

Spectral Method and Regularized MLE Are Both Optimal for Top- K Ranking

Yuxin Chen* Jianqing Fan† Cong Ma† Kaizheng Wang†

August 2017; Revised July 2018

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 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 $\sin \Theta$ 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, eigenvector perturbation analysis, leave-one-out analysis, reversible Markov chain.

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 et al., 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)

Author names are sorted alphabetically.

*Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA; Email: yuxin.chen@princeton.edu.

†Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA; Email: [jqfan, cong, kaizheng@princeton.edu](mailto:{jqfan, cong, kaizheng}@princeton.edu).

model (Bradley and Terry, 1952; Luce, 1959), which assigns a latent preference score $\{w_i^*\}_{1 \leq i \leq n}$ 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,

$$\mathbb{P}\{\text{item } j \text{ is preferred over item } i\} = \frac{w_j^*}{w_i^* + w_j^*} \quad (1)$$

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. **The spectral method.** By connecting the winning probability in (1) with the transition probability of a reversible Markov chain, the spectral method attempts recovery of $\{w_i^*\}$ via the leading left eigenvector of a sample transition matrix. This procedure, also known as Rank Centrality (Negahban et al., 2017a), bears similarity to the PageRank algorithm.
2. **The 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 et al., 2017a).

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 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 $\boldsymbol{\pi} \in \mathbb{R}^n$, we define the inner product space indexed by $\boldsymbol{\pi}$ as a vector space in \mathbb{R}^n endowed with the inner product $\langle \mathbf{x}, \mathbf{y} \rangle_{\boldsymbol{\pi}} = \sum_{i=1}^n \pi_i x_i y_i$. The corresponding vector norm and the induced matrix norm are defined respectively as $\|\mathbf{x}\|_{\boldsymbol{\pi}} = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle_{\boldsymbol{\pi}}}$ and $\|\mathbf{A}\|_{\boldsymbol{\pi}} = \sup_{\|\mathbf{x}\|_{\boldsymbol{\pi}}=1} \|\mathbf{x}^{\top} \mathbf{A}\|_{\boldsymbol{\pi}}$.

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 $\mathbf{L}_{\mathcal{G}} = \sum_{(i,j) \in \mathcal{E}, i > j} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^{\top}$ the (unnormalized) Laplacian matrix (Chung, 1997) associated with it, where $\{\mathbf{e}_i\}_{1 \leq i \leq n}$ are the standard basis vectors in \mathbb{R}^n . For a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with n real eigenvalues, we let $\lambda_1(\mathbf{A}) \geq \lambda_2(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$ 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

$$\mathbf{w}^* = [w_1^*, \dots, w_n^*]^{\top} \quad (2)$$

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

$$\boldsymbol{\theta}^* = [\theta_1^*, \dots, \theta_n^*]^{\top}, \quad \text{where } \theta_i^* = \log w_i^*. \quad (3)$$

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

$$w_i^* \in [w_{\min}, w_{\max}], \quad \text{or} \quad \theta_i^* \in [\theta_{\min}, \theta_{\max}] \quad (4)$$

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

$$\kappa := \frac{w_{\max}}{w_{\min}}. \quad (5)$$

Notably, the current paper primarily focuses on the case with a *fixed* dynamic range (i.e. κ 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

$$w_{\max} \geq w_1^* \geq w_2^* \geq \dots \geq w_n^* \geq w_{\min}, \quad (6)$$

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

Comparison graph. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ 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 \mathcal{E} . 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 $(i, j) \in \mathcal{E}$, 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

$$y_{i,j}^{(l)} \stackrel{\text{ind.}}{=} \begin{cases} 1, & \text{with probability } \frac{w_j^*}{w_i^* + w_j^*} = \frac{e^{\theta_j^*}}{e^{\theta_i^*} + e^{\theta_j^*}}, \\ 0, & \text{else.} \end{cases} \quad (7)$$

By convention, we set $y_{i,j}^{(l)} = 1 - y_{j,i}^{(l)}$ for all $(i, j) \in \mathcal{E}$ 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

$$\mathbf{y} := \{y_{i,j} \mid (i, j) \in \mathcal{E}\}, \quad \text{where } y_{i,j} := \frac{1}{L} \sum_{l=1}^L y_{i,j}^{(l)}. \quad (8)$$

To simplify the notation, we shall also take

$$y_{i,j}^* := \frac{w_j^*}{w_i^* + w_j^*} = \frac{e^{\theta_j^*}}{e^{\theta_i^*} + e^{\theta_j^*}}.$$

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 et al., 2017a), 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

$$P_{i,j} = \begin{cases} \frac{1}{d} y_{i,j}, & \text{if } (i, j) \in \mathcal{E}, \\ 1 - \frac{1}{d} \sum_{k:(i,k) \in \mathcal{E}} y_{i,k}, & \text{if } i = j, \\ 0, & \text{otherwise,} \end{cases} \quad (9)$$

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 et al. (2017a), $\boldsymbol{\pi}$ is a faithful estimate of \mathbf{w}^* 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 \mathbf{w}^* , it is perhaps more convenient to look at the population transition matrix $\mathbf{P}^* = [P_{i,j}^*]_{1 \leq i, j \leq n}$:

$$P_{i,j}^* = \begin{cases} \frac{1}{d} \frac{w_j^*}{w_i^* + w_j^*}, & \text{if } (i, j) \in \mathcal{E}, \\ 1 - \frac{1}{d} \sum_{k:(i,k) \in \mathcal{E}} \frac{w_k^*}{w_i^* + w_k^*}, & \text{if } i = j, \\ 0, & \text{otherwise,} \end{cases}$$

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

$$\boldsymbol{\pi}^* := \frac{1}{\sum_{i=1}^n w_i^*} [w_1^*, w_2^*, \dots, w_n^*]^\top \quad (10)$$

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,

$$\pi_i^* P_{i,j}^* = \pi_j^* P_{j,i}^*, \quad \forall (i, j). \quad (11)$$

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

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 (9).

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

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

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)

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) &:= - \sum_{(i,j) \in \mathcal{E}, i > j} \left\{ y_{j,i} \log \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} + (1 - y_{j,i}) \log \frac{e^{\theta_j}}{e^{\theta_i} + e^{\theta_j}} \right\} \\ &= \sum_{(i,j) \in \mathcal{E}, i > j} \left\{ -y_{j,i} (\theta_i - \theta_j) + \log (1 + e^{\theta_i - \theta_j}) \right\}. \end{aligned} \quad (12)$$

The regularized MLE then amounts to solving the following convex program

$$\text{minimize}_{\boldsymbol{\theta} \in \mathbb{R}^n} \quad \mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y}) := \mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) + \frac{1}{2} \lambda \|\boldsymbol{\theta}\|_2^2, \quad (13)$$

for a regularization parameter $\lambda > 0$. As will be discussed later, we shall adopt the choice $\lambda \asymp \sqrt{\frac{np \log n}{L}}$ 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

$$\Delta_K := \frac{w_K^* - w_{K+1}^*}{w_{\max}}. \quad (14)$$

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,

$$N = (1 + o(1)) \mathbb{E}[N] = (1 + o(1)) n^2 p L / 2 \quad (15)$$

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

Theorem 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*

$$\frac{n^2 p L}{2} \geq \frac{c_1 n \log n}{\Delta_K^2} \quad (16)$$

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 (13). Here, we take $d = c_d np$ in the spectral method and $\lambda = c_\lambda \sqrt{\frac{np \log n}{L}}$ in the regularized MLE, where $c_d \geq 2$ and $c_\lambda > 0$ are some absolute constants.

Remark 1. 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. 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 3. 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 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. $\kappa = O(1)$.

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

$$n^2 p L \leq 2c_2 \frac{(1 - \epsilon)n \log n - 2}{\Delta_K^2}, \quad (17)$$

where $c_2 = w_{\min}^4 / (4w_{\max}^4)$. Then for any ranking procedure ψ , one can find a score vector \mathbf{w}^* 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

$$n^2 p L / 2 \geq n^2 p / 2 \gtrsim \sqrt{n^3 \log n}, \quad (18)$$

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 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 (18). 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 1, the key is to obtain precise control of the ℓ_∞ loss of the score estimates. Our results are as follows.

Theorem 3 (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 $d = c_d n p$ for some constant $c_d \geq 2$ in Algorithm 1. Then the spectral estimate $\boldsymbol{\pi}$ satisfies*

$$\frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty} \lesssim \sqrt{\frac{\log n}{npL}} \quad (19)$$

with probability $1 - O(n^{-5})$, where $\boldsymbol{\pi}^*$ is the normalized score vector (cf. (10)).

Theorem 4 (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 $e^{\boldsymbol{\theta}} := [e^{\theta_1}, \dots, e^{\theta_n}]^\top$.

Theorems 3–4 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 et al. (2017a) (or Theorem 9 of this paper that improves it by removing the logarithmic factor), which obeys

$$\frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_2}{\|\boldsymbol{\pi}^*\|_2} \lesssim \sqrt{\frac{\log n}{npL}} \quad (20)$$

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 et al., 2017a). 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 3–4 immediately lead to exact top- K identification as claimed by Theorem 1.

Proof of Theorem 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

$$\frac{\pi_i - \pi_j}{\|\boldsymbol{\pi}^*\|_\infty} \geq \frac{\pi_i^* - \pi_j^*}{\|\boldsymbol{\pi}^*\|_\infty} - \frac{|\pi_i - \pi_i^*|}{\|\boldsymbol{\pi}^*\|_\infty} - \frac{|\pi_j - \pi_j^*|}{\|\boldsymbol{\pi}^*\|_\infty} \geq \Delta_K - \frac{2\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty}. \quad (21)$$

In addition, it follows from Theorem 3 as well as our sample complexity assumption that

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

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 (21) 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

$$(\boldsymbol{\pi} - \boldsymbol{\pi}^*)^\top = \boldsymbol{\pi}^\top \mathbf{P} - \boldsymbol{\pi}^{*\top} \mathbf{P}^* = (\boldsymbol{\pi} - \boldsymbol{\pi}^*)^\top \mathbf{P} + \underbrace{\boldsymbol{\pi}^{*\top} (\mathbf{P} - \mathbf{P}^*)}_{:= \boldsymbol{\xi}}. \quad (22)$$

When $p = 1$ and $\frac{w_{\max}}{w_{\min}} \asymp 1$, 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

$$\frac{\|\boldsymbol{\xi}\|_\infty}{\|\boldsymbol{\pi}^*\|_\infty} \lesssim \sqrt{\frac{\log n}{npL}} \asymp \frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_2 \sqrt{\log n}}{\|\boldsymbol{\pi}^*\|_2}, \quad (23)$$

which coincides with the optimal rate. Further, if we look at each entry of (22), then for all $1 \leq m \leq n$,

$$\begin{aligned} \pi_m - \pi_m^* &= [(\boldsymbol{\pi} - \boldsymbol{\pi}^*)^\top \mathbf{P}]_m + \xi_m, \\ &= \underbrace{(\pi_m - \pi_m^*) P_{m,m}}_{\text{contraction}} + \underbrace{[P_{1,m}, \dots, P_{m-1,m}, 0, P_{m+1,m}, \dots, P_{n,m}]}_{\text{error averaging}} \begin{bmatrix} \pi_1 - \pi_1^* \\ \vdots \\ \pi_n - \pi_n^* \end{bmatrix} + \xi_m. \end{aligned} \quad (24)$$

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 (24) behaves as an entrywise contraction of the error; (ii) the second term of (24) 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 (24), we are left with

$$(1 - P_{m,m}) |\pi_m - \pi_m^*| \lesssim \frac{1}{n} \sum_{i=1}^n |\pi_i - \pi_i^*| + \xi_m, \quad \forall m \quad (25)$$

which further gives,

$$\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty \lesssim \frac{1}{n} \sum_{i=1}^n |\pi_i - \pi_i^*| + \|\boldsymbol{\xi}\|_\infty. \quad (26)$$

There are two possibilities compatible with this bound (26): (1) $\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty \lesssim \frac{1}{n} \sum_{i=1}^n |\pi_i - \pi_i^*|$, and (2) $\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty \lesssim \|\boldsymbol{\xi}\|_\infty \lesssim \frac{\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_2}{\|\boldsymbol{\pi}^*\|_2} \|\boldsymbol{\pi}^*\|_\infty$ by (23). In either case, the errors are fairly delocalized, revealing that

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

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

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

Write $\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}^*) = \mathbf{D} - \mathbf{A}$, where \mathbf{D} and \mathbf{A} denote respectively the diagonal and off-diagonal parts of $\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}^*)$. 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 (24). 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.




Figure 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^2pL

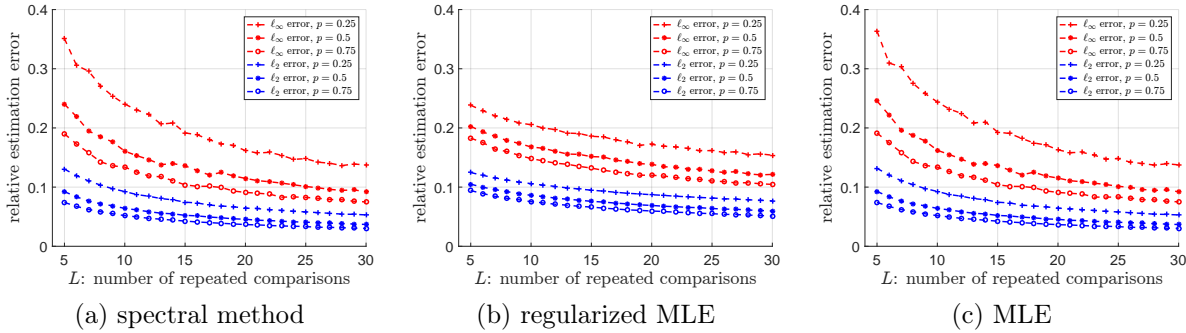


Figure 2: Comparisons between the relative ℓ_∞ error and the relative ℓ_2 error for (a) the spectral method, (b) the regularized MLE and (c) the MLE.

2.6 Numerical experiments

It is worth noting that extensive numerical experiments on both synthetic and real data have already been conducted in [Negahban et al. \(2017a\)](#) 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 w_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 entrywise 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 3-4. Next, we show in Figure 1(c) the relative ℓ_∞ error while fixing the total number of samples (i.e. n^2pL). It can be seen that the performance almost does not change if the sample complexity n^2pL 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 $L = 20$,




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

set $K = 10$, and let $w_i^* = 1$ for all $1 \leq i \leq K$ and $w_j^* = 1 - \Delta$ for all $K + 1 \leq j \leq n$. By construction, the score separation satisfies $\Delta_K = \Delta$. 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 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 et al. (2017a); 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. (2017b); 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 et al., 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 et al., 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 et al., 2016; Eldridge et al., 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 et al., 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 et al., 2010), phase retrieval (Chen and Candès, 2017), graph clustering (Rohe et al., 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 5. *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

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

where $\boldsymbol{\pi}^*$ is the normalized score vector as defined in (10).

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 6. *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 $\bar{\theta}^* := \frac{1}{n} \mathbf{1}^\top \boldsymbol{\theta}^*$ and $e^\theta := [e^{\theta_1}, \dots, e^{\theta_n}]^\top$.

2. the set of top- K ranked items can be recovered exactly by the regularized MLE given in (13), as long as

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

for some sufficiently large constant $c_1 > 0$.

Remark 4. 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 5–6 do not match the lower bound asserted in Theorem 2 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:

$$\Delta_K^* := \frac{w_K^* - w_{K+1}^*}{w_{K+1}^*} \cdot \sqrt{\frac{1}{n} \sum_{i=1}^n \frac{w_{K+1}^* w_i^*}{(w_K^* + w_i^*)^2}}, \quad (27)$$

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

Theorem 7. Fix $\epsilon \in (0, \frac{1}{2})$, and let $\mathcal{G} \sim \mathcal{G}_{n,p}$. Consider any preference score vector \mathbf{w}^* , and let Δ_K^* denote its generalized separation. If

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

then there exists another preference score vector $\tilde{\mathbf{w}}$ 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} .

Proof. See Appendix A. □

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 Appendix B. 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 8), where the spectral gap and matrix perturbation play a crucial rule. However, the current results for controlling these quantities have explicit dependency on κ (Negahban et al. (2017a)). 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 et al., 2017a; 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 11). 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 5 and hence Theorem 3, which characterizes the pointwise 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 8 (Eigenvector perturbation). *Suppose that \mathbf{P} , $\hat{\mathbf{P}}$, and \mathbf{P}^* are probability transition matrices with stationary distributions $\boldsymbol{\pi}$, $\hat{\boldsymbol{\pi}}$, $\boldsymbol{\pi}^*$, respectively. Also, assume that \mathbf{P}^* represents a reversible Markov chain. When $\|\mathbf{P} - \hat{\mathbf{P}}\|_{\boldsymbol{\pi}^*} < 1 - \max\{\lambda_2(\mathbf{P}^*), -\lambda_n(\mathbf{P}^*)\}$, it holds that*

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

Proof. See Appendix C.1. □

Several remarks regarding Theorem 8 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 8 allows one to derive the ℓ_2 estimation error in Negahban et al. (2017a) directly without resorting to the power method; in fact, our ℓ_2 estimation error bound improves upon Negahban et al. (2017a) by some logarithmic factor.

Theorem 9. *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*

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

Proof. See Appendix C.2. □

Notably, Theorem 9 matches the minimax lower bound derived in (Negahban et al., 2017a, 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 1 (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*

$$\mathcal{A}_0 = \left\{ \frac{np}{2} \leq d_{\min} \leq d_{\max} \leq \frac{3np}{2} \right\} \quad (28)$$

obeys

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

Proof. The proof follows from the standard Chernoff bound and is hence omitted. □

Since d is chosen to be $c_d np$ for some constant $c_d \geq 2$, we have, by Lemma 1, that the maximum vertex degree obeys $d_{\max} < d$ with high probability.

5.2 Proof outline of Theorem 5

In this subsection, we outline the proof of Theorem 5.

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

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

Proof. See Appendix C.3. □

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

Lemma 3. *With probability exceeding $1 - O(n^{-5})$, there exists some constant $c > 0$ such that for all $1 \leq m \leq n$,*

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

Proof. See Appendix C.4. □

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 $i \neq j$,

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

with $y_{i,j}^* := \frac{w_j^*}{w_i^* + w_j^*}$. 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 1). 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 4. *Suppose that $n\gamma^2 > c\kappa \log n$ for some sufficiently large constant $c > 0$. With probability at least $1 - O(n^{-5})$,*

$$\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 \leq \frac{16\sqrt{\kappa}}{\gamma} \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_\infty + \|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty, \quad (29)$$

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

Proof. See Appendix C.5. □

Using (Negahban et al., 2017a, Lemmas 3 and 4) and our Lemma 1, we can bound γ from below:

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

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

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

Lemma 6. Suppose that $p > \frac{c_0 \log n}{n}$ for some sufficiently large constant $c_0 > 0$. With probability at least $1 - O(n^{-10})$,

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

Proof. See Appendix C.6. □

The above bound depends on both $\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2$ and $\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty$. We can invoke Lemma 4 and the inequality $\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty \leq \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 + \|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty$ to reach

$$\begin{aligned} |I_4^m| &\lesssim \left(\frac{1}{\sqrt{n}} + \frac{\sqrt{np \log n} + \log n}{d} \right) \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 + \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_\infty \\ &\quad + \frac{\sqrt{np \log n} + \log n}{d} \|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty \\ &\lesssim \left\{ \left(\frac{1}{\sqrt{n}} + \frac{\sqrt{np \log n} + \log n}{d} \right) \frac{\sqrt{\kappa}}{\gamma} + 1 \right\} \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_\infty \\ &\quad + \left(\frac{1}{\sqrt{n}} + \frac{\sqrt{np \log n} + \log n}{d} \right) \|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty. \end{aligned}$$

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

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

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

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

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

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

This finishes the proof of Theorem 5 and Theorem 3.

6 Analysis for the regularized MLE

This section establishes the ℓ_∞ error of the regularized MLE as claimed in Theorem 6 (and also Theorem 4). Recall that in Theorem 6, we compare the regularized MLE $\boldsymbol{\theta}$ with $\boldsymbol{\theta}^* - \bar{\boldsymbol{\theta}}^* \mathbf{1}$. Therefore, without loss of generality we can assume that

$$\mathbf{1}^\top \boldsymbol{\theta}^* = 0. \quad (30)$$

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 (12) can be computed as

$$\nabla \mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) = \sum_{(i,j) \in \mathcal{E}, i > j} \left\{ -y_{j,i} + \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} \right\} (\mathbf{e}_i - \mathbf{e}_j); \quad (31)$$

$$\nabla^2 \mathcal{L}(\boldsymbol{\theta}; \mathbf{y}) = \sum_{(i,j) \in \mathcal{E}, i > j} \frac{e^{\theta_i} e^{\theta_j}}{(e^{\theta_i} + e^{\theta_j})^2} (\mathbf{e}_i - \mathbf{e}_j) (\mathbf{e}_i - \mathbf{e}_j)^\top. \quad (32)$$

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 7. *Let λ be as specified in Theorem 6. The following event*

$$\mathcal{A}_2 := \left\{ \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y})\|_2 \lesssim \sqrt{\frac{n^2 p \log n}{L}} \right\}. \quad (33)$$

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

Proof. See Appendix D.1. □

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 $\mathbf{L}_\mathcal{G} = \sum_{(i,j) \in \mathcal{E}, i > j} (\mathbf{e}_i - \mathbf{e}_j) (\mathbf{e}_i - \mathbf{e}_j)^\top$ the (unnormalized) Laplacian matrix (Chung, 1997) associated with \mathcal{G} . For any matrix \mathbf{A} we let

$$\lambda_{\min, \perp}(\mathbf{A}) := \min \{ \mu \mid \mathbf{z}^\top \mathbf{A} \mathbf{z} \geq \mu \|\mathbf{z}\|_2^2 \text{ for all } \mathbf{z} \text{ with } \mathbf{1}^\top \mathbf{z} = 0 \}, \quad (34)$$

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

Lemma 8. *Suppose that $p > \frac{c_0 \log n}{n}$ for some sufficiently large constant $c_0 > 0$. Then on the event \mathcal{A}_0 as defined in (28), 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 $\frac{e^{\theta_i} e^{\theta_j}}{(e^{\theta_i} + e^{\theta_j})^2} \leq \frac{1}{4}$. It follows immediately from the Hessian in (32) that

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

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. □

Algorithm 2 Gradient descent for computing the regularized MLE.

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

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

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta_t \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^t; \mathbf{y}); \quad (35)$$

end for

Lemma 9. For all $\boldsymbol{\theta} \in \mathbb{R}^n$ such that $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \leq C$ for some $C \geq 0$, we have

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

Proof. See Appendix D.2. □

Lemma 10. 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

$$\mathbb{P}(\lambda_{\min, \perp}(\mathbf{L}_G) \geq np/2) \geq 1 - O(n^{-10}).$$

Proof. Note that $\lambda_{\min, \perp}(\mathbf{L}_G)$ is exactly the spectral gap of the Laplacian matrix. See (Tropp, 2015, Sec 5.3.3) for the derivation of this lemma. □

By combining Lemma 9 with Lemma 10, we reach the following result.

Corollary 1. Under the assumptions of Lemma 10, 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 6

This subsection outlines the main steps for establishing Theorem 6.

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.

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:

$$\eta_t \equiv \eta = \frac{1}{\lambda + np}, \quad t = 0, 1, 2, \dots \quad (36)$$

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,

$$\|\boldsymbol{\theta}^T - \boldsymbol{\theta}\|_\infty \leq \|\boldsymbol{\theta}^T - \boldsymbol{\theta}\|_2 \leq C_0 \kappa^2 \sqrt{\frac{\log n}{npL}} \quad (37)$$

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 (37) 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 $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty$ to $\|e^\boldsymbol{\theta} - e^{\boldsymbol{\theta}^*}\|_\infty$ as claimed in the theorem.

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

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

Proof. See Appendix D.3. □

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

Lemma 11. *On the event \mathcal{A}_0 as defined in (28), one has*

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

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

Proof. This result directly follows from the smoothness property (see Lemma 8), the trivial strong convexity of $\mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y})$ ($\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y}) \succeq \lambda \mathbf{I}_n, \forall \boldsymbol{\theta}$), as well as the convergence property of the gradient descent algorithm (e.g. (Bubeck, 2015, Theorem 3.10)). □

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

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

Lemma 12. *On the event \mathcal{A}_2 as defined in (33), there exists some constant $c_2 > 0$ such that*

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

Proof. See Appendix D.4. □

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

$$\begin{aligned} \|\boldsymbol{\theta}^T - \boldsymbol{\theta}\|_2 &\leq \rho^T \|\boldsymbol{\theta}^0 - \boldsymbol{\theta}\|_2 \leq \rho^T c_2 \sqrt{n} \log \kappa = c_2 \left(1 - \frac{\lambda}{\lambda + np}\right)^T \sqrt{n} \log \kappa \\ &\leq c_2 \exp\left(-\frac{T\lambda}{\lambda + np}\right) \sqrt{n} \kappa^2 \\ &\leq c_2 \exp\left(-\frac{T}{c_3 \log \kappa} \sqrt{\frac{\log n}{npL}}\right) \sqrt{n} \kappa^2 \quad (\text{by } \lambda \asymp \frac{1}{\log \kappa} \sqrt{\frac{np \log n}{L}} < np) \\ &\leq c_2 \exp\left(-\frac{T}{c_4 \log n} \sqrt{\frac{\log n}{npL}}\right) \sqrt{n} \kappa^2 \quad (\text{by } \kappa^4 \lesssim \frac{np}{\log n} \leq \frac{n}{\log n}) \\ &\leq C_0 \kappa^2 \sqrt{\frac{\log n}{npL}}, \end{aligned}$$

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

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

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

$$\boldsymbol{\theta}^{t+1,(m)} = \boldsymbol{\theta}^{t,(m)} - \eta \nabla \mathcal{L}_\lambda^{(m)}(\boldsymbol{\theta}^{t,(m)}), \quad (38)$$

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

$$\begin{aligned} \mathcal{L}_\lambda^{(m)}(\boldsymbol{\theta}; \mathbf{y}) := & \sum_{(i,j) \in \mathcal{E}, i > j, i \neq m, j \neq m} \{-y_{j,i}(\theta_i - \theta_j) + \log(1 + e^{\theta_i - \theta_j})\} \\ & + \sum_{i: i \neq m} p \left\{ -\frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} (\theta_i - \theta_m) + \log(1 + e^{\theta_i - \theta_m}) \right\} + \frac{1}{2} \lambda \|\boldsymbol{\theta}\|_2^2. \end{aligned} \quad (39)$$

Here, the leave-one-out loss function $\mathcal{L}_\lambda^{(m)}(\boldsymbol{\theta}; \mathbf{y})$ 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 $\{y_{i,m} \mid (i,m) \in \mathcal{E}\}$.

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

$$\|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_2 \leq C_1 \kappa \sqrt{\frac{\log n}{pL}}, \quad (40a)$$

$$\max_{1 \leq m \leq n} |\theta_m^{t,(m)} - \theta_m^*| \leq C_2 \kappa^2 \sqrt{\frac{\log n}{npL}}, \quad (40b)$$

$$\max_{1 \leq m \leq n} \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^{t,(m)}\|_2 \leq C_3 \kappa \sqrt{\frac{\log n}{npL}}, \quad (40c)$$

$$\|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_\infty \leq C_4 \kappa^2 \sqrt{\frac{\log n}{npL}}, \quad (40d)$$

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 (40), then the $(t+1)$ -th iterates continue to satisfy these hypotheses. Clearly, it suffices to justify (40) for all $0 \leq t \leq T = n^5$.

Before we dive into the inductive arguments, there are a few direct consequences of (40) that are worth listing. We gather them in the next lemma.

Lemma 13. *Suppose the induction hypotheses (40) 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:*

$$\max_{1 \leq m \leq n} \|\boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^*\|_\infty \leq C_5 \kappa^2 \sqrt{\frac{\log n}{npL}}, \quad (41a)$$

$$\max_{1 \leq m \leq n} \|\boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^*\|_2 \leq C_6 \kappa \sqrt{\frac{\log n}{pL}}. \quad (41b)$$

Proof. See Appendix D.5. □

Note that the base case (i.e. the case for $t = 0$) is trivially true due to the same initial points, namely, $\boldsymbol{\theta}^{0,(m)} = \boldsymbol{\theta}^0 = \boldsymbol{\theta}^*$ for all $1 \leq m \leq n$. We start with the first induction hypothesis (40a), which is supplied below.

Lemma 14. *Suppose the induction hypotheses (40) hold true for the t -th iteration, then with probability at least $1 - O(n^{-10})$, one has*

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

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

Proof. See Appendix D.6. □

The remaining induction steps are provided in the following lemmas.

Lemma 15. *Suppose the induction hypotheses (40) hold true for the t th iteration, then with probability at least $1 - O(n^{-10})$, one has*

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

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

Proof. See Appendix D.7. □

Lemma 16. *Suppose the induction hypotheses (40) hold true for the t -th iteration, then with probability at least $1 - O(n^{-10})$, one has*

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

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

Proof. See Appendix D.8. □

Lemma 17. *Suppose the induction hypotheses (40) hold true for the t -th iteration, then with probability at least $1 - O(n^{-10})$, one has*

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

for any $C_4 \geq C_3 + C_2$.

Proof. See Appendix D.9 □

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

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

which together with the conclusion in Step I results in

$$\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \leq \|\boldsymbol{\theta}^T - \boldsymbol{\theta}^*\|_\infty + \|\boldsymbol{\theta}^T - \boldsymbol{\theta}\|_\infty \leq (C_0 + C_4) \kappa^2 \sqrt{\frac{\log n}{npL}}. \quad (42)$$

6.5 Step III

It remains to show that

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

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

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

where $\tilde{\theta}_m$ is between θ_m and θ_m^* , and θ_{\max} is the largest entry of θ^* . Continuing the derivation and using (42), we arrive at

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

as long as $\kappa^2 \sqrt{\frac{\log n}{npL}}$ is small enough. This completes the proof of Theorem 6.

Acknowledgements

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

A Proof of Theorem 7

As usual, suppose the truth has preference scores w_1^*, \dots, w_n^* . To establish the lower bound, we construct another slightly perturbed scenario where the score of the i th ranked item is \tilde{w}_i as defined by

$$\tilde{w}_i = \begin{cases} w_i^*, & \text{if } i \neq K \text{ and } i \neq K + 1, \\ w_{K+1}^*, & \text{if } i = K, \\ w_K^*, & \text{if } i = K + 1. \end{cases} \quad (43)$$

In words, $\tilde{\mathbf{w}}$ is obtained by swapping the scores of the K th and the $(K + 1)$ th items in \mathbf{w}^* . Clearly, these two score vectors share the same generalized separation measure Δ_K^* , although the top- K items in these two scenarios are not identical. It thus suffices to bound the probability of error in distinguishing these two score vectors given the data.

In the sequel, we denote by \mathbb{P}_0 and \mathbb{P}_1 the probability measures under the scores $\{w_i^*\}$ and $\{\tilde{w}_i\}$, respectively, and let $P_e(\psi)$ represent the probability of error in distinguishing \mathbb{P}_0 and \mathbb{P}_1 using a procedure ψ . In view of (Tsybakov, 2009, Theorem 2.2), if

$$\text{TV}(\mathbb{P}_0, \mathbb{P}_1) \leq \epsilon \quad (44)$$

for some fixed constant $0 < \epsilon < 1$, then

$$\inf_{\psi} P_e(\psi) \geq \frac{1 - \epsilon}{2}. \quad (45)$$

Here, $\text{TV}(\mathbb{P}_0, \mathbb{P}_1)$ represents the total variation distance between \mathbb{P}_0 and \mathbb{P}_1 .

The next step then boils down to characterizing $\text{TV}(\mathbb{P}_0, \mathbb{P}_1)$. To this end, denoting by $\mathbb{P}_0^{i,j}$ (resp. $\mathbb{P}_1^{i,j}$) the distribution of the samples comparing items i and j under \mathbb{P}_0 (resp. \mathbb{P}_1), we obtain that

$$\begin{aligned} \text{TV}(\mathbb{P}_0, \mathbb{P}_1) &= \text{TV}\left(\otimes_{i>j} \mathbb{P}_0^{i,j}, \otimes_{i>j} \mathbb{P}_1^{i,j}\right) \\ &\leq \text{TV}\left(\otimes_{i:i \neq K} \mathbb{P}_0^{i,K}, \otimes_{i:i \neq K} \mathbb{P}_1^{i,K}\right) + \text{TV}\left(\otimes_{i:i \neq K+1} \mathbb{P}_0^{i,K+1}, \otimes_{i:i \neq K+1} \mathbb{P}_1^{i,K+1}\right) \end{aligned} \quad (46)$$

$$\leq \sqrt{\text{KL}\left(\otimes_{i:i \neq K} \mathbb{P}_0^{i,K} \parallel \otimes_{i:i \neq K} \mathbb{P}_1^{i,K}\right)} / 2 \quad (47)$$

$$+ \sqrt{\text{KL}\left(\otimes_{i:i \neq K+1} \mathbb{P}_1^{i,K+1} \parallel \otimes_{i:i \neq K+1} \mathbb{P}_0^{i,K+1}\right) / 2}, \quad (48)$$

where $\text{KL}(P \parallel Q)$ is the KL divergence from Q to P . Here, (46) arises from two facts: (i) \mathbb{P}_0 and \mathbb{P}_1 differ only over locations within $\{(i, K) \mid i \neq K\}$ and $\{(i, K+1) \mid i \neq K+1\}$, and (ii) for any product measure one has

$$\text{TV}(P_1 \otimes Q_1, P_2 \otimes Q_2) \leq \text{TV}(P_1, P_2) + \text{TV}(Q_1, Q_2).$$

Additionally, the inequality (48) comes from Pinsker's inequality (Tsybakov, 2009, Lemma 2.5).

We then look at each term of (48) separately. To begin with, repeating the analysis in (Chen and Suh, 2015, Appendix B), we can demonstrate (using the independence assumption) that

$$\begin{aligned} \text{KL}\left(\otimes_{i:i \neq K} \mathbb{P}_0^{i,K} \parallel \otimes_{i:i \neq K} \mathbb{P}_1^{i,K}\right) &= pL \sum_{i:i \neq K} \text{KL}\left(\mathbb{P}_0(y_{i,K}^{(1)}) \parallel \mathbb{P}_1(y_{i,K}^{(1)})\right) \\ &= pL \sum_{i:i \neq K} \text{KL}\left(\text{Bern}\left(\frac{w_i^*}{w_i^* + w_K^*}\right) \parallel \text{Bern}\left(\frac{w_i^*}{w_i^* + w_{K+1}^*}\right)\right), \end{aligned}$$

where $\text{Bern}(p)$ denotes the Bernoulli distribution with mean p . Upper bounding the KL divergence via χ^2 divergence (see (Tsybakov, 2009, Lemma 2.7)), namely,

$$\begin{aligned} \text{KL}(\text{Bern}(p) \parallel \text{Bern}(q)) &\leq \chi^2(\text{Bern}(p) \parallel \text{Bern}(q)) \\ &= \frac{(p-q)^2}{q} + \frac{(p-q)^2}{1-q} = \frac{(p-q)^2}{q(1-q)}, \end{aligned}$$

we arrive at

$$\begin{aligned} \text{KL}\left(\otimes_{i:i \neq K} \mathbb{P}_0^{i,K} \parallel \otimes_{i:i \neq K} \mathbb{P}_1^{i,K}\right) &\leq pL \sum_{i=1}^n \frac{\left(\frac{w_i^*}{w_i^* + w_{K+1}^*} - \frac{w_i^*}{w_i^* + w_K^*}\right)^2}{\frac{w_i^*}{w_i^* + w_{K+1}^*} \frac{w_{K+1}^*}{w_i^* + w_{K+1}^*}} \\ &= pL \frac{(w_K^* - w_{K+1}^*)^2}{w_{K+1}^*} \sum_{i=1}^n \frac{w_i^*}{(w_i^* + w_K^*)^2}. \end{aligned}$$

Similarly, one can derive

$$\text{KL}\left(\otimes_{i:i \neq K+1} \mathbb{P}_1^{i,K+1} \parallel \otimes_{i:i \neq K+1} \mathbb{P}_0^{i,K+1}\right) \leq pL \frac{(w_K^* - w_{K+1}^*)^2}{w_{K+1}^*} \sum_{i=1}^n \frac{w_i^*}{(w_i^* + w_K^*)^2}.$$

Put together the preceding bounds to reach

$$\text{TV}(\mathbb{P}_0, \mathbb{P}_1) \leq \sqrt{2pL \frac{(w_K^* - w_{K+1}^*)^2}{w_{K+1}^*} \sum_{i=1}^n \frac{w_i^*}{(w_i^* + w_K^*)^2}} = \sqrt{2pLn} \Delta_K^*.$$

As a consequence, if

$$2npL \Delta_K^{*2} \leq \epsilon^2,$$

one necessarily has $\text{TV}(\mathbb{P}_0, \mathbb{P}_1) \leq \epsilon$, which combined with (45) yields $\inf_{\psi} P_e(\psi) \geq \frac{1-\epsilon}{2}$ as claimed.

B Examples for the general κ setting

Recall that in Section 3, we introduce a new metric Δ_K^* . The following three examples shed some light on the potential effectiveness of Δ_K^* as a fundamental information measure.

- *Case 1:* $\kappa = O(1)$. Under this circumstance, it is easy to verify that

$$\Delta_K^* \asymp \Delta_K,$$

and hence all our preceding results for spectral method and regularized MLE for $\kappa = O(1)$ continue to hold with Δ_K replaced by Δ_K^* . In this case, the new lower bound for sample complexity is slightly worse than the previous one (Theorem 2 in the main text) by a factor of $\log n$.

- *Case 2:* Suppose there are 100 items with $w_1^* = \dots = w_5^* = 10$, $w_6^* = \dots = w_{99}^* = 5$ and $w_{100}^* = 10^{-6}$. Our goal is to find the top-5 ranked items. Intuitively, the presence of the 100th item should not affect the hardness of top-5 ranking by much. This intuition is well captured by our new metric in (27) in the main text. Observe that

$$(\Delta_K^*)^2 = \frac{(w_5^* - w_6^*)^2}{w_6^{*2}} \frac{w_6^*}{w_5^*} \left(\frac{1}{100} \sum_{i=1}^{99} \frac{w_i^*/w_5^*}{(1 + w_i^*/w_5^*)^2} + \frac{1}{100} \frac{w_{100}^*/w_5^*}{(1 + w_{100}^*/w_5^*)^2} \right).$$

Since w_{100}^*/w_5^* is exceedingly small (10^{-7} in this example), it is easily seen that $\frac{w_{100}^*/w_5^*}{(1 + w_{100}^*/w_5^*)^2} \lesssim w_{100}^*/w_5^*$ is also extremely small, and hence $(\Delta_K^*)^2$ is not changed by much compared with the case when the 100th item is absent ($(\Delta_K^*)^2 \approx 0.1107$ (resp. 0.1118) for the case when the 100th item is present (resp. absent)). Similarly, consider the case when $w_1^* = 10^6$, $w_2^* = \dots = w_5^* = 1$ and $w_6^* = \dots = w_{100}^* = 0.5$. As one can see, adding the first item will have little influence upon Δ_K^* .

- *Case 3:* Consider finding the top-5 items out of 100 items with $w_1^* = \dots = w_5^* = 10$, $w_6^* = \dots = w_{10}^* = 5$ and $w_{11}^* = \dots = w_{100}^* = 10^{-6}$. The sample complexity needed for exact top-5 recovery will surely increase since the comparisons between $\{1, \dots, 10\}$ and $\{11, \dots, 100\}$ are, with high probability, not useful in determining the relative strength within the group $\{1, \dots, 10\}$. This is also reflected in the generalized separation measure Δ_K^* . Recall that we have

$$(\Delta_K^*)^2 = \frac{(w_5^* - w_6^*)^2}{w_6^{*2}} \frac{w_6^*}{w_5^*} \left(\frac{1}{100} \sum_{i=1}^{10} \frac{w_i^*/w_5^*}{(1 + w_i^*/w_5^*)^2} + \frac{1}{100} \sum_{i=11}^{100} \frac{w_i^*/w_5^*}{(1 + w_i^*/w_5^*)^2} \right).$$

For each $11 \leq i \leq 100$, $\frac{w_i^*/w_5^*}{(1 + w_i^*/w_5^*)^2}$ is exceedingly small. This makes $(\Delta_K^*)^2$ much smaller compared with the case when only $\{1, \dots, 10\}$ are present ($(\Delta_K^*)^2 \approx 0.1118$ (resp. 0.1181) for the case when items $\{11, \dots, 100\}$ are present (resp. absent)). As a result, the required sample size increases accordingly.

C Proofs in Section 5

This section collects proofs of the theorems and lemmas that appear in Section 5.

Before moving on, we note that by Lemma 1 in the main text, the event

$$\mathcal{A}_0 = \left\{ \frac{1}{2}np \leq d_{\min} \leq d_{\max} \leq \frac{3}{2}np \right\}$$

happens with probability at least $1 - O(n^{-10})$. Throughout this section, we shall assume that we are on this event without explicitly referring to it each time. An immediate consequence is that $d_{\max} \leq d$ on this event.

C.1 Proof of Theorem 8

To begin with, we write

$$\boldsymbol{\pi}^\top - \hat{\boldsymbol{\pi}}^\top = \boldsymbol{\pi}^\top \mathbf{P} - \hat{\boldsymbol{\pi}}^\top \hat{\mathbf{P}} = \boldsymbol{\pi}^\top (\mathbf{P} - \hat{\mathbf{P}}) + (\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top \hat{\mathbf{P}}. \quad (49)$$

The last term of the above identity can be further decomposed as

$$(\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top \hat{\mathbf{P}} = (\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top \mathbf{P}^* + (\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top (\hat{\mathbf{P}} - \mathbf{P}^*)$$

$$= (\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top (\mathbf{P}^* - \mathbf{1}\boldsymbol{\pi}^{*\top}) + (\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top (\hat{\mathbf{P}} - \mathbf{P}^*), \quad (50)$$

where we have used the fact that $(\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top \mathbf{1}\boldsymbol{\pi}^{*\top} = \mathbf{0}$. Combining (49) and (50) we get

$$\boldsymbol{\pi}^\top - \hat{\boldsymbol{\pi}}^\top = \boldsymbol{\pi}^\top (\mathbf{P} - \hat{\mathbf{P}}) + (\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top (\mathbf{P}^* - \mathbf{1}\boldsymbol{\pi}^{*\top}) + (\boldsymbol{\pi} - \hat{\boldsymbol{\pi}})^\top (\hat{\mathbf{P}} - \mathbf{P}^*),$$

which together with a little algebra gives

$$\begin{aligned} \|\boldsymbol{\pi} - \hat{\boldsymbol{\pi}}\|_{\boldsymbol{\pi}^*} &\leq \left\| \boldsymbol{\pi}^\top (\mathbf{P} - \hat{\mathbf{P}}) \right\|_{\boldsymbol{\pi}^*} + \|\boldsymbol{\pi} - \hat{\boldsymbol{\pi}}\|_{\boldsymbol{\pi}^*} \|\mathbf{P}^* - \mathbf{1}\boldsymbol{\pi}^{*\top}\|_{\boldsymbol{\pi}^*} + \|\boldsymbol{\pi} - \hat{\boldsymbol{\pi}}\|_{\boldsymbol{\pi}^*} \|\hat{\mathbf{P}} - \mathbf{P}^*\|_{\boldsymbol{\pi}^*} \\ \implies \|\boldsymbol{\pi} - \hat{\boldsymbol{\pi}}\|_{\boldsymbol{\pi}^*} &\leq \frac{\left\| \boldsymbol{\pi}^\top (\mathbf{P} - \hat{\mathbf{P}}) \right\|_{\boldsymbol{\pi}^*}}{1 - \|\mathbf{P}^* - \mathbf{1}\boldsymbol{\pi}^{*\top}\|_{\boldsymbol{\pi}^*} - \|\hat{\mathbf{P}} - \mathbf{P}^*\|_{\boldsymbol{\pi}^*}}. \end{aligned}$$

The theorem follows by recognizing that

$$\|\mathbf{P}^* - \mathbf{1}\boldsymbol{\pi}^{*\top}\|_{\boldsymbol{\pi}^*} = \max\{\lambda_2(\mathbf{P}^*), -\lambda_n(\mathbf{P}^*)\}.$$

C.2 Proof of Theorem 9

By Theorem 8, we obtain

$$\begin{aligned} \|\boldsymbol{\pi}^* - \boldsymbol{\pi}\|_{\boldsymbol{\pi}^*} &\leq \frac{\left\| \boldsymbol{\pi}^{*\top} (\mathbf{P}^* - \mathbf{P}) \right\|_{\boldsymbol{\pi}^*}}{1 - \max\{\lambda_2(\mathbf{P}^*), -\lambda_n(\mathbf{P}^*)\} - \|\mathbf{P}^* - \mathbf{P}\|_{\boldsymbol{\pi}^*}} \stackrel{(i)}{\lesssim} \|\boldsymbol{\pi}^{*\top} (\mathbf{P}^* - \mathbf{P})\|_{\boldsymbol{\pi}^*} \\ &\stackrel{(ii)}{\lesssim} \|\boldsymbol{\pi}^{*\top} (\mathbf{P}^* - \mathbf{P})\|_2 \stackrel{(iii)}{\lesssim} \frac{1}{\sqrt{npL}} \|\boldsymbol{\pi}^*\|_2, \end{aligned}$$

where (i) is a consequence of Lemma 5, (ii) follows from the relationship between $\|\cdot\|_2$ and $\|\cdot\|_{\boldsymbol{\pi}^*}$, and (iii) follows as long as one can justify that

$$\|\boldsymbol{\pi}^{*\top} (\mathbf{P} - \mathbf{P}^*)\|_2 \lesssim \frac{1}{\sqrt{npL}} \|\boldsymbol{\pi}^*\|_2. \quad (51)$$

Therefore, the rest of the proof is devoted to establishing (51). To simplify the notations hereafter, we denote $\boldsymbol{\Delta} := \mathbf{P} - \mathbf{P}^*$. In fact, it is easy to check that for any $i \neq j$,

$$\Delta_{i,j} = P_{i,j} - P_{i,j}^* = \mathbf{1}_{\{(i,j) \in \mathcal{E}\}} \frac{1}{d} \left(\frac{1}{L} \sum_{l=1}^L y_{i,j}^{(l)} - y_{i,j}^* \right) \quad (52)$$

and for $1 \leq i \leq n$, one has

$$\begin{aligned} \Delta_{i,i} &= P_{i,i} - P_{i,i}^* = \left(1 - \sum_{j:j \neq i} P_{i,j} \right) - \left(1 - \sum_{j:j \neq i} P_{i,j}^* \right) \\ &= - \sum_{j:j \neq i} (P_{i,j} - P_{i,j}^*) = - \sum_{j:j \neq i} \Delta_{i,j} \\ &= -\Delta_{i,i}^{\text{lower}} - \Delta_{i,i}^{\text{upper}}. \end{aligned} \quad (53)$$

where

$$\Delta_{i,i}^{\text{lower}} := \sum_{j:j < i} \Delta_{i,j} \quad \text{and} \quad \Delta_{i,i}^{\text{upper}} := \sum_{j:j > i} \Delta_{i,j}.$$

Towards proving (51), we decompose $\boldsymbol{\Delta}$ into four parts $\boldsymbol{\Delta} = \boldsymbol{\Delta}_{\text{lower}} + \boldsymbol{\Delta}_{\text{upper}} + \boldsymbol{\Delta}_{\text{diag,l}} + \boldsymbol{\Delta}_{\text{diag,u}}$, where $\boldsymbol{\Delta}_{\text{lower}}$ is the lower triangular part (excluding the diagonal) of $\boldsymbol{\Delta}$, $\boldsymbol{\Delta}_{\text{upper}}$ is the upper triangular part, and

$$\boldsymbol{\Delta}_{\text{diag,l}} = -\text{diag}\left([\Delta_{i,i}^{\text{lower}}]_{1 \leq i \leq n}\right) \quad \text{and} \quad \boldsymbol{\Delta}_{\text{diag,u}} = -\text{diag}\left([\Delta_{i,i}^{\text{upper}}]_{1 \leq i \leq n}\right).$$

The triangle inequality then gives

$$\|\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}\|_2 \leq \underbrace{\|\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{lower}}\|_2}_{:=I_{\text{lower}}} + \underbrace{\|\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{upper}}\|_2}_{:=I_{\text{upper}}} + \underbrace{\|\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{diag,l}}\|_2}_{:=I_{\text{diag,l}}} + \underbrace{\|\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{diag,u}}\|_2}_{:=I_{\text{diag,u}}}.$$

In what follows, we will focus on controlling the first term I_{lower} . The other three terms can be bounded using nearly identical arguments.

Note that the j component of $\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{lower}}$ can be expressed as

$$[\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{lower}}]_j = \sum_{i:i>j} \pi_i^* \Delta_{i,j}.$$

Recall that for any pair $(i, j) \in \mathcal{E}$, $\Delta_{i,j}$ is a sum of L independent zero-mean random variables, and hence $[\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{lower}}]_j$ is a sum of Ld_j^{lower} independent zero-mean random variables, where

$$d_j^{\text{lower}} := |\{(i, j) \mid (i, j) \in \mathcal{E} \text{ and } i > j\}|.$$

In view of Hoeffding's inequality (Lemma 18), one has, when conditional on \mathcal{G} , that

$$\mathbb{P}\left(\left|[\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{lower}}]_j\right| \geq t\right) \leq 2 \exp\left(-\frac{2t^2}{\frac{1}{(Ld)^2} d_j^{\text{lower}} L (2\|\boldsymbol{\pi}^*\|_\infty)^2}\right).$$

Hence $[\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{lower}}]_j$ can be treated as a sub-Gaussian random variable with variance proxy

$$\sigma^2 \asymp \frac{d_j^{\text{lower}}}{d^2 L} \|\boldsymbol{\pi}^*\|_\infty^2 \lesssim \frac{1}{dL} \|\boldsymbol{\pi}^*\|_\infty^2.$$

Given that the entries of $\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{lower}}$ are independent, we see that

$$(I_{\text{lower}})^2 = \sum_{j=1}^n [\boldsymbol{\pi}^{*\top} \boldsymbol{\Delta}_{\text{lower}}]_j^2,$$

is a quadratic form of a sub-Gaussian vector. On the one hand, $\mathbb{E}[I_{\text{lower}}]^2 \lesssim n\sigma^2$. On the other hand, we invoke (Rudelson et al., 2013, Theorem 1.1) to reach

$$\mathbb{P}\{I_{\text{lower}}^2 - \mathbb{E}[I_{\text{lower}}^2] \geq t\} \leq 2 \exp\left\{-c \min\left\{\frac{t^2}{n\sigma^4}, \frac{t}{\sigma^2}\right\}\right\}$$

for some constant $c > 0$. By choosing $t \asymp \sigma^2 \sqrt{n \log n}$, we see that with probability at least $1 - O(n^{-10})$,

$$I_{\text{lower}}^2 \lesssim \mathbb{E}[I_{\text{lower}}^2] + \sigma^2 \sqrt{n \log n} \lesssim n\sigma^2 \lesssim \frac{n}{dL} \|\boldsymbol{\pi}^*\|_\infty^2 \lesssim \frac{1}{dL} \|\boldsymbol{\pi}^*\|_2^2.$$

The same upper bounds can be derived for other terms using the same arguments. We have thus established (51) by recognizing that $d \gtrsim np$.

C.3 Proof of Lemma 2

Observe that

$$\begin{aligned} I_1^m &= \sum_{j:j \neq m} \pi_j^* (P_{j,m} - P_{j,m}^*) + \pi_m^* (P_{m,m} - P_{m,m}^*) \\ &\stackrel{(i)}{=} \sum_{j:j \neq m} \pi_j^* (P_{j,m} - P_{j,m}^*) + \pi_m^* \left\{ \left(1 - \sum_{j:j \neq m} P_{m,j}\right) - \left(1 - \sum_{j:j \neq m} P_{m,j}^*\right) \right\} \\ &= \sum_{j:j \neq m} (\pi_j^* + \pi_m^*) (P_{j,m} - P_{j,m}^*) \end{aligned}$$

$$= \frac{1}{Ld} \sum_{j:j \neq m} \sum_{l=1}^L (\pi_j^* + \pi_m^*) \mathbb{1}_{(j,m) \in \mathcal{E}} \left(y_{j,m}^{(l)} - y_{j,m}^* \right), \quad (54)$$

where (i) follows from the fact that \mathbf{P} and \mathbf{P}^* are both probability transition matrices. By Lemma 18, one can derive

$$\begin{aligned} \mathbb{P}\left\{ |I_1^m| \geq t \mid \mathcal{G} \right\} &= \mathbb{P}\left\{ \left| \sum_{j:j \neq m} \sum_{l=1}^L (\pi_j^* + \pi_m^*) \mathbb{1}_{(j,m) \in \mathcal{E}} \left(y_{j,m}^{(l)} - y_{j,m}^* \right) \right| \geq Ldt \mid \mathcal{G} \right\} \\ &\leq 2 \exp\left(-\frac{2(Ldt)^2}{Ld_{\max}(2\|\boldsymbol{\pi}^*\|_{\infty})^2} \right). \end{aligned}$$

When $d_{\max} \leq d$, the right hand side is bounded by $2 \exp\left(\frac{-Ldt^2}{2\|\boldsymbol{\pi}^*\|_{\infty}^2} \right)$. Hence

$$\mathbb{P}\left\{ |I_1^m| \geq 4\sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_{\infty} \mid d_{\max} \leq d \right\} \leq 2n^{-8}.$$

The lemma is established by taking the union bounds and using the fact that $\mathbb{P}(d_{\max} \leq d) \geq 1 - O(n^{-10})$ (by Lemma 1 in the main text and the remarks after that).

C.4 Proof of Lemma 3

Similar to the proof of Lemma 2 above, by applying Lemma 18 to the quantity

$$P_{m,m} - P_{m,m}^* = - \sum_{j:j \neq m} (P_{m,j} - P_{m,j}^*) = -\frac{1}{Ld} \sum_{j:j \neq m} \sum_{l=1}^L \mathbb{1}_{(j,m) \in \mathcal{E}} \left(y_{j,m}^{(l)} - y_{j,m}^* \right),$$

we get

$$\mathbb{P}\left\{ \max_m |P_{m,m} - P_{m,m}^*| \geq 2\sqrt{\frac{\log n}{Ld}} \mid d_{\max} \leq d \right\} \leq 2n^{-7}.$$

On the other hand, when the event \mathcal{A}_0 (defined in Lemma 1 in the main text) happens, we have $d_{\min} \geq np/2$ and for all $1 \leq m \leq n$,

$$P_{m,m}^* = 1 - \sum_{j:j \neq m} P_{m,j}^* \leq 1 - \frac{d_{\min}}{d} \cdot \frac{1}{1 + \kappa} \leq 1 - \frac{np}{2d} \cdot \frac{1}{1 + \kappa}.$$

Combining these two pieces completes the proof.

C.5 Proof of Lemma 4

First, by the relationship between $\|\cdot\|_2$ and $\|\cdot\|_{\boldsymbol{\pi}^*}$, we have

$$\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 \leq \frac{1}{\sqrt{\pi_{\min}^*}} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_{\boldsymbol{\pi}^*},$$

where $\pi_{\min}^* := \min_i \pi_i^*$. Invoking Theorem 8, we obtain

$$\begin{aligned} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_{\boldsymbol{\pi}^*} &\leq \frac{\|\boldsymbol{\pi}^{(m)\top} (\mathbf{P}^{(m)} - \mathbf{P})\|_{\boldsymbol{\pi}^*}}{1 - \max\{\lambda_2(\mathbf{P}^*), -\lambda_n(\mathbf{P}^*)\} - \|\mathbf{P} - \mathbf{P}^*\|_{\boldsymbol{\pi}^*}} \\ &\leq \frac{1}{\gamma} \sqrt{\pi_{\max}^*} \|\boldsymbol{\pi}^{(m)\top} (\mathbf{P}^{(m)} - \mathbf{P})\|_2, \end{aligned}$$

where we define $\gamma := 1 - \max\{\lambda_2(\mathbf{P}^*), -\lambda_n(\mathbf{P}^*)\} - \|\mathbf{P} - \mathbf{P}^*\|_{\boldsymbol{\pi}^*}$ and $\pi_{\max}^* := \max_i \pi_i^*$.

To facilitate the analysis of $\|\boldsymbol{\pi}^{(m)\top}(\mathbf{P}^{(m)} - \mathbf{P})\|_2$, we introduce another Markov chain with transition probability matrices $\mathbf{P}^{(m),\mathcal{G}}$, which is also a leave-one-out version of the transition matrix \mathbf{P} . Similar to $\mathbf{P}^{(m)}$, $\mathbf{P}^{(m),\mathcal{G}}$ replaces all the transition probabilities involving the m -th item with their expected values (conditional on \mathcal{G}). Concretely, for $i \neq j$,

$$P_{i,j}^{(m),\mathcal{G}} = \begin{cases} P_{i,j}, & i \neq m, j \neq m, \\ \frac{1}{d} y_{i,j}^* \mathbb{1}_{(i,j) \in \mathcal{E}}, & i = m \text{ or } j = m. \end{cases}$$

And for each $1 \leq i \leq n$, we define

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

to make $\mathbf{P}^{(m),\mathcal{G}}$ a valid probability transition matrix. Hence by the triangle inequality, we see that

$$\|\boldsymbol{\pi}^{(m)\top}(\mathbf{P}^{(m)} - \mathbf{P})\|_2 \leq \underbrace{\|\boldsymbol{\pi}^{(m)\top}(\mathbf{P} - \mathbf{P}^{(m),\mathcal{G}})\|_2}_{:=J_1^m} + \underbrace{\|\boldsymbol{\pi}^{(m)\top}(\mathbf{P}^{(m)} - \mathbf{P}^{(m),\mathcal{G}})\|_2}_{:=J_2^m}.$$

The next step is then to bound J_1^m and J_2^m separately.

For J_1^m , similar to (54), one has

$$\begin{aligned} \left[\boldsymbol{\pi}^{(m)\top}(\mathbf{P} - \mathbf{P}^{(m),\mathcal{G}}) \right]_m &\stackrel{(i)}{=} \left[\boldsymbol{\pi}^{(m)\top}(\mathbf{P} - \mathbf{P}^*) \right]_m \\ &= \frac{1}{Ld} \sum_{j:j \neq m} \sum_{l=1}^L \left(\pi_j^{(m)} + \pi_m^{(m)} \right) \mathbb{1}_{(j,m) \in \mathcal{E}} \left(y_{j,m}^{(l)} - y_{j,m}^* \right), \end{aligned}$$

where (i) comes from the fact that $\mathbf{P}_m^{(m),\mathcal{G}} = \mathbf{P}_m^*$. Recognizing that $\boldsymbol{\pi}^{(m)}$ is statistically independent of $\{y_{j,m}\}_{j \neq m}$, by Hoeffding's inequality in Lemma 18, we get

$$\mathbb{P} \left\{ \left| \left[\boldsymbol{\pi}^{(m)\top}(\mathbf{P} - \mathbf{P}^{(m),\mathcal{G}}) \right]_m \right| \geq 4\sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^{(m)}\|_\infty \right\} \leq 2n^{-8}. \quad (55)$$

And for $j \neq m$, we have

$$\begin{aligned} \left[\boldsymbol{\pi}^{(m)\top}(\mathbf{P} - \mathbf{P}^{(m),\mathcal{G}}) \right]_j &= \sum_i \pi_i^{(m)} (P_{i,j} - P_{i,j}^{(m),\mathcal{G}}) \\ &= \pi_j^{(m)} (P_{j,j} - P_{j,j}^{(m),\mathcal{G}}) + \pi_m^{(m)} (P_{m,j} - P_{m,j}^{(m),\mathcal{G}}) \\ &= \pi_j^{(m)} \frac{1}{d} (y_{j,m}^* - y_{j,m}) \mathbb{1}_{(j,m) \in \mathcal{E}} + \pi_m^{(m)} \frac{1}{d} (y_{m,j} - y_{m,j}^*) \mathbb{1}_{(j,m) \in \mathcal{E}} \\ &= \pi_j^{(m)} (P_{j,m}^* - P_{j,m}) + \pi_m^{(m)} (P_{m,j} - P_{m,j}^*). \end{aligned}$$

In addition by Hoeffding's inequality in Lemma 18, we have

$$\max_{j \neq m} |P_{j,m} - P_{j,m}^*| \leq \frac{2}{d} \sqrt{\frac{\log n}{L}}$$

with probability at least $1 - O(n^{-5})$. As a consequence,

$$\left| \left[\boldsymbol{\pi}^{(m)\top}(\mathbf{P} - \mathbf{P}^{(m),\mathcal{G}}) \right]_j \right| \leq \begin{cases} \frac{4}{d} \sqrt{\frac{\log n}{L}} \|\boldsymbol{\pi}^{(m)}\|_\infty, & \text{if } (j, m) \in \mathcal{E}, \\ 0, & \text{else.} \end{cases} \quad (56)$$

Combining (55) and (56) yields

$$J_1^m \leq 4\sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^{(m)}\|_\infty + \frac{4\sqrt{d_{\max} - 1}}{d} \sqrt{\frac{\log n}{L}} \|\boldsymbol{\pi}^{(m)}\|_\infty \stackrel{(i)}{\leq} 8\sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^{(m)}\|_\infty,$$

where (i) comes from the fact that $d_{\max} \leq d$.

Regarding J_2^m , we invoke the identity $\boldsymbol{\pi}^{*\top}(\mathbf{P}^{(m)} - \mathbf{P}^{(m),\mathcal{G}}) = \mathbf{0}$ to get

$$\boldsymbol{\pi}^{(m)\top}(\mathbf{P}^{(m)} - \mathbf{P}^{(m),\mathcal{G}}) = \left(\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\right)^\top (\mathbf{P}^{(m)} - \mathbf{P}^{(m),\mathcal{G}}).$$

Therefore, for $j \neq m$ we have

$$\begin{aligned} & \left[\left(\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\right)^\top (\mathbf{P}^{(m)} - \mathbf{P}^{(m),\mathcal{G}}) \right]_j \\ &= \sum_i (\pi_i^{(m)} - \pi_i^*) (P_{i,j}^{(m)} - P_{i,j}^{(m),\mathcal{G}}) \\ &= (\pi_j^{(m)} - \pi_j^*) (P_{j,j}^{(m)} - P_{j,j}^{(m),\mathcal{G}}) + (\pi_m^{(m)} - \pi_m^*) (P_{m,j}^{(m)} - P_{m,j}^{(m),\mathcal{G}}) \\ &= -(\pi_j^{(m)} - \pi_j^*) (P_{j,m}^{(m)} - P_{j,m}^{(m),\mathcal{G}}) + (\pi_m^{(m)} - \pi_m^*) (P_{m,j}^{(m)} - P_{m,j}^{(m),\mathcal{G}}). \end{aligned}$$

Recognizing that $|P_{j,m}^{(m)} - P_{j,m}^{(m),\mathcal{G}}| \leq \frac{2}{d}$ for $(j, m) \in \mathcal{E}$ and $|P_{j,m}^{(m)} - P_{j,m}^{(m),\mathcal{G}}| \leq \frac{p}{d}$ for $(j, m) \notin \mathcal{E}$, we have

$$\left| \left[\left(\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\right)^\top (\mathbf{P}^{(m)} - \mathbf{P}^{(m),\mathcal{G}}) \right]_j \right| \leq \begin{cases} \frac{4}{d} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty, & \text{if } (j, m) \in \mathcal{E}, \\ \frac{2p}{d} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty, & \text{else.} \end{cases} \quad (57)$$

And for $j = m$, it holds that

$$\begin{aligned} & \left| \left[\left(\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\right)^\top (\mathbf{P}^{(m)} - \mathbf{P}^{(m),\mathcal{G}}) \right]_m \right| \\ &= \left| (\pi_m^{(m)} - \pi_m^*) (P_{m,m}^{(m)} - P_{m,m}^{(m),\mathcal{G}}) + \sum_{j:j \neq m} (\pi_j^{(m)} - \pi_j^*) (P_{j,m}^{(m)} - P_{j,m}^{(m),\mathcal{G}}) \right| \end{aligned} \quad (58)$$

$$\begin{aligned} & \leq \left| (\pi_m^{(m)} - \pi_m^*) (P_{m,m}^{(m)} - P_{m,m}^{(m),\mathcal{G}}) \right| + \left| \sum_{j:j \neq m} (\pi_j^{(m)} - \pi_j^*) (P_{j,m}^{(m)} - P_{j,m}^{(m),\mathcal{G}}) \right| \\ &= \underbrace{\left| \sum_{j:j \neq m} (\pi_m^{(m)} - \pi_m^*) (P_{m,j}^{(m)} - P_{m,j}^{(m),\mathcal{G}}) \right|}_{:=J_3^m} + \underbrace{\left| \sum_{j:j \neq m} (\pi_j^{(m)} - \pi_j^*) (P_{j,m}^{(m)} - P_{j,m}^{(m),\mathcal{G}}) \right|}_{:=J_4^m}. \end{aligned} \quad (59)$$

Given that $P_{m,j}^{(m)} - P_{m,j}^{(m),\mathcal{G}} = \frac{y_{m,j}^*}{d} (p - \mathbb{1}_{(m,j) \in \mathcal{E}})$, we have

$$J_3^m = \sum_{j:j \neq m} (\pi_m^{(m)} - \pi_m^*) \underbrace{\frac{y_{m,j}^*}{d} (p - \mathbb{1}_{(m,j) \in \mathcal{E}})}_{:=\xi_j^{(m)}}. \quad (60)$$

Since $\|\boldsymbol{\xi}^{(m)}\|_\infty \leq \frac{1}{d} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty$ and $\|\boldsymbol{\xi}^{(m)}\|_2 \leq \frac{1}{d} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_2$, Lemma 19 implies that

$$|J_3^m| \lesssim \frac{\sqrt{np \log n} + \log n}{d} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty$$

with high probability. The same bound holds for J_4^m . Combine (57), (59) and (60) to arrive at

$$J_2^m \lesssim \left(\frac{\sqrt{np \log n} + \log n}{d} + \frac{p\sqrt{n}}{d} + \frac{\sqrt{d}}{d} \right) \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty.$$

Combining all, we deduce that

$$\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 \leq \frac{\sqrt{\kappa}}{\gamma} (J_1^m + J_2^m)$$

$$\begin{aligned}
&\leq \frac{\sqrt{\kappa}}{\gamma} \left(8\sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^{(m)}\|_\infty + C \left(\frac{\sqrt{np \log n} + \log n}{d} + \frac{p\sqrt{n}}{d} + \frac{\sqrt{d}}{d} \right) \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty \right) \\
&\leq \frac{\sqrt{\kappa}}{\gamma} \left(8\sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_\infty + C \left(8\sqrt{\frac{\log n}{Ld}} + \frac{\sqrt{np \log n} + \log n}{d} + \frac{p\sqrt{n}}{d} + \frac{\sqrt{d}}{d} \right) \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty \right) \\
&\stackrel{(i)}{\leq} \frac{8\sqrt{\kappa}}{\gamma} \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_\infty + \frac{1}{2} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty,
\end{aligned}$$

where (i) holds as long as $np\gamma^2 \geq c\kappa \log n$ for c sufficiently large. The triangle inequality

$$\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty \leq \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 + \|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty$$

yields

$$\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 \leq \frac{16\sqrt{\kappa}}{\gamma} \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_\infty + \|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_\infty, \quad (61)$$

which concludes the proof.

C.6 Proof of Lemma 6

For ease of presentation, we define

$$\tilde{y}_{i,j} := \frac{1}{L} \sum_{l=1}^L \tilde{y}_{i,j}^{(l)}$$

for all $i \neq j$, where

$$\tilde{y}_{i,j}^{(l)} \stackrel{\text{ind.}}{=} \begin{cases} 1, & \text{with probability } \frac{w_j^*}{w_i^* + w_j^*}, \\ 0, & \text{else.} \end{cases}$$

This allows us to write \boldsymbol{y} as $y_{i,j} = \tilde{y}_{i,j} \mathbf{1}_{(i,j) \in \mathcal{E}}$ for all $i \neq j$. With this notation in place, we can obtain

$$I_4^m = \sum_{j:j \neq m} (\pi_j^{(m)} - \pi_j^*) \left(\frac{1}{Ld} \sum_{l=1}^L \tilde{y}_{j,m}^{(l)} \right) \mathbf{1}_{(j,m) \in \mathcal{E}}.$$

We can further decompose I_4^m into

$$I_4^m = \mathbb{E}[I_4^m | \mathcal{G}_{-m}, \tilde{\boldsymbol{y}}] + I_4^m - \mathbb{E}[I_4^m | \mathcal{G}_{-m}, \tilde{\boldsymbol{y}}],$$

where \mathcal{G}_{-m} represent the graph without the m -th node, and $\tilde{\boldsymbol{y}} = \{\tilde{y}_{i,j} | i \neq j\}$ represents all the binary outcomes.

We start with the expectation term

$$\begin{aligned}
\mathbb{E}[I_4^m | \mathcal{G}_{-m}, \tilde{\boldsymbol{y}}] &= \sum_{j:j \neq m} (\pi_j^{(m)} - \pi_j^*) \left(\frac{1}{Ld} \sum_{l=1}^L \tilde{y}_{j,m}^{(l)} \right) \mathbb{P}\{(j,m) \in \mathcal{E}\} \\
&\stackrel{(i)}{\leq} p \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_2 \frac{\sqrt{n}}{d} \stackrel{(ii)}{\leq} \frac{1}{2\sqrt{n}} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_2 \\
&\stackrel{(iii)}{\leq} \frac{1}{2\sqrt{n}} (\|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 + \|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_2),
\end{aligned}$$

where (i) comes from the Cauchy-Schwarz inequality, (ii) follows from the choice $d = c_d np \geq 2np$ and (iii) results from the triangle inequality. By Theorem 9, with high probability we have

$$\|\boldsymbol{\pi} - \boldsymbol{\pi}^*\|_2 \leq C_N \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_2 \leq C_N \sqrt{\frac{\log n}{Ld}} \sqrt{n} \|\boldsymbol{\pi}^*\|_\infty,$$

thus indicating that

$$\mathbb{E}[I_4^m \mid \mathcal{G}_{-m}, \tilde{\mathbf{y}}] \leq \frac{1}{2\sqrt{n}} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}\|_2 + \frac{C_N}{2} \sqrt{\frac{\log n}{Ld}} \|\boldsymbol{\pi}^*\|_\infty. \quad (62)$$

When it comes to the fluctuation term, one can write

$$I_4^m - \mathbb{E}[I_4^m \mid \mathcal{G}_{-m}, \tilde{\mathbf{y}}] = \sum_{j:j \neq m} \underbrace{(\pi_j^{(m)} - \pi_j^*)}_{:=\beta_j^{(m)}} \left(\frac{1}{Ld} \sum_{l=1}^L \tilde{y}_{j,m}^{(l)} \right) (\mathbf{1}_{(j,m) \in \mathcal{E}} - \mathbb{P}\{(j,m) \in \mathcal{E}\}).$$

Since $\|\beta^{(m)}\|_2 \leq \frac{1}{d} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_2$ and $\|\beta^{(m)}\|_\infty \leq \frac{1}{d} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty$, one can apply Lemma 19 to derive

$$|I_4^m - \mathbb{E}[I_4^m \mid \mathcal{G}_{-m}, \tilde{\mathbf{y}}]| \lesssim \frac{\sqrt{np \log n} + \log n}{d} \|\boldsymbol{\pi}^{(m)} - \boldsymbol{\pi}^*\|_\infty \quad (63)$$

with high enough probability. The bounds (62) and (63) taken together complete the proof.

D Proofs in Section 6

This section gathers the proofs of the lemmas in Section 6.

D.1 Proof of Lemma 7

Observe that

$$\begin{aligned} \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) &= \lambda \boldsymbol{\theta}^* + \sum_{(i,j) \in \mathcal{E}, i > j} \left(-y_{j,i} + \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_j^*}} \right) (\mathbf{e}_i - \mathbf{e}_j) \\ &= \lambda \boldsymbol{\theta}^* + \frac{1}{L} \sum_{(i,j) \in \mathcal{E}, i > j} \sum_{l=1}^L \underbrace{\left(-y_{j,i}^{(l)} + \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_j^*}} \right)}_{:=\mathbf{z}_{i,j}^{(l)}} (\mathbf{e}_i - \mathbf{e}_j). \end{aligned}$$

It is seen that $\mathbb{E}[\mathbf{z}_{i,j}^{(l)}] = \mathbf{0}$, $\|\mathbf{z}_{i,j}^{(l)}\|_2 \leq \sqrt{2}$,

$$\mathbb{E}[\mathbf{z}_{i,j}^{(l)} \mathbf{z}_{i,j}^{(l)\top}] = \text{Var}[y_{i,j}^{(l)}] (\mathbf{e}_i - \mathbf{e}_j) (\mathbf{e}_i - \mathbf{e}_j)^\top \preceq (\mathbf{e}_i - \mathbf{e}_j) (\mathbf{e}_i - \mathbf{e}_j)^\top$$

$$\text{and} \quad \mathbb{E}[\mathbf{z}_{i,j}^{(l)\top} \mathbf{z}_{i,j}^{(l)}] = \text{Tr} \left(\mathbb{E}[\mathbf{z}_{i,j}^{(l)} \mathbf{z}_{i,j}^{(l)\top}] \right) \leq 2.$$

This implies that with high probability (note that the randomness comes from \mathcal{G}),

$$\left\| \sum_{(i,j) \in \mathcal{E}, i > j} \sum_{l=1}^L \mathbb{E}[\mathbf{z}_{i,j}^{(l)} \mathbf{z}_{i,j}^{(l)\top}] \right\| \leq L \left\| \sum_{(i,j) \in \mathcal{E}, i > j} (\mathbf{e}_i - \mathbf{e}_j) (\mathbf{e}_i - \mathbf{e}_j)^\top \right\| = L \|\mathbf{L}_\mathcal{G}\| \lesssim Lnp$$

and

$$\left| \sum_{(i,j) \in \mathcal{E}, i > j} \sum_{l=1}^L \mathbb{E}[\mathbf{z}_{i,j}^{(l)\top} \mathbf{z}_{i,j}^{(l)}] \right| \leq 2L \left| \sum_{(i,j) \in \mathcal{E}, i > j} 1 \right| \lesssim Ln^2p.$$

Letting $V := \frac{1}{L^2} \max \left\{ \left\| \sum_{(i,j) \in \mathcal{E}} \sum_{l=1}^L \mathbb{E}[\mathbf{z}_{i,j}^{(l)} \mathbf{z}_{i,j}^{(l)\top}] \right\|, \left| \sum_{(i,j) \in \mathcal{E}} \sum_{l=1}^L \mathbb{E}[\mathbf{z}_{i,j}^{(l)\top} \mathbf{z}_{i,j}^{(l)}] \right| \right\}$ and $B := \max_{i,j,l} \|\mathbf{z}_{i,j}^{(l)}\|$, we can invoke the matrix Bernstein inequality (Tropp, 2012, Theorem 1.6) to reach

$$\|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) - \mathbb{E}[\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) \mid \mathcal{G}]\|_2 \lesssim \sqrt{V \log n} + B \log n \lesssim \sqrt{\frac{n^2 p \log n}{L}} + \frac{\log n}{L}$$

$$\lesssim \sqrt{\frac{n^2 p \log n}{L}}$$

with probability at least $1 - O(n^{-10})$. Combining this with the identity $\mathbb{E}[\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) \mid \mathcal{G}] = \lambda \boldsymbol{\theta}^*$ yields

$$\begin{aligned} \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y})\|_2 &\leq \|\mathbb{E}[\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) \mid \mathcal{G}]\|_2 + \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) - \mathbb{E}[\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) \mid \mathcal{G}]\|_2 \\ &\lesssim \lambda \|\boldsymbol{\theta}^*\|_2 + \sqrt{\frac{n^2 p \log n}{L}} \asymp \sqrt{\frac{n^2 p \log n}{L}}, \end{aligned}$$

where the last relation holds because of the facts that $\|\boldsymbol{\theta}^*\|_2 \leq \sqrt{n} \log \kappa$, $\lambda \asymp \frac{1}{\log \kappa} \sqrt{\frac{np \log n}{L}}$, and

$$\lambda \|\boldsymbol{\theta}^*\|_2 \lesssim \frac{\|\boldsymbol{\theta}^*\|_2}{\log \kappa} \sqrt{\frac{np \log n}{L}} \lesssim \sqrt{\frac{n^2 p \log n}{L}}.$$

This concludes the proof.

D.2 Proof of Lemma 9

It suffices to prove that

$$\min_{1 \leq i, j \leq n} \frac{e^{\theta_i} e^{\theta_j}}{(e^{\theta_i} + e^{\theta_j})^2} \geq \frac{1}{4\kappa e^{2C}}$$

for all $\boldsymbol{\theta} \in \mathbb{R}^n$ obeying $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \leq C$. Without loss of generality, suppose $\theta_i \leq \theta_j$. One can divide both the denominator and the numerator by $e^{2\theta_j}$ to obtain

$$\frac{e^{\theta_i} e^{\theta_j}}{(e^{\theta_i} + e^{\theta_j})^2} = \frac{e^{\theta_i - \theta_j}}{(1 + e^{\theta_i - \theta_j})^2} = \frac{e^{-|\theta_i - \theta_j|}}{(1 + e^{-|\theta_i - \theta_j|})^2} \geq \frac{1}{4} e^{-|\theta_i - \theta_j|},$$

where the last relation holds since $(1 + e^{-x})^2 \leq 4$ for all $x \geq 0$. From our assumption, we see that for all $1 \leq i, j \leq n$,

$$|\theta_i - \theta_j| \leq \theta_{\max}^* - \theta_{\min}^* + 2\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_\infty \leq \log \kappa + 2C,$$

which relies on the fact that $\theta_{\max}^* - \theta_{\min}^* \leq \log \kappa$. This allows one to justify that

$$\frac{e^{\theta_i} e^{\theta_j}}{(e^{\theta_i} + e^{\theta_j})^2} \geq \frac{1}{4} e^{-|\theta_i - \theta_j|} \geq \frac{1}{4\kappa e^{2C}}.$$

D.3 Proof of Fact 1

By $\boldsymbol{\theta}^0 = \boldsymbol{\theta}^*$ and $\mathbf{1}^\top \boldsymbol{\theta}^* = 0$, the statement trivially holds true for $t = 0$. Suppose it is true for some $t \geq 0$. Then

$$\begin{aligned} \mathbf{1}^\top \boldsymbol{\theta}^{t+1} &= \mathbf{1}^\top (\boldsymbol{\theta}^t - \eta_t \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^t; \mathbf{y})) \stackrel{(i)}{=} -\eta_t \mathbf{1}^\top \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^t; \mathbf{y}) \\ &= -\eta_t \mathbf{1}^\top (\nabla \mathcal{L}(\boldsymbol{\theta}^t; \mathbf{y}) + \lambda \boldsymbol{\theta}^t) \stackrel{(ii)}{=} -\eta_t \mathbf{1}^\top \nabla \mathcal{L}(\boldsymbol{\theta}^t; \mathbf{y}) \stackrel{(iii)}{=} 0, \end{aligned}$$

where the equalities (i) and (ii) follow from the fact that $\mathbf{1}^\top \boldsymbol{\theta}^t = 0$, whereas the last identity (iii) arises from the gradient expression (31) and the simple fact that $\mathbf{1}^\top (\mathbf{e}_i - \mathbf{e}_j) = 0$ for any i and j . This completes the whole proof.

D.4 Proof of Lemma 12

It follows from the optimality of $\boldsymbol{\theta}$ as well as the mean value theorem that

$$\mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) \geq \mathcal{L}_\lambda(\boldsymbol{\theta}; \mathbf{y}) = \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}) + \langle \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}), \boldsymbol{\theta} - \boldsymbol{\theta}^* \rangle$$

$$+ \frac{1}{2} (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^\top \nabla^2 \mathcal{L}_\lambda(\tilde{\boldsymbol{\theta}}; \mathbf{y}) (\boldsymbol{\theta} - \boldsymbol{\theta}^*),$$

where $\tilde{\boldsymbol{\theta}}$ is between $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^*$. This together with the Cauchy-Schwarz inequality gives

$$\begin{aligned} \frac{1}{2} (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^\top \nabla^2 \mathcal{L}_\lambda(\tilde{\boldsymbol{\theta}}; \mathbf{y}) (\boldsymbol{\theta} - \boldsymbol{\theta}^*) &\leq -\langle \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y}), \boldsymbol{\theta} - \boldsymbol{\theta}^* \rangle \\ &\leq \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y})\|_2 \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2. \end{aligned}$$

The above inequality gives

$$\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq \frac{2 \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y})\|_2}{\lambda_{\min}(\nabla^2 \mathcal{L}_\lambda(\tilde{\boldsymbol{\theta}}; \mathbf{y}))}. \quad (64)$$

From the trivial lower bound $\lambda_{\min}(\nabla^2 \mathcal{L}_\lambda(\tilde{\boldsymbol{\theta}}; \mathbf{y})) \geq \lambda$, the preceding inequality gives

$$\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq \frac{2 \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y})\|_2}{\lambda}. \quad (65)$$

On the event $\mathcal{A}_2 = \left\{ \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*; \mathbf{y})\|_2 \lesssim \sqrt{\frac{n^2 p \log n}{L}} \right\}$ and in the presence of the choice $\lambda \asymp \frac{1}{\log \kappa} \sqrt{\frac{np \log n}{L}}$, we obtain $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq c_2 \log \kappa \sqrt{n}$ for some constant $c_2 > 0$.

D.5 Proof of Lemma 13

In regard to the first consequence, one can apply the triangle inequality to show

$$\begin{aligned} \max_{1 \leq m \leq n} \|\boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^*\|_\infty &\leq \max_{1 \leq m \leq n} \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^{t,(m)}\|_2 + \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_\infty \\ &\leq C_3 \kappa \sqrt{\frac{\log n}{npL}} + C_4 \kappa^2 \sqrt{\frac{\log n}{npL}} \\ &\leq C_5 \kappa^2 \sqrt{\frac{\log n}{npL}}, \end{aligned}$$

as long as $C_5 \geq C_4 + C_3$. Similarly, for the second one, we have

$$\begin{aligned} \max_{1 \leq m \leq n} \|\boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^*\|_2 &\leq \max_{1 \leq m \leq n} \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^{t,(m)}\|_2 + \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_2 \\ &\leq C_3 \kappa \sqrt{\frac{\log n}{npL}} + C_1 \kappa \sqrt{\frac{\log n}{pL}} \\ &\leq C_6 \kappa \sqrt{\frac{\log n}{pL}}, \end{aligned}$$

as soon as $C_6 \geq C_1 + C_3$.

D.6 Proof of Lemma 14

In view of the gradient update rule (35), we have

$$\begin{aligned} \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^* &= \boldsymbol{\theta}^t - \eta \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^t) - \boldsymbol{\theta}^* \\ &= \boldsymbol{\theta}^t - \eta \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^t) - [\boldsymbol{\theta}^* - \eta \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*)] - \eta \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*) \\ &= \left\{ \mathbf{I}_n - \eta \int_0^1 \nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}(\tau)) d\tau \right\} (\boldsymbol{\theta}^t - \boldsymbol{\theta}^*) - \eta \nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*), \end{aligned} \quad (66)$$

where we denote $\boldsymbol{\theta}(\tau) := \boldsymbol{\theta}^* + \tau(\boldsymbol{\theta}^t - \boldsymbol{\theta}^*)$. Here, the last identity results from the fundamental theorem of calculus (Lang, 1993, Chapter XIII, Theorem 4.2). Let $\theta_{\max}(\tau) := \max_i \theta_i(\tau)$ and $\theta_{\min}(\tau) := \min_i \theta_i(\tau)$. Combining the induction hypothesis (40d) with the definition of $\boldsymbol{\theta}(\tau)$, one can see that for all $0 \leq \tau \leq 1$,

$$\theta_{\max}(\tau) - \theta_{\min}(\tau) \leq \theta_{\max}^* - \theta_{\min}^* + 2 \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_{\infty} \leq \log \kappa + \epsilon,$$

for any sufficiently small $\epsilon > 0$, as long as

$$2C_4\kappa^2 \sqrt{\frac{\log n}{npL}} \leq \epsilon.$$

This together with Lemma 8 and Corollary 1 reveals that for any $0 \leq \tau \leq 1$,

$$\frac{1}{10\kappa}np + \lambda \leq \frac{1}{8\kappa e^\epsilon}np + \lambda \leq \lambda_{\min, \perp}(\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}(\tau))) \leq \lambda_{\max}(\nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}(\tau))) \leq \lambda + np, \quad (67)$$

where the first inequality holds as long as $\epsilon > 0$ is small enough. Denoting $\mathbf{A} = \int_0^1 \nabla^2 \mathcal{L}_\lambda(\boldsymbol{\theta}(\tau)) d\tau$ and using the triangle inequality, we can derive from (66) that

$$\|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^*\|_2 \leq \|(\mathbf{I}_n - \eta\mathbf{A})(\boldsymbol{\theta}^t - \boldsymbol{\theta}^*)\|_2 + \eta \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*)\|_2. \quad (68)$$

Since $\mathbf{1}^\top(\boldsymbol{\theta}^t - \boldsymbol{\theta}^*) = 0$, the first term on the right hand side of (68) is controlled by

$$\|(\mathbf{I}_n - \eta\mathbf{A})(\boldsymbol{\theta}^t - \boldsymbol{\theta}^*)\|_2 \leq \max\{|1 - \eta\lambda_{\min, \perp}(\mathbf{A})|, |1 - \eta\lambda_{\max}(\mathbf{A})|\} \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_2.$$

Substitute (67) into the above inequality to reach

$$\|(\mathbf{I}_n - \eta\mathbf{A})(\boldsymbol{\theta}^t - \boldsymbol{\theta}^*)\|_2 \leq \left(1 - \frac{1}{10\kappa}\eta np\right) \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_2. \quad (69)$$

Substitute (69) back to (68) and use the induction hypothesis (40d) to conclude that

$$\begin{aligned} \|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^*\|_2 &\leq \left(1 - \frac{1}{10\kappa}\eta np\right) \|\boldsymbol{\theta}^t - \boldsymbol{\theta}^*\|_2 + \eta \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*)\|_2 \\ &\leq \left(1 - \frac{1}{10\kappa}\eta np\right) C_1\kappa \sqrt{\frac{\log n}{pL}} + C\eta \sqrt{\frac{n^2 p \log n}{L}} \\ &\leq C_1\kappa \sqrt{\frac{\log n}{pL}} \end{aligned}$$

for some constants $C, C_1 > 0$. Here, the second inequality makes use of the facts that $\|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}^*)\|_2 \lesssim \sqrt{\frac{n^2 p \log n}{L}}$ (see Lemma 7). The last line holds with the proviso that $C_1 > 0$ is sufficiently large.

D.7 Proof of Lemma 15

Consider any m ($1 \leq m \leq n$). According to the gradient update rule (38), one has

$$\begin{aligned} \theta_m^{t+1, (m)} - \theta_m^* &= \theta_m^{t, (m)} - \eta \left[\nabla \mathcal{L}_\lambda^{(m)}(\boldsymbol{\theta}^{t, (m)}) \right]_m - \theta_m^* \\ &= \theta_m^{t, (m)} - \theta_m^* - \eta \left[p \sum_{i: i \neq m} \left\{ \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} - \frac{e^{\theta_i^{t, (m)}}}{e^{\theta_i^{t, (m)}} + e^{\theta_m^{t, (m)}}} \right\} \right] - \eta \lambda \theta_m^{t, (m)}, \end{aligned} \quad (70)$$

where the last line follows by the construction of $\mathcal{L}_\lambda^{(m)}$. Apply the mean value theorem to obtain

$$\frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} - \frac{e^{\theta_i^{t, (m)}}}{e^{\theta_i^{t, (m)}} + e^{\theta_m^{t, (m)}}} = \frac{1}{1 + e^{\theta_m^* - \theta_i^*}} - \frac{1}{1 + e^{\theta_m^{t, (m)} - \theta_i^{t, (m)}}}$$

$$= -\frac{e^{c_i}}{(1+e^{c_i})^2} \left[\theta_m^* - \theta_i^* - \left(\theta_m^{t,(m)} - \theta_i^{t,(m)} \right) \right], \quad (71)$$

where c_i is some real number lying between $\theta_m^* - \theta_i^*$ and $\theta_m^{t,(m)} - \theta_i^{t,(m)}$. Substituting (71) back into (70) and rearranging terms yield

$$\begin{aligned} \theta_m^{t+1,(m)} - \theta_m^* &= \left(1 - \eta\lambda - \eta p \sum_{i:i \neq m} \frac{e^{c_i}}{(1+e^{c_i})^2} \right) \left(\theta_m^{t,(m)} - \theta_m^* \right) \\ &\quad + \eta p \sum_{i:i \neq m} \frac{e^{c_i}}{(1+e^{c_i})^2} \left(\theta_i^{t,(m)} - \theta_i^* \right) - \eta\lambda \theta_m^* \end{aligned}$$

and hence

$$\begin{aligned} \left| \theta_m^{t+1,(m)} - \theta_m^* \right| &\leq \left| 1 - \eta\lambda - \eta p \sum_{i:i \neq m} \frac{e^{c_i}}{(1+e^{c_i})^2} \right| \left| \theta_m^{t,(m)} - \theta_m^* \right| \\ &\quad + \frac{\eta p}{4} \sum_{i:i \neq m} \left| \theta_i^{t,(m)} - \theta_i^* \right| + \eta\lambda \|\boldsymbol{\theta}^*\|_\infty \\ &\leq \left| 1 - \eta\lambda - \eta p \sum_{i:i \neq m} \frac{e^{c_i}}{(1+e^{c_i})^2} \right| \left| \theta_m^{t,(m)} - \theta_m^* \right| \\ &\quad + \frac{\eta p}{4} \sqrt{n} \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_2 + \eta\lambda \|\boldsymbol{\theta}^*\|_\infty. \end{aligned}$$

Here, the first inequality comes from the triangle inequality and the elementary inequality $\frac{e^c}{(1+e^c)^2} \leq \frac{1}{4}$ for any $c \in \mathbb{R}$, whereas the second relation holds owing to the Cauchy-Schwarz inequality, namely

$$\sum_{i:i \neq m} \left| \theta_i^{t,(m)} - \theta_i^* \right| \leq \sqrt{n} \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_2.$$

From $\frac{e^c}{(1+e^c)^2} \leq \frac{1}{4}$ we also obtain that

$$1 - \eta\lambda - \eta p \sum_{i:i \neq m} \frac{e^{c_i}}{(1+e^{c_i})^2} \geq 1 - \eta\lambda - \eta p \frac{n}{4} \geq 1 - \eta(np + \lambda) \geq 0$$

and

$$\begin{aligned} \left| 1 - \eta\lambda - \eta p \sum_{i:i \neq m} \frac{e^{c_i}}{(1+e^{c_i})^2} \right| &= 1 - \eta\lambda - \eta p \sum_{i:i \neq m} \frac{e^{c_i}}{(1+e^{c_i})^2} \\ &\leq 1 - \eta p(n-1) \min_{i:i \neq m} \frac{e^{c_i}}{(1+e^{c_i})^2}. \end{aligned}$$

To further upper bound $\left| \theta_m^{t+1,(m)} - \theta_m^* \right|$, it suffices to obtain a lower bound on $\frac{e^{c_i}}{(1+e^{c_i})^2}$. Toward this, it is easy to see from (41a) that

$$\begin{aligned} \max_{i:i \neq m} |c_i| &\leq \max_{i:i \neq m} |\theta_m^* - \theta_i^*| + \max_{i:i \neq m} \left| \theta_m^* - \theta_i^* - \left(\theta_m^{t,(m)} - \theta_i^{t,(m)} \right) \right| \\ &\leq \log \kappa + 2 \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_\infty \\ &\leq \log \kappa + \epsilon, \end{aligned}$$

as long as

$$2C_5 \kappa^2 \sqrt{\frac{\log n}{npL}} \leq \epsilon.$$

This further reveals that for $\epsilon > 0$ small enough, one has

$$\frac{e^{c_i}}{(1 + e^{c_i})^2} = \frac{e^{-|c_i|}}{(1 + e^{-|c_i|})^2} \geq \frac{e^{-|c_i|}}{4} \geq \frac{1}{4e^\epsilon \kappa} \geq \frac{1}{5\kappa}.$$

Taking the previous bounds collectively, we arrive at

$$\begin{aligned} \left| \theta_m^{t+1, (m)} - \theta_m^* \right| &\leq \left(1 - \frac{1}{10\kappa} \eta mp \right) \left| \theta_m^{t, (m)} - \theta_m^* \right| + \frac{\eta p \sqrt{n}}{4} \left\| \theta^{t, (m)} - \theta^* \right\|_2 + \eta \lambda \|\theta^*\|_\infty \\ &\leq \left(1 - \frac{1}{10\kappa} \eta mp \right) C_2 \kappa^2 \sqrt{\frac{\log n}{npL}} + \frac{\eta p \sqrt{n}}{4} C_6 \kappa \sqrt{\frac{\log n}{pL}} + c_\lambda \eta \sqrt{\frac{np \log n}{L}} \\ &\leq C_2 \kappa^2 \sqrt{\frac{\log n}{npL}}, \end{aligned}$$

as long as $C_2 \gg \max\{C_6, c_\lambda\}$. Here the second line comes from the setting of λ , namely,

$$\lambda \|\theta^*\|_\infty = c_\lambda \frac{\|\theta^*\|_\infty}{\log \kappa} \sqrt{\frac{np \log n}{L}} \leq c_\lambda \sqrt{\frac{np \log n}{L}}$$

since $\|\theta^*\|_\infty \leq \log \kappa$.

D.8 Proof of Lemma 16

Consider any $1 \leq m \leq n$. Apply the update rules (35) and (38) to obtain

$$\begin{aligned} \theta^{t+1} - \theta^{t+1, (m)} &= \theta^t - \eta \nabla \mathcal{L}_\lambda(\theta^t) - \left[\theta^{t, (m)} - \eta \nabla \mathcal{L}_\lambda^{(m)}(\theta^{t, (m)}) \right] \\ &= \theta^t - \eta \nabla \mathcal{L}_\lambda(\theta^t) - \left[\theta^{t, (m)} - \eta \nabla \mathcal{L}_\lambda(\theta^{t, (m)}) \right] \\ &\quad - \eta \left(\nabla \mathcal{L}_\lambda(\theta^{t, (m)}) - \nabla \mathcal{L}_\lambda^{(m)}(\theta^{t, (m)}) \right) \\ &= \underbrace{\left(\mathbf{I}_n - \eta \int_0^1 \nabla^2 \mathcal{L}_\lambda(\theta(\tau)) d\tau \right)}_{:= \mathbf{v}_1} (\theta^t - \theta^{t, (m)}) \\ &\quad - \underbrace{\eta \left(\nabla \mathcal{L}_\lambda(\theta^{t, (m)}) - \nabla \mathcal{L}_\lambda^{(m)}(\theta^{t, (m)}) \right)}_{:= \mathbf{v}_2}, \end{aligned}$$

where we abuse the notation and denote $\theta(\tau) = \theta^{t, (m)} + \tau(\theta^t - \theta^{t, (m)})$, and the last identity results from the fundamental theorem of calculus (Lang, 1993, Chapter XIII, Theorem 4.2). In what follows, we control \mathbf{v}_1 and \mathbf{v}_2 separately.

- Regarding the term \mathbf{v}_1 , repeating the same argument as in Appendix D.6 yields

$$\|\mathbf{v}_1\|_2 \leq \left(1 - \frac{1}{10\kappa} \eta mp \right) \left\| \theta^t - \theta^{t, (m)} \right\|_2$$

as long as $\eta \leq \frac{1}{\lambda + np}$.

- When it comes to \mathbf{v}_2 , one can use the gradient definitions to reach

$$\frac{1}{\eta} \mathbf{v}_2 = \sum_{i: i \neq m} \left\{ \left(-y_{m, i} + \frac{e^{\theta_i^{t, (m)}}}{e^{\theta_i^{t, (m)}} + e^{\theta_m^{t, (m)}}} \right) \mathbb{1}_{\{(i, m) \in \mathcal{E}\}} - p \left(-y_{m, i}^* + \frac{e^{\theta_i^{t, (m)}}}{e^{\theta_i^{t, (m)}} + e^{\theta_m^{t, (m)}}} \right) \right\} (\mathbf{e}_i - \mathbf{e}_m)$$

$$\begin{aligned}
&= \frac{1}{L} \sum_{i:(i,m) \in \mathcal{E}} \underbrace{\sum_{l=1}^L \left(-y_{m,i}^{(l)} + \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} \right)}_{:= \mathbf{u}^m} (\mathbf{e}_i - \mathbf{e}_m) \\
&+ \underbrace{\sum_{i:i \neq m} \left\{ \left(-\frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} + \frac{e^{\theta_i^{t,(m)}}}{e^{\theta_i^{t,(m)}} + e^{\theta_m^{t,(m)}}} \right) (\mathbb{1}_{\{(i,m) \in \mathcal{E}\}} - p) \right\}}_{:= \mathbf{v}^m} (\mathbf{e}_i - \mathbf{e}_m). \tag{72}
\end{aligned}$$

In the sequel, we control the two terms of (72) separately.

– For the first term \mathbf{u}^m in (72), we make the observation that

$$u_i^m = \begin{cases} \frac{1}{L} \sum_{l=1}^L \left(-y_{m,i}^{(l)} + \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} \right), & \text{if } (i, m) \in \mathcal{E}; \\ \frac{1}{L} \sum_{i:(i,m) \in \mathcal{E}} \sum_{l=1}^L \left(y_{m,i}^{(l)} - \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} \right), & \text{if } i = m; \\ 0, & \text{else.} \end{cases}$$

Since $\left| y_{m,i}^{(l)} - \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} \right| \leq 1$ and $\text{card}(\{i : (i, m) \in \mathcal{E}\}) \asymp np$, we can apply Hoeffding's inequality and union bounds to get for all $1 \leq m \leq n$,

$$|u_m^m| \lesssim \sqrt{\frac{np \log n}{L}} \quad \text{and} \quad |u_i^m| \lesssim \sqrt{\frac{\log n}{L}} \quad \forall i \text{ obeying } (i, m) \in \mathcal{E},$$

which further gives

$$\|\mathbf{u}^m\|_2 \leq |u_m^m| + \sqrt{\sum_{i:(i,m) \in \mathcal{E}} (u_i^m)^2} \lesssim \sqrt{\frac{np \log n}{L}}, \quad \forall 1 \leq m \leq n.$$

– We then turn to the second term \mathbf{v}^m of (72). This is a zero-mean random vector that satisfies

$$v_i^m = \begin{cases} \xi_i (1 - p), & \text{if } (i, m) \in \mathcal{E}, \\ -\sum_{i:i \neq m} \xi_i (\mathbb{1}_{\{(i,m) \in \mathcal{E}\}} - p), & \text{if } i = m, \\ -\xi_i p, & \text{else,} \end{cases}$$

where

$$\xi_i := -\frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_m^*}} + \frac{e^{\theta_i^{t,(m)}}}{e^{\theta_i^{t,(m)}} + e^{\theta_m^{t,(m)}}} = -\frac{1}{1 + e^{\theta_m^* - \theta_i^*}} + \frac{1}{1 + e^{\theta_m^{t,(m)} - \theta_i^{t,(m)}}}.$$

The first step is to bound the size of the coefficient ξ_i . Define $g(x) = (1 + e^x)^{-1}$ for $x \in \mathbb{R}$. We have $|g'(x)| \leq 1$ and thus

$$\begin{aligned}
|\xi_i| &= \left| g\left(\theta_m^{t,(m)} - \theta_i^{t,(m)}\right) - g\left(\theta_m^* - \theta_i^*\right) \right| \\
&\leq \left| \left(\theta_m^{t,(m)} - \theta_i^{t,(m)}\right) - \left(\theta_m^* - \theta_i^*\right) \right| \leq \left| \theta_m^* - \theta_i^{t,(m)} \right| + \left| \theta_m^* - \theta_m^{t,(m)} \right|.
\end{aligned}$$

This indicates that

$$|\xi_i| \leq 2 \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_\infty \quad \text{and} \quad \sum_{i=1}^n \xi_i^2 \leq 4n \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_\infty^2.$$

Applying the Bernstein inequality in Lemma 19 we obtain

$$|v_m^m| \lesssim \sqrt{\left(p \sum_{i=1}^n \xi_i^2 \right) \log n} + \max_{1 \leq i \leq n} |\xi_i| \log n$$

$$\lesssim \left(\sqrt{np \log n} + \log n \right) \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_{\infty}$$

with high probability. As a consequence,

$$\begin{aligned} \|\mathbf{v}^m\|_2 &\leq |v_m^m| + \sqrt{\sum_{i:(i,m) \in \mathcal{E}} (v_i^m)^2} + \sqrt{\sum_{i:(i,m) \notin \mathcal{E} \text{ and } i \neq m} (v_i^m)^2} \\ &\lesssim \left(\sqrt{np \log n} + \log n \right) \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_{\infty} \\ &\quad + \sqrt{np} \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_{\infty} + p\sqrt{n} \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_{\infty} \\ &\lesssim \left(\sqrt{np \log n} + \log n \right) \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_{\infty} \lesssim \left(\sqrt{np \log n} \right) \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_{\infty}, \end{aligned}$$

as long as $np \gg \log n$.

Putting the above results together, we see that

$$\|\mathbf{v}_2\|_2 \lesssim \eta \left(\sqrt{\frac{np \log n}{L}} + \sqrt{np \log n} \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_{\infty} \right).$$

- Combine the above two bounds to deduce for some $C > 0$

$$\begin{aligned} \left\| \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^{t+1,(m)} \right\|_2 &\leq \left(1 - \frac{1}{10\kappa} \eta np \right) \left\| \boldsymbol{\theta}^t - \boldsymbol{\theta}^{t,(m)} \right\|_2 \\ &\quad + C\eta \left(\sqrt{\frac{np \log n}{L}} + \sqrt{np \log n} \left\| \boldsymbol{\theta}^{t,(m)} - \boldsymbol{\theta}^* \right\|_{\infty} \right) \\ &\leq \left(1 - \frac{1}{10\kappa} \eta np \right) C_3 \kappa \sqrt{\frac{\log n}{npL}} \\ &\quad + C\eta \left(\sqrt{\frac{np \log n}{L}} + \sqrt{np \log n} C_5 \kappa^2 \sqrt{\frac{\log n}{npL}} \right) \\ &\leq C_3 \kappa \sqrt{\frac{\log n}{npL}}, \end{aligned}$$

as soon as C_3 is sufficiently large and

$$\kappa^2 \sqrt{\frac{\log n}{np}} \ll 1.$$

D.9 Proof of Lemma 17

Consider any m ($1 \leq m \leq n$). It is easily seen from the triangle inequality that

$$\begin{aligned} \left| \boldsymbol{\theta}_m^{t+1} - \boldsymbol{\theta}_m^* \right| &\leq \left| \boldsymbol{\theta}_m^{t+1} - \boldsymbol{\theta}_m^{t+1,(m)} \right| + \left| \boldsymbol{\theta}_m^{t+1,(m)} - \boldsymbol{\theta}_m^* \right| \\ &\leq \left\| \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^{t+1,(m)} \right\|_2 + \left| \boldsymbol{\theta}_m^{t+1,(m)} - \boldsymbol{\theta}_m^* \right| \\ &\leq C_3 \kappa \sqrt{\frac{\log n}{npL}} + C_2 \kappa^2 \sqrt{\frac{\log n}{npL}} \\ &\leq C_4 \kappa^2 \sqrt{\frac{\log n}{npL}}, \end{aligned}$$

with the proviso that $C_4 \geq C_3 + C_2$.

E Hoeffding’s and Bernstein’s inequalities

This section collects two standard concentration inequalities used throughout the paper, which can be easily found in textbooks such as [Boucheron et al. \(2013\)](#). The proofs are omitted.

Lemma 18 (Hoeffding’s inequality). *Let $\{X_i\}_{1 \leq i \leq n}$ be a sequence of independent random variables where $X_i \in [a_i, b_i]$ for each $1 \leq i \leq n$, and $S_n = \sum_{i=1}^n X_i$. Then*

$$\mathbb{P}(|S_n - \mathbb{E}[S_n]| \geq t) \leq 2e^{-2t^2 / \sum_{i=1}^n (b_i - a_i)^2}.$$

The next lemma is about a user-friendly version of the Bernstein inequality.

Lemma 19 (Bernstein’s inequality). *Consider n independent random variables z_l ($1 \leq l \leq n$), each satisfying $|z_l| \leq B$. For any $a \geq 2$, one has*

$$\left| \sum_{l=1}^n z_l - \mathbb{E} \left[\sum_{l=1}^n z_l \right] \right| \leq \sqrt{2a \log n \sum_{l=1}^n \mathbb{E}[z_l^2]} + \frac{2a}{3} B \log n$$

with probability at least $1 - 2n^{-a}$.

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*.
- AMMAR, A. and SHAH, D. (2011). Ranking: Compare, don’t score. In *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*.
- AMMAR, A. and SHAH, D. (2012). Efficient rank aggregation using partial data. In *SIGMETRICS*, vol. 40. 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*, ACM, New York, NY, USA.
- BOUCHERON, S., LUGOSI, G. and MASSART, P. (2013). *Concentration inequalities*. Oxford University Press, Oxford. A nonasymptotic theory of independence, With a foreword by Michel Ledoux.
- BRADLEY, R. A. and TERRY, M. E. (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 *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28. ICML’13*, JMLR.org.
- CHEN, X., BENNETT, P. N., COLLINS-THOMPSON, K. and HORVITZ, E. (2013). Pairwise ranking aggregation in a crowdsourced setting. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining. WSDM ’13*, ACM, New York, NY, USA.
- CHEN, X., GOPI, S., MAO, J. and SCHNEIDER, J. (2017). Competitive analysis of the top-k ranking problem. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms. SODA ’17*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.

- CHEN, Y. and CANDÈS, E. (2016). The projected power method: An efficient algorithm for joint alignment from pairwise differences. *arXiv preprint arXiv:1609.05820* .
- CHEN, Y. and CANDÈS, E. J. (2017). Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Comm. Pure Appl. Math.* **70** 822–883.
- 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* .
- CHEN, Y. and SUH, C. (2015). Spectral mle: Top-k rank aggregation from pairwise comparisons. In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37*. ICML'15, JMLR.org.
- CHUNG, F. R. K. (1997). *Spectral graph theory*, vol. 92 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI.
- DAVIS, C. and KAHAN, W. M. (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 *Proceedings of the 10th International Conference on World Wide Web*. WWW '01, ACM, New York, NY, USA.
- EL KAROUI, 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. (2016). An ℓ_∞ eigenvector perturbation bound and its application to robust covariance estimation. *arXiv preprint arXiv:1603.03516* .
- FORD, L. R., JR. (1957). Solution of a ranking problem from binary comparisons. *Amer. Math. Monthly* **64** 28–33.
- HAJEK, B., OH, S. and XU, J. (2014). Minimax-optimal inference from partial rankings. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*. NIPS'14, MIT Press, Cambridge, MA, USA.
- HECKEL, R., SHAH, N. B., RAMCHANDRAN, K. and WAINWRIGHT, M. J. (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.
- JAMIESON, K. G. and NOWAK, R. D. (2011). Active ranking using pairwise comparisons. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*. NIPS'11, Curran Associates Inc., USA.
- 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.
- KESHAVAN, R. H., MONTANARI, A. and OH, S. (2010). Matrix completion from noisy entries. *J. Mach. Learn. Res.* **11** 2057–2078.

- 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*, vol. 52. Institut Henri Poincaré.
- KOLTCHINSKII, V. and XIA, D. (2016). Perturbation of linear forms of singular vectors under gaussian noise. In *High Dimensional Probability VII*. Springer, 397–423.
- LANG, S. (1993). Real and functional analysis. *Springer-Verlag, New York*, **10** 11–13.
- LU, Y. and NEGAHBAN, S. N. (2014). Individualized rank aggregation using nuclear norm regularization. *arXiv preprint arXiv:1410.0860* .
- LUCE, R. D. (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. (2017a). Rank centrality: ranking from pairwise comparisons. *Oper. Res.* **65** 266–287.
- NEGAHBAN, S., OH, S., THEKUMPARAMPIL, K. K. and XU, J. (2017b). Learning from comparisons and choices. *arXiv preprint arXiv:1704.07228* .
- PANANJADY, A., MAO, C., MUTHUKUMAR, V., WAINWRIGHT, M. J. and COURTADE, T. A. (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 *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32. ICML'14, JMLR.org*.
- RAJKUMAR, A. and AGARWAL, S. (2016). When can we rank well from comparisons of $o(n \log n)$ non-actively chosen pairs? In *29th Annual Conference on Learning Theory*.
- ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* **39** 1878–1915.
- RUDELSON, M., VERSHYNIN, R. ET AL. (2013). Hanson-wright inequality and sub-gaussian concentration. *Electronic Communications in Probability* **18**.
- SHAH, N. B., BALAKRISHNAN, S., GUNTUBOYINA, A. and WAINWRIGHT, M. J. (2017). Stochastically transitive models for pairwise comparisons: statistical and computational issues. *IEEE Trans. Inform. Theory* **63** 934–959.
- SHAH, N. B. and WAINWRIGHT, M. J. (2015). Simple, robust and optimal ranking from pairwise comparisons. *arXiv preprint arXiv:1512.08949* .
- SOUFIANI, H. A., CHEN, W. Z., PARKES, D. C. 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, Curran Associates Inc., USA.
- SUH, C., TAN, V. Y. F. and ZHAO, R. (2017). Adversarial top- K ranking. *IEEE Trans. Inform. Theory* **63** 2201–2225.
- SUR, P., CHEN, Y. and CANDÈS, E. J. (2017). The likelihood ratio test in high-dimensional logistic regression is asymptotically a rescaled Chi-square. *arXiv preprint arXiv:1706.01191* .

- TROPP, J. A. (2012). User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.* **12** 389–434.
- TROPP, J. A. (2015). An introduction to matrix concentration inequalities. *Found. Trends Mach. Learn.* **8** 1–230.
- TSYBAKOV, A. B. (2009). Introduction to nonparametric estimation.
- ZHONG, Y. and BOUMAL, N. (2017). Near-optimal bounds for phase synchronization. *arXiv preprint arXiv:1703.06605* .