Published in final edited form as:

Ann Stat. 2019 ; 47(4): 2204–2235. doi:10.1214/18-AOS1745.

SPECTRAL METHOD AND REGULARIZED MLE ARE BOTH OPTIMAL FOR TOP- K RANKING*

YUXIN CHEN[†],

DEPARTMENT OF ELECTRICAL ENGINEERING, PRINCETON UNIVERSITY PRINCETON, NEW JERSEY 08544,
yuxin.chen@princeton.edu

JIANQING FAN^{‡,†},

SCHOOL OF DATA SCIENCE, FUDAN UNIVERSITY, SHANGHAI, CHINA, AND, DEPARTMENT OF OPERATIONS
RESEARCH &, FINANCIAL ENGINEERING, PRINCETON, UNIVERSITY, PRINCETON NEW JERSEY 08544,
jqfan@princeton.edu

CONG MA[†],

DEPARTMENT OF OPERATIONS RESEARCH &, FINANCIAL ENGINEERING, PRINCETON UNIVERSITY, PRINCETON, NEW
JERSEY 08544, congma@princeton.edu

KAIZHENG WANG[†]

DEPARTMENT OF OPERATIONS RESEARCH &, FINANCIAL ENGINEERING, PRINCETON UNIVERSITY, PRINCETON, NEW
JERSEY 08544, kaizheng@princeton.edu

Abstract

This paper is concerned with the problem of top- K ranking from pairwise comparisons. Given a collection of n items and a few pairwise comparisons across them, one wishes to identify the set of K items that receive the highest ranks. To tackle this problem, we adopt the logistic parametric model — the Bradley-Terry-Luce model, where each item is assigned a latent preference score, and where the outcome of each pairwise comparison depends solely on the relative scores of the two items involved. Recent works have made significant progress towards characterizing the performance (e.g. the mean square error for estimating the scores) of several classical methods, including the spectral method and the maximum likelihood estimator (MLE). However, where they stand regarding top- K ranking remains unsettled.

We demonstrate that under a natural random sampling model, the spectral method alone, or the regularized MLE alone, is minimax optimal in terms of the sample complexity — the number of paired comparisons needed to ensure exact top- K identification, for the fixed dynamic range regime. This is accomplished via optimal control of the entrywise error of the score estimates. We complement our theoretical studies by numerical experiments, confirming that both methods yield

* Author names are sorted alphabetically.

[†]Princeton University

[‡]Fudan University

MSC 2010 subject classifications: Primary 62F07; secondary 62B10

SUPPLEMENTARY MATERIAL

Supplement: Additional Proofs

(<http://www.e-publications.org/ims/support/download/imsart-ims.zip>). Additional proofs of the results in the paper can be found in the Supplementary Material.

low entrywise errors for estimating the underlying scores. Our theory is established via a novel leave-one-out trick, which proves effective for analyzing both iterative and non-iterative procedures. Along the way, we derive an elementary eigenvector perturbation bound for probability transition matrices, which parallels the Davis-Kahan Θ theorem for symmetric matrices. This also allows us to close the gap between the ℓ_2 error upper bound for the spectral method and the minimax lower limit.

Keywords

top- K ranking; pairwise comparisons; spectral method; regularized MLE; entrywise perturbation; leave-one-out analysis; reversible Markov chains

1. Introduction.

Imagine we have a large collection of n items, and we are given partially revealed comparisons between pairs of items. These paired comparisons are collected in a non-adaptive fashion, and could be highly noisy and incomplete. The aim is to aggregate these partial preferences so as to identify the K items that receive the highest ranks. This problem, which is called *top- K rank aggregation*, finds applications in numerous contexts, including web search (Dwork et al., 2001), recommendation systems (Baltrunas, Makcinskas and Ricci, 2010), sports competition (Masse, 1997), to name just a few. The challenge is both statistical and computational: how can one achieve reliable top- K ranking from a minimal number of pairwise comparisons, while retaining computational efficiency?

1.1. Popular approaches.

To address the aforementioned challenge, many prior approaches have been put forward based on certain statistical models. Arguably one of the most widely used parametric models is the Bradley-Terry-Luce (BTL) model (Bradley and Terry, 1952; Luce, 1959), which assigns a latent preference score $\{\omega_i^*\}_{1 \leq i \leq n}$ to each of the n items. The BTL model posits that: the chance of each item winning a paired comparison is determined by the relative scores of the two items involved, or more precisely,

$$\mathbb{P}\left\{\text{item } j \text{ is preferred over item } i\right\} = \frac{\omega_j^*}{\omega_i^* + \omega_j^*} \quad (1.1)$$

in each comparison of item i against item j . The items are repeatedly compared in pairs according to this parametric model. The task then boils down to identifying the K items with the highest preference scores, given these pairwise comparisons.

Among the ranking algorithms tailored to the BTL model, the following two procedures have received particular attention, both of which rank the items based on appropriate estimates of the latent preference scores.

1. **spectral method.** By connecting the winning probability in (1.1) with the transition probability of a reversible Markov chain, the spectral method attempts recovery of $\{\omega_i^*\}$ via the leading left eigenvector of a sample transition matrix. This procedure, also known as Rank Centrality (Negahban, Oh and Shah, 2017), bears similarity to the PageRank algorithm.
2. **maximum likelihood estimator (MLE).** This approach proceeds by finding the score assignment that maximizes the likelihood function (Ford, 1957). When parameterized appropriately, solving the MLE becomes a convex program, and hence is computationally feasible. There are also important variants of the MLE that enforce additional regularization.

Details are postponed to Section 2.2. In addition to their remarkable practical applicability, these two ranking paradigms are appealing in theory as well. For instance, both of them provably achieve intriguing ℓ_2 accuracy when estimating the latent preference scores (Negahban, Oh and Shah, 2017).

Nevertheless, the ℓ_2 error for estimating the latent scores merely serves as a “meta-metric” for the ranking task, which does not necessarily reveal the accuracy of top- K identification. In fact, given that the ℓ_2 loss only reflects the estimation error in some average sense, it is certainly possible that an algorithm obtains minimal ℓ_2 estimation loss but incurs (relatively) large errors when estimating the scores of the highest ranked items. Interestingly, a recent work Chen and Suh (2015) demonstrates that: a careful combination of the spectral method and the coordinate-wise MLE is optimal for top- K ranking. This leaves open the following natural questions: *where does the spectral alone, or the MLE alone, stand in top- K ranking? Are they capable of attaining exact top- K ranking? Are they capable of attaining exact top- K recovery from minimal samples?* These questions form the primary objectives of our study.

As we will elaborate later, the spectral method part of the preceding questions was recently explored by (Jang et al., 2016), for a regime where a relatively large fraction of item pairs have been compared. However, it remains unclear how well the spectral method can perform in a much broader — and often much more challenging — regime, where the fraction of item pairs being compared may be vanishingly small. Additionally, the ranking accuracy of the MLE (and its variants) remains unknown.

1.2. Main contributions.

The central focal point of the current paper is to assess the accuracy of both the spectral method and the regularized MLE in top- K identification. Assuming that the pairs of items being compared are randomly selected and that the preference scores fall within a *fixed* dynamic range, our paper delivers a somewhat surprising message:

Both the spectral method and the regularized MLE achieve perfect identification of top- K ranked items under optimal sample complexity (up to some constant factor)!

It is worth emphasizing that these two algorithms succeed even under the sparsest possible regime, a scenario where only an exceedingly small fraction of pairs of items have been

compared. This calls for precise control of the entrywise error — as opposed to the ℓ_2 loss — for estimating the scores. To this end, our theory is established upon a novel *leave-one-out argument*, which might shed light on how to analyze the entrywise error for more general optimization problems.

As a byproduct of the analysis, we derive an elementary eigenvector perturbation bound for (asymmetric) probability transition matrices, which parallels Davis-Kahan's sin Θ theorem for symmetric matrices. This simple perturbation bound immediately leads to an improved ℓ_2 error bound for the spectral method, which allows to close the gap between the theoretical performance of the spectral method and the minimax lower limit.

1.3. Notation.

Before proceeding, we introduce a few notations that will be useful throughout. To begin with, for any strictly positive probability vector $\pi \in \mathbb{R}^n$, we define the inner product space indexed by π as a vector space in $\mathbb{R}^n$ endowed with the inner product $\langle \mathbf{x}, \mathbf{y} \rangle_\pi = \sum_{i=1}^n \pi_i x_i y_i$.

The corresponding vector norm and the induced matrix norm are defined respectively as

$$\|\mathbf{x}\|_\pi = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle_\pi} \text{ and } \|A\|_\pi = \sup_{\|\mathbf{x}\|_\pi = 1} \|\mathbf{x}^\top A\|_\pi.$$

Additionally, the notation $f(n) = O(g(n))$ or $f(n) \lesssim g(n)$ means there is a constant $c > 0$ such that $|f(n)| \leq c|g(n)|$, $f(n) = \Omega(g(n))$ or $f(n) \gtrsim g(n)$ means there is a constant $c > 0$ such that $|f(n)| \geq c|g(n)|$, $f(n) = \Theta(g(n))$ or $f(n) \asymp g(n)$ means that there exist constants $c_1, c_2 > 0$ such that $c_1|g(n)| \leq |f(n)| \leq c_2|g(n)|$, and $f(n) = o(g(n))$ means $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0$.

Given a graph $\mathcal{G}$ with vertex set $\{1, 2, \dots, n\}$ and edge set $\mathcal{E}$, we denote by

$$L_{\mathcal{G}} = \sum_{(i,j) \in \mathcal{E}, i > j} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top \text{ the (unnormalized) Laplacian matrix (Chung, 1997)}$$

associated with it, where $\{\mathbf{e}_i\}_{1 \leq i \leq n}$ are the standard basis vectors in $\mathbb{R}^n$. For a matrix

$A \in \mathbb{R}^{n \times n}$ with n real eigenvalues, we let $\lambda_1(A) \geq \lambda_2(A) \geq \dots \geq \lambda_n(A)$ be the eigenvalues sorted in descending order.

2. Statistical models and main results.

2.1. Problem setup.

We begin with a formal introduction of the Bradley-Terry-Luce parametric model for binary comparisons.

Preference scores.—As introduced earlier, we assume the existence of a positive latent score vector

$$\boldsymbol{\omega}^* = [\omega_1^*, \dots, \omega_n^*]^\top \quad (2.1)$$

that comprises the underlying preference scores $\{\omega_i^* > 0\}_{1 \leq i \leq n}$ assigned to each of the n items. Alternatively, it is sometimes more convenient to reparameterize the score vector by

$$\boldsymbol{\theta}^* = [\theta_1^*, \dots, \theta_n^*]^\top, \text{ where } \theta_i^* = \log \omega_i^*. \quad (2.2)$$

These scores are assumed to fall within a *dynamic range* given by

$$\omega_i^* \in [\omega_{\min}, \omega_{\max}], \text{ or } \theta_i^* \in [\theta_{\min}, \theta_{\max}] \quad (2.3)$$

for all $1 \leq i \leq n$ and for some $\omega_{\min} > 0$, $\omega_{\max} > 0$, $\theta_{\min} = \log \omega_{\min}$, and $\theta_{\max} = \log \omega_{\max}$. We also introduce the *condition number* as

$$\kappa := \omega_{\max} / \omega_{\min}. \quad (2.4)$$

Notably, the current paper primarily focuses on the case with a *fixed* dynamic range (i.e. k is a fixed constant independent of n), although we will also discuss extensions to the large dynamic range regime in Section 3. Without loss of generality, it is assumed that

$$\omega_{\max} \geq \omega_1^* \geq \omega_2^* \geq \dots \geq \omega_n^* \geq \omega_{\min}, \quad (2.5)$$

meaning that items 1 through K are the desired top- K ranked items.

Comparison graph.—Let $\mathcal{G} = (\mathcal{V}, \varepsilon)$ stand for a comparison graph, where the vertex set $\mathcal{V} = \{1, 2, \dots, n\}$ represents the n items of interest. The items i and j are compared if and only if (i, j) falls within the edge set ε . Unless otherwise noted, we assume that $\mathcal{G}$ is drawn from the Erdős-Rényi random graph $\mathcal{G}_{n, p}$, such that an edge between any pair of vertices is present independently with some probability p . In words, p captures the fraction of item pairs being compared.

Pairwise comparisons.—For each $(i, j) \in \varepsilon$, we obtain L independent paired comparisons between items i and j . Let $y_{i,j}^{(l)}$ be the outcome of the l -th comparison, which is independently drawn as

$$y_{i,j}^{(l)} \stackrel{\text{ind.}}{=} \begin{cases} 1, & \text{with probability } \frac{\omega_j^*}{\omega_i^* + \omega_j^*} = \frac{e^{\theta_j^*}}{e^{\theta_i^*} + e^{\theta_j^*}}, \\ 0, & \text{else.} \end{cases} \quad (2.6)$$

By convention, we set $y_{i,j}^{(l)} = 1 - y_{j,i}^{(l)}$ for all $(i, j) \in \varepsilon$ throughout the paper. This is also known as the *logistic* pairwise comparison model, due to its strong resemblance to logistic regression. It is self-evident that the sufficient statistics under this model are given by

$$\mathbf{y} := \{y_{i,j} | (i, j) \in \varepsilon\}, \quad \text{where} \quad y_{i,j} := \frac{1}{L} \sum_{l=1}^L y_{i,j}^{(l)}. \quad (2.7)$$

To simplify the notation, we shall also take $y_{i,j}^* := \frac{\omega_j^*}{\omega_i^* + \omega_j^*} = \frac{e^{\theta_j^*}}{e^{\theta_i^*} + e^{\theta_j^*}}$.

Goal.—The goal is to identify the *set* of top- K ranked items — that is, the *set* of K items that enjoy the largest preference scores — from the pairwise comparison data $\mathbf{y}$.

2.2. Algorithms.

2.2.1. The spectral method: Rank centrality.—The spectral ranking algorithm, or Rank Centrality (Negahban, Oh and Shah, 2017), is motivated by the connection between the pairwise comparisons and a random walk over a directed graph. The algorithm starts by converting the pairwise comparison data $\mathbf{y}$ into a transition matrix $\mathbf{P} = [P_{i,j}]_{1 \leq i, j \leq n}$ in such a way that

$$P_{i,j} = \begin{cases} \frac{1}{d} y_{i,j} & \text{if } (i, j) \in \varepsilon, \\ 1 - \frac{1}{d} \sum_{k: (i,k) \in \varepsilon} y_{i,k} & \text{if } i = j, \\ 0, & \text{otherwise,} \end{cases} \quad (2.8)$$

for some given normalization factor $d > 0$, and then proceeds by computing the stationary distribution $\boldsymbol{\pi} \in \mathbb{R}^n$ of the Markov chain induced by $\mathbf{P}$. As we shall see later, the parameter d is taken to be on the same order of the maximum vertex degree of $\mathcal{G}$ while ensuring the non-negativity of $\mathbf{P}$. As asserted by Negahban, Oh and Shah (2017), $\boldsymbol{\pi}$ is a faithful estimate of $\boldsymbol{\omega}^*$ up to some global scaling. The algorithm is summarized in Algorithm 1.

To develop some intuition regarding why this spectral algorithm gives a reasonable estimate of $\boldsymbol{\omega}^*$, it is perhaps more convenient to look at the population transition matrix

$$\mathbf{P}^* = [P_{i,j}^*]_{1 \leq i, j \leq n} :$$

$$P_{i,j}^* = \begin{cases} \frac{1}{d} \frac{\omega_j^*}{\omega_i^* + \omega_j^*}, & \text{if } (i, j) \in \varepsilon, \\ 1 - \frac{1}{d} \sum_{k: (i,k) \in \varepsilon} \frac{\omega_k^*}{\omega_i^* + \omega_k^*}, & \text{if } i = j, \\ 0, & \text{otherwise,} \end{cases}$$

which coincides with $\mathbf{P}$ by taking $L \rightarrow \infty$. It can be seen that the normalized score vector

$$\boldsymbol{\pi}^* = \frac{1}{\sum_{i=1}^n \omega_i^*} [\omega_1^*, \omega_2^*, \dots, \omega_n^*]^\top \quad (2.9)$$

is the stationary distribution of the Markov chain induced by the transition matrix $\mathbf{P}^*$, since $\mathbf{P}^*$ and $\boldsymbol{\pi}^*$ are in detailed balance, namely,

$$\pi_i^* P_{i,j}^* = \pi_j^* P_{j,i}^*, \quad \forall (i, j). \quad (2.10)$$

As a result, one expects the stationary distribution of the sample version $\mathbf{P}$ to form a good estimate of $\boldsymbol{\omega}^*$, provided the sample size is sufficiently large.

2.2.2. The regularized MLE.—Under the BTL model, the negative log-likelihood function conditioned on $\mathcal{G}$ is given by (up to some global scaling)

$$\mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) = - \sum_{(i,j) \in \mathcal{E}, i > j} \left\{ y_{j,i} \log \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} + (1 - y_{j,i}) \log \frac{e^{\theta_j}}{e^{\theta_i} + e^{\theta_j}} \right\}$$

Algorithm 1

Spectral method (Rank Centrality).

Input the comparison graph $\mathcal{G}$, sufficient statistics $\mathbf{y}$, and the normalization factor d .

Define the probability transition matrix $\mathbf{P} = [P_{i,j}]_{1 \leq i, j \leq n}$ as in (2.8).

Compute the leading left eigenvector $\boldsymbol{\pi}$ of $\mathbf{P}$.

Output the K items that correspond to the K largest entries of $\boldsymbol{\pi}$.

$$= \sum_{(i,j) \in \mathcal{E}, i > j} \left\{ -y_{j,i} (\theta_i - \theta_j) + \log \left(1 + e^{\theta_i - \theta_j} \right) \right\}. \quad (2.11)$$

The regularized MLE then amounts to solving the following convex program

$$\underset{\boldsymbol{\theta} \in \mathbb{R}^n}{\text{minimize}} \quad \mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y}) = \mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) + \frac{1}{2} \lambda \|\boldsymbol{\theta}\|_2^2, \quad (2.12)$$

for a regularization parameter $\lambda > 0$. As will be discussed later, we shall adopt the choice $\lambda \asymp \sqrt{\frac{n p \log n}{L}}$ throughout this paper. For the sake of brevity, we let $\boldsymbol{\theta}$ represent the resulting penalized maximum likelihood estimate whenever it is clear from the context. Similar to the spectral method, one reports the K items associated with the K largest entries of $\boldsymbol{\theta}$.

2.3. Main results.

The most challenging part of top- K ranking is to distinguish the K -th and the $(K+1)$ -th items. In fact, the score difference of these two items captures the distance between the item sets $\{1, \dots, K\}$ and $\{K+1, \dots, n\}$. Unless their latent scores are sufficiently separated, the finite-sample nature of the model would make it infeasible to distinguish these two critical items. With this consideration in mind, we define the following separation measure

$$\Delta_K := \frac{\omega_K^* - \omega_{K+1}^*}{\omega_{\max}}. \quad (2.13)$$

This metric turns out to play a crucial role in determining the minimal sample complexity for perfect top- K identification.

The main finding of this paper concerns the optimality of both the spectral method and the regularized MLE in the presence of a fixed dynamic range (i.e. $\kappa = O(1)$). Recall that under the BTL model, the total number N of samples we collect concentrates sharply around its mean, namely,

$$N = (1 + o(1)) \mathbb{E}[N] = (1 + o(1)) n^2 p L / 2 \quad (2.14)$$

occurs with high probability. Our main result is stated in terms of the sample complexity required for exact top- K identification.

THEOREM 2.1. *Consider the pairwise comparison model specified in section 2.1 with $\kappa = O(1)$.*

Suppose that $p > \frac{c_0 \log n}{n}$ and that

$$\frac{n^2 p L}{2} \geq \frac{c_1 n \log n}{\Delta_K^2} \quad (2.15)$$

for some sufficiently large positive constants c_0 and c_1 . Further assume $L \leq c_2 \cdot n^{c_3}$ for any absolute constants $c_2, c_3 > 0$. With probability exceeding $1 - O(n^{-5})$, the set of top- K ranked items can be recovered exactly by the spectral method given in Algorithm 1, and by the regularized MLE given in (2.12). Here, we take $d = c_d n p$ in the spectral method and $\lambda = c_\lambda \sqrt{\frac{n p \log n}{L}}$ in the regularized MLE, where $c_d \geq 2$ and $c_\lambda > 0$ are some absolute constants.

Remark 2.2. We emphasize that $p \geq \frac{c_0 \log n}{n}$ for $c_0 \geq 1$ is a fundamental requirement for the ranking task. In fact, if $p < (1 - \epsilon) \frac{\log n}{n}$ for any constant $\epsilon > 0$, then the comparison graph

$\mathcal{G} \sim \mathcal{G}_{n,p}$ is disconnected with high probability. This means that there exists at least one isolated item (which has not been compared with any other item) and cannot be ranked.

Remark 2.3. In fact, the assumption that $L \leq c_2 \cdot n^{c_3}$ for any absolute constants $c_2, c_3 > 0$ is not needed for the spectral method.

Remark 2.4. Here, we assume the same number of comparisons L to simplify the presentation as well as the proof. The result still holds true if we have distinct $L_{i,j}$'s for each $i \neq j$, as long as $n^2 p \min_i \neq_j L_{i,j} \gtrsim \frac{n \log n}{\Delta_K^2}$.

Theorem 2.1 asserts that both the spectral method and the regularized MLE achieve a sample complexity on the order of $\frac{n \log n}{\Delta_K^2}$. Encouragingly, this sample complexity coincides with the minimax limit identified in (Chen and Suh, 2015, Theorem 2) in the fixed dynamic range, i.e. $\epsilon := \mathcal{O}(1)$.

THEOREM 2.5 (Chen and Suh (2015)). Fix $\epsilon \in (0, \frac{1}{2})$, and suppose that

$$n^2 p L \leq 2c_2 \frac{(1-\epsilon)n \log n - 2}{\Delta_K^2}, \quad (2.16)$$

where $c_2 = \omega_{\min}^4 / (4\omega_{\max}^4)$. Then for any ranking procedure ψ , one can find a score vector ω^* with separation Δ_K such that ψ fails to retrieve the top- K items with probability at least ϵ .

We are now positioned to compare our results with Jang et al. (2016), which also investigates the accuracy of the spectral method for top- K ranking. Specifically, Theorem 3 in Jang et al. (2016) establishes the optimality of the spectral method for the relatively dense regime where $p \gtrsim \sqrt{\frac{\log n}{n}}$. In this regime, however, the total sample size necessarily exceeds

$$n^2 p L / 2 \geq n^2 p / 2 \gtrsim \sqrt{n^3 \log n}, \quad (2.17)$$

which rules out the possibility of achieving minimal sample complexity if Δ_K is sufficiently large. For instance, consider the case where $\Delta_K \asymp 1$, then the optimal sample size — as revealed by Theorem 2.1 or (Chen and Suh, 2015, Theorem 1) — is on the order of

$$(n \log n) / \Delta_K^2 \asymp n \log n,$$

which is a factor of $\sqrt{\frac{n}{\log n}}$ lower than the bound in (2.17). By contrast, our results hold all the way down to the sparsest possible regime where $p \asymp \frac{\log n}{n}$, confirming the optimality of the

spectral method even for the most challenging scenario. Furthermore, we establish that the regularized MLE shares the same optimality guarantee as the spectral method, which was previously out of reach.

2.4. Optimal control of entrywise estimation errors.

In order to establish the ranking accuracy as asserted by Theorem 2.1, the key is to obtain precise control of the ℓ_∞ loss of the score estimates. Our results are as follows.

THEOREM 2.6 (Entrywise error of the spectral method). *Consider the pairwise comparison model in Section 2.1 with $\kappa = O(1)$. Suppose $p > \frac{c_0 \log n}{n}$ for some sufficiently large constant $c_0 > 0$. Choose $d = c_d np$ for some constant $c_d \geq 2$ in Algorithm 1. Then the spectral estimate $\boldsymbol{\pi}$ satisfies*

$$\frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty} \lesssim \sqrt{\frac{\log n}{npL}} \quad (2.18)$$

with probability $1 - O(n^{-5})$, where $\boldsymbol{\pi}^*$ is the normalized score vector (cf. (2.9)).

THEOREM 2.7 (Entrywise error of the regularized MLE). *Consider the pairwise comparison model specified in Section 2.1 with $\kappa = O(1)$. Suppose that $p \geq \frac{c_0 \log n}{n}$ for some sufficiently large constant $c_0 > 0$ and that $L \leq c_2 \cdot n^{c_3}$ for any absolute constants $c_2, c_3 > 0$. Set the regularization parameter to be $\lambda = c_\lambda \sqrt{\frac{np \log n}{L}}$ for some absolute constant $c_\lambda > 0$. Then the regularized MLE $\boldsymbol{\theta}$ satisfies*

$$\frac{\|e^{\boldsymbol{\theta}} - e^{\boldsymbol{\theta}^*} - \bar{\boldsymbol{\theta}}^* \mathbf{1}\|_\infty}{\|e^{\boldsymbol{\theta}^*} - \bar{\boldsymbol{\theta}}^* \mathbf{1}\|_\infty} \lesssim \sqrt{\frac{\log n}{npL}}$$

with probability exceeding $1 - O(n^{-5})$, where $\bar{\boldsymbol{\theta}}^* := \frac{1}{n} \mathbf{1}^\top \boldsymbol{\theta}^*$ and $e^{\boldsymbol{\theta}} := \left[e^{\theta_1}, \dots, e^{\theta_n} \right]^\top$.

Theorems 2.6–2.7 indicate that if the number of comparisons associated with each item — which concentrates around npL — exceeds the order of $\log n$, then both methods are able to achieve a small ℓ_∞ error when estimating the scores.

Recall that the ℓ_2 estimation error of the spectral method has been characterized by Negahban, Oh and Shah (2017) (or Theorem 5.2 of this paper that improves it by removing the logarithmic factor), which obeys

$$\frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_2}{\|\boldsymbol{\pi}^*\|_2} \lesssim \sqrt{\frac{\log n}{npL}} \quad (2.19)$$

with high probability. Similar theoretical guarantees have been derived for another variant of the MLE (the constrained version) under a uniform sampling model as well (Negahban, Oh and Shah, 2017). In comparison, our results indicate that the estimation errors for both algorithms are almost evenly spread out across all coordinates rather than being localized or clustered. Notably, the pointwise errors revealed by Theorems 2.6–2.7 immediately lead to exact top- K identification as claimed by Theorem 2.1.

Proof of Theorem 2.1. In what follows, we prove the theorem for the spectral method part. The regularized MLE part follows from an almost identical argument and hence is omitted.

Since the spectral algorithm ranks the items in accordance with the score estimate $\boldsymbol{\pi}$, it suffices to demonstrate that

$$\pi_i - \pi_j > 0, \quad \forall 1 \leq i \leq K, K+1 \leq j \leq n$$

To this end, we first apply the triangle inequality to get

$$\frac{\pi_i - \pi_j}{\|\boldsymbol{\pi}^*\|_\infty} \geq \frac{\pi_i^* - \pi_j^*}{\|\boldsymbol{\pi}^*\|_\infty} - \frac{|\pi_i - \pi_i^*|}{\|\boldsymbol{\pi}^*\|_\infty} - \frac{|\pi_j - \pi_j^*|}{\|\boldsymbol{\pi}^*\|_\infty} \geq \Delta_K - \frac{2\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty}. \quad (2.20)$$

In addition, it follows from Theorem 2.6 as well as our sample complexity assumption that

$$\frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty} \lesssim \sqrt{\frac{\log n}{npL}} \text{ and } n^2 pL \gtrsim \frac{n \log n}{\Delta_K^2}.$$

These conditions taken collectively imply that $\frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty} < \frac{1}{2}\Delta_K$ as long as $\frac{npL\Delta_K^2}{\log n}$ exceeds some sufficiently large constant. Substitution into (2.20) reveals that $\pi_i - \pi_j > 0$, as claimed. $\square$

2.5. Heuristic arguments.

We pause to develop some heuristic explanation as to why the estimation errors are expected to be spread out across all entries. For simplicity, we focus on the case where $p = 1$ and L is sufficiently large, so that $\mathbf{y}$ and $\mathbf{P}$ sharply concentrate around $\mathbf{y}^*$ and $\mathbf{P}^*$, respectively.

We begin with the spectral algorithm. Since $\boldsymbol{\pi}$ and $\boldsymbol{\pi}^*$ are respectively the invariant distributions of the Markov chains induced by $\mathbf{P}$ and $\mathbf{P}^*$, we can decompose

$$(\boldsymbol{\pi} - \boldsymbol{\pi}^*)^\top = \boldsymbol{\pi}^\top \mathbf{P} - \boldsymbol{\pi}^{*\top} \mathbf{P}^* = (\boldsymbol{\pi} - \boldsymbol{\pi}^*)^\top \mathbf{P} + \underbrace{\boldsymbol{\pi}^{*\top} (\mathbf{P} - \mathbf{P}^*)}_{:= \boldsymbol{\xi}} \quad (2.21)$$

When $p = 1$ and $\frac{\omega_{\max}}{\omega_{\min}} \asymp 1$, the entries of $\boldsymbol{\pi}^*$ (resp. the off-diagonal entries of $\mathbf{P}^*$ and $\mathbf{P} - \mathbf{P}^*$) are all of the same order and, as a result, the energy of the uncertainty term $\boldsymbol{\xi}$ is spread out (using standard concentration inequalities). In fact, we will demonstrate in Section 5.2 that

$$\frac{\|\boldsymbol{\xi}\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty} \lesssim \sqrt{\frac{\log n}{npL}} \asymp \frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_2 \sqrt{\log n}}{\|\boldsymbol{\pi}^*\|_2}, \quad (2.22)$$

which coincides with the optimal rate. Further, if we look at each entry of (2.21), then for all $1 \leq m \leq n$,

$$\pi_m - \pi_m^* = [(\boldsymbol{\pi} - \boldsymbol{\pi}^*)^\top \mathbf{P}]_m + \xi_m, \quad (2.23)$$

$$= \underbrace{(\pi_m - \pi_m^*) P_{m,m}}_{\text{contraction}} + \underbrace{[P_{1,m}, \dots, P_{m-1,m}, 0, P_{m+1,m}, \dots, P_{n,m}]}_{\text{error averaging}} \begin{bmatrix} \pi_1 - \pi_1^* \\ \vdots \\ \pi_n - \pi_n^* \end{bmatrix} + \xi_m.$$

By construction of the transition matrix, one can easily verify that $P_{m,m}$ is bounded away from 1 and $P_{j,m} \asymp \frac{1}{n}$ for all $j \neq m$. As a consequence, the identity $\boldsymbol{\pi}^\top \mathbf{P} = \boldsymbol{\pi}^\top$ allows one to treat each $\pi_m - \pi_m^*$ as a mixture of three effects: (i) the first term of (2.23) behaves as an entrywise contraction of the error; (ii) the second term of (2.23) is a (nearly uniformly weighted) average of the errors over all coordinates, which can essentially be treated as a *smoothing* operator applied to the error components; and (iii) the uncertainty term ξ_m .

Rearranging terms in (2.23), we are left with

$$(1 - P_{m,m})|\pi_m - \pi_m^*| \lesssim \frac{1}{n} \sum_{i=1}^n |\pi_i - \pi_i^*| + \xi_m, \quad \forall m \quad (2.24)$$

which further gives,

$$\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty \lesssim \frac{1}{n} \sum_{i=1}^n |\pi_i - \pi_i^*| + \|\boldsymbol{\xi}\|_\infty. \quad (2.25)$$

There are two possibilities compatible with this bound (2.25): (1)

$\|\pi - \pi^*\|_\infty \lesssim \frac{1}{n} \sum_{i=1}^n |\pi_i - \pi_i^*|$, and (2) $\|\pi - \pi^*\|_\infty \lesssim \|\xi\|_\infty \lesssim \frac{\|\pi - \pi^*\|_2}{\|\pi^*\|_2} \|\pi^*\|_\infty$ by (2.22). In

either case, the errors are fairly delocalized, revealing that

$$\frac{\|\pi - \pi^*\|_\infty}{\|\pi^*\|_\infty} \lesssim \left\{ \frac{1}{n} \frac{\|\pi - \pi^*\|_1}{\|\pi^*\|_\infty}, \frac{\|\pi - \pi^*\|_2 \sqrt{\log n}}{\|\pi^*\|_2} \right\}.$$

We now move on to the regularized MLE, following a very similar argument. By the optimality condition that $\nabla \mathcal{L}_\lambda(\theta) = 0$, one can derive (for some η to be specified later)

$$\begin{aligned} \theta - \theta^* &= \theta - \eta \nabla \mathcal{L}_\lambda(\theta) - \theta^* \\ &= \theta - \eta \nabla \mathcal{L}_\lambda(\theta) - (\theta^* - \eta \nabla \mathcal{L}_\lambda(\theta^*)) - \underbrace{\eta \nabla \mathcal{L}_\lambda(\theta^*)}_{:= \zeta} \\ &\approx (I - \eta \nabla^2 \mathcal{L}_\lambda(\theta^*))(\theta - \theta^*) - \zeta \end{aligned}$$

Write $\nabla^2 \mathcal{L}_\lambda(\theta^*) = D - A$, where D and A denote respectively the diagonal and off-diagonal parts of $\nabla^2 \mathcal{L}_\lambda(\theta^*)$. Under our assumptions, one can check that $D_{m,m} \asymp n$ for all $1 \leq m \leq n$ and $A_{j,m} \asymp 1$ for any $j \neq m$. With these notations in place, one can write the entrywise error as follows

$$\theta_m - \theta_m^* = (1 - \eta D_{m,m})(\theta_m - \theta_m^*) + \sum_{j: j \neq m} \eta A_{j,m}(\theta_j - \theta_j^*) - \zeta_m.$$

By choosing $\eta = c_2/n$ for some sufficiently small constant $c_2 > 0$, we get $1 - \eta D_{m,m} < 1$ and $\eta A_{j,m} \asymp 1/n$. Therefore, the right-hand side of the above relation also comprises a contraction term as well as an error smoothing term, similar to (2.23). Carrying out the same argument as for the spectral method, we see that the estimation errors of the regularized MLE are expected to be spread out.

2.6. Numerical experiments.

It is worth noting that extensive numerical experiments on both synthetic and real data have already been conducted in Negahban, Oh and Shah (2017) to confirm the practicability of both the spectral method and the regularized MLE. See also Chen and Suh (2015) for the experiments on the Spectral-MLE algorithm. This section provides some additional simulations to complement their experimental results as well as our theory. Throughout the experiments, we set the number of items n to be 200, while the number of repeated

comparisons L and the edge probability p can vary with the experiments. Regarding the tuning parameters, we choose $d = 2d_{\max}$ in the spectral method where $d_{\max}$ is the maximum degree of the graph and $\lambda = 2\sqrt{\frac{np \log n}{L}}$ in the regularized MLE, which are consistent with the configurations considered in the main theorems. Additionally, we also display the experimental results for the unregularized MLE, i.e. $\lambda = 0$. All of the results are averaged over 100 Monte Carlo simulations.

We first investigate the ℓ_∞ error of the spectral method and the (regularized) MLE when estimating the preference scores. To this end, we generate the latent scores ω_i^* ($1 \leq i \leq n$) independently and uniformly at random over the interval $[0.5, 1]$. Figure 1(a) (resp. Figure 1(b)) displays the entry wise error in the spectral score estimation as the number of repeated comparisons L (resp. the edge probability p) varies. As is seen from the plots, the ℓ_∞ error of all methods gets smaller as p and L increase, confirming our results in Theorems 2.6–2.7. Next, we show in Figure 1(c) the relative ℓ_∞ error while fixing the total number of samples (i.e. $n^2 p L$). It can be seen that the performance almost does not change if the sample complexity $n^2 p L$ remains the same. It is also interesting to see that the ℓ_∞ error of the spectral method and the MLE are very similar. In addition, Figure 2 illustrates the relative ℓ_∞ error and the relative ℓ_2 error in score estimation for all three methods. As we can see, the relative ℓ_∞ errors are not much larger than the relative ℓ_2 errors (recall that $n = 200$), thus offering empirical evidence that the errors in the score estimates are spread out across all entries.

Further, we examine the top- K ranking accuracy of all three methods. Here, we fix $p = 0.25$ and $L = 20$, set $K = 10$, and let $\omega_i^* = 1$ for all $1 \leq i \leq K$ and $\omega_j^* = 1 - \Delta$ for all $K + 1 \leq j \leq n$.

By construction, the score separation satisfies $\kappa = \frac{1}{\Delta}$. Figure 3 illustrates the accuracy in identifying the top- K ranked items. The performance of them improves when the score separation becomes larger, which matches our theory in Theorem 2.1.

2.7. Other related works.

The problem of ranking based on partial preferences has received much attention during the past decade. Two types of observation models have been considered: (1) *the cardinal-based model*, where users provide explicit numerical ratings of the items, (2) *the ordinal-based model*, where users are asked to make comparative measurements. See Ammar and Shah (2011) for detailed comparisons between them.

In terms of the ordinal-based model — and in particular, ranking from pairwise comparisons — both parametric and nonparametric models have been extensively studied. For example, Hunter (2004) examined variants of the parametric BTL model, and established the convergence properties of the minorization-maximization algorithm for computing the MLE. Moreover, the BTL model falls under the category of *low-rank parametric models*, since the preference matrix is generated by passing a rank-2 matrix through the logistic link function (Rajkumar and Agarwal, 2016). Additionally, the work Jiang et al. (2011) proposed a least-

squares type method to estimate the full ranking, which generalizes the simple Borda count algorithm (Ammar and Shah, 2011). For many of these algorithms, the sample complexities needed for perfect total ranking were determined by Rajkumar and Agarwal (2014), although the top- K ranking accuracy was not considered there.

Going beyond the parametric models, a recent line of works Shah et al. (2017); Shah and Wainwright (2015); Chen et al. (2017); Pananjady et al. (2017) considered the nonparametric *stochastically transitive model*, where the only model assumption is that the comparison probability matrix follows certain transitivity rules. This type of models subsumes the BTL model as a special case. For instance, Shah and Wainwright (2015) suggested a simple counting-based algorithm which can reliably recover the top- K ranked items for various models. However, the sampling paradigm considered therein is quite different from ours in the sparse regime; for instance, their model does not come close to the setting where p is small but L is large, which is the most challenging regime of the model adopted in our paper and Negahban, Oh and Shah (2017); Chen and Suh (2015).

All of the aforementioned papers concentrate on the case where there is a single ground-truth ordering. It would also be interesting to investigate the scenarios where different users might have different preference scores. To this end, Negahban et al. (2017); Lu and Negahban (2014) imposed the low-rank structure on the underlying preference matrix and adopted the nuclear-norm relaxation approach to recover the users' preferences. Additionally, several papers explored the ranking problem for the more general Plackett-Luce model (Hajek, Oh and Xu, 2014; Soufiani et al., 2013), in the presence of adaptive sampling (Jamieson and Nowak, 2011; Busa-Fekete et al., 2013; Heckel et al., 2016; Agarwal et al., 2017), for the crowdsourcing scenario (Chen et al., 2013), and in the adversarial setting (Suh, Tan and Zhao, 2017). These are beyond the scope of the present paper.

Speaking of the error metric, the ℓ_∞ norm is appropriate for top- K ranking problem and other learning problems as well. In particular, ℓ_∞ perturbation bounds for eigenvectors of symmetric matrices (Koltchinskii and Lounici, 2016; Fan, Wang and Zhong, 2018; Eldridge, Belkin and Wang, 2017; Abbe et al., 2017) and singular vectors of general matrices (Koltchinskii and Xia, 2016) have been studied. In stark contrast, we study the ℓ_∞ norm errors of the leading eigenvector of a class of asymmetric matrices (probability transition matrix) and the regularized MLE. Furthermore, most existing results require the expectations of data matrices to have low rank, at least approximately. We do not impose such assumptions.

When it comes to the technical tools, it is worth noting that the leave-one-out idea has been invoked to analyze random designs for other high-dimensional problems, e.g. robust M-estimators (El Karoui, 2017), confidence intervals for Lasso (Javanmard and Montanari, 2015), likelihood ratio test (Sur, Chen and Candes, 2017), and nonconvex statistical learning (Ma et al., 2017; Chen et al., 2018). In particular, Zhong and Boumal (2017) and Abbe et al. (2017) use it to precisely characterize entrywise behavior of eigenvectors of a large class of symmetric random matrices, which improves upon prior ℓ_∞ eigenvector analysis.

Consequently, they are able to show the sharpness of spectral methods in many popular models. Our introduction of leave-one-out auxiliary quantities is similar in spirit to these papers.

Finally, the family of spectral methods has been successfully applied in numerous applications, e.g. matrix completion (Keshavan, Montanari and Oh, 2010), phase retrieval (Chen and Candès, 2017), graph clustering (Rohe, Chatterjee and Yu, 2011; Abbe et al., 2017), joint alignment (Chen and Candès, 2016). All of them are designed based on the eigenvectors of some symmetric matrix, or the singular vectors if the matrix of interest is asymmetric. Our paper contributes to this growing literature by establishing a sharp *eigenvector* perturbation analysis framework for an important class of *asymmetric* matrices — the probability transition matrices.

3. Extension: general dynamic range.

All of the preceding results concern the regime with a fixed dynamic range (i.e. $\kappa = O(1)$). This section moves on to discussing the case with large κ .

To start with, by going through the same proof technique, we can readily obtain — in the general κ setting — the following performance guarantees for both the spectral estimate $\boldsymbol{\pi}$ and the regularized MLE $\boldsymbol{\theta}$.

THEOREM 3.1. *Consider the pairwise comparison model in Section 2.1. Suppose that*

$p > \frac{c_0 \kappa^5 \log n}{n}$ for some sufficiently large constant $c_0 > 0$, and choose $d = c_d np$ for some constant $c_d \geq 2$ in Algorithm 1. Then with probability exceeding $1 - O(n^{-5})$,

1. *the spectral estimate $\boldsymbol{\pi}$ satisfies*

$$\frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty} \lesssim K \sqrt{\frac{\log n}{npL}},$$

where $\boldsymbol{\pi}^$ is the normalized score vector as defined in (2.9).*

2. *the set of top- K ranked items can be recovered exactly by the spectral method given in Algorithm 1, as long as*

$$\frac{n^2 p L}{2} \geq c_1 \frac{\kappa^2 n \log n}{\Delta_K^2}$$

for some sufficiently large constant $c_1 > 0$.

THEOREM 3.2. *Consider the pairwise comparison model in Section 2.1. Suppose that*

$p \geq \frac{c_0 \kappa^4 \log n}{n}$ for some sufficiently large constant $c_0 > 0$ and that $L \leq c_2 \cdot n^{c_3}$ for any absolute

constants $c_2, c_3 > 0$. Set the regularization parameter to be $\lambda = c\lambda \frac{1}{\log \kappa} \sqrt{\frac{np \log n}{L}}$ for some absolute constant $c_\lambda > 0$. Then with probability exceeding $1 - O(n^{-5})$,

1. the regularized MLE $\boldsymbol{\theta}$ satisfies

$$\frac{\|e^{\boldsymbol{\theta}} - e^{\boldsymbol{\theta}^*} - \bar{\boldsymbol{\theta}}^* \mathbf{1}\|_\infty}{\|e^{\boldsymbol{\theta}^*} - \bar{\boldsymbol{\theta}}^* \mathbf{1}\|_\infty} \lesssim \kappa^2 \sqrt{\frac{\log n}{npL}}$$

where $\bar{\boldsymbol{\theta}}^* := \frac{1}{n} \mathbf{1}^\top \boldsymbol{\theta}^*$ and $e^{\boldsymbol{\theta}} := \begin{bmatrix} e^{\theta_1}, \dots, e^{\theta_n} \end{bmatrix}^\top$.

2. the set of top- K ranked items can be recovered exactly by the regularized MLE given in (2.12), as long as

$$\frac{n^2 pL}{2} \geq c_1 \frac{\kappa^4 n \log n}{\Delta_K^2}$$

for some sufficiently large constant $c_1 > 0$.

Remark 3.3. The guarantees on exact top- K recovery for both the spectral method and the regularized MLE are immediate consequences of their ℓ_∞ error bound, as we have argued in Section 2.4. Hence we will focus on proving the ℓ_∞ error bound in Sections 5–6.

Notably, the achievability bounds for top- K ranking in Theorems 3.1–3.2 do not match the lower bound asserted in Theorem 2.5 in terms of κ . This is partly because the separation measure Δ_K fails to capture the information bottleneck for the general κ setting. In light of this, we introduce the following new measure that seems to be a more suitable metric to reflect the hardness of the top- K ranking problem:

$$\Delta_K^* := \frac{\omega_K^* - \omega_{K+1}^*}{\omega_{K+1}^*} \cdot \sqrt{\frac{1}{n} \sum_{i=1}^n \frac{\omega_{K+1}^* \omega_i^*}{(\omega_K^* + \omega_i^*)^2}}, \quad (3.1)$$

which will be termed the *generalized separation measure*. Informally, $(\Delta_K^*)^2$ is a reasonably tight upper bound on certain normalized KL divergence metric (see the proof of Theorem 3.4 for details in the Supplemental Material). With this metric in place, we derive another lower bound as follows.

THEOREM 3.4. Fix $\epsilon \in (0, \frac{1}{2})$, and let $\mathcal{G} \sim \mathcal{G}_{n,p}$. Consider any preference score vector $\boldsymbol{\omega}^*$, and let Δ_K^* denote its generalized separation. If

$$n^2 pL \leq \frac{\epsilon^2}{2} \frac{n}{(\Delta_K^*)^2},$$

then there exists another preference score vector $\tilde{\omega}$ with the same generalized separation Δ_K^* and different top- K items such that $P_e(\psi) \geq \frac{1-\epsilon}{2}$ for any ranking scheme ψ . Here, $P_e(\psi)$ represents the probability of error in distinguishing these two vectors given $\mathbf{y}$.

The preceding sample complexity lower bound scales inversely proportionally to $(\Delta_K^*)^2$. To see why this generalized measure may be more suitable compared to the original separation metric, we single out three examples in the Supplementary Material. Unfortunately, our current analyses do not yield a matching upper bound with respect to Δ_K^* unless κ is a constant. For instance, the analysis of the spectral method relies on the eigenvector perturbation bound (Theorem 5.1), where the spectral gap and matrix perturbation play a crucial role. However, the current results for controlling these quantities have explicit dependency on κ (Negahban, Oh and Shah (2017)). It is not clear whether we could incorporate the new measure to eliminate such dependency on κ . This calls for more refined analysis techniques, which we leave for future investigation.

Moreover, it is not obvious whether the spectral method alone or the regularized MLE alone can achieve the minimal sample complexity in the general κ regime. It is possible that one needs to first screen out those items with extremely high or low scores using methods like Borda count (Ammar and Shah, 2012), as advocated by (Negahban, Oh and Shah, 2017; Chen and Suh, 2015; Jang et al., 2016). All in all, finding tight upper bounds for general κ remains an open question.

4. Discussion.

This paper justifies the optimality of both the spectral method and the regularized MLE for top- K rank aggregation for the fixed dynamic range case. Our theoretical studies are by no means exhaustive, and there are numerous directions that would be of interest for future investigations. We point out a few possibilities as follows.

General condition number κ .

As mentioned before, our current theory is optimal in the presence of a fixed dynamic range with $\kappa = O(1)$. We have also made a first attempt in considering the large κ regime. It is desirable to characterize the statistical and computational limits for more general κ .

Goodness-of-fit.

Throughout this paper, we have assumed the BTL model captures the randomness underlying the data we collect. A practical question is whether the real data actually follows the BTL model. It would be interesting to investigate how to test the goodness-of-fit of this model.

Unregularized MLE.

We have studied the optimality of the regularized MLE with the regularization parameter $\lambda \asymp \sqrt{\frac{np \log n}{L}}$. Our analysis relies on the regularization term to obtain convergence of the gradient descent algorithm (see Lemma 6.7). It is natural to ask whether such a regularization term is necessary or not. This question remains open.

More general comparison graphs.

So far we have focused on a tractable but somewhat restrictive comparison graph, namely, the Erdős-Rényi random graph. It would certainly be important to understand the performance of both methods under a broader family of comparison graphs, and to see which algorithms would enable optimal sample complexities under general sampling patterns.

Entrywise perturbation analysis for convex optimization.

This paper provides the ℓ_∞ perturbation analysis for the regularized MLE using the leave-one-out trick as well as an inductive argument along the algorithmic updates. We expect this analysis framework to carry over to a much broader family of convex optimization problems, which may in turn offer a powerful tool for showing the stability of optimization procedures in an entrywise fashion.

5. Analysis for the spectral method.

This section is devoted to proving Theorem 3.1 and hence Theorem 2.6, which characterizes the point-wise error of the spectral estimate.

5.1. Preliminaries.

Here, we gather some preliminary facts about reversible Markov chains as well as the Erdős-Rényi random graph.

The first important result concerns the eigenvector perturbation for probability transition matrices, which can be treated as the analogue of the celebrated Davis-Kahan $\sin \Theta$ theorem (Davis and Kahan, 1970). Due to its potential importance for other problems, we promote it to a theorem as follows.

THEOREM 5.1 (Eigenvector perturbation). *Suppose that P , $\hat{P}$, and P^* are probability transition matrices with stationary distributions π , $\hat{\pi}$, π^* , respectively. Also, assume that P^* represents a reversible Markov chain.*

When $\|P - \hat{P}\|_{\pi^} < 1 - \max\{\lambda_2(P^*), -\lambda_n(P^*)\}$, it holds that*

$$\|\pi - \hat{\pi}\|_{\pi^*} \leq \frac{\|\pi^\top (P - \hat{P})\|_{\pi^*}}{1 - \max\{\lambda_2(P^*), -\lambda_n(P^*)\} - \|P - \hat{P}\|_{\pi^*}}.$$

Several remarks regarding Theorem 5.1 are in order. First, in contrast to standard perturbation results like Davis-Kahan's $\sin \Theta$ theorem, our theorem involves three matrices in total, where $\mathbf{P}$, $\hat{\mathbf{P}}$, and $\mathbf{P}^*$ can all be arbitrary. For example, one may choose $\mathbf{P}^*$ to be the population transition matrix, and $\mathbf{P}$ and $\hat{\mathbf{P}}$ as two finite-sample versions associated with $\mathbf{P}^*$. Second, we only impose reversibility on $\mathbf{P}^*$, whereas $\mathbf{P}$ and $\hat{\mathbf{P}}$ need not induce reversible Markov Chains. Third, Theorem 5.1 allows one to derive the ℓ_2 estimation error in Negahban, Oh and Shah (2017) directly without resorting to the power method; in fact, our ℓ_2 estimation error bound improves upon Negahban, Oh and Shah (2017) by some logarithmic factor.

THEOREM 5.2. *Consider the pairwise comparison model specified in Section 2.1 with $\kappa = O(1)$. Suppose $p \geq c_0 \frac{\log n}{n}$ for some sufficiently large constant $c_0 > 0$ and $d \geq c_d np$ for $c_d \geq 2$ in Algorithm 1. With probability exceeding $1 - O(n^{-5})$, one has*

$$\frac{\|\pi - \pi^*\|_2}{\|\pi^*\|_2} \lesssim \frac{1}{\sqrt{npL}}.$$

Notably, Theorem 5.2 matches the minimax lower bound derived in (Negahban, Oh and Shah, 2017, Theorem 3). As far as we know, this is the first result that demonstrates the orderwise optimality of the spectral method when measured by the ℓ_2 loss.

The next result is concerned with the concentration of the vertex degrees in an Erdős-Rényi random graph.

Lemma 5.3 (Degree concentration). Suppose that $\mathcal{G} \sim \mathcal{G}_{n,p}$. Let d_i be the degree of node i , $d_{\min} = \min_{1 \leq i \leq n} d_i$ and $d_{\max} = \max_{1 \leq i \leq n} d_i$. If $p \geq \frac{c_0 \log n}{n}$ for some sufficiently large constant $c_0 > 0$, then the following event

$$\mathcal{A}_0 = \left\{ \frac{np}{2} \leq d_{\min} \leq d_{\max} \leq \frac{3np}{2} \right\} \quad (5.1)$$

obeys

$$\mathbb{P}(\mathcal{A}_0) \geq 1 - O(n^{-10}).$$

PROOF. The proof follows from the standard Chernoff bound and is hence omitted. $\square$

Since d is chosen to be $c_d np$ for some constant $c_d \geq 2$, we have, by Lemma 5.3, that the maximum vertex degree obeys $d_{\max} < d$ with high probability.

5.2. Proof outline of Theorem 3.1.

In this subsection, we outline the proof of Theorem 3.1.

Recall that $\boldsymbol{\pi} = [\pi_1, \dots, \pi_n]^\top$ and $\boldsymbol{\pi}^* = [\pi_1^*, \dots, \pi_n^*]^\top$ are the stationary distributions associated with $\mathbf{P}$ and $\mathbf{P}^*$, respectively. This gives

$$\boldsymbol{\pi}^\top \mathbf{P} = \boldsymbol{\pi}^\top \quad \text{and} \quad \boldsymbol{\pi}^{*\top} \mathbf{P}^* = \boldsymbol{\pi}^{*\top}$$

For each $1 \leq m \leq n$, one can decompose

$$\begin{aligned} \pi_m - \pi_m^* &= \boldsymbol{\pi}^\top \mathbf{P}_{\cdot m} - \boldsymbol{\pi}^{*\top} \mathbf{P}_{\cdot m}^* = \boldsymbol{\pi}^{*\top} (\mathbf{P}_{\cdot m} - \mathbf{P}_{\cdot m}^*) + (\boldsymbol{\pi} - \boldsymbol{\pi}^*)^\top \mathbf{P}_{\cdot m} \\ &= \underbrace{\sum_j \pi_j^* (P_{j,m} - P_{j,m}^*)}_{:= I_1^m} + \underbrace{(\pi_m - \pi_m^*) P_{m,m}}_{:= I_2^m} + \sum_{j: j \neq m} (\pi_j - \pi_j^*) P_{j,m}, \end{aligned}$$

where $\mathbf{P}_{\cdot m}$ (resp. $\mathbf{P}_{\cdot m}^*$) denotes the m -th column of $\mathbf{P}$ (resp. $\mathbf{P}^*$). Then it boils down to controlling I_1^m, I_2^m and $\sum_{j: j \neq m} (\pi_j - \pi_j^*) P_{j,m}$.

1. Since $\boldsymbol{\pi}^*$ is deterministic while $\mathbf{P}$ is random, we can easily control I_1^m using Hoeffding's inequality. The bound is the following.

Lemma 5.4. With probability exceeding $1 - O(n^{-5})$, one has

$$\max_m |I_1^m| \lesssim \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_\infty.$$

2. Next, we show the term I_2^m behaves as a contraction of $|\pi_m - \pi_m^*|$.

Lemma 5.5. With probability exceeding $1 - O(n^{-5})$, there exists some constant $c > 0$ such that for all $1 \leq m \leq n$,

$$|I_2^m| \leq \left(1 - \frac{\text{np}}{2(1+\kappa)d} + c\sqrt{\frac{\log n}{Ld}}\right) |\pi_m - \pi_m^*|.$$

3. The statistical dependency between $\boldsymbol{\pi}$ and $\mathbf{P}$ introduces difficulty in obtaining a sharp estimate of the third term $\sum_{j: j \neq m} (\pi_j - \pi_j^*) P_{j,m}$. Nevertheless, the leave-one-out technique helps us decouple the dependency and obtain effective control of this term. The key component of the analysis is the introduction of a new probability transition matrix $\mathbf{P}^{(m)}$, which is a leave-one-out version of the original matrix $\mathbf{P}$. More precisely, $\mathbf{P}^{(m)}$ replaces all of the transition probabilities

involving the m -th item with their expected values (unconditional on $\mathcal{G}$); that is, for any $i \neq j$,

$$P_{i,i}^{(m)} = \begin{cases} P_{i,j} & i \neq j, j \neq m \\ \frac{P_{i,j}}{d} y_{i,j}^* & i = m \text{ or } j = m \end{cases}$$

with $y_{i,j}^* = \frac{\omega_j^*}{\omega_j^* + \omega_i^*}$. For any $1 \leq i \leq n$, set

$$P_{i,i}^{(m)} = 1 - \sum_{j: j \neq i} P_{i,j}^{(m)}$$

in order to ensure that $\mathbf{P}^{(m)}$ is a probability transition matrix. In addition, we let $\boldsymbol{\pi}^{(m)}$ be the stationary distribution of the Markov chain induced by $\mathbf{P}^{(m)}$. As will be demonstrated later, the main advantages of introducing $\boldsymbol{\pi}^{(m)}$ are two-fold: (1) the original spectral estimate $\boldsymbol{\pi}$ is very well approximated by $\boldsymbol{\pi}^{(m)}$, and (2) $\boldsymbol{\pi}^{(m)}$ is statistically independent of the connectivity of the m -th node and the comparisons with regards to the m -th item. Now we further decompose $\sum_{j: j \neq m} (\pi_j - \pi_j^*) P_{j,m}$:

$$\sum_{j: j \neq m} (\pi_j - \pi_j^*) P_{j,m} = \underbrace{\sum_{j: j \neq m} (\pi_j - \pi_j^{(m)}) P_{j,m}}_{:= I_3^m} + \underbrace{\sum_{j: j \neq m} (\pi_j^{(m)} - \pi_j^*) P_{j,m}}_{:= I_4^m}.$$

4. For I_3^m , we apply the Cauchy-Schwarz inequality to obtain that with probability at least $1 - O(n^{-10})$,

$$|I_3^m| \leq \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 \left(\sum_{j: j \neq m} P_{j,m}^2 \right)^{\frac{1}{2}} \stackrel{(i)}{\leq} \frac{1}{\sqrt{d}} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2,$$

where (i) follows from the fact that $P_{j,m} \leq \frac{1}{d}$ for all $j \neq m$ and $d_{\max} \leq d$ on the event $\mathcal{A}_0$ (defined in Lemma 5.3). Consequently, it suffices to control the ℓ_2 difference between the original spectral estimate $\boldsymbol{\pi}$ and its leave-one-out version $\boldsymbol{\pi}^{(m)}$. This is accomplished in the following lemma.

Lemma 5.6. Suppose that $n\gamma^2 > c\kappa \log n$ for some sufficiently large constant $c > 0$. With probability at least $1 - O(n^{-5})$,

$$\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 \leq \frac{16\sqrt{\kappa}}{\gamma} \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_{\infty} + \|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_{\infty}, \quad (5.2)$$

where $\kappa = \omega_{\max}/\omega_{\min}$ and $\gamma = 1 - \max\{\lambda_2(\mathbf{P}^*), -\lambda_n(\mathbf{P}^*)\} - \|\mathbf{P} - \mathbf{P}^*\|_{\pi^*}$.

Using (Negahban, Oh and Shah, 2017, Lemmas 3 and 4) and our Lemma 5.3, we can bound γ from below:

Lemma 5.7 (Spectral gap, Negahban, Oh and Shah (2017)). Under the model specified in Section 2.1, if $p \geq c_0 \frac{\log n}{n} \max\left\{1, \frac{\kappa^5}{L}\right\}$ for some sufficiently large constant $c_0 > 0$, then with probability at least $1 - O(n^{-5})$,

$$\gamma = 1 - \max\{\lambda_2(\mathbf{P}^*), -\lambda_n(\mathbf{P}^*)\} - \|\mathbf{P} - \mathbf{P}^*\|_{\pi^*} \geq \frac{1}{2\kappa^2} \frac{d_{\min}}{d_{\max}} \geq \frac{1}{6\kappa^2}.$$

5. In order to control I_4^m , we exploit the statistical independence between $\pi^{(m)}$ and $\mathbf{P}_m$. Specifically, we demonstrate that:

Lemma 5.8. Suppose that $p > \frac{c_0 \log n}{n}$ for some sufficiently large constant $c_0 > 0$. With probability at least $1 - O(n^{-10})$,

$$|I_4^m| \lesssim \frac{1}{\sqrt{n}} \|\pi^{(m)} - \pi\|_2 + \sqrt{\frac{\log n}{Ld}} \|\pi^*\|_{\infty} + \frac{\sqrt{np \log n} + \log n}{d} \|\pi^{(m)} - \pi^*\|_{\infty}.$$

The above bound depends on both $\|\pi^{(m)} - \pi\|_2$ and $\|\pi^{(m)} - \pi^*\|_{\infty}$. We can invoke Lemma 5.6 and the inequality $\|\pi^{(m)} - \pi^*\|_{\infty} \leq \|\pi^{(m)} - \pi\|_2 + \|\pi - \pi^*\|_{\infty}$ to reach

$$|I_4^m| \lesssim \left(\frac{1}{\sqrt{n}} + \frac{\sqrt{np \log n} + \log n}{d}\right) \|\pi^{(m)} - \pi\|_2 + \sqrt{\frac{\log n}{Ld}} \|\pi^*\|_{\infty}$$

$$+ \frac{\sqrt{np \log n} + \log n}{d} \|\pi - \pi^*\|_{\infty}$$

$$\lesssim \left\{ \left(\frac{1}{\sqrt{n}} + \frac{\sqrt{np \log n} + \log n}{d} \right) \frac{\sqrt{\kappa}}{\gamma} + 1 \right\} \sqrt{\frac{\log n}{Ld}} \|\pi^*\|_{\infty}$$

$$+ \left(\frac{1}{\sqrt{n}} + \frac{\sqrt{np \log n} + \log n}{d} \right) \|\pi - \pi^*\|_{\infty}.$$

6. Finally we put the preceding bounds together. When $\frac{np}{\kappa^5 \log n}$ is large enough, with high probability, for some absolute constants $c_1, c_2, c_3 > 0$ one has

$$\begin{aligned}
& \left(\frac{np}{2(1+\kappa)d} - c_1 \sqrt{\frac{\log n}{Ld}} \right) |\pi_m - \pi_m^*| \\
& \leq \left\{ c_2 + \left(\frac{1}{\sqrt{d}} + \frac{c_3}{\sqrt{n}} + c_3 \frac{\sqrt{np \log n} + \log n}{d} \right) \frac{16\sqrt{\kappa}}{\gamma} \right\} \sqrt{\frac{\log n}{Ld}} \|\pi^*\|_\infty \\
& \quad + \left(\frac{1}{\sqrt{d}} + \frac{c_3}{\sqrt{n}} + 2c_3 \frac{\sqrt{np \log n} + \log n}{d} \right) \|\pi - \pi^*\|_\infty.
\end{aligned}$$

simultaneously for all $1 \leq m \leq n$. By taking the maximum over m on the left-hand side and combining terms, we get

$$\begin{aligned}
& \underbrace{\left(\frac{np}{2(1+\kappa)d} - c_1 \sqrt{\frac{\log n}{Ld}} - \frac{1}{\sqrt{d}} - \frac{c_3}{\sqrt{n}} - 2c_3 \frac{\sqrt{np \log n} + \log n}{d} \right)}_{:= \alpha_1} \|\pi - \pi^*\|_\infty \\
& \leq \underbrace{\left\{ c_2 + \left(\frac{1}{\sqrt{d}} + \frac{c_3}{\sqrt{n}} + c_3 \frac{\sqrt{np \log n} + \log n}{d} \right) \frac{16\sqrt{\kappa}}{\gamma} \right\} \sqrt{\frac{\log n}{Ld}}}_{:= \alpha_2} \|\pi^*\|_\infty.
\end{aligned}$$

Hence, as long as $\frac{np}{\kappa^5 \log n}$ is sufficiently large, one has

$$c_1 \sqrt{\frac{\log n}{Ld}} + \frac{1}{\sqrt{d}} + \frac{c_3}{\sqrt{n}} + 2c_3 \frac{\sqrt{np \log n} + \log n}{d} \lesssim \sqrt{\frac{\log n}{np}} \lesssim \frac{1}{\sqrt{\kappa^5}},$$

which further leads to $\alpha_1 \gtrsim 1/\kappa$, $\alpha_2 \lesssim 1$, and

$$\|\pi - \pi^*\|_\infty \lesssim \kappa \sqrt{\frac{\log n}{Ld}} \|\pi^*\|_\infty \asymp \kappa \sqrt{\frac{\log n}{npL}} \|\pi^*\|_\infty.$$

This finishes the proof of Theorem 3.1 and Theorem 2.6.

6. Analysis for the regularized MLE.

This section establishes the ℓ_∞ error of the regularized MLE as claimed in Theorem 3.2 (and also Theorem 2.7). Recall that in Theorem 3.2, we compare the regularized MLE $\boldsymbol{\theta}$ with $\boldsymbol{\theta}^* - \bar{\boldsymbol{\theta}}^* \mathbf{1}$. Therefore, without loss of generality we can assume that

$$\mathbf{1}^\top \boldsymbol{\theta}^* = 0 \quad (6.1)$$

This combined with the fact that $\theta_{\max} - \theta_{\min} = \log \kappa$ reveals that

$$\|\boldsymbol{\theta}^*\|_\infty \leq \log \kappa \quad \text{and} \quad \|\boldsymbol{\theta}^*\|_2 \leq \sqrt{n \log \kappa}.$$

In addition, we assume that $L = O(n^5)$ in this section. It is straightforward to extend the proof to cover $L \leq c_2 \cdot n^{c_3}$ for any constants $c_2, c_3 > 0$.

6.1. Preliminaries and notation.

Before proceeding to the proof, we gather some basic facts. To begin with, the gradient and the Hessian of $\mathcal{L}(\cdot; \mathbf{y})$ in (2.11) can be computed as

$$\nabla \mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) = \sum_{(i,j) \in \mathcal{E}, i > j} \left\{ -y_{j,i} + \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} \right\} (\mathbf{e}_i - \mathbf{e}_j); \quad (6.2)$$

$$\nabla^2 \mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) = \sum_{(i,j) \in \mathcal{E}, i > j} \frac{e^{\theta_i} e^{\theta_j}}{\left(e^{\theta_i} + e^{\theta_j} \right)^2} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top. \quad (6.3)$$

Here $\mathbf{e}_1, \dots, \mathbf{e}_n$ stand for the canonical basis vectors in $\mathbb{R}^n$. When evaluated at the truth $\boldsymbol{\theta}^*$, the size of the gradient can be controlled as follows.

Lemma 6.1. Let λ be as specified in Theorem 3.2. The following event

$$\mathcal{A}_2 := \left\{ \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y})\|_2 \lesssim \sqrt{\frac{n^2 p \log n}{L}} \right\}. \quad (6.4)$$

occurs with probability exceeding $1 - O(n^{-10})$.

The following lemmas characterize the smoothness and the strong convexity of the function $\mathcal{L}_\lambda(\cdot; \mathbf{y})$. In the sequel, we denote by $L_{\mathcal{G}} = \sum_{(i,j) \in \mathcal{E}, i > j} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top$ the (unnormalized) Laplacian matrix (Chung, 1997) associated with $\mathcal{G}$. For any matrix $\mathbf{A}$ we let

$$\lambda_{\min, \perp}(\mathbf{A}) := \min \left\{ \mu \mid \mathbf{z}^\top \mathbf{A} \mathbf{z} \geq \mu \|\mathbf{z}\|_2^2 \text{ for all } \mathbf{z} \text{ with } \mathbf{1}^\top \mathbf{z} = 0 \right\}, \quad (6.5)$$

namely, the smallest eigenvalue when restricted to vectors orthogonal to $\mathbf{1}$.

Lemma 6.2. Suppose that $p > \frac{c_0 \log n}{n}$ for some sufficiently large constant $c_0 > 0$. Then on the event A_0 as defined in (5.1), one has

$$\lambda_{\max}(\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y})) \leq \lambda + np, \quad \forall \boldsymbol{\theta} \in \mathbb{R}^n.$$

PROOF. Note that, $\frac{\theta_i \theta_j}{\left(\theta_i + \theta_j\right)^2} \leq \frac{1}{4}$. It follows immediately from the Hessian in (6.3) that

$$\lambda_{\max}(\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y})) \leq \lambda + \frac{1}{4} \|L_{\mathcal{G}}\| \leq \lambda + \frac{1}{2} d_{\max},$$

where $d_{\max}$ is the maximum vertex degree in the graph $\mathcal{G}$. In addition, on the event $\mathcal{A}_0$ we have $d_{\max} \leq 2np$, which completes the proof. $\square$

Lemma 6.3. For all $\boldsymbol{\theta} \in \mathbb{R}^n$ such that $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \leq C$ for some $C \geq 0$, we have

$$\lambda_{\min, \perp}(\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y})) \geq \lambda + \frac{1}{4\kappa e^{2C}} \lambda_{\min, \perp}(L_{\mathcal{G}}).$$

Lemma 6.4. Let $\mathcal{G} \sim \mathcal{G}_{n,p}$, and suppose that $p > \frac{c_0 \log n}{n}$ for some sufficiently large constant $c_0 > 0$. Then one has

$$\mathbb{P}(\lambda_{\min, \perp}(L_{\mathcal{G}}) \geq np/2) \geq 1 - O(n^{-10}).$$

PROOF. Note that $\lambda_{\min, \perp}(L_{\mathcal{G}})$ is exactly the spectral gap of the Laplacian matrix. See (Tropp, 2015, Sec 5.3.3) for the derivation of this lemma. $\square$

By combining Lemma 6.3 with Lemma 6.4, we reach the following result.

Corollary 6.5. Under the assumptions of Lemma 6.4, with probability exceeding $1 - O(n^{-10})$ one has

$$\lambda_{\min, \perp}(\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y})) \geq \lambda + \frac{1}{8\kappa e^{2C}} np$$

simultaneously for all $\boldsymbol{\theta}$ obeying $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \leq C$ for some $C \geq 0$.

6.2. Proof outline of Theorem 3.2.

This subsection outlines the main steps for establishing Theorem 3.2.

Rather than directly resorting to the optimality condition, we adopt an algorithmic perspective to analyze the regularized MLE $\boldsymbol{\theta}$. Specifically, we consider the standard gradient descent algorithm that is expected to converge to the minimizer $\boldsymbol{\theta}$, and analyze the trajectory of this iterative algorithm instead. The algorithm is stated in Algorithm 2.

Algorithm 2

Gradient descent for computing the regularized MLE.

Initialize $\boldsymbol{\theta}^0 = \boldsymbol{\theta}^*$.

for $t = 0, 1, 2, \dots, T-1$ **do**

$$(6.6) \quad \boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta_t \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^t; \mathbf{y});$$

Notably, this gradient descent algorithm is not practical since the initial point is set to be $\boldsymbol{\theta}^*$. Nevertheless, it is helpful for analyzing the statistical accuracy of the regularized MLE $\boldsymbol{\theta}$. In what follows, we shall adopt a time-invariant step size rule:

$$\eta_t \equiv \eta = 1/(\lambda + np), \quad t = 0, 1, 2, \dots \quad (6.7)$$

Our proof can be divided into three steps:

- I. establish — via standard optimization theory — that the output $\boldsymbol{\theta}^T$ of Algorithm 2 is sufficiently close to the regularized MLE $\boldsymbol{\theta}$, namely,

$$\|\boldsymbol{\theta}^T - \boldsymbol{\theta}\|_\infty \leq \|\boldsymbol{\theta}^T - \boldsymbol{\theta}\|_2 \leq C_0 \kappa^2 \sqrt{\frac{\log n}{npL}} \quad (6.8)$$

For $T = n^5$, where $C_0 > 0$ is some absolute constant;

- II. use the leave-one-out argument to demonstrate that: the output $\boldsymbol{\theta}^T$ is close to the truth $\boldsymbol{\theta}^*$ in an entrywise fashion, i.e.

$$\|\boldsymbol{\theta}^T - \boldsymbol{\theta}^*\|_\infty \leq C_4 \kappa^2 \sqrt{\frac{\log n}{npL}}$$

for some universal constant $C_4 > 0$. Combining this with (6.8) yields

$$\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \lesssim \kappa^2 \sqrt{\frac{\log n}{npL}};$$

- III. the final step is to translate the perturbation bound on $\|\theta - \theta^*\|_\infty$ to $\|e^\theta - e^{\theta^*}\|_\infty$ as claimed in the theorem.

Before continuing, we single out an important fact that will be used throughout the proof.

Fact 6.6. Suppose $\mathbf{1}^\top \theta^* = 0$. Then we have $\mathbf{I}^\top \theta^t = 0$ for all $t \geq 0$.

6.3. Step I.

The first step relies heavily on optimization theory, namely the theory of gradient descent on strongly convex and smooth functions.

1. It is seen that the sequence $\{\theta^t\}_{t=1}^\infty$ converges geometrically fast to the regularized MLE θ , a property that is standard in convex optimization literature. This claim is summarized in the following lemma.

Lemma 6.7. On the event $\mathcal{A}_0$ as defined in (5.1), one has

$$\|\theta^t - \theta\|_2 \leq \rho^t \|\theta^0 - \theta\|_2,$$

where $\rho = 1 - \frac{\lambda}{\lambda + np}$.

PROOF. This result directly follows from the smoothness property (see Lemma 6.2), the trivial strong convexity of $\mathcal{L}_\lambda(\theta; \mathbf{y})$ ($\nabla^2 \mathcal{L}_\lambda(\theta; \mathbf{y}) \geq \lambda \mathbf{I}_n, \forall \theta$), as well as the convergence property of the gradient descent algorithm (e.g. (Bubeck, 2015, Theorem 3.10)). $\square$

A direct consequence of this convergence result and Fact 6.6 is that $\mathbf{1}^\top \theta = 0$ for the regularized MLE θ .

2. We then control $\|\theta^0 - \theta\|_2$. Recall that $\theta^0 = \theta^*$ and we have:

Lemma 6.8. On the event $\mathcal{A}_2$ as defined in (6.4), there exists some constant $c_2 > 0$ such that

$$\|\theta^0 - \theta\|_2 = \|\theta - \theta^*\|_2 \leq c_2 \sqrt{n \log \kappa}.$$

3. The previous two claims taken together lead us to conclude that

$$\begin{aligned} \|\theta^T - \theta\|_2 &\leq \rho^T \|\theta^0 - \theta\|_2 \leq \rho^T c_2 \sqrt{n \log \kappa} = c_2 \left(1 - \frac{\lambda}{\lambda + np}\right)^T \sqrt{n \log \kappa} \\ &\leq c_2 \exp\left(-\frac{T\lambda}{\lambda + np}\right) \sqrt{n \log \kappa} \end{aligned}$$

$$\begin{aligned}
&\leq c_2 \exp\left(-\frac{T}{c_3 \log \kappa} \sqrt{\frac{\log n}{npL}}\right) \sqrt{n} \kappa^2 \quad \left(\text{by } \lambda \asymp \frac{1}{\log \kappa} \sqrt{\frac{np \log n}{L}} < np\right) \\
&\leq c_2 \exp\left(-\frac{T}{c_4 \log n} \sqrt{\frac{\log n}{npL}}\right) \sqrt{n} \kappa^2 \quad \left(\text{by } \kappa^4 \lesssim \frac{np}{\log n} \leq \frac{n}{\log n}\right) \\
&\leq C_0 \kappa^2 \sqrt{\frac{\log n}{npL}},
\end{aligned}$$

for some constants $c_3, c_4 C_0 > 0, L \lesssim n^5$, and sufficiently large T (recall that $T = n^5$). The above bounds are somewhat loose, but they suffice for our purpose. We then naturally obtain

$$\|\theta^T - \theta\|_\infty \leq \|\theta^T - \theta\|_2 \leq C_0 \kappa^2 \sqrt{\frac{\log n}{npL}}$$

as claimed. This finishes the first step of the proof.

6.4. Step II.

The purpose of this step is to show that all iterates $\{\theta^t\}_{0 \leq t \leq T}$ are sufficiently close to θ^* in terms of the ℓ_∞ -norm distance. To facilitate analysis, for each $1 \leq m \leq n$, we introduce a leave-one-out sequence $\{\theta^{t,(m)}\}$ constructed via the following update rule

$$\theta^{t+1,(m)} = \theta^{t,(m)} - \eta \nabla \mathcal{L}_\lambda^{(m)}(\theta^{t,(m)}), \quad (6.9)$$

where $\theta^{0,(m)} = \theta^0 = \theta^*$ and

$$\begin{aligned}
\mathcal{L}_\lambda^{(m)}(\theta; \mathbf{y}) &:= \sum_{(i,j) \in \mathcal{E}, i > j, i \neq m, j \neq m} \left\{ -y_{j,i} (\theta_i - \theta_j) + \log \left(1 + e^{\theta_i - \theta_j} \right) \right\} \\
&+ \sum_{i: i \neq m} p \left\{ -\frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} (\theta_i - \theta_m) + \log \left(1 + e^{\theta_i - \theta_m} \right) \right\} + \frac{1}{2} \lambda \|\theta\|_2^2. \quad (6.10)
\end{aligned}$$

Here, the leave-one-out loss function $\mathcal{L}_\lambda^{(m)}(\theta; \mathbf{y})$ replaces all log-likelihood components involving the m -th item with their expected values (unconditional on $\mathcal{G}$). For any $1 \leq m \leq n$,

the auxiliary sequence $\{\boldsymbol{\theta}^{t,(m)}\}$ serves as a reasonably good proxy for $\{\boldsymbol{\theta}^t\}$, while remaining statistically independent of $\{y_{i,m} | (i,m) \in \varepsilon\}$.

Our proof in this step is inductive in nature. For the sake of clarity, we first list all induction hypotheses needed in our analysis:

$$\|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_2 \leq C_1 \kappa \sqrt{\frac{\log n}{pL}}, \quad (6.11a)$$

$$\max_{1 \leq m \leq n} |\theta_m^{t,(m)} - \theta_m^*| \leq C_2 \kappa^2 \sqrt{\frac{\log n}{npL}}, \quad (6.11b)$$

$$\max_{1 \leq m \leq n} \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^{t,(m)}\|_2 \leq C_3 \kappa \sqrt{\frac{\log n}{npL}}, \quad (6.11c)$$

$$\|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_\infty \leq C_4 \kappa^2 \sqrt{\frac{\log n}{npL}} \quad (6.11d)$$

where $C_1, \dots, C_4 > 0$ are some absolute constants. We aim to show that if the iterates at the t -th iteration — i.e. $\boldsymbol{\theta}^t$ and $\{\boldsymbol{\theta}^{t,(m)}\}_{1 \leq m \leq n}$ — satisfy the induction hypotheses (6.11), then the $(t+1)$ -th iterates continue to satisfy these hypotheses. Clearly, it suffices to justify (6.11) for all $0 \leq t \leq T = n^5$.

Before we dive into the inductive arguments, there are a few direct consequences of (6.11) that are worth listing. We gather them in the next lemma.

Lemma 6.9. Suppose the induction hypotheses (6.11) hold true for the t -th iteration, then there exist some universal constants $C_5, C_6 > 0$ such that the following two bounds hold:

$$\max_{1 \leq m \leq n} \|\boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^*\|_\infty \leq C_5 \kappa^2 \sqrt{\frac{\log n}{npL}}, \quad (6.12a)$$

$$\max_{1 \leq m \leq n} \|\boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^*\|_2 \leq C_6 \kappa \sqrt{\frac{\log n}{pL}}. \quad (6.12b)$$

Note that the base case (i.e. the case for $t = 0$) is trivially true due to the same initial points, namely, $\theta^{0,(m)} = \theta^0 = \theta^*$ for all $1 \leq m \leq n$. We start with the first induction hypothesis (6.11a), which is supplied below.

Lemma 6.10. Suppose the induction hypotheses (6.11) hold true for the t -th iteration, then with probability at least $1 - O(n^{-10})$, one has

$$\|\theta^{t+1} - \theta^*\|_2 \leq C_1 \kappa \sqrt{\frac{\log n}{pL}},$$

as long as the step size obeys $0 < \eta \leq \frac{1}{\lambda + np}$ and $C_1 > 0$ is sufficiently large.

The remaining induction steps are provided in the following lemmas.

Lemma 6.11. Suppose the induction hypotheses (6.11) hold true for the t -th iteration, then with probability at least $1 - O(n^{-10})$, one has

$$\max_{1 \leq m \leq n} |\theta_m^{t+1,(m)} - \theta_m^*| \leq C_2 \kappa^2 \sqrt{\frac{\log n}{npL}},$$

with the proviso that $0 < \eta \leq \frac{1}{\lambda + np}$ and $C_2 \gtrsim C_6 + c_\lambda$.

Lemma 6.12. Suppose the induction hypotheses (6.11) hold true for the t -th iteration, then with probability at least $1 - O(n^{-10})$, one has

$$\max_{1 \leq m \leq n} \|\theta^{t+1} - \theta^{t+1,(m)}\|_2 \leq C_3 \kappa \sqrt{\frac{\log n}{npL}},$$

as long as the step size obeys $0 < \eta \leq \frac{1}{\lambda + np}$ and $C_3 > 0$ is sufficiently large.

Lemma 6.13. Suppose the induction hypotheses (6.11) hold true for the t -th iteration, then with probability at least $1 - O(n^{-10})$, one has

$$\|\theta^{t+1} - \theta^*\|_\infty \leq C_4 \kappa^2 \sqrt{\frac{\log n}{npL}}$$

for any $C_4 = C_3 + C_2$.

Taking the union bound over $T = n^5$ iterations yields that with probability at least $1 - O(n^{-5})$,

$$\|\theta^T - \theta^*\|_\infty \leq C_4 \kappa^2 \sqrt{\frac{\log n}{npL}},$$

which together with the conclusion in Step I results in

$$\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \leq \|\boldsymbol{\theta}^T - \boldsymbol{\theta}^*\|_\infty + \|\boldsymbol{\theta}^T - \boldsymbol{\theta}\|_\infty \leq (C_0 + C_4)\kappa^2 \sqrt{\frac{\log n}{npL}}. \quad (6.13)$$

6.5. Step III.

It remains to show that

$$\frac{\|e^{\boldsymbol{\theta}} - e^{\boldsymbol{\theta}^*}\|_\infty}{\|e^{\boldsymbol{\theta}}\|_\infty} \lesssim \kappa^2 \sqrt{\frac{\log n}{npL}}.$$

Toward this end, we observe that for each $1 \leq m \leq n$

$$\frac{|e^{\theta_m} - e^{\theta_m^*}|}{e^{\theta_{\max}}} = \frac{e^{\tilde{\theta}_m}(\theta_m - \theta_m^*)}{e^{\theta_{\max}}} \leq \frac{e^{\theta_{\max} + \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty} \cdot |\theta_m - \theta_m^*|}{e^{\theta_{\max}}},$$

where $\tilde{\theta}_m$ is between θ_m and θ_m^* , and $\theta_{\max}$ is the largest entry of $\boldsymbol{\theta}^*$. Continuing the derivation and using (6.13), we arrive at

$$\max_{1 \leq m \leq n} \frac{|e^{\theta_m} - e^{\theta_m^*}|}{e^{\theta_{\max}}} \leq \frac{e^{\theta_{\max} + \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty}}{e^{\theta_{\max}}} \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \lesssim \kappa^2 \sqrt{\frac{\log n}{npL}}$$

as long as $\kappa^2 \sqrt{\frac{\log n}{npL}}$ is small enough. This completes the proof of Theorem 3.2.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements.

Y. C. is supported by ARO grant W911NF-18-1-0303 and by Princeton SEAS innovation award. J. F. is supported by NSF grants DMS-1662139 and DMS-1712591 and NIH grant 2R01-GM072611-13.

References.

- ABBE E, Fan J, Wang K and ZHONG Y (2017). Entrywise Eigenvector Analysis of Random Matrices with Low Expected Rank. arXiv preprint arXiv:1709.09565.
- AGARWAL A, AGARWAL S, ASSADI S and KHANNA S (2017). Learning with Limited Rounds of Adaptivity: Coin Tossing, Multi-Armed Bandits, and Ranking from Pairwise Comparisons. In Conference on Learning Theory 39–75.
- AMMAR A and SHAH D (2011). Ranking: Compare, don't score. In 2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton) 776–783.

- AMMAR A and SHAH D (2012). Efficient rank aggregation using partial data In SIG-METRICS 40 355–366. ACM.
- BALTRUNAS L, MAKCINSKAS T and RICCI F (2010). Group Recommendations with Rank Aggregation and Collaborative Filtering. In Proceedings of the Fourth ACM Conference on Recommender Systems RecSys '10 119–126 ACM, New York, NY, USA.
- BRADLEY RA and TERRY ME (1952). Rank analysis of incomplete block designs. I. The method of paired comparisons. *Biometrika* 39 324–345.
- BUBECK S (2015). Convex Optimization: Algorithms and Complexity. *Found. Trends Mach. Learn.* 8 231–357.
- BUSA-FEKETE R, SZÖRÉNYI B, WENG P, CHENG W and HÜLLERMEIER E (2013). Top-k Selection Based on Adaptive Sampling of Noisy Preferences. In International Conference on Machine Learning.
- CHEN, Y. and CANDES, E. (2016). The Projected Power Method: An Efficient Algorithm for Joint Alignment from Pairwise Differences. accepted to Communications on Pure and Applied Mathematics.
- CHEN Y and CANDÈS EJ (2017). Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Comm. Pure Appl. Math.* 70 822–883.
- CHEN Y and SUH C (2015). Spectral MLE: Top- K Rank Aggregation from Pairwise Comparisons. In International Conference on Machine Learning 371–380.
- CHEN X, BENNETT PN, COLLINS-THOMPSON K and HORVITZ E (2013). Pairwise Ranking Aggregation in a Crowdsourced Setting. In ACM International Conference on Web Search and Data Mining 193–202 ACM, New York, NY, USA.
- CHEN X, GOPI S, MAO J and SSCHNEIDER J (2017). Competitive Analysis of the top- K Ranking Problem. In Symposium on Discrete Algorithms 1245–1264.
- Chen Y, Chi Y, Fan J and Ma c. (2018). Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. arXiv:1803.07726.
- CHUNG FRK (1997). Spectral graph theory 92. American Mathematical Society.
- DAVIS C and KAHAN WM (1970). The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.* 7 1–46.
- DWORK C, KUMAR R, NAOR M and SIVAKUMAR D (2001). Rank Aggregation Methods for the Web. In International Conference on World Wide Web 613–622.
- EI KAROUÏ N (2017). On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*.
- ELDRIDGE J, BELKIN M and WANG Y (2017). Unperturbed: spectral analysis beyond Davis-Kahan. arXiv preprint arXiv:1706.06516.
- FAN J, WANG W and ZHONG Y (2018). An ℓ_∞ Eigenvector Perturbation Bound and Its Application. *Journal of Machine Learning Research* 18 1–42.
- FORD LR Jr (1957). Solution of a ranking problem from binary comparisons. *Amer. Math. Monthly* 64 28–33. MR0097876
- HAJEK B, Oh S and XU J (2014). Minimax-optimal Inference from Partial Rankings. In Neural Information Processing Systems 1475–1483.
- HECKEL R, SHAH NB, RAMCHANDRAN K and WAINWRIGHT MJ (2016). Active Ranking from Pairwise Comparisons and when Parametric Assumptions Don't Help. arXiv preprint arXiv:1606.08842.
- HUNTER d. R(2004). MM algorithms for generalized Bradley-Terry models. *Ann. Statist.* 32 384–406. MR2051012
- JAMIESON KG and NOWAK RD (2011). Active Ranking Using Pairwise Comparisons. In Neural Information Processing Systems 2240–2248.
- JANG M, KIM S, SUH C and OH S (2016). Top- K Ranking from Pairwise Comparisons: When Spectral Ranking is Optimal. arXiv preprint arXiv:1603.04153
- JAVANMARD A and MONTANARI A (2015). De-biasing the lasso: Optimal sample size for Gaussian designs. arXiv preprint arXiv:1508-02757.

- JIANG X, LIM L-H, YAO Y and YE Y (2011). Statistical ranking and combinatorial Hodge theory. *Math. Program.* 127 203–244. MR2776715
- KESHAVAN RH, MONTANARI A and OH S (2010). Matrix completion from noisy entries. *J. Mach. Learn. Res.* 11 2057–2078. MR2678022
- KOLTCHINSKII V and LOUNICI K (2016). Asymptotics and concentration bounds for bilinear forms of spectral projectors of sample covariance In *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* 52 1976–2013. Institut Henri Poincaré.
- Koltchinskii V and XIA, D. (2016). Perturbation of linear forms of singular vectors under gaussian noise In *High Dimensional Probability VII* 397–423. Springer.
- LU Y and NEGAHBAN SN (2014). Individualized Rank Aggregation using Nuclear Norm Regularization. *arXiv preprint arXiv:1410.0860*.
- LUCE RD (1959). *Individual choice behavior: A theoretical analysis*. John Wiley & Sons, Inc., New York; Chapman & Hall, Ltd, London.
- Ma C, Wang K, Chi Y and Chen Y (2017). Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. *arXiv preprint arXiv:1711.10467*.
- MASSE K (1997). Statistical models applied to the rating of sports teams. Technical report Bluefield College.
- NEGAHBAN S, OH S and SHAH D (2017). Rank centrality: ranking from pairwise comparisons. *Oper. Res.* 65 266–287.
- NEGAHBAN S, OH S, THEKUMPARAMPIL KK and XU J (2017). Learning from Comparisons and Choices. *arXiv preprint arXiv:1704.07228*.
- PANANJADY A, MAO C, MUTHUKUMAR V, WAINWRIGHT MJ and COUR-TADE TA (2017). Worst-case vs Average-case Design for Estimation from Fixed Pairwise Comparisons. *arXiv preprint arXiv:1707.06217*.
- RAJKUMAR A and AGARWAL S (2014). A Statistical Convergence Perspective of Algorithms for Rank Aggregation from Pairwise Data. In *International Conference on Machine Learning* I-118–I-126.
- RAJKUMAR A and AGARWAL S (2016). When Can We Rank Well from Comparisons of $O(n \log n)$ Non-Actively Chosen Pairs? In *Conference on Learning Theory* 1376–1401.
- ROHE K, CHATERJEE S and YU B (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* 39 1878–1915.
- SHAH NB and WAINWRIGHT MJ (2015). Simple, robust and optimal ranking from pairwise comparisons. *arXiv preprint arXiv:1512.08949*.
- SHAH NB, BALAKRISHNAN S, GUNTUBOYINA A and WAINWRIGHT MJ (2017). Stochastically transitive models for pairwise comparisons: statistical and computational issues. *IEEE Trans. Inform. Theory* 63 934–959.
- SOUFIANI HA, CHEN WZ, PARKES DC and XIA L (2013). Generalized Method-of-moments for Rank Aggregation. In *Proceedings of the 26th International Conference on Neural Information Processing Systems*. NIPS'13 2706–2714.
- SUH C, TAN VYF and ZHAO R (2017). Adversarial top-K ranking. *IEEE Trans. Inform. Theory* 63 2201–2225.
- SUR P, CHEN Y and CANDES EJ (2017). The Likelihood Ratio Test in High-Dimensional Logistic Regression Is Asymptotically a Rescaled Chi-Square. *arXiv preprint arXiv:1706.01191*.
- TROPP JA (2015). An Introduction to Matrix Concentration Inequalities. *Found. Trends Mach. Learn.* 8 1–230.
- ZHONG Y and BOUMAL N (2017). Near-optimal bounds for phase synchronization. *arXiv preprint arXiv:1703.06605*.

FIG 1.
Empirical performance of the spectral method and the (regularized) MLE: (a) ℓ_∞ error vs. L .
(b) ℓ_∞ error vs. p and (c) ℓ_∞ error vs. $n^2 pL$

FIG 2.

Comparisons between the relative ℓ_∞ error and the relative ℓ_2 error for (a) the spectral method, (b) the regularized MLE and (c) the MLE.

FIG 3.
The top-K ranking accuracy of both the spectral method and the regularized MLE.