s disease patients.

7 DISCUSSION

In this paper, partial recovery was studied under the normalized Kendall’s tau distance. Another commonly used metric is the normalized Spearman’s footrule distance defined by

$$\rho(\pi_1, \pi_2) = \frac{2}{p(p-1)} \sum_{i=1}^p |\pi_1(i) - \pi_2(i)|, \quad \pi_1, \pi_2 \in \mathcal{S}_p.$$

A celebrated result by Diaconis and Graham (1977) shows that $\tau_K(\pi_1, \pi_2) \leq \rho(\pi_1, \pi_2) \leq 2\tau_K(\pi_1, \pi_2)$, which means the two distances are equivalent. As a consequence, all the theoretical results presented in this paper concerning the Kendall’s tau distance also hold for the Spearman’s footrule distance without any change.

The minimax optimality of the proposed estimator $\hat{\pi}$ was investigated in Section 5 by examining the asymptotic sharpness of the MSG condition for exact recovery, and by obtaining the matching minimax risk lower bound for partial recovery. There are a few issues that deserve further investigation. For both exact and partial recovery, it is unclear to what extent the GSS condition is necessary. In our risk analysis, the perturbation bound for the left singular subspace (Cai and Zhang 2018) was used. In fact, similar results can be obtained using the concentration bound for the linear functionals of singular vectors (Koltchinskii and Xia 2016). Nevertheless, it remains to show whether the GSS condition is also asymptotically sharp. In addition, in Theorem 6, the matching minimax lower bound was obtained only for nonvanishing Γ/σ . It remains to show whether the rate $\sigma/(p\Gamma)$ is minimax optimal when $\Gamma/\sigma \rightarrow 0$. The difficulty lies in finding a $p^{1+\delta}$ -sphere packing of the group \mathcal{S}_p equipped with the Kendall’s tau distance for any $0 < \delta < 1$, while the pairwise ℓ_2 distances of the packing elements are also well controlled. Some initial steps have been made in the so-called rank modulation theory (Barg and Mazumdar 2010; Mazumdar, Barg, and Zemor 2013).

There are several related problems that are also of significant theoretical and practical interest. Firstly, although we used the Kendall’s tau distance or the equivalent Spearman’s footrule distance as the metric for partial recovery, other distances such the Hamming distance, Spearman’s rank correlation distance, and Ulam’s distance have also been used as the performance metrics for par-

tial recovery in other permutation estimation problems (Göloğlu et al. 2015; Mukherjee 2016). It is therefore of interest to see how $\hat{\pi}$ performs under these losses. Secondly, our proposed estimator $\hat{\pi}$ implicitly performs a (linear) dimension reduction technique and only uses the information contained in the first eigenvector of A in (4). A natural extension is to consider the eigen-subspace spanned by the first k eigenvectors and to estimate the permutation in a sequential manner.

The present paper focuses on the estimation of the permutation matrix Π . It is also of interest to estimate the underlying signal matrix Θ or some functionals of it. For example, in microbiome growth dynamics studies, it is of significant interest to estimate the peak-to-trough ratio $\exp(\theta_{kp} - \theta_{k1})$ for $k = 1, \dots, n$, which measures the microbial growth rate for the k th sample, and to identify the samples with peak-to-trough ratio of 1. It is also interesting to identify the bacteria that show differential growth dynamics between disease and normal individuals. Finally, robust permutation recovery methods that can relax the Gaussian or sub-Gaussian assumption of the noise in the permuted monotone matrix model are needed. For example, in some applications, the columns of the noise matrix are not independent, or the variance levels across the noise matrix are not identical. In these cases, we argue that, as long as the marginal distributions of the noise matrix entries remain sub-Gaussian, the analytical framework of the current paper can still be applied, but with more efforts to control the underlying heteroskedasticity. Toward this end, results from the recent work of Zhang, Cai, and Wu (2018) can be very useful, in terms of the new technical tools that parallel the ones used in the current paper to analyse the homoskedastic PCA (cf. Lemma 2 and 3). Finally, to account for non-informative samples, sparse PCA (Cai, Ma, and Wu 2013; Yuan and Zhang 2013) can be considered. These are interesting problems left for future research.

8 PROOFS OF THE MAIN THEOREMS

In this section, we prove Theorems 1 and 2 in detail and briefly sketch the proofs of Theorems 3 and 4. We also prove the minimax lower bounds in Theorems 5 and 6. Proofs of other results including the technical lemmas can be found in the online Supplementary Materials.

Proof of Theorem 1. Let $X = \Theta + Z$. It follows that $Y = X\Pi$. By right invariance of the 0-1 loss with respect to permutation composition, we have

$$\ell((\mathbf{r}(\hat{w}^\top Y))^{-1}, \pi) = \ell((\mathbf{r}(\hat{w}^\top X\Pi))^{-1}, \pi) = \ell((\mathbf{r}(\hat{w}^\top X))^{-1} \circ \pi, \pi) = \ell((\mathbf{r}(\hat{w}^\top X))^{-1}, id).$$

Thus it suffices to study the risk $\mathbb{E}\ell((\mathbf{r}(\hat{w}^\top X))^{-1}, id) = P((\mathbf{r}(\hat{w}^\top X))^{-1} \neq id)$. In fact,

$$\begin{aligned} P((\mathbf{r}(\hat{w}^\top X))^{-1} \neq id) &\leq P\left(\bigcup_{i=1}^{p-1} \left\{ \sum_{k=1}^n \hat{w}_k X_{ki} \geq \sum_{k=1}^n \hat{w}_k X_{k,i+1} \right\}\right) \\ &\leq \sum_{i=1}^{p-1} P\left(\sum_{k=1}^n \hat{w}_k X_{ki} \geq \sum_{k=1}^n \hat{w}_k X_{k,i+1}\right) = \sum_{i=1}^{p-1} P\left(\sum_{k=1}^n \hat{w}_k (X_{ki} - X_{k,i+1}) \geq 0\right), \end{aligned} \quad (13)$$

which further reduces to obtaining an upper bound for $P_i = P(\sum_{k=1}^n \hat{w}_k (X_{ki} - X_{k,i+1}) \geq 0)$. By definition, \hat{w} is the first eigenvector of $A = Y(I - \frac{1}{p}ee^\top)(I - \frac{1}{p}ee^\top)Y^\top$. Simple calculation yields $\Pi(I - \frac{1}{p}ee^\top)(I - \frac{1}{p}ee^\top)\Pi^\top = (I - \frac{1}{p}ee^\top)(I - \frac{1}{p}ee^\top)$ for any $\Pi \in \mathcal{S}_p$. So \hat{w} is also the first eigenvector of $A = X(I - \frac{1}{p}ee^\top)(I - \frac{1}{p}ee^\top)X^\top \equiv TT^\top$, where $T \in \mathbb{R}^{n \times p}$. Note that T admits the decomposition $T = \Theta' + E \in \mathbb{R}^{n \times p}$ where $E_{ij} \sim N(0, (p-1)\sigma^2/p)$ and $\Theta' = a\eta'^\top, \eta'_j = \eta_j - \frac{1}{p}\sum_{i=1}^p \eta_i$. In particular, $T_{\cdot i} = X_{\cdot i} - \bar{X}_{row}$ where $\bar{X}_{row} = p^{-1}\sum_{i=1}^p X_{\cdot i} \in \mathbb{R}^n$ is the vector of row means of X . We denote $\phi_{ij} = T_{\cdot i} - T_{\cdot j} = X_{\cdot i} - X_{\cdot j}$ and denote $w = a/\|a\|_2 \in \mathbb{R}^n$ as the first eigenvector of the rank-one matrix $\Theta'\Theta'^\top$. Now following (13), we have

$$\begin{aligned} P_i &= P\left(\sum_{k=1}^n \hat{w}_k (X_{ki} - X_{k,i+1}) \geq 0\right) = P\left(w^\top \phi_{ij} + (\hat{w} - w)^\top \phi_{ij} \geq 0\right) \\ &= P\left(w^\top \phi_{i,i+1} + (\hat{w} - w)^\top \phi_{ij} \geq 0, |1 - (\hat{w}^\top w)^2| \leq \delta\right) \\ &\quad + P\left(w^\top \phi_{i,i+1} + (\hat{w} - w)^\top \phi_{ij} \geq 0, |1 - (\hat{w}^\top w)^2| > \delta\right) \end{aligned}$$

for some $\delta > 0$. By definition, up to a change of sign for \hat{w} , we have $0 \leq \hat{w}^\top w \leq 1$. Then $|1 - (\hat{w}^\top w)^2| \leq \delta$ implies $|(\hat{w} - w)^\top \phi_{i,i+1}| \leq \|\hat{w} - w\|_2 \|\phi_{i,i+1}\|_2 \leq \sqrt{2\delta} \|\phi_{i,i+1}\|_2$, where the first inequality follows from Cauchy-Schwartz and the second inequality used $\|\hat{w} - w\|_2 = \sqrt{2(1 - \hat{w}^\top w)} \leq \sqrt{2(1 - (\hat{w}^\top w)^2)}$. Thus

$$P_i \leq P\left(w^\top \phi_{i,i+1} \geq -\sqrt{2\delta} \|\phi_{i,i+1}\|_2\right) + P\left(|1 - (\hat{w}^\top w)^2| > \delta\right). \quad (14)$$

The following lemmas provide upper bounds for the two probability events in the last expression.

Lemma 1. *Under the conditions of Theorem 1, denote $\Gamma_i = \|a\|_2(\eta_i - \eta_{i+1})$, then for any $\delta > 0$, we have*

$$P\left(w^\top \phi_{i,i+1} \geq -\sqrt{2\delta}\|\phi_{i,i+1}\|\right) \leq \Phi\left(C_1\sqrt{\delta}\Psi_i^{1/2} + \frac{\Gamma_i}{\sigma}\right) + \frac{C_2}{p^c} \quad (15)$$

for $\Psi_i = (\sqrt{n} + \sqrt{\log p})^2 + \frac{\Gamma_i^2}{\sigma^2} + \frac{|\Gamma_i|}{\sigma}\sqrt{\log p}$ and some constants $C_1, C_2, c > 0$.

Lemma 2. *Suppose $\lambda_1^2(\Theta') \geq C\sigma^2(n + \sqrt{np})$ for some $C > 0$, it follows that*

$$P\left(|1 - (\hat{w}^\top w)^2| \leq C_1 \frac{\sigma^2(\lambda_1^2(\Theta') + \sigma^2 p)(n + \log p)}{\lambda_1^4(\Theta')}\right) \geq 1 - \frac{C_2}{p^c}$$

for some $C_1, C_2, c > 0$.

Now since $\frac{1}{p} \sum_{1 \leq i < j \leq p} (\eta_i - \eta_j)^2 = \sum_{j=1}^p (\eta_j - \frac{1}{p} \sum_{i=1}^p \eta_i)^2 = \sum_{j=1}^p \eta_j'^2$, we have $\lambda_1^2(\Theta') = \Lambda > C_0\sigma^2(n + \sqrt{np})$ for some $C_0 > 0$. Set $\delta = C_0\sigma^2 \frac{(\lambda_1^2(\Theta') + \sigma^2 p)(n + \log p)}{\lambda_1^4(\Theta')}$. It follows that $\delta = o(1)$.

Combining Lemma 1 and Lemma 2, we have

$$P_i \leq \Phi\left(C\sqrt{\delta}\left[(\sqrt{n} + \sqrt{\log p})^2 + \frac{\Gamma_i^2}{\sigma^2} + \frac{|\Gamma_i|}{\sigma}\sqrt{\log p}\right]^{1/2} + \frac{\Gamma_i}{\sigma}\right) + \frac{C}{p^c} \quad (16)$$

for some $C, c > 0$. The rest of the analysis is divided into several cases.

Case 1. $\log p \lesssim n$. In this case, we have $P_i \leq \Phi\left(C\sqrt{\delta}\left[n + \frac{\Gamma_i^2}{\sigma^2} + \frac{|\Gamma_i|}{\sigma}\sqrt{\log p}\right]^{1/2} + \frac{\Gamma_i}{\sigma}\right) + \frac{C}{p^c}$. In addition, if $|\Gamma_i|/\sigma \lesssim \sqrt{n}$, we have $P_i \leq \Phi\left(C\sqrt{\delta n} + \frac{\Gamma_i}{\sigma}\right) + \frac{C}{p^c} \leq \frac{C'}{p^c}$, where the last inequality follows from $\sqrt{\log p} \lesssim \Gamma/\sigma \leq |\Gamma_i|/\sigma \lesssim \sqrt{n}$ and $\Lambda \gtrsim \sigma^2 n \left(\frac{\sigma^2 n}{\Gamma^2} + \frac{\sigma\sqrt{p}}{\Gamma}\right)$. If instead $|\Gamma_i|/\sigma \gtrsim \sqrt{n}$, we have $P_i \leq \Phi\left(C\sqrt{\delta} \frac{|\Gamma_i|}{\sigma} + \frac{\Gamma_i}{\sigma}\right) + \frac{C}{p^c} \leq \frac{C'}{p^c}$, where the last inequality follows from $|\Gamma_i|/\sigma \gtrsim \sqrt{n} \gtrsim \sqrt{\log p}$ and $\delta = o(1)$. Hence, in Case 1, (16) can be bounded by $O(p^{-c})$.

Case 2. $\log p \gtrsim n$. In this case, we have $P_i \leq \Phi\left(C\sqrt{\delta}\left[\log p + \frac{\Gamma_i^2}{\sigma^2} + \frac{|\Gamma_i|}{\sigma}\sqrt{\log p}\right]^{1/2} + \frac{\Gamma_i}{\sigma}\right) + \frac{C}{p^c}$. In addition, since $|\Gamma_i| \geq \Gamma \gtrsim \sigma\sqrt{\log p}$ and $\delta = o(1)$, we have $P\left(\sum_{k=1}^n \hat{u}_k(X_{ki} - X_{k,i+1}) \geq 0\right) \leq \Phi\left(\frac{C\sqrt{\delta}}{\sigma}|\Gamma_i| + \frac{\Gamma_i}{\sigma}\right) + \frac{C}{p^c} \leq \frac{C'}{p^c}$. This shows that, in Case 2, (16) can also be bounded by $O(p^{-c})$.

As a result, it follows that, up to a change of sign for \hat{w} , $P((\tau(\hat{w}^\top X))^{-1} \neq id) = O(p^{-c})$ for some constant $c > 0$. \square

Proof of Theorem 2. Firstly, by invariance property of Kendall's tau distance, $\mathbb{E}[\tau_K(\hat{\pi}, \pi)] = \mathbb{E}[\tau_K((\mathbf{r}(\hat{w}^\top X))^{-1}, id)] = \mathbb{E}[\tau_K((\mathbf{r}(\hat{w}^\top X)), id)]$. It then follows

$$\begin{aligned}\mathbb{E}[\tau_K(\hat{\pi}, \pi)] &= \frac{2}{p(p-1)} \sum_{i < j} P([\mathbf{r}(\hat{w}^\top X)]_i \geq [\mathbf{r}(\hat{w}^\top X)]_j) \\ &= \frac{2}{p(p-1)} \sum_{i < j} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right).\end{aligned}$$

The summation in the last expression can be divided into two parts, namely, the consecutive differences and non-consecutive differences, i.e.,

$$\begin{aligned}\sum_{i < j} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) &= \sum_{(i,j):j=i+1} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) \\ &\quad + \sum_{(i,j):j>i+1} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right).\end{aligned}$$

In the following, we first show

$$P_i = P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{k,i+1}) \geq 0\right) \leq \frac{ce^{-\Gamma^2/2\sigma^2}}{\Gamma/\sigma + \sqrt{\Gamma^2/\sigma^2 + 8/\pi}} + \frac{C}{p^c} \quad (17)$$

so that

$$\sum_{(i,j):j=i+1} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) \leq \frac{cpe^{-\Gamma^2/2\sigma^2}}{\Gamma/\sigma + \sqrt{\Gamma^2/\sigma^2 + 8/\pi}} + \frac{C}{p^c}. \quad (18)$$

for some $C, c > 0$. Then we show that

$$\sum_{(i,j):j>i+1} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) \leq C \frac{p\sigma}{\Gamma} \min\left\{1, e^{-\Gamma^2/2\sigma^2} \log\left(1 + \frac{2\sigma^2}{\Gamma^2}\right)\right\} + \frac{C}{p^c}. \quad (19)$$

Combining (18) and (19), we conclude that

$$\mathbb{E}[\tau_K(\hat{w}^\top Y, \pi)] \leq \frac{C\sigma}{p\Gamma} \min\left\{1, e^{-\Gamma^2/2\sigma^2} \log\left(1 + \frac{2\sigma^2}{\Gamma^2}\right)\right\} + \frac{Ce^{-\Gamma^2/2\sigma^2}}{p(\Gamma/\sigma + \sqrt{8/\pi})} + \frac{C}{p^{c+2}},$$

which completes the proof, as the bound $\mathbb{E}[\tau_K(\hat{w}^\top Y, \pi)] \leq 1$ is trivial.

Proof of (17). Following the same argument as the proof of Theorem 1, we have for $1 \leq i \leq p-1$ and $\delta = \frac{\sigma^2(n+\log p)(\lambda_1^2 + \sigma^2 p)}{\lambda_1^4}$, $P_i \leq P(w^\top \phi_{ij} \geq -\sqrt{2\delta} \|\phi_{i,i+1}\|_2) + P(|1 - (\hat{w}^\top w)^2| > \delta)$, where the second term can be bounded using Lemma 2. For the first term, by Lemma 1, for

$\lambda_1^4(\Theta') \geq \sigma^2(p + \sigma^2\lambda_1^2(\Theta'))(\log p + n)$, we have

$$P_i \leq \Phi\left(C\sqrt{\delta}\left[(\sqrt{n} + \sqrt{\log p})^2 + \frac{\Gamma_i^2}{\sigma^2} + \frac{|\Gamma_i|}{\sigma}\sqrt{\log p}\right]^{1/2} + \frac{\Gamma_i}{\sigma}\right) + \frac{C}{p^c}.$$

Using same argument as the proof of Theorem 1, it holds that $P_i \leq \Phi\left(\frac{\Gamma_i}{\sigma}\right) + \frac{C}{p^c}$. Equation (17) then follows by using formula 7.1.13 of Abramowitz and Stegun (1965) that $\Phi(-t) < \frac{2}{t+\sqrt{t^2+8/\pi}}\phi(t)$ for $t \geq 0$.

Proof of (19). For the set of indices $S = \{(i, j) : 1 \leq i < j \leq p, j > i + 1\}$, we further divide it into two subsets $S_1 = \{(i, j) : 1 \leq i < j \leq p, j > i + \lfloor \sigma\sqrt{C\log p}/\Gamma \rfloor\}$ and $S_2 = \{(i, j) : 1 \leq i < j \leq p, i + 1 < j \leq i + \lfloor \sigma\sqrt{C\log p}/\Gamma \rfloor\}$ for some constant $C > 0$. Apparently we have the decomposition

$$\begin{aligned} \sum_{(i,j):j>i+1} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) &= \sum_{(i,j) \in S_1} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) \\ &+ \sum_{(i,j) \in S_2} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) \end{aligned} \quad (20)$$

For the first term, by construction, it can be shown using the same argument (see supplementary materials) in Theorem 1 that

$$\sum_{(i,j) \in S_1} P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) \leq \frac{C|S_1|}{p^c} \leq \frac{C}{p^{c_0}}. \quad (21)$$

Now for the second term in (20), similar argument yields, for $(i, j) \in S_2$,

$$P\left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0\right) \leq \frac{c \exp(-|i-j|^2\Gamma^2/(2\sigma^2))}{|i-j|\Gamma/\sigma + \sqrt{|i-j|^2\Gamma^2/\sigma^2 + 8/\pi}} + \frac{C}{p^c}.$$

Note that, on the one hand,

$$\frac{e^{-|i-j|^2\Gamma^2/(2\sigma^2)}}{|i-j|\Gamma/\sigma + \sqrt{|i-j|^2\Gamma^2/\sigma^2 + 8/\pi}} \leq \frac{\sigma e^{-|i-j|^2\Gamma^2/2\sigma^2}}{|i-j|\Gamma}.$$

We have

$$\sum_{(i,j) \in S_2} \frac{\sigma e^{-|i-j|^2\Gamma^2/2\sigma^2}}{|i-j|\Gamma} = \frac{\sigma}{\Gamma} \sum_{k=2}^{p \wedge \lfloor \sqrt{\log p}/\Gamma \rfloor} \left(\frac{e^{-k^2\Gamma^2/2\sigma^2} p}{k} - e^{-k^2\Gamma^2/2\sigma^2} \right) = T_1 - T_2.$$

For the rest of the proof, we assume $\Gamma \leq \sigma\sqrt{\log p}/2$, otherwise the set S_2 will vanish. Then

$$T_1 = \frac{\sigma p}{\Gamma} \sum_{k=2}^{p \wedge \lfloor \sigma\sqrt{\log p}/\Gamma \rfloor} \frac{e^{-k^2\Gamma^2/2\sigma^2}}{k} \leq \frac{\sigma p}{\Gamma} \int_1^{p \wedge \lfloor \sigma\sqrt{\log p}/\Gamma \rfloor} \frac{e^{-x^2\Gamma^2/2\sigma^2}}{x} dx$$

where the last inequality used monotonicity of the integrand. The integral in the last inequality, after change of variable, can be bounded by an exponential integral $\text{Ei}(\Gamma^2/2\sigma^2)$, which has an upper bound

$$\int_1^{p \wedge \lfloor \sigma\sqrt{\log p}/\Gamma \rfloor} \frac{e^{-x^2\Gamma^2/2\sigma^2}}{x} dx = \frac{1}{2} \int_{\Gamma^2/2\sigma^2}^{(p^2\Gamma^2/2\sigma^2) \wedge (\log p/2)} \frac{e^{-t}}{t} dt \leq \frac{1}{2} \int_{\Gamma^2/2\sigma^2}^{\infty} \frac{e^{-t}}{t} dt \leq e^{-\Gamma^2/2\sigma^2} \log \left(1 + \frac{2\sigma^2}{\Gamma^2} \right)$$

so that $T_1 \leq \frac{\sigma p}{\Gamma} e^{-\Gamma^2/2\sigma^2} \log \left(1 + \frac{2\sigma^2}{\Gamma^2} \right)$. For T_2 , we have $T_2 = \frac{\sigma}{\Gamma} \sum_{k=2}^{p \wedge \lfloor \sigma\sqrt{\log p}/\Gamma \rfloor} e^{-k^2\Gamma^2/2\sigma^2} \geq \frac{\sigma}{\Gamma} e^{-2\Gamma^2/\sigma^2}$. Therefore,

$$\sum_{(i,j) \in S_2} P \left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0 \right) \leq \frac{C\sigma p}{\Gamma} e^{-\Gamma^2/2\sigma^2} \log \left(1 + \frac{2\sigma^2}{\Gamma^2} \right) + \frac{C}{p^c}. \quad (22)$$

On the other hand, note that

$$\frac{e^{-|i-j|^2\Gamma^2/(2\sigma^2)}}{|i-j|\Gamma/\sigma + \sqrt{|i-j|^2\Gamma^2/\sigma^2 + 8/\pi}} \leq ce^{-|i-j|^2\Gamma^2/(2\sigma^2)}.$$

We have

$$\begin{aligned} \sum_{(i,j) \in S_2} e^{-|i-j|^2\Gamma^2/(2\sigma^2)} &= \sum_{k=2}^{p \wedge \lfloor \sigma\sqrt{\log p}/\Gamma \rfloor} pe^{-k^2\Gamma^2/(2\sigma^2)} - \sum_{k=2}^{p \wedge \lfloor \sigma\sqrt{\log p}/\Gamma \rfloor} ke^{-k^2\Gamma^2/(2\sigma^2)} \\ &\leq p \int_1^{\infty} e^{-k^2\Gamma^2/2\sigma^2} dk - 2e^{-2\Gamma^2/\sigma^2} \leq Cp\sigma/\Gamma. \end{aligned}$$

Thus

$$\sum_{(i,j) \in S_2} P \left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0 \right) \leq Cp\sigma/\Gamma + \frac{C'}{p^c}. \quad (23)$$

Combining (22) and (23), we have

$$\sum_{(i,j) \in S_2} P \left(\sum_{k=1}^n \hat{w}_k(X_{ki} - X_{kj}) \geq 0 \right) \leq C \frac{p\sigma}{\Gamma} \min \left\{ 1, e^{-\Gamma^2/2\sigma^2} \log \left(1 + \frac{2\sigma^2}{\Gamma^2} \right) \right\} + \frac{C}{p^c} \quad (24)$$

Combining (21) and (24), we have (19). \square

Proof of Theorem 3 and Theorem 4. Here we only provide a sketch of the proofs. We refer the readers to our Supplementary Material for detailed proofs. The proofs follow essentially from

the same argument as the proofs of Theorem 1 and Theorem 2, respectively. However, in place of Lemma 2 used therein, we need the following lemma that provides a perturbation bound for the leading eigenvector of approximate rank-one matrices, which could be of independent interest.

Lemma 3. *Suppose $p \gtrsim n$ and $\lambda_1^2(\Theta') \geq \lambda_2^2(\Theta') + C\sigma^2(n + \sqrt{np})$ for some $C > 0$. Let $w = u'_1$ be the first left singular vector of Θ' , it follows that,*

$$P\left(|1 - (\hat{w}^\top w)^2| \leq \frac{C\sigma^2(\lambda_1^2(\Theta') + \sigma^2 p)(n + \log p)}{(\lambda_1^2(\Theta') - \lambda_2^2(\Theta'))^2}\right) \geq 1 - \frac{C}{p^c}.$$

The proof of Lemma 3 is nontrivial, which depends on a combination of the generic perturbation bound obtained by Cai and Zhang (2018) and new concentration inequalities of approximate rank-one matrices (see Supplementary Materials).

Proof of Theorem 5. The proof relies on the following lemma adapted from (Tsybakov 2009).

Lemma 4. *Assume that for some integer $M \geq 2$ there exist distinct parameters $\theta_0, \dots, \theta_M$ from the parameter space Θ and mutually absolutely continuous probability measures P_0, \dots, P_M with $P_j = P_{\theta_j}$ for $j = 0, 1, \dots, M$, defined on a common probability space (Ω, \mathcal{F}) such that the averaged K-L divergence $\frac{1}{M} \sum_{j=1}^M D(P_j, P_0) \leq \frac{1}{8} \log M$. Then, for every measurable mapping $\hat{\theta} : \Omega \rightarrow \Theta$,*

$$\max_{j=0, \dots, M} P_j(\hat{\theta} \neq \theta_j) \geq \frac{\sqrt{M}}{\sqrt{M} + 1} \left(\frac{3}{4} - \frac{1}{2\sqrt{\log M}} \right).$$

We construct the $(M + 1) = p$ points parameter space as follows. We define p permutations from \mathcal{S}_p as an identity plus $(p-1)$ consecutive swaps, i.e., $\pi_0 = id$, $\pi_k = (k, k+1)$ for $k = 1, \dots, p-1$. The signal matrix $\Theta_0 = a\eta^\top$ where $a = (1, \dots, 1)^\top \in \mathbb{R}^n$ and $\eta = (0, \delta, \dots, (p-1)\delta)^\top \in \mathbb{R}^p$, $\delta = \frac{\sigma}{4}\sqrt{\log p/n}$. In this way, we have $\Gamma = \|a\|_2 \cdot \min_{1 \leq i \leq p-1} |\eta_i - \eta_{i+1}| = \frac{\sigma}{4}\sqrt{\log p}$. Let P_k corresponds to the joint probability measure of Y under (Θ_0, π_k) for $k = 0, 1, \dots, p-1$, and let p_k be the pdf of P_k , we have $p_0(x) = \prod_{i=1}^n \prod_{j=1}^p \phi_{\eta_j}(x_{ij})$, $p_k(x) = \prod_{i=1}^n \prod_{j=1}^p \phi_{\eta_{\pi_k(j)}}(x_{ij})$ for $k = 1, \dots, p-1$, where ϕ_μ is the pdf of Gaussian distribution $N(\mu, \sigma^2)$. Now we calculate the

KL-divergence

$$\begin{aligned} D(P_k, P_0) &= \int \log \left(\frac{p_k(x)}{p_0(x)} \right) p_0(x) dx = \int \frac{n}{2\sigma^2} \sum_{i=1}^p [(x_{1j} - \eta_{\pi_k(j)})^2 - (x_{1j} - \eta_j)^2] p_0(x) dx \\ &= \frac{n\delta^2}{\sigma^2} = \frac{\log p}{16}. \end{aligned}$$

Then, we have for $p \geq 10$, $\frac{1}{p-1} \sum_{k=1}^{p-1} D(P_k, P_0) = \frac{\log p}{16} \leq \frac{1}{8} \log(p-1)$. It follows from Lemma 4 that, $\inf_{\hat{\pi}} \sup_{(\pi, \Theta) \in \mathcal{D}_1} P(\hat{\pi} \neq \pi) \geq \inf_{\hat{\pi}} \max_{j=0, \dots, p-1} P_j(\hat{\pi} \neq \pi_j) \geq 0.3$ as long as $p \geq 10$. In addition, $\inf_{\hat{\pi}} \sup_{(\pi, \Theta) \in \mathcal{D}'_1} P(\hat{\pi} \neq \pi) \geq \inf_{\hat{\pi}} \sup_{(\pi, \Theta) \in \mathcal{D}_1} P(\hat{\pi} \neq \pi)$ as $\mathcal{D}_1 \subset \mathcal{D}'_1$. \square

Proof of Theorem 6. The proof relies on the following lemma from Tsybakov (2009).

Lemma 5. Assume that $M \geq 2$ and suppose that Θ contains elements $\theta_0, \theta_1, \dots, \theta_M$ such that: (i) $d(\theta_j, \theta_k) \geq 2s > 0$ for any $0 \leq j < k \leq M$; (ii) for any $j = 1, \dots, M$, $\frac{1}{M} \sum_{j=1}^M D(P_j, P_0) \leq \alpha \log M$ with $0 < \alpha < 1/8$ and $P_j = P_{\theta_j}$ for $j = 0, 1, \dots, M$. Then

$$\inf_{\hat{\theta}} \sup_{\theta \in \Theta} P_{\theta}(d(\hat{\theta}, \theta) \geq s) \geq \frac{\sqrt{M}}{1 + \sqrt{M}} \left(1 - 2\alpha - \sqrt{\frac{2\alpha}{\log M}} \right) > 0.$$

We also need the following sphere packing lemma proved by Mao, Weed, and Rigollet (2017), which is a direct consequence of the well-celebrated Varshamov-Gilbert bound.

Lemma 6. For any $r < p/2$, there exists a subset \mathcal{Q}_r of \mathcal{S}_p such that (i) $\log |\mathcal{Q}_r| \geq \frac{r}{5} \log(p/r)$, (ii) for any elements $\pi_1, \pi_2 \in \mathcal{Q}_r$, we have $\binom{p}{2} \cdot \tau_K(\pi_1, \pi_2) \geq r$, and (iii) for any $\pi \in \mathcal{Q}_r$, we have $\|\pi - id\|_2^2 \leq 2r$.

For $t/\sigma \geq 2$, we set $r = \frac{p\sigma}{t} e^{-t^2/2\sigma^2} < p/2$. Let $\pi_0 = id$ and $\pi_1, \dots, \pi_{|\mathcal{Q}_r|}$ be the elements of \mathcal{Q}_r . The signal matrix $\Theta_0 = a\eta^\top$ where $a = (1/\sqrt{320n}, \dots, 1/\sqrt{320n})^\top \in \mathbb{R}^n$ and $\eta = (t, \dots, pt)^\top \in \mathbb{R}^p$. Let P_k be the joint probability measure of Y under (Θ_0, π_k) for $k = 0, 1, \dots, |\mathcal{Q}_r|$, and let p_k be the pdf of P_k . By Lemma 6, the KL-divergence

$$D(P_k, P_0) = \int \log \left(\frac{p_k(x)}{p_0(x)} \right) p_0(x) dx = \frac{t^2}{320\sigma^2} \|\pi_k - id\|_2^2 \leq \frac{pt}{160\sigma} e^{-t^2/2\sigma^2}.$$

and therefore

$$\frac{1}{p-1} \sum_{k=1}^{p-1} D(P_k, P_0) \leq \frac{pt}{160\sigma} e^{-t^2/2\sigma^2} \leq \frac{p\sigma}{80t} e^{-t^2/2\sigma^2} \log \left(\frac{t}{\sigma} e^{t^2/2\sigma^2} \right) \leq \frac{1}{16} \log |\mathcal{Q}_r|.$$

Without loss of generality, we assume $|\mathcal{Q}_r| \geq 2$. By Lemma 5, it then follows that,

$$\inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}_2(t)} P\left(\tau_K(\hat{\pi}, \pi) \geq \frac{\sigma}{2pt} e^{-t^2/2\sigma^2}\right) \geq C_1,$$

for some absolute constant $C_1 > 0$. By Markov's inequality, we have

$$\inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}_2(t)} \mathbb{E}[\tau_K(\hat{\pi}, \pi)] \geq \frac{\sigma}{2pt} e^{-t^2/2\sigma^2} \inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}_2(t)} P\left(\tau_K(\hat{\pi}, \pi) \geq \frac{\sigma}{2pt} e^{-t^2/2\sigma^2}\right) \geq \frac{C_1 \sigma}{pt} e^{-t^2/2\sigma^2}.$$

The relationship $\inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}_2(t)} \mathbb{E}[\tau_K(\hat{\pi}, \pi)] \geq \inf_{\hat{\pi}} \sup_{(\Theta, \pi) \in \mathcal{D}_2(t)} \mathbb{E}[\tau_K(\hat{\pi}, \pi)]$ follows from $\mathcal{D}_L \subset \mathcal{D}$. The rate $1/p^2$ follows by setting $t = C_2 \sigma \sqrt{\log p}$ for some $C_2 > 0$. \square

ACKNOWLEDGEMENT

We would like to thank the Associate Editor and the anonymous referees for many helpful suggestions that significantly improved the paper. R. M. would also like to thank Rui Duan, Anru Zhang, Yuan Gao and Shulei Wang for stimulating discussions at various stages of this project.

FUNDING

SUPPLEMENTARY MATERIALS

In our online Supplemental Materials, we prove Theorem 3-4, Proposition 1-4, as well as the technical lemmas. Some supplementary simulations, figures and tables are included in the appendix.

References

- Abel, S., Zur Wiesch, P. A., Chang, H.-H., Davis, B. M., Lipsitch, M., and Waldor, M. K. (2015), "Sequence tag-based analysis of microbial population dynamics," *Nature Methods*, 12, 223.
- Abramowitz, M., and Stegun, I. A. (1965), *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*, vol. 55, Courier Corporation.
- Barg, A., and Mazumdar, A. (2010), "Codes in permutations and error correction for rank modulation," in *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, IEEE, pp. 854–858.
- Boulund, F., Pereira, M. B., Jonsson, V., and Kristiansson, E. (2018), "Computational and statistical considerations in the analysis of metagenomic data," in *Metagenomics*, Elsevier, pp. 81–102.

- Bremer, H., and 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., and Banfield, J. F. (2016), “Measurement of bacterial replication rates in microbial communities,” *Nature Biotechnology*, 34, 1256.
- Cai, T. T., Ma, Z., and Wu, Y. (2013), “Sparse PCA: Optimal rates and adaptive estimation,” *The Annals of Statistics*, 41, 3074–3110.
- Cai, T. T., and Zhang, A. (2018), “Rate-optimal perturbation bounds for singular subspaces with applications to high-dimensional statistics,” *The Annals of Statistics*, 46, 60–89.
- Chatterjee, S., Guntuboyina, A., and Sen, B. (2015), “On risk bounds in isotonic and other shape restricted regression problems,” *The Annals of Statistics*, 43, 1774–1800.
- (2018), “On matrix estimation under monotonicity constraints,” *Bernoulli*, 24, 1072–1100.
- Collier, O., and Dalalyan, A. S. (2016), “Minimax rates in permutation estimation for feature matching,” *The Journal of Machine Learning Research*, 17, 162–192.
- Cooper, S., and Helmstetter, C. E. (1968), “Chromosome replication and the division cycle of *Escherichia coli* B/r,” *Journal of Molecular Biology*, 31, 519–540.
- Cullina, D., and Kiyavash, N. (2016), “Improved achievability and converse bounds for Erdős-Renyi graph matching,” in *ACM SIGMETRICS Performance Evaluation Review*, ACM, vol. 44, pp. 63–72.
- Currie, R. R., and Pandher, G. S. (2011), “Finance journal rankings and tiers: An active scholar assessment methodology,” *Journal of Banking & Finance*, 35, 7–20.
- Deshpande, S. K., and Jensen, S. T. (2016), “Estimating an NBA player’s impact on his team’s chances of winning,” *Journal of Quantitative Analysis in Sports*, 12, 51–72.
- Diaconis, P. (1988), *Group Representations in Probability and Statistics.*, Institute of Mathematical Statistics Lecture Notes–Monograph Series (11).
- Diaconis, P., and Graham, R. L. (1977), “Spearman’s footrule as a measure of disarray,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 262–268.
- Flammarion, N., Mao, C., and Rigollet, P. (2019), “Optimal rates of statistical seriation,” *Bernoulli*, 25, 623–653.
- Gao, F., Luo, H., and 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., and Li, H. (2018), “Quantifying and comparing bacterial growth dynamics in multiple metagenomic samples,” *Nature Methods*, 15, 1041–1044.
- Göloğlu, F., Lember, J., Riet, A.-E., and Skachek, V. (2015), “New bounds for permutation codes in Ulam metric,” in *Information Theory (ISIT), 2015 IEEE International Symposium on*, IEEE,

- pp. 1726–1730.
- Kendall, M. G. (1938), “A new measure of rank correlation,” *Biometrika*, 30, 81–93.
- Koltchinskii, V., and Xia, D. (2016), “Perturbation of linear forms of singular vectors under gaussian noise,” in *High Dimensional Probability VII*, Springer, pp. 397–423.
- Korem, T., Zeevi, D., Suez, J., Weinberger, A., Avnit-Sagi, T., Pompan-Lotan, M., Matot, E., Jona, G., Harmelin, A., and Cohen, N. (2015), “Growth dynamics of gut microbiota in health and disease inferred from single metagenomic samples,” *Science*, aac4812.
- Lewis, J. D., Chen, E. Z., Baldassano, R. N., Otley, A. R., Griffiths, A. M., Lee, D., Bittinger, K., Bailey, A., Friedman, E. S., Hoffmann, C., et al. (2015), “Inflammation, antibiotics, and diet as environmental stressors of the gut microbiome in pediatric Crohn’s disease,” *Cell Host & Microbe*, 18, 489–500.
- Li, D., Liu, C., Luo, R., Sadakane, K., and Lam, T. (2015), “MEGAHIT: an ultra-fast single-node solution for large and complex metagenomics assembly via succinct de Bruijn graph,” *Bioinformatics*, 15, 1674–1676.
- Mao, C., Weed, J., and Rigollet, P. (2017), “Minimax rates and efficient algorithms for noisy sorting,” *arXiv preprint arXiv:1710.10388*.
- Mazumdar, A., Barg, A., and Zemor, G. (2013), “Constructions of rank modulation codes,” *IEEE Transactions on Information Theory*, 59, 1018–1029.
- Mukherjee, S. (2016), “Estimation in exponential families on permutations,” *The Annals of Statistics*, 44, 853–875.
- Myhrvold, C., Kotula, J. W., Hicks, W. M., Conway, N. J., and Silver, P. A. (2015), “A distributed cell division counter reveals growth dynamics in the gut microbiota,” *Nature Communications*, 6, 10039.
- Pananjady, A., Wainwright, M. J., and Courtade, T. A. (2016), “Linear regression with an unknown permutation: Statistical and computational limits,” *arXiv preprint arXiv:1608.02902*.
- (2017), “Denoising linear models with permuted data,” in *Information Theory (ISIT), 2017 IEEE International Symposium on*, IEEE, pp. 446–450.
- Rendle, S., Balby Marinho, L., Nanopoulos, A., and Schmidt-Thieme, L. (2009), “Learning optimal ranking with tensor factorization for tag recommendation,” in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, pp. 727–736.
- Rigollet, P., and Weed, J. (2018), “Uncoupled isotonic regression via minimum Wasserstein deconvolution,” *arXiv preprint arXiv:1806.10648*.
- Slawski, M., and Ben-David, E. (2017), “Linear Regression with Sparsely Permuted Data,” *arXiv preprint arXiv:1710.06030*.

- Tsybakov, A. B. (2009), *Introduction to Nonparametric Estimation*, Springer Series in Statistics. Springer, New York.
- von Meijenfeldt, F. B., Arkhipova, K., Cambuy, D. D., Coutinho, F. H., and Dutilh, B. E. (2019), “Robust taxonomic classification of uncharted microbial sequences and bins with CAT and BAT,” *bioRxiv*, 530188.
- Wu, Y., Tang, Y., Tringe, S., Simmons, B., and Singer, S. (2014), “MaxBin: an automated binning method to recover individual genomes from metagenomes using an expectation-maximization algorithm,” *Microbiome*, 2, 26.
- Yuan, X., and Zhang, T. (2013), “Truncated power method for sparse eigenvalue problems,” *Journal of Machine Learning Research*, 14, 899–925.
- Zhang, A., Cai, T. T., and Wu, Y. (2018), “Heteroskedastic PCA: Algorithm, optimality, and applications,” *arXiv preprint arXiv:1810.08316*.

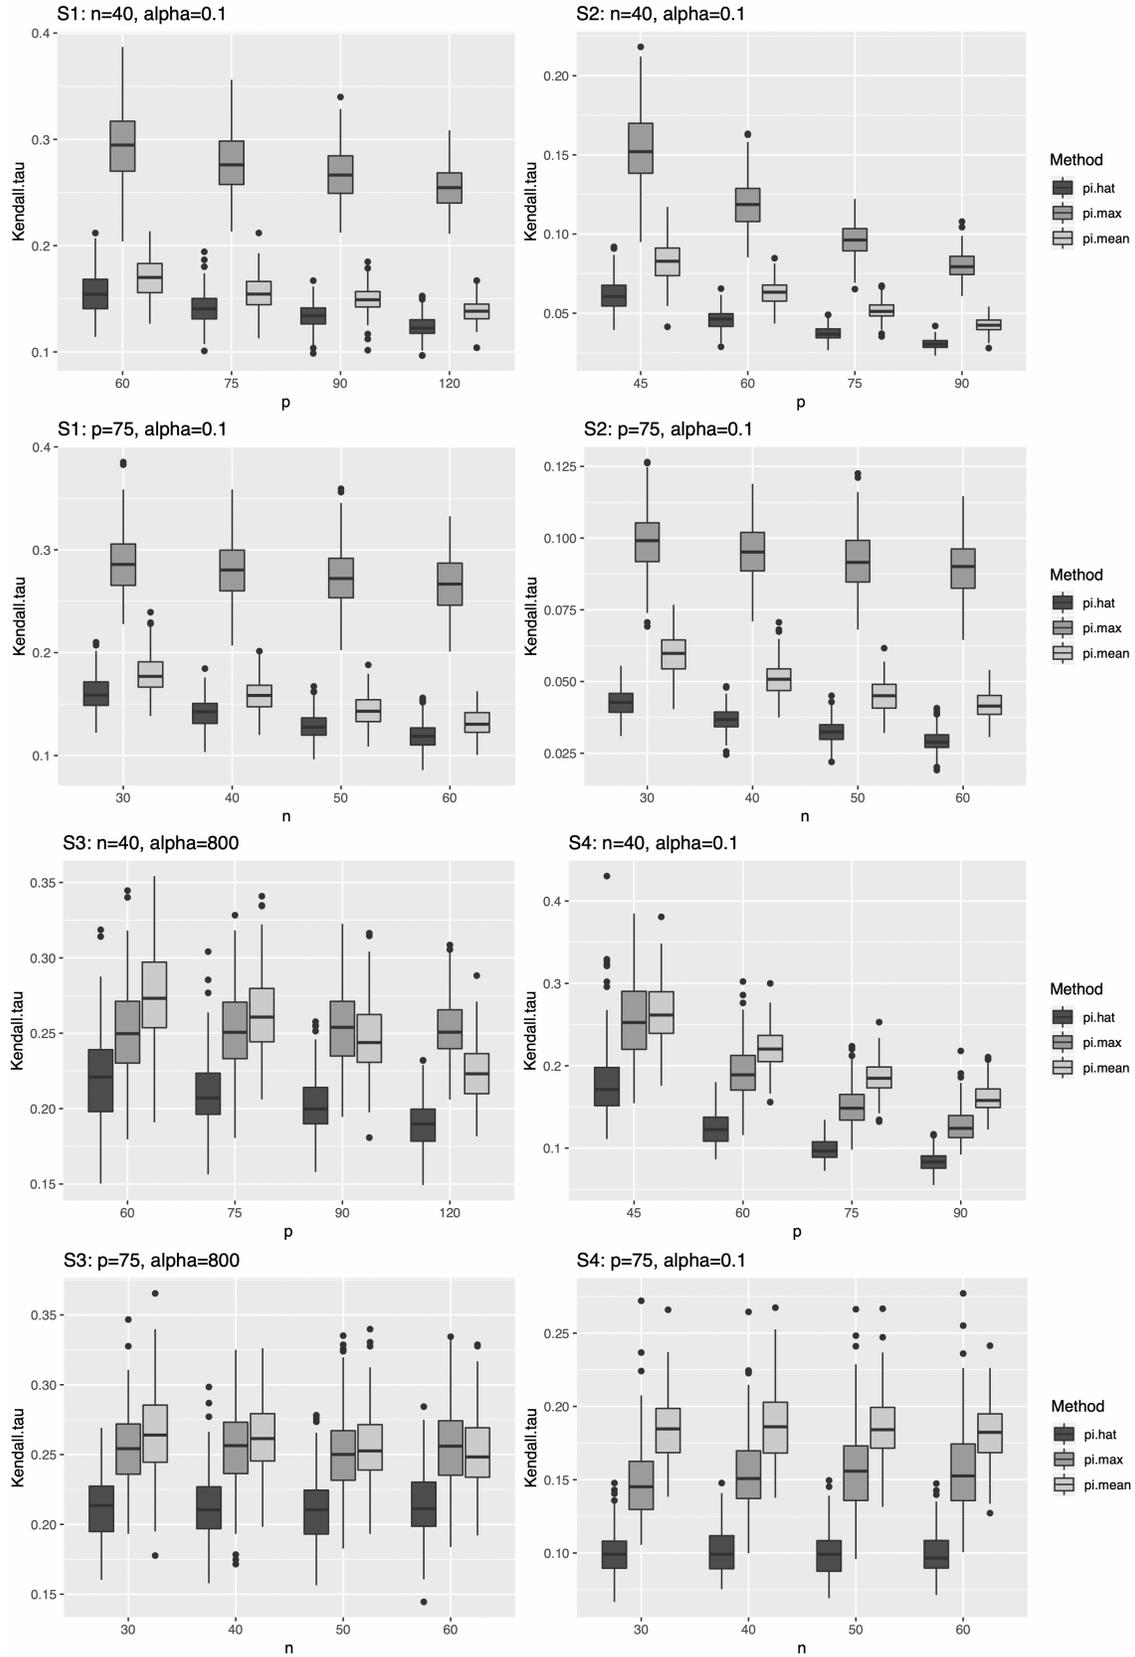


Figure 4: Boxplots of the empirical normalized Kendall's distance between the estimated and true permutations under different models. $\hat{\pi}$: proposed estimator; π_{mean} : mean-based estimator; π_{max} : max-based estimator.

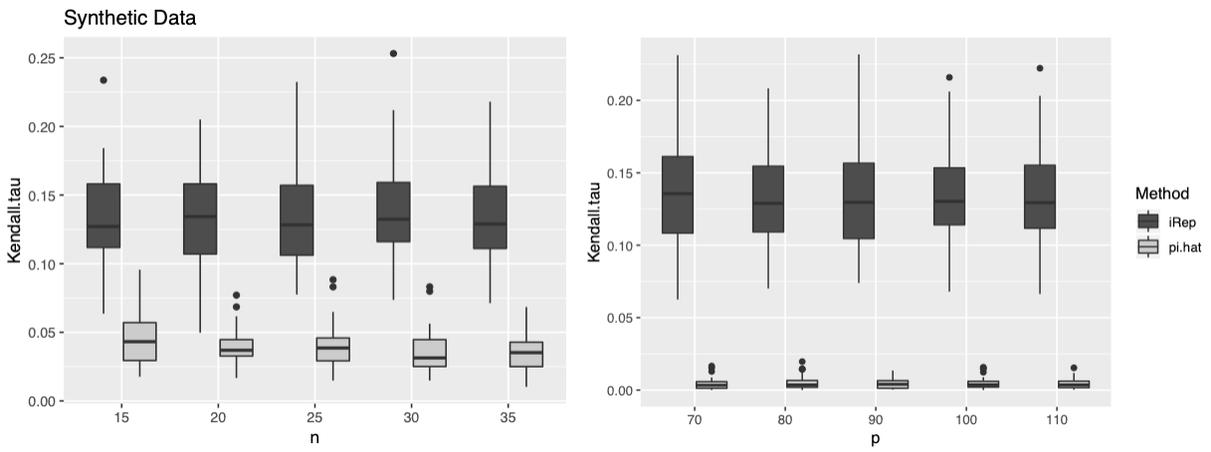


Figure 5: Boxplots of the normalized Kendall's distance between the estimated contig orders and the true orders for different sample sizes n and different numbers of contigs p . The lighter ones correspond to our proposed method and the darker ones correspond to the iRep estimation method.