Random pairing MLE for estimation of item parameters in Rasch model

Yuepeng Yang[1] and Cong Ma[1]

[1]Department of Statistics, University of Chicago

June 19, 2024

Abstract

The Rasch model, a classical model in the item response theory, is widely used in psychometrics to model the relationship between individuals' latent traits and their binary responses on assessments or questionnaires. In this paper, we introduce a new likelihood-based estimator—random pairing maximum likelihood estimator (RP-MLE) and its bootstrapped variant multiple random pairing MLE (MRP-MLE) that faithfully estimate the item parameters in the Rasch model. The new estimators have several appealing features compared to existing ones. First, both work for sparse observations, an increasingly important scenario in the big data era. Second, both estimators are provably minimax optimal in terms of finite sample ℓ_∞ estimation error. Lastly, RP-MLE admits precise distributional characterization that allows uncertainty quantification on the item parameters, e.g., construction of confidence intervals of the item parameters. The main idea underlying RP-MLE and MRP-MLE is to randomly pair user-item responses to form item-item comparisons. This is carefully designed to reduce the problem size while retaining statistical independence. We also provide empirical evidence of the efficacy of the two new estimators using both simulated and real data.

1 Introduction

The item response theory (IRT) [ER13] is a framework widely used in psychometrics to model the relationship between individuals' latent traits (such as ability or personality) and their responses on assessments or questionnaires. It is particularly useful in the development, analysis, and scoring of tests and assessments; see a recent survey [CLLY21] for a statistical account of IRT.

Among statistical models in IRT, the Rasch model [Ras60] is a simple but fundamental one for modeling binary responses. Specifically, for a user t (e.g., test-taker) and an item i (e.g., test problem), Rasch model assumes that the response user t has to item i is binary and obeys

$$\mathbb{P}\left[\text{user } t \text{ ``beats'' item } i\right] = \frac{e^{\zeta_t^\star}}{e^{\zeta_t^\star} + e^{\theta_i^\star}},$$

where $\zeta_t^\star, \theta_i^\star \in \mathbb{R}$ are latent traits of user t and item i, respectively. The term "beats" here refers to positive responses, such as answering an exam question correctly, writing a favorable review of a product, etc. We also call this response a comparison between user t and item i.

In this paper, we focus on estimating the item parameters $\boldsymbol{\theta}^\star$, which is one of the four main statistical tasks surrounding the Rasch model (or IRT more generally) listed in [CLLY21]. Estimating the item parameters is also quite useful in practice. For instance, in education testing, $\boldsymbol{\theta}^\star$ could reveal the difficulty of the exam questions, while in product reviews, $\boldsymbol{\theta}^\star$ could reveal the popularity of the products. Various methods have been proposed for estimating the item parameters $\boldsymbol{\theta}^\star$, including the joint maximum likelihood estimator (JMLE), the marginal maximum likelihood estimator, the conditional maximum likelihood estimator (CMLE), as well as the spectral estimator recently proposed in [NZ22, NZ23]. We refer readers to a recent article [Rob21] for comparisons between different item parameter estimation methods. However, three main gaps remain in tackling item estimation in the Rasch model:

- **Non-asymptotic guarantee**. Apart from the recently proposed spectral estimator [NZ22, NZ23], most theoretical guarantees for the likelihood-based estimators are asymptotic. Since all the estimation procedures are necessarily applied with finite samples, the asymptotic guarantee alone fails to inform practitioners about the performance of different estimators when working with a limited number of samples.

- **Sparse observations**. It is not uncommon to encounter situations where each user only responds to a handful of questions (or items). This brings the challenge of incomplete and sparse observations. Among the likelihood based estimators, CMLE is known to have numerical issues with incomplete observations [Mol95]. While the spectral estimator [NZ22, NZ23] is capable of handling incomplete observations, it still requires the observations to be relatively dense. We will elaborate on this point later.

- **Uncertainty quantification**. Beyond estimation, uncertainty quantification on the item parameters is central to realizing the full potential of the Rasch model. However, existing results do not address this problem under sparse observations. An exception is the recent work [CLOX23], which is based on joint estimation and inference on the item parameters $\boldsymbol{\theta}^\star$ and the user parameters $\boldsymbol{\zeta}^\star$. Their sampling scheme is restrictive and requires a relatively dense sampling rate.

In light of the above gaps, we raise the following natural question:

Can we develop an estimator for the item parameters $\boldsymbol{\theta}^\star$ that (1) enjoys optimal estimation guarantee in finite sample, and (2) is amenable to tight uncertainty quantification, when the observations are sparse?

1.1 Main contributions

The main contribution of our work is the proposal of a novel estimator named random pairing maximum likelihood estimator (RP-MLE in short) that achieves the two desiderata listed above.

In essence, RP-MLE compiles user-item comparisons to item-item comparisons by randomly pairing responses of the same user to different items. This pairing procedure is carefully designed to extract information of the item parameters while retaining statistical independence. After this compilation step, item parameters $\boldsymbol{\theta}^\star$ are estimated by the MLE $\widehat{\boldsymbol{\theta}}$ given the item-item comparisons.

Even when the observations are extremely sparse, RP-MLE achieves the following:

- Regarding estimation, we show that both RP-MLE and its bootstrapped version enjoy optimal finite sample ℓ_∞ error guarantee. Compared to the conventional ℓ_2 error guarantee, the ℓ_∞ guarantee, as an entrywise guarantee is more fine-grained. Consequently, we also show that RP-MLE can recover the top-K items with minimal sample complexity.

- While the optimal ℓ_∞ error guarantee directly yields optimal ℓ_2 guarantee, such guarantee is only correct in an order-wise sense. We provide a refined finite-sample ℓ_2 error guarantee of RP-MLE that is precise even in the leading constant.

- Supplementing the estimation guarantee, we also build an inferential framework based on RP-MLE $\widehat{\boldsymbol{\theta}}$. More specifically, we precisely characterize the asymptotic distribution of $\widehat{\boldsymbol{\theta}}$. This result facilitates several inferential tasks such as hypothesis testing and construction of confidence regions of $\boldsymbol{\theta}^\star$.

We test our methods on both synthetic and real data, which clearly show competitive empirical estimation performance. The inferential result on synthetic data also closely matches our theoretical predictions.

1.2 Prior art

Item response theory. The item response theory is a popular statistical framework for modeling response data. It often involves a probabilistic model that links categorical responses to latent traits of both users and items. Early endeavors include [Ras60] that introduces the Rasch model studied herein and [LNB68] that describes a more general framework using parametric models. Popular IRT models include the Rasch model, the two-parameter model (2PL), and the three-parameter logistic model (3PL). As response data widely appears in real life, IRT finds application in numerous fields including educational assessment [DC10], psychometrics [LNB68], political science [VHSA20], and medical assessment [FBC05]. See also [CLLY21] for an overview of IRT from statisticians' perspective.

Latent score estimation for Rasch model. An important statistical question in the Rasch model is to estimate the latent parameters of the items. As the Rasch model is an explicit probabilistic model, many methods are based on the principle of maximum likelihood estimation. For instance, marginal MLE assumes a prior on the user parameters that is either given or optimized within a parametric distribution family. The item parameter is then estimated by maximizing the marginal likelihood. A drawback is that MMLE relies on a good prior. On the other hand, joint MLE (JMLE) makes no distributional assumption and maximizes the joint likelihood w.r.t. both the item and user parameters. However, it is not consistent for estimating the item parameters when the number m of items is fixed [Gho95]. Interested readers may also consult [Lin99] for an overview of other classical estimators.

Several methods are more relevant to our proposed estimator RP-MLE as they follow a similar philosophy to form item-item pairs from user-item responses. Pseudo MLE (PMLE) [Zwi95] maximizes the sum of the log-likelihood of all pairs of responses from the same users to different items. However, due to the dependency issue, no satisfying finite sample performance guarantee has been established. Another related approach is the spectral method, in which a Markov chain on the space of items is formed and the item parameters are estimated via the stationary distribution of the Markov chain. The most recent works in this category are [NZ22, NZ23], which essentially use the same idea as pseudo MLE in forming item-item comparisons.

The Bradley-Terry-Luce model with sparse comparisons. An informed reader may realize that the Rasch model resembles the Bradley-Terry-Luce (BTL) model [Luc59, BT52] in the ranking literature. Indeed, one can view the Rasch model as a special case of the BTL model that distinguishes the two groups of users and items, and only allows inter-group comparisons. There has been a recent surge in interest in studying top-K ranking in the BTL model [CS15, JKSO16, CFMW19, CGZ22, GSZ23] and its extensions [FHY22, FLWY22, FLWY23], especially under sparse observations of the pairwise responses. Most notably, under a uniform sampling scheme, [CFMW19] shows that (regularized) MLE and spectral methods are both optimal in top-K ranking and [GSZ23] provides inference results for both methods. Going beyond uniform sampling, [Che23, LSR22] investigate the performance of MLE in the BTL model with a general comparison graph and later [YCOM24] improves the analysis to show the optimality of MLE for the BTL model in both uniform and semi-random sampling.

Notation. For a positive integer n, we denote $[n] = \{1, 2, \dots, n\}$. For any $a, b \in \mathbb{R}$, $a \wedge b$ means the minimum of a, b and $a \vee b$ means the maximum of a, b. We use $\xrightarrow{\mathrm{d}}$ to denote convergence in distribution. For a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$, we use $\lambda_1(\boldsymbol{A}) \geq \lambda_2(\boldsymbol{A}) \geq \cdots \geq \lambda_n(\boldsymbol{A})$ to denote its eigenvalues and $\boldsymbol{A}^\dagger$ to denote its Moore-Penrose pseudo-inverse. For symmetric matrices $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{n\times n}$, $\boldsymbol{A} \preceq \boldsymbol{B}$ means $\boldsymbol{B} - \boldsymbol{A}$ is positive semidefinite, i.e., $\boldsymbol{v}^\top(\boldsymbol{B} - \boldsymbol{A})\boldsymbol{v} \geq 0$ for any $\boldsymbol{v} \in \mathbb{R}^n$. We use $\boldsymbol{e}_i$ to denote the standard unit vector with 1 at i-th coordinate and 0 elsewhere. Unless specified otherwise, $\log(\cdot)$ denotes the natural log.

2 Problem setup and new estimator RP-MLE

In this section, we first introduce the formal setup of the item parameter estimation problem in the Rasch model. Then we present the newly proposed estimator RP-MLE along with the rationale behind its development.

2.1 Problem setup

The Rasch model considers pairwise comparisons between elements of two groups: users and items. Let n (resp. m) be the number of users (resp. items). Rasch assumes a user parameter $\boldsymbol{\zeta}^\star \in \mathbb{R}^n$ and an item parameter $\boldsymbol{\theta}^\star \in \mathbb{R}^m$ that measures the latent traits (e.g., difficulty of a problem) of users and items,

Algorithm 1 Random Pairing Maximum Likelihood Estimator (RP-MLE)

1. For each tester t,
   (a) Randomly split the n_t problems taken by tester t into $\lfloor n_t/2 \rfloor$ pairs of problems.
   (b) For each $(i,j) \in [m] \times [m]$, do the following:
      i. If (i,j) is selected as a pair in Step 1(a), $R_{ij}^t = 1$. Furthermore, if $X_{ti} \neq X_{tj}$, let $Y_{ij}^t = \mathbb{1}\{X_{ti} < X_{tj}\}$ and $L_{ij}^t = 1$; if $X_{ti} = X_{tj}$, let $L_{ij}^t = 0$.
      ii. If (i,j) is not selected as a pair in Step 1(a), let $L_{ij}^t = 0$ and $R_{ij}^t = 0$.
2. Let $\mathcal{E}_Y$ be a set of edges defined by $\mathcal{E}_Y \coloneqq \{(i,j) : \sum_{t=1} L_{ij}^t \geq 1\}$ and let $\mathcal{G}_Y = ([m], \mathcal{E}_Y)$. For each $(i,j) \in \mathcal{E}_Y$, let $L_{ij} \coloneqq \sum_{t=1}^n L_{ij}^t$ and $Y_{ij} \coloneqq (1/L_{ij}) \sum_{\{t : L_{ij}^t = 1\}} Y_{ij}^t$.
3. Compute MLE on Y_{ij}, i.e.,
$$\widehat{\boldsymbol{\theta}} \coloneqq \arg \min_{\boldsymbol{\theta}: \mathbf{1}_m^\top \boldsymbol{\theta} = 0} \mathcal{L}(\boldsymbol{\theta})$$
   where
$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{(i,j) \in \mathcal{E}_Y, i > j} L_{ij} \left(-Y_{ji}(\theta_i - \theta_j) + \log(1 + e^{\theta_i - \theta_j}) \right). \tag{2}$$
4. Return the top-K items by selecting the top-K entries of $\widehat{\boldsymbol{\theta}}$.

respectively. For a subset of possible user-item pairs $\mathcal{E}_X \subset [n] \times [m]$, we observe binary responses $\{X_{ti}\}_{(t,i) \in \mathcal{E}_X}$ obeying
$$\mathbb{P}[X_{ti} = 1] = \frac{e^{\theta_i^\star}}{e^{\zeta_t^\star} + e^{\theta_i^\star}}. \tag{1}$$
Here $X_{ti} = 1$ means user t has negative response against item i (e.g., unable to solve a problem). The goal is to estimate $\boldsymbol{\theta}^\star$, the item parameters.

To model sparse observations/comparisons, we assume that $\mathbb{P}[(t,i) \text{ is compared}] = p$ independently for every $(t,i) \in [n] \times [m]$. To put it in the language of graph theory, we denote the associated bipartite comparison graph to be $\mathcal{G}_X = (\mathcal{V}_X, \mathcal{E}_X)$, where $\mathcal{V}_X$ consists of n users and m items. Then essentially, we are assuming that the bipartite graph follows an Erdős-Rényi random model.[1]

Before moving on, we introduce the condition numbers that characterize the range of the latent traits. Let κ_1, κ_2, and κ be defined by $\log(\kappa_1) = \max_{ij}\{|\theta_i^\star - \theta_j^\star|\}$, $\log(\kappa_2) = \max_{ti}\{|\zeta_t^\star - \theta_i^\star|\}$, and $\kappa = \max\{\kappa_1, \kappa_2\}$, respectively.

2.2 Random pairing maximum likelihood estimator

In this section, we present our main method RP-MLE; see Algorithm 1. The algorithm can be divided into two parts. The first part—Steps 1 and 2—uses random pairing to compile the observed user-item responses $\boldsymbol{X} \in \mathbb{R}^{n \times m}$ to item-item comparisons $\boldsymbol{Y} \in \mathbb{R}^{m \times m}$. The second part—Steps 3 and 4—computes a standard MLE on the item-item comparisons. Some intuitions regarding the development of RP-MLE are in order.

Random pairing to construct item-item comparisons. The idea of pairing is that by matching the responses X_{ti} with X_{tj}, we form a comparison between items i and j to directly extract information of item parameters $\theta_i^\star$ and $\theta_j^\star$. More specifically, the item-item comparisons $\boldsymbol{Y}$ follow the Bradley-Terry-Luce model [BT52, Luc59],i.e., $\mathbb{P}[Y_{ij}^t] = e^{\theta_j^\star}/(e^{\theta_i^\star} + e^{\theta_j^\star})$; see Section 4.1 for a formal argument.

By compiling user-item responses to item-item comparisons, we reduce the size of the data matrix from $n \times m$ to $m \times m$, and also the number of intrinsic parameters from $n + m$ to m, since the likelihood function

[1] Alternatively, we can assume each user responds to mp items uniformly at random. Our estimator and performance guarantee continue to work in this sampling scheme.

Algorithm 2 Multiple Random Pairing Maximum Likelihood Estimator (MRP-MLE)

1. Let n_{split} be the number of runs. For $i = 1, \dots, n_{\text{split}}$, run RP-MLE (Algorithm 1), each time with an independent random splitting in Step 1. Let the estimated latent scores be $\widehat{\boldsymbol{\theta}}^{(i)}$.
2. Estimate the latent score with
$$\widehat{\boldsymbol{\theta}} = \frac{1}{n_{\text{split}}} \sum_{i=1}^{n_{\text{split}}} \widehat{\boldsymbol{\theta}}^{(i)}.$$
3. Estimate the top-K items by selecting the top-K entries of $\widehat{\boldsymbol{\theta}}$.

(2) of Y_{ij}^t is completely independent with the user parameter $\zeta_t^\star$.

More importantly, the pairing is performed in a disjoint fashion. This ensures that all constructed item-item comparisons Y_{ij}^t are independent with each other; see Section 4.1 for a formal statement. This is the key ingredient that enables us to improve over previous implementation of item-item comparisons, such as pseudo-likelihood [Cho82, Zwi95] and spectral methods [NZ22, NZ23].

A variant via bootstrapping. A drawback of this random pairing is that it potentially induces a loss of information since not every possible pairing is considered. Once X_{ti} is paired with X_{tj}, we cannot pair X_{ti} with another response X_{tl}. That being said, we will later show that the ℓ_∞ error of RP-MLE is still rate-optimal up to logarithmic factors. Hence the loss of information can at most incur a small constant factor in terms of estimation error. Nevertheless, we provide a remedy to this phenomenon in MRP-MLE (Algorithm 2) by running (in other words, bootstrapping) the RP-MLE multiple times with different random data splitting and averaging the resulting estimates. MRP-MLE trivially enjoys the same theoretical guarantee (cf. Theorem 1) while improving the estimation error in practice over RP-MLE. See Figure 3 in Section 5.1 for the empirical evidence.

3 Main results

In this section, we collect the main theoretical guarantees for RP-MLE and its variant MRP-MLE. Section 3.1 focuses on the finite sample ℓ_∞ error bound. While one can translate ℓ_∞ error into an ℓ_2 bound, in Section 3.2, we present a much sharper characterization of the ℓ_2 error of RP-MLE. Lastly in Section 3.3, we provide a distributional characterization of RP-MLE.

3.1 ℓ_∞ estimation error and top-K recovery

Without loss of generality, we assume that the scores of the items are ordered, i.e., $\theta_1^\star \geq \theta_2^\star \geq \cdots \geq \theta_m^\star$, and denote $\Delta_K := \theta_K^\star - \theta_{K+1}^\star$. In words, Δ_K measures the difference between the difficulty levels of items K and $K+1$. The following theorem provides ℓ_∞ error bounds and top-K recovery guarantee for both RP-MLE and MRP-MLE. We defer its proof to Section 4.2.

Theorem 1. *Suppose that $mp \geq 2$ and $np \geq C_1 \kappa_1^4 \kappa_2^5 \log^3(n)$ for some sufficiently large constant $C_1 > 0$. Suppose that there exists some constant $\alpha > 0$ such that $m \leq n^\alpha$. Let $\widehat{\boldsymbol{\theta}}$ be the* RP-MLE *estimator. With probability at least $1 - O(n^{-10})$, $\widehat{\boldsymbol{\theta}}$ satisfies*

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|_\infty \leq C_2 \kappa_1 \kappa_2^{1/2} \sqrt{\frac{\log(n)}{np}}.$$

Consequently, the estimator is able to exactly recover the top-K items as soon as

$$np \geq \frac{C_3 \kappa_1^2 \kappa_2^1 \log(n)}{\Delta_K^2}.$$

Here $C_2, C_3 > 0$ are some universal constants. All the claims continue to hold for MRP-MLE *as long as there exists some constant $\beta > 0$ such that $n_{\text{split}} \leq n^{\beta}$.*

Some remarks are in order.

Finite sample minimax optimality. Intuitively, as we have $\Theta(mnp)$ sampled comparisons and m parameters to estimate, the optimal ℓ_∞ error should be $O(\sqrt{m/(mnp)}) = O(1/\sqrt{np})$. This is formalized in the following result from [NZ23].

Proposition 1 (Minimax lower bound, Theorems 3.3 and 3.4 in [NZ23])**.** *Assume that $np \geq C_1$ for some suffiently large constant $C_1 > 0$. For any n and m, there exists a class of user and item parameters Θ such that*

$$\inf_{\hat{\boldsymbol{\theta}}} \sup_{(\boldsymbol{\zeta}^\star, \boldsymbol{\theta}^\star) \in \Theta} \mathbb{E} \left\| \widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star \right\|_2^2 \geq \frac{C_2 m}{np}, \qquad \textit{and} \qquad \inf_{\hat{\boldsymbol{\theta}}} \sup_{(\boldsymbol{\zeta}^\star, \boldsymbol{\theta}^\star) \in \Theta} \mathbb{E} \left\| \widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star \right\|_\infty^2 \geq \frac{C_2}{np},$$

where $C_2 > 0$ is some constant. Moreover if $np \leq C_K \log(m)/\Delta_K^2$ for some constant $C_K > 0$, we have

$$\inf_{\hat{\boldsymbol{\theta}}} \sup_{(\boldsymbol{\zeta}^\star, \boldsymbol{\theta}^\star) \in \Theta} \mathbb{P} \left[\widehat{\boldsymbol{\theta}} \textit{ fails to identify all top-}K \textit{ items} \right] \geq \frac{1}{2}.$$

Comparing our upper bounds with the lower bound in the proposition, we can see that both RP-MLE and MRP-MLE are rate-optimal in ℓ_∞ estimation error and top-K recovery sample complexity, up to logarithmic and κ factors.

Sample size requirement. While the rates are optimal, it is worth noting that in Theorem 1 we have made several sample size requirements. We now elaborate on them.

First, the assumption $mp \geq 2$ is a mild requirement on the expected number of items compared by each user. This is required as we need user t to compare at least two items to form a comparison between items. In fact, if a user only responds to one item, it is clear that this data point is not useful at all for item parameter estimation.

Second, it is a standard and necessary requirement to have $np \gtrsim \log(n)$ to make sure that each item is compared to at least one user with high probability. In Theorem 1 we require an extra $\log^2(n)$ factor to suppress a quadratic error term that comes up in the analysis. This cubic log factor can possibly be loose, but it is a minor issue and we leave it as a problem for future research.

Lastly, $m \leq n^{\alpha}$ and $n_{\text{split}} \leq n^{\beta}$ are both minor as we only need these to allow union bounds over m and n_{split}.

Comparison with [NZ23]. The closest result to our paper in terms of ℓ_∞ guarantee for the Rasch model appears in the recent work [NZ23]. Their spectral method uses a similar construction of the item-item comparisons but without disjoint pairing. To provide detailed comparisons, we restate their results below.

Proposition 2 (Informal, Theorem 3.1 in [NZ23])**.** *Assume that $p \gtrsim \log(m)/\sqrt{n}$ and $mp \gtrsim \log(n)$, with probability at least $1 - O(m^{-10} + n^{-10})$, spectral estimator $\widehat{\boldsymbol{\theta}}_{\text{spectral}}$ satisfies*

$$\|\widehat{\boldsymbol{\theta}}_{\text{spectral}} - \widehat{\boldsymbol{\theta}}\|_\infty \lesssim \kappa^9 \sqrt{\frac{\log(m)}{np}}.$$

This error rate is similar to ours. However, the required sample size is much larger as they require $p \gtrsim \log(m)/\sqrt{n}$. Our result makes a significant improvement by allowing a much smaller sampling rate p, cf. $mp \geq 2$ and $np \gtrsim \log^3(n)$. In fact, as we have argued earlier, it is nearly the sparsest possible. In addition, our methods enjoy a significantly better error rate dependency on κ. In Section 5.1, we provide empirical evidence for this improvement: when κ is large, RP-MLE and MRP-MLE outperform the spectral methods in [NZ23].

3.2 Refined ℓ_2 error characterization

The ℓ_∞ error guarantee in Theorem 1 immediately implies an ℓ_2 error bound

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\| \le C\kappa_1\kappa_2^{1/2}\sqrt{\frac{m\log(n)}{np}}, \tag{3}$$

which is rate-optimal compared to the minimax lower bound in Proposition 1. However, this guarantee is only correct in an order-wise sense. In this section, we present a refined characterization of $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|$ that is precise in the leading constant.

We start with necessary notation. Let $z_{ij} := e^{\theta_i^\star}e^{\theta_j^\star}/(e^{\theta_i^\star}+e^{\theta_j^\star})^2$ and $\widehat{z}_{ij} := e^{\hat{\theta}_i}e^{\hat{\theta}_j}/(e^{\hat{\theta}_i}+e^{\hat{\theta}_j})^2$. Define a matrix

$$\boldsymbol{L}_{L\tilde{z}} = \sum_{(i,j)\in\mathcal{E}_Y, i>j} L_{ij}\widetilde{z}_{ij}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top \tag{4}$$

for $\tilde{z} \in \{z, \widehat{z}\}$ and let $\boldsymbol{L}_{L\tilde{z}}^\dagger$ be its pseudo-inverse. One can view $\boldsymbol{L}_{L\tilde{z}}$ as a weighted graph Laplacian. In the following theorem, we show that $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|$ concentrates tightly around $[\mathrm{Trace}(\boldsymbol{L}_{L\tilde{z}}^\dagger)]^{1/2}$. We defer the analysis to Section 4.3 and the complete proof to Section D.

Theorem 2. *Instate the assumptions of Theorem 1. Then for some constants $C_1, C_2 > 0$, with probability at least $1 - O(n^{-10})$, we have*

$$\left|\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\| - \sqrt{\mathrm{Trace}(\boldsymbol{L}_{Lz}^\dagger)}\right| \le C_1\kappa_1^3\kappa_2\sqrt{\frac{\log(n)}{np}} + \frac{C_2\kappa_1^6\sqrt{m}\log^2(n)}{np}, \tag{5}$$

$$\left|\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\| - \sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\widehat{z}}^\dagger)}\right| \le C_1\kappa_1^3\kappa_2\sqrt{\frac{\log(n)}{np}} + \frac{C_2(\kappa_1^6 + \kappa_1^{7/2}\kappa_2^2)\sqrt{m}\log^2(n)}{np}. \tag{6}$$

Theorem 2 is more refined compared to (3). First, it provides both upper and lower bounds for $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|$. Second, there is no hidden constant in front of the leading term $[\mathrm{Trace}(\boldsymbol{L}_{Lz}^\dagger)]^{1/2}$. In addition, inspecting the proof of Theorem 2, we see that

$$\sqrt{\frac{m-1}{np}} \le \sqrt{\mathrm{Trace}(\boldsymbol{L}_{Lz}^\dagger)} \le 4\kappa_1^{1/2}\kappa_2^{1/2}\sqrt{\frac{m}{np}},$$

and the same holds for $[\mathrm{Trace}(\boldsymbol{L}_{L\widehat{z}}^\dagger)]^{1/2}$. Consequently, the right hand sides of both (5) and (6) are lower order terms compared to $[\mathrm{Trace}(\boldsymbol{L}_{Lz}^\dagger)]^{1/2}$ when $n, m \to \infty$. Indeed this recovers the naive ℓ_2 bound (3) under appropriate sample size assumptions.

3.3 Uncertainty quantification

Now we move on to uncertainty quantification about the RP-MLE estimator $\widehat{\boldsymbol{\theta}}$ in order to perform, say, hypothesis testing on the item parameters $\boldsymbol{\theta}^\star$. Our starting point is the following characterization of the limiting distribution of $\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star$. The analysis is deferred to Section 4.3 and the complete proof is deferred to Section E.

Theorem 3. *Instate the assumptions of Theorem 1 and fix m. As $np/\log^4(n) \to \infty$, for $\tilde{z} \in \{z, \widehat{z}\}$, we have*

$$\boldsymbol{L}_{L\tilde{z}}^{1/2}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star) \xrightarrow{\mathrm{d}} \mathcal{N}\left(0, \boldsymbol{I}_m - \frac{1}{m}\boldsymbol{1}_m\boldsymbol{1}_m^\top\right).$$

Though simple, we record the following immediate consequence on the asymptotic normality of $\widehat{\theta}_i - \theta_i^\star$ for any single coordinate i.

Corollary 1. *Instate the assumptions of Theorem 1. As $np/\log^4(n) \to \infty$, for $\tilde{z} \in \{z, \widehat{z}\}$, we have*

$$\frac{\widehat{\theta}_i - \theta_i^\star}{\sqrt{[\boldsymbol{L}_{L\tilde{z}}^\dagger]_{ii}}} \xrightarrow{\mathrm{d}} \mathcal{N}(0, 1). \tag{7}$$

We compare our result with that in [CLOX23]. Theorem 8 therein studies the Rasch model and provides an inference result for a joint estimator of $\boldsymbol{\theta}^\star$ and $\boldsymbol{\zeta}^\star$. However, it requires a dense sampling scheme with $p \gtrsim 1/\sqrt{\log(n)/n}$ and $p \gtrsim n \log^2(n)/m^2$. It also requires that both m and n tend to infinity. These two assumptions are more restrictive than ours.

Theorem 3 is quite general as it allows various applications to inference. As an illustrative example, we demonstrate how to construct confidence intervals for a particular item parameter $\theta_i^\star$.

Fix any $\alpha > 0$, the level of significance. Define

$$\mathcal{C}_i^+(\alpha) := \widehat{\theta}_i + \Phi^{-1}(1-\alpha) \left[\boldsymbol{L}_{L\hat{z}}^\dagger\right]_{ii}^{1/2},$$
$$\mathcal{C}_i^-(\alpha) := \widehat{\theta}_i - \Phi^{-1}(1-\alpha) \left[\boldsymbol{L}_{L\hat{z}}^\dagger\right]_{ii}^{1/2},$$

where Φ is the cdf for the standard normal distribution. Corollary 1 implies the following proposition.

Proposition 3. *Suppose the statement in Theorem 3 holds. Then for any (i,j), we have*

$$\lim_{np\to\infty} \mathbb{P}\left[\theta_i^\star \geq \mathcal{C}_i^+(\alpha)\right] = \alpha, \tag{8a}$$
$$\lim_{np\to\infty} \mathbb{P}\left[\theta_i^\star \leq \mathcal{C}_i^-(\alpha)\right] = \alpha, \tag{8b}$$

and

$$\lim_{np\to\infty} \mathbb{P}\left[\mathcal{C}_i^-(\alpha/2) \leq \theta_i^\star \leq \mathcal{C}_i^+(\alpha/2)\right] = \alpha. \tag{9}$$

This proposition can be extended to confidence intervals for any linear combination of $\boldsymbol{\theta}^\star$, such as $\theta_i^\star - \theta_j^\star$. These confidence intervals naturally enable hypothesis testings such as

$$\mathcal{H}_0 : \theta_j^\star \geq \theta_i^\star, \qquad \mathcal{H}_1 : \theta_j^\star < \theta_i^\star$$

and

$$\mathcal{H}_0 : \text{item } i \text{ is a top-}K \text{ item}, \qquad \mathcal{H}_1 : \text{item } i \text{ is not a top-}K \text{ item}.$$

Various other inference tasks have been considered in the ranking literature [LFL23][FLWY22, FLWY23]. We expect our general distributional characterization in Theorem 3 to be readily applied for these tasks.

4 Analysis

In this section, we present the main steps in the analysis to obtain theoretical results in the previous section. Section 4.1 gives a complete argument on the reduction to the BTL model we alluded in Section 2.2, Section 4.2 provides the analysis of the ℓ_∞ error, and Section 4.3 provides the analysis of both the ℓ_2 error and the asymptotic distribution as they share similar components in the analysis.

4.1 Reduction to Bradley-Terry-Luce model

A key idea in RP-MLE is the randomized pairing in Steps 1 and 2 of Algorithm 1. It compiles the user-item responses $\boldsymbol{X}$ to item-item comparisons $\boldsymbol{Y}$. In this section, we make a detailed argument that $\boldsymbol{Y}$ follows the Bradley-Terry-Luce model with a non-uniform sampling scheme.

Recall that $L_{ij}^t := \mathbb{1}\{X_{ti} \neq X_{tj}\}$ and $Y_{ij}^t := \mathbb{1}\{X_{ti} < X_{tj}\}$. The following fact provides the distribution of Y_{ij}^t conditional on $L_{ij}^t = 1$. We defer its proof to Section A.

Fact 1. *Let i, j be two items and t be a user. Suppose that user t has responded to both items i and j. Let X_{ti} and X_{tj} be the responses sampled from the probability model (1). Then we have*

$$\mathbb{P}[X_{ti} < X_{tj} \mid L_{ij}^t = 1] = \frac{e^{\theta_j^\star}}{e^{\theta_i^\star} + e^{\theta_j^\star}},$$

and

$$\mathbb{P}[L_{ij}^t = 1] \geq \frac{2\kappa_2}{(1+\kappa_2)^2}. \tag{10}$$

Fact 1 shows that conditional on $L_{ij}^t = 1$, Y_{ij}^t follows the BTL model with parameters $\boldsymbol{\theta}^\star$. More importantly, as we deploy random pairing (cf. Step 1a), each response X_{ti} is used at most once. As a result, conditional on $\{L_{ij}^t\}_{ijt}$, Y_{ij}^t's are jointly independent across users and items. In light of these, we can equivalently describe the data generating process of $\boldsymbol{Y}$ as follows:

1. For each user-item pair (t, i), there is a comparison between them with probability p independently.
2. Randomly split the n_t problems taken by tester t into $\lfloor n_t/2 \rfloor$ pairs of problems. (Step 1(a) of Algorithm 1)
3. For all (t, i, j), items i and j are compared by user t if $L_{ij}^t := \mathbb{1}\{X_{ti} \neq X_{tj}\} = 1$.
4. Conditioned on $L_{ij}^t = 1$, one observes the outcome $Y_{ij}^t := \mathbb{1}\{X_{ti} < X_{tj}\}$.

Steps 1–3 generates a non-uniform comparison graph $\mathcal{E}_Y$ between items. Step 4 reveals the independent outcomes of these comparisons following the BTL model, conditional on the graph $\mathcal{E}_Y$. This justifies that (2) is truly the likelihood function of the BTL model conditional on the comparison graph $\mathcal{E}_Y$.

4.2 Analysis of ℓ_∞ error

We have seen that analyzing RP-MLE under the Rasch model can be reduced to analyzing the MLE under the BTL model. This reduction allows us to invoke the general theory of MLE in the BTL model established in the recent work [YCOM24].

To facilitate the presentation, we introduce the necessary notation. For any $i \in [m]$, let $d_i := \sum_{j:j \neq i} L_{ij}$ be the weighted degree of item i in $\mathcal{G}_Y$ and $d_{\max} = \max_{i\in[m]} d_i$. Let the weighted graph Laplacian $\boldsymbol{L}_L$ be

$$\boldsymbol{L}_L := \sum_{i,j:i>j} L_{ij}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top.$$

The following lemma adapts Theorem 3 of the recent work [YCOM24] to our setting.

Lemma 1 (Theorem 3 in [YCOM24])**.** *Assume that $\mathcal{G}_Y$ is connected, and that*

$$[\lambda_{m-1}(\boldsymbol{L}_L)]^5 \geq C_1 \kappa_1^4 (d_{\max})^4 \log^2(n) \tag{11}$$

for some large enough constant $C_1 > 0$. Then with probability at least $1 - n^{-10}$, we have

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|_\infty \leq C_2 \kappa_1 \sqrt{\frac{\log(n)}{\lambda_{m-1}(\boldsymbol{L}_L)}}$$

for some constant $C_2 > 0$.

To leverage this general result, we need to characterize the spectral and degree properties of the comparison graph $\mathcal{G}_Y$, which is achieved in the following two lemmas. The proofs are deferred to Section B.

Lemma 2 (Degree bound in $\mathcal{G}_Y$)**.** *Suppose that $np \geq C\kappa_2^2 \log(n)$ for some large enough constant $C > 0$ and $m \leq n^\alpha$ for some sufficiently large constant $\alpha > 0$. With probability at least $1 - 2n^{-10}$, for all $i \in [m]$,*

$$\frac{1}{24\kappa_2} np \leq d_i \leq \frac{3}{2} np. \tag{12}$$

Lemma 3. *Suppose $mp \geq 2$, $np \geq C\kappa_2^2 \log(n)$ for some large enough constant C, and $m \leq n^\alpha$ for some constant $\alpha > 0$. With probability at least $1 - 10n^{-10}$, we have*

$$\frac{np}{4\kappa_2} \leq \lambda_{m-1}(\boldsymbol{L}_L) \leq \lambda_1(\boldsymbol{L}_L) \leq 3np \tag{13}$$

and

$$\frac{np}{16\kappa_1\kappa_2} \leq \lambda_{m-1}(\boldsymbol{L}_{Lz}) \leq \lambda_1(\boldsymbol{L}_{Lz}) \leq np. \tag{14}$$

4.2.1 Proof of Theorem 1

Now we are ready to prove Theorem 1. We focus on analyzing RP-MLE, as the analysis of MRP-MLE follows immediately from the union bound of the different data splitting and the triangular inequality:

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|_\infty \le \frac{1}{n_{\text{split}}} \sum_{i=1}^{n_{\text{split}}} \|\widehat{\boldsymbol{\theta}}^{(i)} - \boldsymbol{\theta}^\star\|_\infty.$$

By assumption we have $mp \ge 2$ and $np \ge C_1 \kappa_1^4 \kappa_2^5 \log^3(n)$ for some constant $C_1 > 0$. Then we can apply Lemmas 3 and 2 to see that

$$\frac{np}{4\kappa_2} \le \lambda_{m-1}(\boldsymbol{L}_L), \qquad \text{and} \qquad d_{\max} \le \frac{3}{2} np.$$

We observe that (11) is satisfied as long as $np \ge C_1 \kappa_1^4 \kappa_2^5 \log^3(n)$ for some constant C_1 that is large enough. Invoking Lemma 1, we conclude that

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|_\infty \le C_2 \kappa_1 \sqrt{\frac{\log(n)}{\lambda_{m-1}(\boldsymbol{L}_L)}} \le 2C_2 \kappa_1 \kappa_2^{1/2} \sqrt{\frac{\log(n)}{np}}.$$

It remains to show the top-K recovery sample complexity. As $\theta_1^\star \ge \cdots \ge \theta_K^\star > \theta_{K+1}^\star \ge \cdots \ge \theta_m^\star$ by assumption, it suffices to show $\widehat{\theta}_i - \widehat{\theta}_j > 0$ for any $i \le K$ and $j > K$. Using the ℓ_∞ error bound, we have that

$$\widehat{\theta}_i - \widehat{\theta}_j \ge \left(\theta_i^\star - \theta_j^\star\right) - \left|\widehat{\theta}_i - \theta_i^\star\right| - \left|\widehat{\theta}_j - \theta_j^\star\right| \ge \Delta_K - 4C_2 \kappa_1 \kappa_2^{1/2} \sqrt{\frac{\log(n)}{np}}.$$

Then $\widehat{\theta}_i - \widehat{\theta}_j > 0$ as long as

$$np \ge \frac{16 C_2^2 \kappa_1^2 \kappa_2 \log(n)}{\Delta_K^2}.$$

4.3 Analysis for refined ℓ_2 error and uncertainty quantification

To make the main text concise, we sketch the key ideas that lead to both Theorems 2 and 3.

We start with some necessary notation. For convenience, for any (i,j), renumber the set of comparisons $\{Y_{ji}^t\}_{t: L_{ij}^t = 1}$ as $\{Y_{ji}^{(l)}\}_{l=1,\dots,L_{ij}^t}$ in an arbitrary order. Let $\epsilon_{ij}^{(l)} := Y_{ji}^{(l)} - \sigma(\theta_i^\star - \theta_j^\star)$ where $\sigma(x) = e^x/(1+e^x)$ is the sigmoid function. Let $\boldsymbol{B} := \left[\cdots, \sqrt{z_{ij}}(\boldsymbol{e}_i - \boldsymbol{e}_j), \cdots\right]_{i>j:(i,j)\in\mathcal{E}_Y} \in \mathbb{R}^{m \times L_{\text{total}}}$ (repeat L_{ij} times for edge (i,j)), where $L_{\text{total}} := \sum_{i>j:(i,j)\in\mathcal{E}_Y} L_{ij}$ is the total number of observed effective comparisons in $\mathcal{G}_Y$. Define

$$\widehat{\boldsymbol{\epsilon}} := \left[\cdots, \epsilon_{ij}^{(l)}/\sqrt{z_{ij}}, \cdots\right]_{i>j:(i,j)\in\mathcal{E}_Y, l=1,\dots,L_{ij}^t} \in \mathbb{R}^{L_{\text{total}}}.$$

The key step in establishing Theorems 2 and 3 is the following refined characterization of the estimation error $\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star$.

Lemma 4. *Instate the assumptions of Theorem 1. With probability at least $1 - O(n^{-10})$, the estimator $\widehat{\boldsymbol{\theta}}$ given by the Algorithm (1) can be written as*

$$\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star = -\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B} \widehat{\boldsymbol{\epsilon}} + \boldsymbol{r}, \tag{15}$$

where $\boldsymbol{r} \in \mathbb{R}^m$ is a random vector obeying $\|\boldsymbol{r}\|_\infty \le C\kappa_1^6 \log^2(n)/(np)$ for some constant $C > 0$.

As the residual term $\boldsymbol{r}$ has small magnitude, we may analyze the properties of $\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star$ by focusing on the leading term $-\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B} \widehat{\boldsymbol{\epsilon}}$.

Proof outline of Theorem 2. Using the approximation $\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star \approx -\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}$, we can turn our attention to the concentration of

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|^2 \approx \|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2 = \widehat{\boldsymbol{\epsilon}}^\top \boldsymbol{B}^\top \boldsymbol{L}_{Lz}^\dagger \boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}.$$

This is a quadratic form of sub-Gaussian random variables. The Hanson-Wright inequality tells us that $\|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2$ concentrates around its expectation—Trace($\boldsymbol{L}_{Lz}^\dagger$). The complete proof is deferred to Section D for Theorem 2.

Proof outline of Theorem 3. With Lemma 4 in hand, we can expand $\boldsymbol{L}_{Lz}^{1/2}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star)$ to reach

$$\begin{aligned}\boldsymbol{L}_{Lz}^{1/2}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star) &\approx -(\boldsymbol{L}_{Lz}^\dagger)^{1/2}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}} \\ &= -\sum_{i>j:(i,j)\in\mathcal{E}_Y} \sum_{l=1}^{L_{ij}} \underbrace{\left[Y_{ji}^{(l)} - \sigma(\theta_i^\star - \theta_j^\star)\right](\boldsymbol{L}_{Lz}^\dagger)^{1/2}(\boldsymbol{e}_i - \boldsymbol{e}_j)}_{=:\boldsymbol{u}_{ij}^{(l)}}.\end{aligned}$$

The main component of the last expression is a sum of bounded random variables $\boldsymbol{u}_{ij}^{(l)}$ that are independent conditional on $\mathcal{G}_Y$. This sum has zero mean and variance

$$\sum_{i>j:(i,j)\in\mathcal{E}_Y} \sum_{l=1}^{L_{ij}} z_{ij}(\boldsymbol{L}_{Lz}^\dagger)^{1/2}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top(\boldsymbol{L}_{Lz}^\dagger)^{1/2} = \boldsymbol{I}_m - \frac{1}{m}\boldsymbol{1}_m\boldsymbol{1}_m^\top.$$

Therefore applying multivariate CLT, we obtain

$$-(\boldsymbol{L}_{Lz}^\dagger)^{1/2}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}} \xrightarrow{\mathrm{d}} \mathcal{N}\left(0, \boldsymbol{I}_m - \frac{1}{m}\boldsymbol{1}_m\boldsymbol{1}_m^\top\right).$$

See the complete proof in Section E.

4.4 Proof outline of Lemma 4

Due to the importance and usefulness of Lemma 4, we provide a sketch of its proof and leave the full one to Section C.

The proof is inspired by the proof of Theorem 1 in [Che23], which analyzes MLE via the trajectory of the preconditioned gradient descent (PGD) dynamic starting from ground truth. More precisely, letting $\boldsymbol{\theta}^0 = \boldsymbol{\theta}^\star$, we consider the PGD iterates defined by

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta\boldsymbol{L}_{Lz}^\dagger \nabla\mathcal{L}(\boldsymbol{\theta}^t),$$

where $\eta > 0$ is the step size of PGD. [Che23] shows that this dynamic converges to $\widehat{\boldsymbol{\theta}}$. We proceed one step further by establishing precise distributional characterization of $\hat{\boldsymbol{\theta}}$ via analyzing PGD. With Taylor expansion, the gradient can be decomposed into

$$\nabla\mathcal{L}(\boldsymbol{\theta}^t) = \boldsymbol{L}_{Lz}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star) - \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} + \boldsymbol{r}^t,$$

where $\boldsymbol{r}^t$ is a residual vector with small magnitude. Then the PGD update becomes

$$\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^\star = (1-\eta)(\boldsymbol{\theta}^t - \boldsymbol{\theta}^\star) - \eta\left(\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} - \boldsymbol{L}_{Lz}^\dagger \boldsymbol{r}^t\right).$$

We establish Lemma 4 by solving this recursive relation. More specifically, as $\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}$ does not depend on t and $\|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{r}^t\|_\infty$ can be controlled for each step t, taking $t \to \infty$, we see that

$$\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star = \lim_{t\to\infty} \boldsymbol{\theta}^t - \boldsymbol{\theta}^\star = -\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} + \boldsymbol{r}$$

for some residual term $\boldsymbol{r}$ that is well controlled in ℓ_∞ norm.

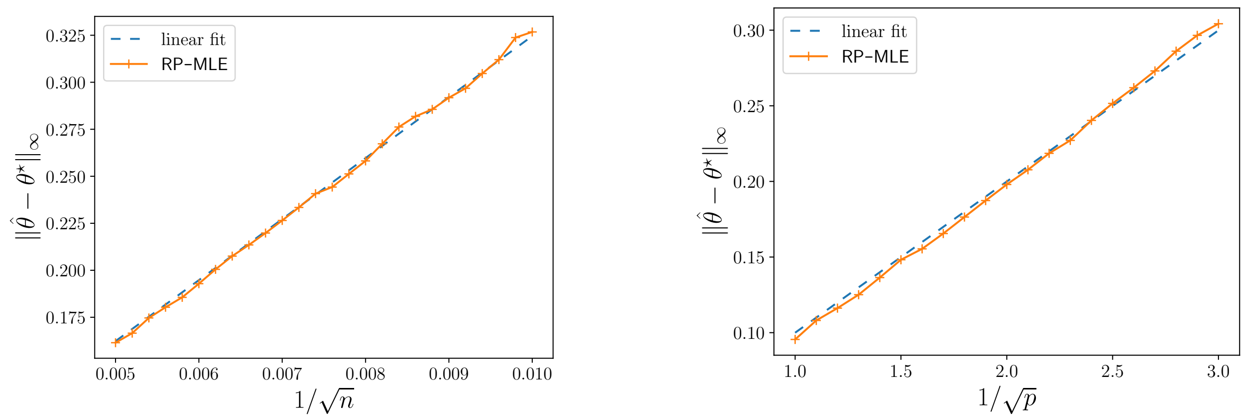

(a) $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|_\infty$ with v.s. $1/\sqrt{n}$. The parameter is chosen to be $m = 50, p = 0.1$ and n varies from 10000 to 40000. Each point represents the average of 1000 trials.

(b) $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|_\infty$ v.s. $1/\sqrt{p}$. The parameter is chosen to be $m = 50, n = 10000$ and p varies from 1/9 to 1. Each point represents the average of 1000 trials.

Figure 1: Estimation error $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|_\infty$ of RP-MLE with varying n and p.

5 Experiment

In this section, we demonstrate the empirical performance of RP-MLE and its variant MRP-MLE using both simulated and real data.

5.1 Simulations

We begin by using simulated data to validate our theoretical results and compare our estimators with existing ones for the Rasch model.

The data generating process follows the model specified in Section 2.1. Unless specified otherwise, in each trial, the user and item parameters are randomly drawn from

$$\widetilde{\boldsymbol{\zeta}}^{\star} \sim \mathcal{N}(0, \boldsymbol{I}_n), \qquad \text{and} \qquad \widetilde{\boldsymbol{\theta}}^{\star} \sim \mathcal{N}(0, \boldsymbol{I}_m).$$

Afterwards, $\boldsymbol{\zeta}^{\star}$ and $\boldsymbol{\theta}^{\star}$ is computed by shifting $\widetilde{\boldsymbol{\zeta}}^{\star}$ and $\widetilde{\boldsymbol{\theta}}^{\star}$ to zero mean.

5.1.1 ℓ_∞ estimation error

Here, we investigate the ℓ_∞ estimation error. First, we validate the theoretical result in ℓ_∞ estimation error of RP-MLE. Second, we show how much advantage the MRP-MLE brings through multiple runs of data splitting. Lastly, we compare our methods with existing comparison-based algorithms, including the case where κ_1, κ_2 are large.

Validating the theoretical result. Theorem 1 tells us that the ℓ_∞ error scales as $1/\sqrt{np}$. As shown in Figure 1, $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|_\infty$ is very close to being linear w.r.t. both $1/\sqrt{n}$ and $1/\sqrt{p}$. This is consistent with our theoretical predictions.

Multiple runs in MRP-MLE. As we have discussed after Theorem 1, the random data splitting could incur a small loss of information. We have provided a remedy MRP-MLE (Algorithm 2) by averaging over multiple runs with independent data splitting. In Figure 2, we show the empirical evidence of its effectiveness. For each trial, we run 50 independent data splitting and obtain the 50 estimates $\widehat{\boldsymbol{\theta}}^{(i)}$. We then report the ℓ_∞ error of the running average, i.e., $\|\frac{1}{k}\sum_{i=1}^{k}\widehat{\boldsymbol{\theta}}^{(i)} - \boldsymbol{\theta}^{\star}\|$ for $k = 1, \ldots, 50$. As a baseline for comparison, we also run pseudo MLE on the same datasets, which uses all possible pairs of comparison. We can see that by averaging over more runs of data splittings, MRP-MLE gives a better estimation. In particular, one can see

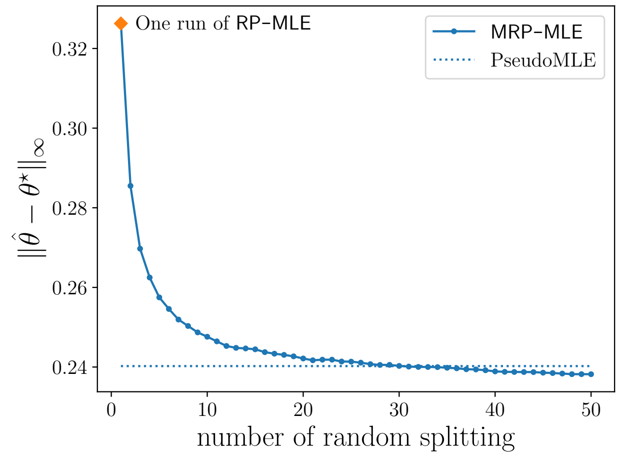

Figure 2: The ℓ_∞ estimation error of MRP-MLE with different numbers of data splitting. The dashed line is the performance of Pseudo MLE. The parameters are chosen to be $m = 50, p = 0.1, n = 10000$. Each point is averaged over 1000 trials.

that with sufficient number of independent data splittings, the average estimation error matches and even beats pseudo MLE.

Comparison with existing estimator. We then compare our algorithms with two other comparison-based algorithms: pseudo MLE (PMLE) and the spectral method in [NZ23]. We plot the ℓ_∞ error $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|_\infty$ against n in Figure 3. We can see that all four algorithms follow a similar trend. The performance of MRP-MLE is on par with PMLE and better than the spectral method. Our proposed algorithm not only enjoys an improved theoretical guarantee but is also competitive in practice.

Performance with large κ_1, κ_2. While we assume $\kappa = \max\{\kappa_1, \kappa_2\} = O(1)$ is most of this article, the scenario with large κ can be practically relevant. Here we compare the ℓ_∞ error of different methods when the condition number is large. Fixing κ, we draw the user and item parameters with

$$\widetilde{\boldsymbol{\zeta}}^{\star} \sim \mathrm{Unif}(0, \log(\kappa)) \qquad \text{and} \qquad \widetilde{\boldsymbol{\theta}}^{\star} \sim \mathrm{Unif}(0, \log(\kappa))$$

and compute $\boldsymbol{\zeta}^{\star}$ and $\boldsymbol{\theta}^{\star}$ by shifting $\widetilde{\boldsymbol{\zeta}}^{\star}$ and $\widetilde{\boldsymbol{\theta}}^{\star}$ to zero mean. Figure 4 shows the performance of different estimators as κ varies. The MLE-based approaches including RP-MLE and MRP-MLE achieve better ℓ_∞ error than the spectral method when κ is large.

5.1.2 Top-K recovery

In this part, we investigate the performance of different algorithms in top-K recovery. We set $\theta_i^{\star} = (1 - K/m)\Delta_K$ for $i \le K$ and $\theta_i^{\star} = (-K/m)\Delta_K$ otherwise. For any estimator $\widehat{\boldsymbol{\theta}}$, we define top-$K$ recovery rate to be

$$\frac{1}{K} \left| \{ i \le K : i \in A_K \} \right|,$$

where A_K is an arbitrary K-element set such that $\widehat{\theta}_i \ge \widehat{\theta}_j$ for any $i \in A_K, j \notin A_K$. We compare the top-K recovery rate of PMLE and the spectral method in [NZ23] with RP-MLE and MRP-MLE in Figure 5. The recovery rate of PMLE, spectral method and MRP-MLE is similar, indicating again that our algorithm performs well in practice.

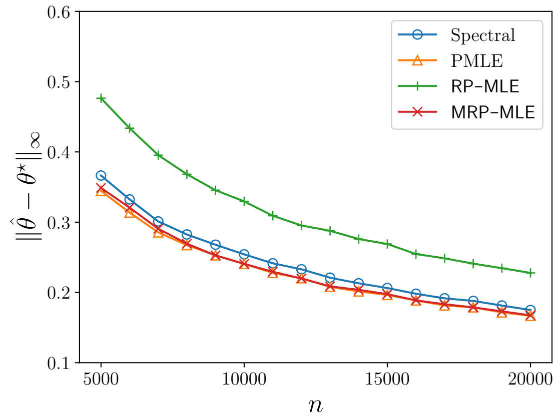

Figure 3: $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|_\infty$ v.s. n using Spectral method, PMLE, RP-MLE, and MRP-MLE using 20 data splittings. The parameter is chosen to be $m = 50, p = 0.1$ and n varies from 5000 to 20000. The result is averaged over 1000 trials.

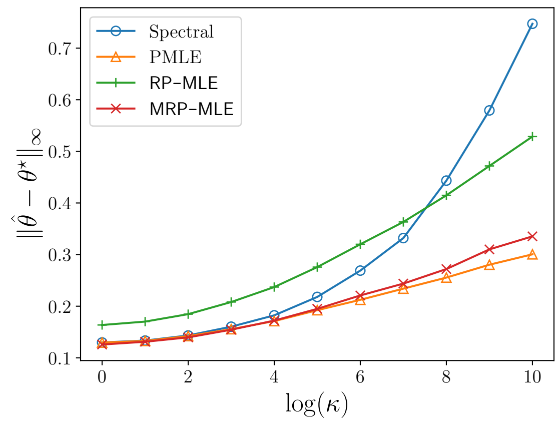

Figure 4: $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|_\infty$ v.s. $\log(\kappa)$ using Spectral method, PMLE, RP-MLE, and MRP-MLE using 20 data splittings. The parameter is chosen to be $m = 50, p = 0.1, n = 20000$ and κ varies from 1 to e^{10}. The result is averaged over 1000 trials.

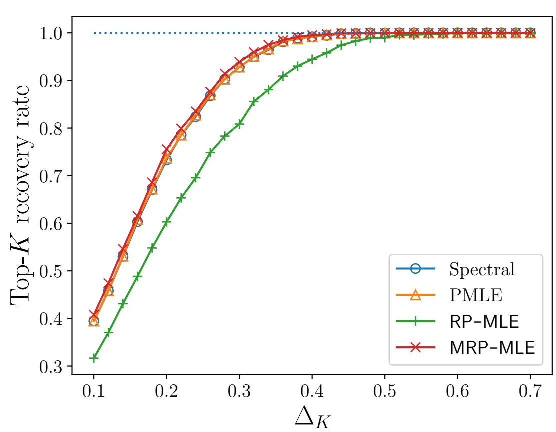

Figure 5: Top-K recovery rate using spectral method, PMLE, RP-MLE, and MRP-MLE using 20 data splittings. The parameter is chosen to be $m = 10000, m = 50, p = 0.1, K = 5$ and Δ_K varies from 0.1 to 0.7. The result is averaged over 1000 trials.

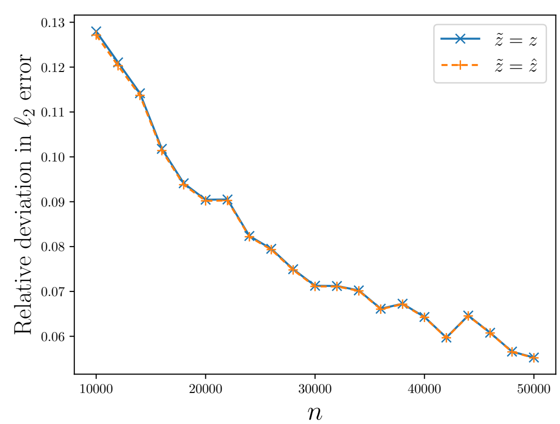

Figure 6: The relative deviation of $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\|$ from $\sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\tilde{z}}^\dagger)}$ v.s. n for both $\widetilde{z} = \widehat{z}$ and $\widetilde{z} = z$. The parameter is chosen to be $p = 0.1$, n varies from 10000 to 50000 and $m = n/500$. The result is averaged over 1000 trials.

5.1.3 Refined ℓ_2 estimation error

In this part of the section, we show the empirical evidence of Theorem 2. In Theorem 2 we have shown that the ℓ_2 error concentrate around $\sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\widetilde{z}}^{\dagger})}$ for $\widetilde{z} \in \{\widehat{z}, z\}$. In each trial we compute the following quantity

$$\frac{\left| \|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\| - \sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\widetilde{z}}^{\dagger})} \right|}{\sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\widetilde{z}}^{\dagger})}}$$

for both $\widetilde{z} = \widehat{z}$ and $\widetilde{z} = z$. This measures the relative deviation of $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|$ from $\sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\widetilde{z}}^{\dagger})}$. In Figure 6 we consider the regime where p and n/m is fixed. In this case, Theorem 2 implies that

$$\frac{\left| \|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\| - \sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\widetilde{z}}^{\dagger})} \right|}{\sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\widetilde{z}}^{\dagger})}} \lesssim \frac{\sqrt{\frac{1}{np} + \frac{\sqrt{m}}{np}}}{\sqrt{\frac{m}{np}}} \lesssim \frac{1}{\sqrt{n}}.$$

In other words, $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|$ concentrate tightly around $\mathrm{Trace}(\boldsymbol{L}_{L\widetilde{z}}^{\dagger})$. We can see that the deviation is very small between $\widetilde{z} = \widehat{z}$ and $\widetilde{z} = z$. In both cases, the relative deviation of $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|$ from $\sqrt{\mathrm{Trace}(\boldsymbol{L}_{L\widetilde{z}}^{\dagger})}$ decreases as n and m increase as expected.

5.1.4 Inference

Lastly, we validate our inferential results in Section 3.3.

Asymptotic normality. We first check the asymptotic normality presented in Theorem 3 and Corollary 1. For simplicity, we focus on the first coordinate. Theorem 3 and Corollary 1 have claimed that

$$\left[\boldsymbol{L}_{Lz}^{1/2}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star})\right]_1 \xrightarrow{\mathrm{d}} \mathcal{N}\left(0, 1 - \frac{1}{m}\right)$$

and

$$\frac{\widehat{\theta}_1 - \theta_1^{\star}}{\sqrt{[\boldsymbol{L}_{Lz}^{\dagger}]_{11}}} \xrightarrow{\mathrm{d}} \mathcal{N}(0, 1).$$

For simplicity we set both $\boldsymbol{\zeta}^{\star}$ and $\boldsymbol{\theta}^{\star}$ to be all zero vectors. In Figure 7, we present the QQ-plot comparing the generated instances and the theoretical distributions. In both cases, the spread of the term of interest matches the normality claimed in Theorem 3 and Corollary 1.

Coverage of the confidence intervals. We compare the empirical coverage rate of the two-sided confidence intervals in Proposition 3 with the theoretical ones. Without loss of generality we consider the confidence interval for the first coordinate. In each trial, we compute the empirical coverage rate

$$1 - \widehat{\alpha} = \frac{1}{m} \sum_{i=1}^{m} \mathbb{1}\{\mathcal{C}_i^{-}(\alpha/2) \le \theta_i^{\star} \le \mathcal{C}_i^{+}(\alpha/2)\}$$

and report the average $1 - \widehat{\alpha}$ over the trials. We do this for multiple α's and compare $1 - \alpha$ with $1 - \widehat{\alpha}$ in Table 1. The empirical coverage rate $1 - \widehat{\alpha}$ is very close to the theoretical coverage rate $1 - \alpha$.

5.2 LSAT real data

We study a real dataset (LSAT) on the Law School Admissions Test from [DBL70]. LSAT has full observation of 1000 people answering 5 problems. Foro each person-item pair it records whether the answer is correct or not. The second row in Table 2 lists how many people answer each problem correctly. From the first

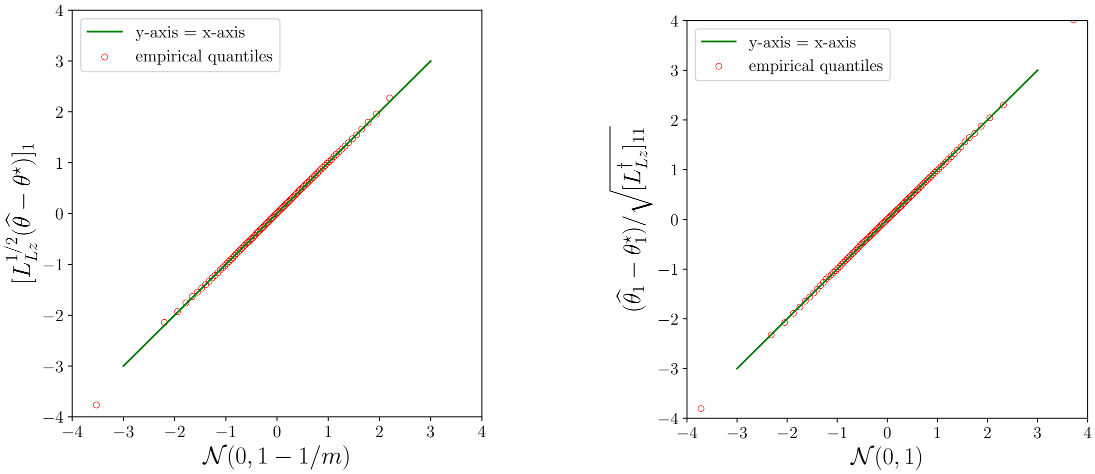

Figure 7: Quantile-quantile plots of $\left[\boldsymbol{L}_{Lz}^{1/2}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star)\right]_1$ and $(\widehat{\theta}_1 - \theta_1^\star)/\sqrt{[\boldsymbol{L}_{Lz}^\dagger]_{11}}$ v.s. zero-mean normal distributions. The estimate $\widehat{\boldsymbol{\theta}}$ is computed with RP-MLE. The parameters are set to be $n = 10000, m = 10, p = 0.5$. Data is collected over 10000 trials.

$1-\alpha$	0.8	0.9	0.95	0.98	0.99	0.995	0.999
$1-\widehat{\alpha}$	0.801145	0.900635	0.95108	0.980395	0.99025	0.99523	0.99904

Table 1: Theoretical coverage rate $1-\alpha$ and empirical coverage rate $1-\widehat{\alpha}$ of the two-sided confidence intervals $[\mathcal{C}_i^-(\alpha/2), \mathcal{C}_i^+(\alpha/2)]$. The parameters are set to be $n = 10000, m = 20, p = 0.1$. The result is averaged over 10000 trials.

look, Problem 3 seems to be the hardest question. We proceed to quantify the hardness of these problems under the Rasch model and infer how confident we are in claiming it is the hardest. We use RP-MLE to get an estimate of the latent scores and calculate two-sided confidence intervals at level $\alpha = 0.01$ for each coordinate with the framework in Section 3.3. See Table 2 for the result. Higher latent score here means higher difficulty. The estimated parameters agree with the total number of correct answers in reverse order. We can also see that the confidence interval lower bound of $\theta_3^\star$ is larger than the largest confidence interval upper bound of $\theta_1^\star, \theta_2^\star, \theta_4^\star, \theta_5^\star$. With Bonferroni correction, we have a 0.95 confidence level to say Problem 3 is the most difficult problem.

Now we assume Problem 3 is the top-1 item and investigate the top-1 recovery rate of different algorithms on LSAT under incomplete observation. In each trial we randomly select $\widetilde{n}$ people, and for each of them we randomly select their outcome on $\widetilde{m}$ problems. We then estimate $\boldsymbol{\theta}^\star$ using the subsampled data with different methods and compare their top-1 recovery rate, i.e., the proportion of trials where the top-1 item is correctly identified. In Figure 8 we see results similar to the simulation. Our algorithm MRP-MLE has a similar recovery rate compared to PMLE and spectral method.

Problem	1	2	3	4	5
Total correct	924	709	553	763	870
θ estimate	-1.2139	0.4093	1.1123	0.2977	-0.6053
CI lower bound	-1.5957	0.1449	0.8321	0.0234	-0.9050
CI upper bound	-0.8322	0.6736	1.3926	0.5720	-0.3057

Table 2: Latent score estimate calculated using RP-MLE and confidence interval calculated with the framework in Section 3.3. The significance level is chosen to be $\alpha = 0.01$ for each coordinate.

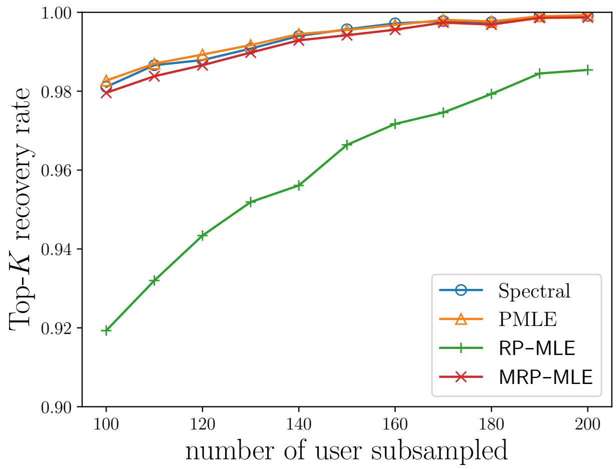

Figure 8: Top-1 recovery rate using Spectral method, PMLE, RP-MLE, and MRP-MLE using 20 data splittings. The parameters are chosen to be $\widetilde{m} = 4$ and $\widetilde{n}$ varies from 100 to 200. The result is averaged over 10000 trials.

6 Discussion

In this paper, we propose two new likelihood-based estimators RP-MLE and MRP-MLE for item parameter estimation in the Rasch model. Both enjoy optimal finite sample estimation guarantee and RP-MLE is further amenable to tight uncertainty quantification. All this is achieved even when the user-item response data are extremely sparse (cf. [NZ23]). Below, we identify several questions that are interesting for further investigation:

- **Does PMLE achieve optimal theoretical guarantee?** In our experiments, pseudo MLE has shown a similar performance to MRP-MLE. This naturally leads to the question of whether PMLE can enjoy the same theoretical guarantee. This is relevant to our work because our methods can be viewed as a modification of pseudo MLE by incorporating random disjoint pairing to decouple statistical dependency among paired Y_{ij}'s. It remains unclear whether such dependency is a fundamental bottleneck.

- **Inference for multiple data splitting.** In MRP-MLE (Algorithm 2), we average the estimates over multiple runs of data splitting to improve the empirical result over RP-MLE. While this is effective for estimation, we do not have an immediate analog for inference as the estimations $\{\widehat{\boldsymbol{\theta}}^{(i)}\}_{i\in[n_{\text{split}}]}$ are dependent on each other. It is an interesting question whether we can characterize the distribution of $\widehat{\boldsymbol{\theta}}$ in MRP-MLE and whether it leads to better inference, e.g., tighter confidence intervals.

- **Extending random pairing to other models in IRT.** Some IRT models parameterize the latent score of users and items differently from the Rasch model. For instance, the two-parameter logistic model (2PL) assumes that X_{ti}, the response of user t to item i, follows the law

$$\mathbb{P}[X_{ti} = 1] = \frac{\exp(a_t^\star + b_t^\star \theta_i^\star)}{1 + \exp(a_t^\star + b_t^\star \theta_i^\star)},$$

where $\boldsymbol{\theta}^\star$ is the latent score of the items while $\boldsymbol{a}^\star$ and $\boldsymbol{b}^\star$ are two parameter vectors for the users. Unlike the Rasch model, in the 2PL model,

$$\mathbb{P}[X_{ti} > X_{tj} \mid X_{ti} \neq X_{tj}] = \frac{\exp(b_t^\star \theta_i^\star)}{\exp(b_t^\star \theta_i^\star) + \exp(b_t^\star \theta_j^\star)}$$

is not independent of the user parameter $\boldsymbol{b}^\star$. Therefore reduction to the BTL model is no longer true in this case. It is interesting and non-trivial to extend the idea of random pairing to the 2PL model.

Acknowledgement

YY and CM were partially supported by the National Science Foundation via grant DMS-2311127.

References

[BT52] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.

[CFMW19] Yuxin Chen, Jianqing Fan, Cong Ma, and Kaizheng Wang. Spectral method and regularized mle are both optimal for top-k ranking. *Annals of statistics*, 47(4):2204, 2019.

[CGZ22] Pinhan Chen, Chao Gao, and Anderson Y Zhang. Partial recovery for top-k ranking: Optimality of mle and suboptimality of the spectral method. *The Annals of Statistics*, 50(3):1618–1652, 2022.

[Che23] Yanxi Chen. Ranking from pairwise comparisons in general graphs and graphs with locality. *arXiv preprint arXiv:2304.06821*, 2023.

[Cho82] Bruce Choppin. A fully conditional estimation procedure for rasch model parameters. 1982.

[CLLY21] Yunxiao Chen, Xiaoou Li, Jingchen Liu, and Zhiliang Ying. Item response theory–a statistical framework for educational and psychological measurement. *arXiv preprint arXiv:2108.08604*, 2021.

[CLOX23] Yunxiao Chen, Chengcheng Li, Jing Ouyang, and Gongjun Xu. Statistical inference for noisy incomplete binary matrix. *Journal of Machine Learning Research*, 24(95):1–66, 2023.

[CS15] Yuxin Chen and Changho Suh. Spectral mle: Top-k rank aggregation from pairwise comparisons. In *International Conference on Machine Learning*, pages 371–380. PMLR, 2015.

[DBL70] R Darrell Bock and Marcus Lieberman. Fitting a response model for n dichotomously scored items. *Psychometrika*, 35(2):179–197, 1970.

[DC10] André F De Champlain. A primer on classical test theory and item response theory for assessments in medical education. *Medical education*, 44(1):109–117, 2010.

[ER13] Susan E Embretson and Steven P Reise. *Item response theory*. Psychology Press, 2013.

[FBC05] James F Fries, Bonnie Bruce, and David Cella. The promise of promis: using item response theory to improve assessment of patient-reported outcomes. *Clinical and experimental rheumatology*, 23(5):S53, 2005.

[FHY22] Jianqing Fan, Jikai Hou, and Mengxin Yu. Uncertainty quantification of mle for entity ranking with covariates. *arXiv preprint arXiv:2212.09961*, 2022.

[FLWY22] Jianqing Fan, Zhipeng Lou, Weichen Wang, and Mengxin Yu. Ranking inferences based on the top choice of multiway comparisons. *arXiv preprint arXiv:2211.11957*, 2022.

[FLWY23] Jianqing Fan, Zhipeng Lou, Weichen Wang, and Mengxin Yu. Spectral ranking inferences based on general multiway comparisons. *arXiv preprint arXiv:2308.02918*, 2023.

[Gho95] Malay Ghosh. Inconsistent maximum likelihood estimators for the rasch model. *Statistics & Probability Letters*, 23(2):165–170, 1995.

[GSZ23] Chao Gao, Yandi Shen, and Anderson Y Zhang. Uncertainty quantification in the bradley–terry–luce model. *Information and Inference: A Journal of the IMA*, 12(2):1073–1140, 2023.

[JKSO16] Minje Jang, Sunghyun Kim, Changho Suh, and Sewoong Oh. Top-k ranking from pairwise comparisons: When spectral ranking is optimal. *arXiv preprint arXiv:1603.04153*, 2016.

[LFL23] Yue Liu, Ethan X Fang, and Junwei Lu. Lagrangian inference for ranking problems. *Operations research*, 71(1):202–223, 2023.

[Lin99] John M Linacre. Understanding rasch measurement: estimation methods for rasch measures. *Journal of outcome measurement*, 3:382–405, 1999.

[LNB68] FM Lord, MR Novick, and Allan Birnbaum. Statistical theories of mental test scores. 1968.

[LSR22] Wanshan Li, Shamindra Shrotriya, and Alessandro Rinaldo. ℓ_∞ bounds of the MLE in the BTL model under general comparison graphs. In *Uncertainty in Artificial Intelligence*, pages 1178–1187. PMLR, 2022.

[Luc59] R Duncan Luce. Individual choice behavior. 1959.

[Mol95] Ivo W. Molenaar. *Estimation of Item Parameters*, pages 39–51. Springer New York, New York, NY, 1995.

[NZ22] Duc Nguyen and Anderson Ye Zhang. A spectral approach to item response theory. *Advances in neural information processing systems*, 35:38818–38830, 2022.

[NZ23] Duc Nguyen and Anderson Ye Zhang. Optimal and private learning from human response data. In *International Conference on Artificial Intelligence and Statistics*, pages 922–958. PMLR, 2023.

[Ost59] Alexander M Ostrowski. A quantitative formulation of sylvester's law of inertia. *Proceedings of the National Academy of Sciences*, 45(5):740–744, 1959.

[Rai19] Martin Raič. A multivariate berry–esseen theorem with explicit constants. 2019.

[Ras60] Georg Rasch. Studies in mathematical psychology: I. probabilistic models for some intelligence and attainment tests. 1960.

[Rob21] Alexander Robitzsch. A comparison of estimation methods for the rasch model. 2021.

[RV13] Mark Rudelson and Roman Vershynin. Hanson-wright inequality and sub-gaussian concentration. 2013.

[Sch92] Bernhard A Schmitt. Perturbation bounds for matrix square roots and pythagorean sums. *Linear algebra and its applications*, 174:215–227, 1992.

[Spi07] Daniel A Spielman. Spectral graph theory and its applications. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 29–38. IEEE, 2007.

[Tro15] Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.

[Ver18] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

[VHSA20] Steven M Van Hauwaert, Christian H Schimpf, and Flavio Azevedo. The measurement of populist attitudes: Testing cross-national scales using item response theory. *Politics*, 40(1):3–21, 2020.

[Wed73] Per-Åke Wedin. Perturbation theory for pseudo-inverses. *BIT Numerical Mathematics*, 13:217–232, 1973.

[YCOM24] Yuepeng Yang, Antares Chen, Lorenzo Orecchia, and Cong Ma. Top-K ranking with a monotone adversary. *Conference on learning theory*, 2024.

[Zwi95] Aeilko H Zwinderman. Pairwise parameter estimation in rasch models. *Applied Psychological Measurement*, 19(4):369–375, 1995.

A Proof of Fact 1

Expanding the probability of the random events, we have that

$$\begin{aligned}
\mathbb{P}[X_{ti} < X_{tj} \mid X_{ti} \neq X_{tj}] &= \frac{\mathbb{P}[X_{ti} = 0, X_{tj} = 1]}{\mathbb{P}[X_{ti} = 0, X_{tj} = 1 \text{ or } X_{ti} = 1, X_{tj} = 0]} \\
&= \frac{e^{\zeta_t^\star} e^{\theta_j^\star}}{(e^{\zeta_t^\star} + e^{\theta_i^\star})(e^{\zeta_t^\star} + e^{\theta_j^\star})} \Bigg/ \left(\frac{e^{\zeta_t^\star} e^{\theta_j^\star}}{(e^{\zeta_t^\star} + e^{\theta_i^\star})(e^{\zeta_t^\star} + e^{\theta_j^\star})} + \frac{e^{\theta_i^\star} e^{\zeta_t^\star}}{(e^{\zeta_t^\star} + e^{\theta_i^\star})(e^{\zeta_t^\star} + e^{\theta_j^\star})} \right) \\
&= \frac{e^{\zeta_t^\star} e^{\theta_j^\star}}{e^{\zeta_t^\star} (e^{\theta_i^\star} + e^{\theta_j^\star})} \\
&= \frac{e^{\theta_j^\star}}{e^{\theta_i^\star} + e^{\theta_j^\star}}.
\end{aligned}$$

Now consider $\mathbb{P}[X_{ti} \neq X_{tj}]$, we have

$$\begin{aligned}
\mathbb{P}[X_{ti} \neq X_{tj}] &= \frac{e^{\zeta_t^\star} e^{\theta_j^\star}}{(e^{\zeta_t^\star} + e^{\theta_i^\star})(e^{\zeta_t^\star} + e^{\theta_j^\star})} + \frac{e^{\theta_i^\star} e^{\zeta_t^\star}}{(e^{\zeta_t^\star} + e^{\theta_i^\star})(e^{\zeta_t^\star} + e^{\theta_j^\star})} \\
&= \frac{e^{\theta_j^\star - \zeta_t^\star} + e^{\theta_i^\star - \zeta_t^\star}}{(1 + e^{\theta_i^\star - \zeta_t^\star})(1 + e^{\theta_j^\star - \zeta_t^\star})}. \qquad (16)
\end{aligned}$$

Let $f : [1/\kappa_2, \kappa_2]^2 \to \mathbb{R}$ defined by

$$f(a, b) \coloneqq \frac{a + b}{(1 + a)(1 + b)}.$$

Its partial derivatives are

$$\frac{\partial}{\partial a} f(a, b) = \frac{b^2 - 1}{(1 + a)^2 (1 + b)^2} \quad \text{and} \quad \frac{\partial}{\partial b} f(a, b) = \frac{a^2 - 1}{(1 + a)^2 (1 + b)^2}.$$

It is now easy to see that the minimum or maximum of f can only happen if $(a, b) = (1, 1)$ or $(a, b) \in \{1/\kappa_2, \kappa_2\}^2$. After comparing the value of f at these points, we conclude that f achieves minimum at

$$f(1/\kappa_2, 1/\kappa_2) = f(\kappa_2, \kappa_2) = \frac{2\kappa_2}{(1 + \kappa_2)^2}.$$

By the definition of κ_2, $|\theta_l^\star - \zeta_t^\star| \leq \log(\kappa_2)$ for any $l \in [m]$. Then (16) fits the definition of f and the proof is completed.

B Degree and spectral properties of the comparison graphs

In this section, we present the analysis for lemmas that characterize the degree and spectral properties of the comparison graphs. We start with a lemma that controls the degrees in $\mathcal{G}_X$, and then prove Lemmas 2 and 3.

B.1 Degree range of $\mathcal{G}_X$

Recall that n_t is the number of neighbors of user t in $\mathcal{G}_X$. Furthermore, we denote m_i as the number of users that is compared with problem i and at least another item, i.e.,

$$m_i \coloneqq |\{t : (t, i) \in \mathcal{E}_X, n_t \geq 2\}|. \qquad (17)$$

The following lemma controls the size of n_t and m_i.

Lemma 5 (Degree bounds in $\mathcal{G}_X$)**.** *Suppose that $np \geq C\log(n)$ for some large enough constant $C > 0$ and that $m \leq n^\alpha$ for some constant $\alpha > 0$. Then with probability at least $1 - 2n^{-10}$, for all $i \in [m]$, we have*

$$\frac{1}{4}np \leq m_i \leq \frac{3}{2}np. \tag{18}$$

Moreover, with probability at least $1 - n^{-10}$, for all $t \in [n]$, we have

$$n_t \leq \left(\frac{3}{2}mp\right) \vee 165\log(n). \tag{19}$$

Proof. We prove the two claims in the lemma sequentially.

Fix any t, i. One has

$$\begin{aligned}\mathbb{P}\left[t : (t,i) \in \mathcal{E}_X, n_t \geq 2\right] &= \mathbb{P}\left[(t,i) \in \mathcal{E}_X\right] - \mathbb{P}\left[(t,i) \in \mathcal{E}_X \text{ and } n_t = 1\right] \\ &= p - p(1-p)^{m-1} \\ &\geq p(1 - e^{-(m-1)p}) \\ &\geq \frac{1}{2}p, \end{aligned} \tag{20}$$

as long as $mp \geq 2$. Let $\mu_i := \mathbb{E}[m_i]$. By the linearity of expectation, we have

$$np/2 \leq \mu_i \leq \sum_t \mathbb{P}\left[(t,i) \in \mathcal{E}_X\right] = np. \tag{21}$$

Fix $i \in [m]$. Since the sampling is independent with different t, by the Chernoff bound,

$$\mathbb{P}[|m_i - \mu_i| \leq (1/2)\mu_i] \leq 2e^{-\frac{1}{12}\mu_i} \leq 2e^{-\frac{1}{24}np} \leq m^{-1}n^{-10}$$

as long as $np \geq C\log(n)$ for large enough constant C. Applying (21) and union bound on $i \in [m]$ yields (18).

Moving on to (19), we first consider the case where $mp \geq 110\log(n)$. By Chernoff bound,

$$\mathbb{P}[n_t \geq (3/2)mp] \leq e^{-\frac{1}{10}mp} \leq n^{-11}.$$

In the case of $mp < 110\log(n)$, the quantity $\mathbb{P}\left[n_t \geq 165\log(n)\right]$ clear decreases as p decreases. So we may use the case $mp = 110\log(n)$ to bound this quantity and conclude that

$$\mathbb{P}\left[n_t \geq 165\log(n)\right] \leq n^{-11}.$$

Finally we apply union bound on $t \in [n]$ to reach (19).

□

B.2 Proof of Lemma 2

The assumption of this lemma allows us to invoke Lemma 5. For the upper bound of d_i, since (18) is true,

$$\begin{aligned} d_i &= \sum_{j:j\neq i} L_{ij} \\ &= \sum_{t:(t,i)\in\mathcal{E}_X} \sum_{j:j\neq i} L_{ij}^t \\ &\overset{\text{(i)}}{=} \sum_{t:(t,i)\in\mathcal{E}_X, n_t\geq 2} \sum_{j:j\neq i} L_{ij}^t \\ &\overset{\text{(ii)}}{\leq} \sum_{t:(t,i)\in\mathcal{E}_X, n_t\geq 2} \sum_{j:j\neq i} R_{ij}^t \leq m_i \leq \frac{3}{2}np. \end{aligned}$$

Here (i) holds since L_{ij}^t can only be 0 when $n_t \le 1$, and (ii) holds since $L_{ij}^t \le R_{ij}^t$ by definition. For the lower bound of d_i, notice that for any (t,i), $\sum_j L_{ij}^t$ is either 0 or 1. Fix $\mathcal{E}_X$ and only consider randomness on L_{ij}^t. By Hoeffding's inequality,

$$\begin{aligned}
\mathbb{P}\left\{ d_i - \sum_{t:(t,i)\in\mathcal{E}_X, n_t\ge 2} \mathbb{E}\left[\sum_{j:j\neq i} L_{ij}^t \mid (t,i)\in\mathcal{E}_X\right] \le -\frac{1}{12\kappa_2} np \right\} &\le \exp\left(-\frac{(1/72)\kappa_2^{-2} n^2 p^2}{m_i}\right) \\
&\le \exp\left(-\frac{\kappa_2^{-2} np}{108}\right) \\
&\le m^{-1} n^{-10} \qquad (22)
\end{aligned}$$

as long as $np \ge 1200\kappa_2^2 \log(n)$. The second to last inequality uses (18). For each $(t,i)\in\mathcal{E}_X$,

$$\begin{aligned}
\mathbb{E}\left[\sum_j L_{ij}^t \mid (t,i)\in\mathcal{E}_X\right] &= \sum_j \mathbb{P}\left[R_{ij}^t = 1 \mid (t,i)\in\mathcal{E}_X\right] \mathbb{P}\left[\sum_j L_{ij}^t = 1 \mid R_{ij}^t = 1\right] \\
&\ge \sum_j \mathbb{P}\left[R_{ij}^t = 1 \mid (t,i)\in\mathcal{E}_X\right] \frac{2\kappa_2}{(1+\kappa_2)^2} \\
&\ge \frac{2\lfloor n_t/2 \rfloor}{n_t} \cdot \frac{2\kappa_2}{(1+\kappa_2)^2} \ge \frac{1}{3\kappa_2}
\end{aligned}$$

as long as $n_t \ge 2$. The first inequality here uses Fact 1. Then by definition of m_i in (17),

$$\sum_{t:(t,i)\in\mathcal{E}_X, n_t\ge 2} \mathbb{E}\left[\sum_j L_{ij}^t \mid (t,i)\in\mathcal{E}_X\right] \ge \frac{1}{3\kappa_2} m_i. \qquad (23)$$

Combining (22), (23) and (18),

$$d_i \ge \frac{1}{3\kappa_2} m_i - \frac{1}{12\kappa_2} np \ge \frac{1}{24\kappa_2} np.$$

Applying union bound over $i \in [m]$ yields the desired result.

B.3 Proof of Lemma 3

We first consider the spectrum of $\boldsymbol{L}_L$. Recall that

$$\begin{aligned}
\boldsymbol{L}_L &= \sum_{(i,j)\in\mathcal{E}_Y, i>j} L_{ij} (\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top \\
&= \sum_{t=1}^n \underbrace{\sum_{(i,j)\in\mathcal{E}_Y, i>j} L_{ij}^t (\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top}_{\boldsymbol{L}_L^t}.
\end{aligned}$$

For the upper bound, it is clear from Lemmas 10 and 2 that $\lambda_1(\boldsymbol{L}_L) \le 2\max_i d_i \le 3np$. For the lower bound, we use the matrix Chernoff inequality (see Section 5 of [Tro15]). Let $\boldsymbol{R} \in \mathbb{R}^{(n-1)\times n}$ be a partial isometry such that $\boldsymbol{R}\boldsymbol{R}^\top = \boldsymbol{I}_{m-1}$ and $\boldsymbol{R}\boldsymbol{1} = \boldsymbol{0}$. Then $\lambda_{m-1}(\boldsymbol{L}_L) = \lambda_{m-1}(\boldsymbol{R}\boldsymbol{L}_L\boldsymbol{R}^\top)$. For any $t\in[n]$, by (19),

$$0 \le \lambda_{m-1}(\boldsymbol{R}\boldsymbol{L}_L^t\boldsymbol{R}^\top) \le \lambda_1(\boldsymbol{R}\boldsymbol{L}_L^t\boldsymbol{R}^\top) = \lambda_1(\boldsymbol{L}_L^t) \le 2.$$

The last inequality follows from Lemma 10 since $\sum_j \boldsymbol{L}_{ij}^t \le 1$ for any $t\in[n]$ and $i\in[m]$. By Fact 1,

$$\mathbb{P}[L_{ij}^t = 1 \mid R_{ij}^t = 1] \ge \frac{2\kappa_2}{(1+\kappa_2)^2} \ge 1/(2\kappa_2).$$

Then

$$\lambda_{m-1}(\mathbb{E}\boldsymbol{R}\boldsymbol{L}_L\boldsymbol{R}^\top) = \lambda_{m-1}\left(\boldsymbol{R}\sum_{t=1}^{n}\sum_{i>j}\mathbb{E}L_{ij}^t(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^\top\boldsymbol{R}^\top\right)$$
$$\geq \frac{1}{2\kappa_2}\lambda_{m-1}\left(\boldsymbol{R}\sum_{t=1}^{n}\sum_{i>j}\mathbb{E}R_{ij}^t(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^\top\boldsymbol{R}^\top\right)$$

Moreover $\sum_{i>j}\mathbb{E}R_{ij}^t \geq \frac{1}{2}(\mathbb{E}[n_t]-1) = (mp-1)/2$, where the -1 accounts for possible unpaired X_{ti}. By symmetry $\mathbb{E}R_{ij}^t$ is the same for any (i,j). Then for any (i,j),

$$\mathbb{E}R_{ij}^t \geq \frac{mp-1}{2}\Big/\binom{m}{2} \geq \frac{p}{2m}$$

as long as $mp\geq 2$. Thus

$$\lambda_{m-1}(\mathbb{E}\boldsymbol{R}\boldsymbol{L}_L\boldsymbol{R}^\top) \geq \frac{1}{2\kappa_2}\lambda_{m-1}\left(\boldsymbol{R}\sum_{t=1}^{n}\sum_{i>j}\frac{p}{2m}(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^\top\boldsymbol{R}^\top\right)$$
$$= \frac{1}{2\kappa_2}\cdot n\cdot\frac{p}{2m}\cdot m$$
$$= \frac{np}{4\kappa_2}.$$

Now invoke the matrix Chernoff inequality, we have

$$\mathbb{P}\left\{[\lambda_{m-1}(\boldsymbol{R}\boldsymbol{L}_L\boldsymbol{R}^\top)] \leq \frac{np}{8\kappa_2}\right\} \leq m\cdot\left[\frac{e^{-1/2}}{(1/2)^{1/2}}\right]^{\frac{np}{4\kappa_2}/2} \leq n^{-10}$$

as long as $np\geq C\kappa_2\log(n)$ for some large enough constant C.

The spectrum of $\boldsymbol{L}_{Lz}$ comes directly from the spectrum of $\boldsymbol{L}_L$. Recall

$$\boldsymbol{L}_{Lz} = \sum_{(i,j)\in\mathcal{E}_Y, i>j} L_{ij}z_{ij}(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^\top.$$

By Lemma 9,

$$\lambda_1(\boldsymbol{L}_{Lz}) = \max_{\boldsymbol{v}:\|\boldsymbol{v}\|=1}\boldsymbol{v}^\top\sum_{(i,j)\in\mathcal{E}_Y,i>j}L_{ij}z_{ij}(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^\top\boldsymbol{v}$$
$$\leq \frac{1}{4}\max_{\boldsymbol{v}:\|\boldsymbol{v}\|=1}\boldsymbol{v}^\top\sum_{(i,j)\in\mathcal{E}_Y,i>j}L_{ij}(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^\top\boldsymbol{v}$$
$$= \frac{1}{4}\lambda_1(\boldsymbol{L}_L) \leq np,$$

and

$$\lambda_{m-1}(\boldsymbol{L}_{Lz}) = \min_{\boldsymbol{v}:\|\boldsymbol{v}\|=1,\boldsymbol{v}^\top\boldsymbol{1}_m=0}\boldsymbol{v}^\top\sum_{(i,j)\in\mathcal{E}_Y,i>j}L_{ij}z_{ij}(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^\top\boldsymbol{v}$$
$$\geq \frac{1}{4\kappa_1}\min_{\boldsymbol{v}:\|\boldsymbol{v}\|=1,\boldsymbol{v}^\top\boldsymbol{1}_m=0}\boldsymbol{v}^\top\sum_{(i,j)\in\mathcal{E}_Y,i>j}L_{ij}(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^\top\boldsymbol{v}$$
$$= \frac{1}{4\kappa_1}\lambda_{m-1}(\boldsymbol{L}_L) \geq \frac{np}{16\kappa_1\kappa_2}.$$

C Proof of Lemma 4

We study the MLE $\widehat{\boldsymbol{\theta}}$ by analyzing the iterates of preconditioned gradient descent starting from the ground truth. Let $\boldsymbol{\theta}_0 = \boldsymbol{\theta}^\star$ be the starting point and $\eta > 0$ be the stepsize that is small enough. The preconditioned gradient descent iterates are given by

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta \boldsymbol{L}_{Lz}^\dagger \nabla \mathcal{L}(\boldsymbol{\theta}^t). \tag{24}$$

Consider the Taylor expansion of $\nabla \mathcal{L}(\boldsymbol{\theta}^t)$, we have

$$\begin{aligned}
\nabla \mathcal{L}(\boldsymbol{\theta}^t) &= \sum_{i>j:(i,j)\in\mathcal{E}_Y} \sum_{l=1}^{L_{ij}} \left(\left(-Y_{ji}^{(l)} + \sigma(\theta_i^t - \theta_j^t) \right) (\boldsymbol{e}_i - \boldsymbol{e}_j) \right) \\
&= \sum_{i>j:(i,j)\in\mathcal{E}_Y} \sum_{l=1}^{L_{ij}} \left(\left(-\epsilon_{ji}^{(l)} - \sigma(\theta_i^\star - \theta_j^\star) + \sigma(\theta_i^t - \theta_j^t) \right) (\boldsymbol{e}_i - \boldsymbol{e}_j) \right) \\
&= \sum_{i>j:(i,j)\in\mathcal{E}_Y} \sum_{l=1}^{L_{ij}} \left(\left(-\epsilon_{ji}^{(l)} + \sigma'(\theta_i^\star - \theta_j^\star)(\delta_i^t - \delta_j^t) + \frac{1}{2}\sigma''(\xi_{ij}^t)(\delta_i^t - \delta_j^t)^2 \right) (\boldsymbol{e}_i - \boldsymbol{e}_j) \right).
\end{aligned}$$

Here $\boldsymbol{\delta}^t := \boldsymbol{\theta}^t - \boldsymbol{\theta}^\star$ and for all (i,j), $\xi_{ij}^t \in \mathbb{R}$ is some number that lies between $\theta_i^\star - \theta_j^\star$ and $\theta_i^t - \theta_j^t$. As $\sigma'(\theta_i^\star - \theta_j^\star) = z_{ij}$ and $\delta_i^t - \delta_j^t = (\boldsymbol{e}_i - \boldsymbol{e}_j)^\top \boldsymbol{\delta}^t$, we have

$$\nabla \mathcal{L}(\boldsymbol{\theta}^t) = \boldsymbol{L}_{Lz}\boldsymbol{\delta}^t - \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} + \boldsymbol{r}^t,$$

where $\boldsymbol{r}^t = \sum_{i>j:(i,j)\in\mathcal{E}_Y} \sum_{l=1}^{L_{ij}} [\frac{1}{2}\sigma''(\xi_{ij}^t)(\delta_i^t - \delta_j^t)^2(\boldsymbol{e}_i - \boldsymbol{e}_j)]$. Feeding this into (24), we have

$$\boldsymbol{\delta}^{t+1} = (1-\eta)\,\boldsymbol{\delta}^t - \eta \left(\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} - \boldsymbol{L}_{Lz}^\dagger \boldsymbol{r}^t \right) \tag{25}$$

By definition $\boldsymbol{\delta}^0 = \boldsymbol{0}$. Applying this recursive relation $t-1$ times we obtain

$$\begin{aligned}
\boldsymbol{\delta}^t &= -\eta \sum_{i=0}^{t-1} (1-\eta)^i \boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} + \sum_{i=0}^{t-1} (1-\eta)^{t-1-i} \boldsymbol{L}_{Lz}^\dagger \boldsymbol{r}^i \\
&= -\left[1 - (1-\eta)^t\right] \boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} + \sum_{i=0}^{t-1} (1-\eta)^{t-1-i} \boldsymbol{L}_{Lz}^\dagger \boldsymbol{r}^i.
\end{aligned}$$

At this point, we invoke an existing result on these terms that have been studied in [YCOM24] for a more general setting. The proof of Lemma 2 combined with Lemmas 3 and 4 in [YCOM24] reveal the following properties of (24).

Lemma 6. *Instate the assumptions of Theorem 1. Suppose that*

$$\kappa_1^3 \frac{(d_{\max})^2 \log^2(n)}{(\lambda_{m-1}(\boldsymbol{L}_{Lz}))^3} \leq C_1 \kappa_1 \sqrt{\frac{\log(n)}{\lambda_{m-1}(\boldsymbol{L}_{Lz})}}, \tag{26}$$

for some constant $C_1 > 0$. Then with probability at least $1 - n^{-10}$, the precondition gradient descent dynamic satisfies the following properties:

1. *There exists a unique minimizer $\widehat{\boldsymbol{\theta}}$ of (2).*
2. *There exist some α_1, α_2 obeying $0 < \alpha_1 \leq \alpha_2$ such that any $t \in \mathbb{N}$,*

$$\|\boldsymbol{\theta}^t - \widehat{\boldsymbol{\theta}}\|_{\boldsymbol{L}_{Lz}} \leq (1 - \eta\alpha_1)^t \|\boldsymbol{\theta}^0 - \widehat{\boldsymbol{\theta}}\|_{\boldsymbol{L}_{Lz}},$$

provided that $0 < \eta \leq 1/\alpha_2$.

3. *For any k, l and iteration $t \geq 0$,*

$$\left|\left(\theta_k^t - \theta_l^t\right) - \left(\theta_k^\star - \theta_l^\star\right)\right| \leq C_2 \kappa_1 \sqrt{\frac{\log(n)}{\lambda_{m-1}(\boldsymbol{L}_L)}} \tag{27}$$

for some constant $C_2 > 0$.

4. *For any k, l and iteration $t \geq 0$,*

$$\left|(\boldsymbol{e}_k - \boldsymbol{e}_l)^\top \boldsymbol{L}_{Lz}^\dagger \boldsymbol{r}^t\right| \leq C_3 \kappa_1^3 \frac{(d_{\max})^2 \log^2(n)}{(\lambda_{m-1}(\boldsymbol{L}_L))^3}$$

for some constant $C_3 > 0$.

Lemmas 3 and 2 imply that

$$\frac{np}{4\kappa_1\kappa_2} \leq \lambda_{m-1}(\boldsymbol{L}_{Lz}) \qquad \text{and} \qquad d_{\max} \leq \frac{3}{2} np.$$

Then the condition (26) holds as long as $np \geq C_4 \kappa_1^2 \kappa_2 \log^3(n)$ for some large enough constant C_4. Invoke Lemma 6 to see that for any (k, l, t),

$$\left|(\boldsymbol{e}_k - \boldsymbol{e}_l)^\top \boldsymbol{L}_{Lz}^\dagger \boldsymbol{r}^t\right| \leq C_5 \kappa_1^6 \frac{\log^2(n)}{np}, \tag{28}$$

where $C_5 > 0$ is some constant. Furthermore, the convergence given by Lemma 6 implies that

$$\widehat{\boldsymbol{\delta}} = \lim_{t\to\infty} \boldsymbol{\delta}^t = -\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B} \widehat{\boldsymbol{\epsilon}} + \underbrace{\eta \lim_{t\to\infty} \sum_{i=0}^{t-1} (1-\eta)^{t-1-i} \boldsymbol{L}_{Lz}^\dagger \boldsymbol{r}^i}_{=:\boldsymbol{r}}. \tag{29}$$

It remains to control the ℓ_∞ norm of $\boldsymbol{r}$. For any (k, l), (28) shows that

$$|r_k - r_j| \leq \eta \lim_{t\to 0} \sum_{i=0}^{t-1} (1-\eta)^{t-1-i} C_5 \kappa_1^6 \frac{\log^2(n)}{np} = C_5 \kappa_1^6 \frac{\log^2(n)}{np}.$$

As $\boldsymbol{1}^\top \boldsymbol{L}_{Lz}^\dagger = 0$, the above inequality implies that

$$\begin{aligned} |r_k| &= \left|\frac{1}{m} \cdot m \boldsymbol{e}_k^\top \boldsymbol{r}\right| = \left|\frac{1}{m} \sum_{l=1}^m (\boldsymbol{e}_k - \boldsymbol{e}_l)^\top \boldsymbol{r}\right| \\ &= \left|\frac{1}{m} \sum_{l=1}^m (r_k - r_l)\right| \leq C_5 \kappa_1^6 \frac{\log^2(n)}{np}. \end{aligned}$$

The proof is now completed.

D Proof of Theorem 2

We start with the proof of (5). By Lemma 4 we can express the MLE estimation error $\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star$ as

$$\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star = -\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B} \widehat{\boldsymbol{\epsilon}} + \boldsymbol{r} \tag{30}$$

for $\boldsymbol{B}, \widehat{\boldsymbol{\epsilon}}$ defined in Section 4.3 and $\boldsymbol{r} \in \mathbb{R}^m$ is a residual term obeying $\|\boldsymbol{r}\|_\infty \leq C_1 \kappa_1^6 \log^2 n/(np)$ for some constant C_1.

We first focus on the main term $\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}}$. Expanding $\boldsymbol{B}$ and $\widehat{\boldsymbol{\epsilon}}$, we rewrite it as

$$\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}} = \sum_{i>j:(i,j)\in\mathcal{E}_Y}\sum_{l=1}^{L_{ij}}\underbrace{\epsilon_{ij}^{(l)}\boldsymbol{L}_{Lz}^{\dagger}(\boldsymbol{e}_i-\boldsymbol{e}_j)}_{=:\boldsymbol{u}_{ij}^{(l)}}.$$

It is easy to see that conditional on $\mathcal{G}_Y$, $\{\boldsymbol{u}_{ij}^{(l)}\}_{i,j,l}$ is a set of independent zero-mean random variables. Thus we can expand $\mathbb{E}[\|\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2]$ to be

$$\begin{aligned}
\mathbb{E}\left[\|\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2\right] &= \mathbb{E}\sum_{i>j:(i,j)\in\mathcal{E}_Y}\sum_{l=1}^{L_{ij}}\boldsymbol{u}_{ij}^{(l)\top}\boldsymbol{u}_{ij}^{(l)} \\
&= \mathbb{E}\sum_{i>j:(i,j)\in\mathcal{E}_Y}\sum_{l=1}^{L_{ij}}\mathrm{Trace}\left(\boldsymbol{u}_{ij}^{(l)}\boldsymbol{u}_{ij}^{(l)\top}\right) \\
&\overset{\text{(i)}}{=} \mathrm{Trace}\left(\sum_{i>j:(i,j)\in\mathcal{E}_Y}\sum_{l=1}^{L_{ij}}\boldsymbol{L}_{Lz}^{\dagger}z_{ij}(\boldsymbol{e}_i-\boldsymbol{e}_j)(\boldsymbol{e}_i-\boldsymbol{e}_j)^{\top}\boldsymbol{L}_{Lz}^{\dagger}\right) \\
&\overset{\text{(ii)}}{=} \mathrm{Trace}(\boldsymbol{L}_{Lz}^{\dagger}). \qquad (31)
\end{aligned}$$

Here (i) follows from the equality

$$z_{ij} = \mathrm{Var}(\epsilon_{ji}^{(l)}) = \sigma'(\theta_i^\star-\theta_j^\star) = \frac{e^{\theta_i^\star}e^{\theta_j^\star}}{(e^{\theta_i^\star}+e^{\theta_j^\star})^2},$$

and (ii) follows from the definition of $\boldsymbol{L}_{Lz}$. By Lemma 3,

$$\frac{m}{2np} \le (m-1)\lambda_{m-1}(\boldsymbol{L}_{Lz}^{\dagger}) \le \mathrm{Trace}(\boldsymbol{L}_{Lz}^{\dagger}) \le m\|\boldsymbol{L}_{Lz}^{\dagger}\| \le \frac{16\kappa_1\kappa_2 m}{np}. \qquad (32)$$

Moreover, $\{\epsilon_{ij}^{(l)}\}_{i,j,l}$ is a set of sub-Gaussian random variable with variance proxy $1/z_{ij}$ (see, e.g., Section 2.5 in [Ver18]) and $1/z_{ij} \le 4\kappa_1$ by Lemma 9. Applying Hanson-Wright inequality (see [RV13]), for any scalar $a > 0$ we have

$$\begin{aligned}
&\mathbb{P}\left[\left|\|\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2 - \mathbb{E}\left[\|\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2\right]\right| > a\right] \qquad (33) \\
&\quad\le 2\exp\left[-C_2\left(\frac{a^2}{(4\kappa_1)^4\|\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\|_{\mathrm{F}}^2}\wedge\frac{a}{(4\kappa_1)^2\|\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\|}\right)\right]
\end{aligned}$$

for some constant $C_2 > 0$. For $\|\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\|_{\mathrm{F}}$ and $\|\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\|$ we have

$$\begin{aligned}
\|\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\| &\le \|\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\|_{\mathrm{F}} \\
&= \sqrt{\mathrm{Trace}\left(\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\right)} \\
&= \sqrt{\mathrm{Trace}\left(\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\boldsymbol{B}^{\top}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{B}\boldsymbol{B}^{\top}\right)} \\
&\overset{\text{(i)}}{=} \sqrt{\mathrm{Trace}\left(\boldsymbol{L}_{Lz}^{\dagger}\boldsymbol{L}_{Lz}^{\dagger}\right)} \\
&\overset{\text{(ii)}}{\le} \frac{\sqrt{m}}{\lambda_{m-1}(\boldsymbol{L}_{Lz})} \overset{\text{(iii)}}{\le} \frac{16\kappa_1\kappa_2\sqrt{m}}{np}.
\end{aligned}$$

Here (i) follows from the fact that $\boldsymbol{B}\boldsymbol{B}^\top = \boldsymbol{L}_{Lz}$, (ii) follows from Lemma 3 and the fact that $\text{Trace}(\boldsymbol{M}) \le m\|\boldsymbol{M}\|$ for any $m\times m$ matrix $\boldsymbol{M}$, and (iii) follows from Lemma 3. Now substitute $a = C_3\kappa_1^3\kappa_2\sqrt{m}\log(n)/(np)$ in (33) for some large enough constant C_3. We have that with probability at least $1-2n^{-10}$,

$$\left|\|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2 - \mathbb{E}\left[\|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2\right]\right| \le \frac{C_3\kappa_1^3\kappa_2\sqrt{m}\log(n)}{np}. \tag{34}$$

Combining this with (31) and (32),

$$\begin{aligned}\left|\|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\| - \sqrt{\text{Trace}(\boldsymbol{L}_{Lz}^\dagger)}\right| &= \frac{\left|\|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\|^2 - \text{Trace}(\boldsymbol{L}_{Lz}^\dagger)\right|}{\|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\| + \sqrt{\text{Trace}(\boldsymbol{L}_{Lz}^\dagger)}} \\ &\le \frac{C_3\kappa_1^3\kappa_2\sqrt{m}\log(n)/(np)}{\sqrt{m/(2np)}} \\ &\le C_4\kappa_1^3\kappa_2\sqrt{\frac{\log(n)}{np}}\end{aligned}$$

for some constant $C_4 > 0$.

Substituting (34) and (31) into (30), we have that for some constant C_1, C_2,

$$\begin{aligned}\left|\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\| - \sqrt{\text{Trace}(\boldsymbol{L}_{Lz}^\dagger)}\right| &\le \left|\|\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}\| - \sqrt{\text{Trace}(\boldsymbol{L}_{Lz}^\dagger)}\right| + \|\boldsymbol{r}\| \\ &\le C_4\kappa_1^3\kappa_2\sqrt{\frac{\log(n)}{np}} + \frac{C_1\kappa_1^6\sqrt{m}\log^2(n)}{np}.\end{aligned} \tag{35}$$

The proof of (5) is now completed. For (6), the following lemma connects $\widehat{z}$ and z. The proof is deferred to the end of this section.

Lemma 7. *Instate the assumptions of Theorem 1, then with probability at least $1-C_5n^{-10}$ for some constant $C_5 > 0$,*

$$\|\boldsymbol{L}_{Lz}^\dagger - \boldsymbol{L}_{L\hat{z}}^\dagger\| \le \frac{C_6\kappa_1^{7/2}\kappa_2^2}{(np)^{3/2}}$$

for some large enough constant C_6.

Combining this lemma with Weryl's inequality, we have that

$$\left|\text{Trace}(\boldsymbol{L}_{Lz}^\dagger) - \text{Trace}(\boldsymbol{L}_{L\hat{z}}^\dagger)\right| \le \frac{C_6 m\kappa_1^{7/2}\kappa_2^2}{(np)^{3/2}}.$$

Then by (32),

$$\begin{aligned}\left|\sqrt{\text{Trace}(\boldsymbol{L}_{L\hat{z}}^\dagger)} - \sqrt{\text{Trace}(\boldsymbol{L}_{Lz}^\dagger)}\right| &= \frac{\left|\text{Trace}(\boldsymbol{L}_{Lz}^\dagger) - \text{Trace}(\boldsymbol{L}_{L\hat{z}}^\dagger)\right|}{\sqrt{\text{Trace}(\boldsymbol{L}_{Lz}^\dagger)} + \sqrt{\text{Trace}(\boldsymbol{L}_{L\hat{z}}^\dagger)}} \\ &= \frac{C_6 m\kappa_1^{7/2}\kappa_2^2/(np)^{3/2}}{\sqrt{m/(2np)}} \le \frac{\sqrt{2}C_6\kappa_1^{7/2}\kappa_2^2\sqrt{m}}{np}.\end{aligned}$$

Using triangular inequality, we conclude that

$$\left|\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star\| - \sqrt{\text{Trace}(\boldsymbol{L}_{L\hat{z}}^\dagger)}\right| \le C_4\kappa_1^3\kappa_2\sqrt{\frac{\log(n)}{np}} + \frac{(C_1\kappa_1^6 + \sqrt{2}C_6\kappa_1^{7/2}\kappa_2^2)\sqrt{m}\log^2(n)}{np}.$$

Proof of Lemma 7. Recall σ is the sigmoid function and its derivative σ' is 1-Lipschitz. By Theorem 1, for all (i,j),

$$\begin{aligned}|z_{ij} - \widehat{z}_{ij}| &= \left|\sigma'(\theta_i^\star - \theta_j^\star) - \sigma'(\theta_i^\star - \theta_j^\star)\right| \\ &\le \left|\left(\widehat{\theta}_i - \widehat{\theta}_j\right) - \left(\theta_i^\star - \theta_j^\star\right)\right| \le C_7 \kappa_1 \kappa_2^{1/2}\sqrt{\frac{\log(n)}{np}}\end{aligned}$$

for some constant $C_7 > 0$. Then

$$\begin{aligned}\|\boldsymbol{L}_{Lz} - \boldsymbol{L}_{L\hat{z}}\| &= \max_{\boldsymbol{v}\in\mathbb{R}^m:\|\boldsymbol{v}\|=1}\left|\boldsymbol{v}^\top \sum_{i>j:(i,j)\in\mathcal{E}_Y} L_{ij}\left(z_{ij} - \widehat{z}_{ij}\right)(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top \boldsymbol{v}\right| \\ &\le \max_{\boldsymbol{v}\in\mathbb{R}^m:\|\boldsymbol{v}\|=1} \sum_{i>j:(i,j)\in\mathcal{E}_Y} |z_{ij} - \widehat{z}_{ij}|\, \boldsymbol{v}^\top L_{ij}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top \boldsymbol{v} \\ &\le C_7\kappa_1\kappa_2^{1/2}\sqrt{\frac{\log(n)}{np}} \cdot \|\boldsymbol{L}_L\| \\ &\le 3C_7\kappa_1\kappa_2^{1/2}\sqrt{np\log(n)},\end{aligned}$$

where the last line follows from Lemma 3. As $np \ge C_8\kappa_1^4\kappa_2^3\log^2(n)$ for some large enough constant C_8, $\|\boldsymbol{L}_{Lz} - \boldsymbol{L}_{L\hat{z}}\| \le np/(32\kappa_1\kappa_2)$. By Weryl's inequality and Lemma 3,

$$\lambda_{m-1}(\boldsymbol{L}_{L\hat{z}}) \ge \lambda_{m-1}(\boldsymbol{L}_{Lz}) - \|\boldsymbol{L}_{Lz} - \boldsymbol{L}_{L\hat{z}}\| \ge \frac{np}{16\kappa_1\kappa_2} - \frac{np}{32\kappa_1\kappa_2} = \frac{np}{32\kappa_1\kappa_2}.$$

This implies that $\|\boldsymbol{L}_{Lz}^\dagger\| \le 16\kappa_1\kappa_2/(np)$ and $\|\boldsymbol{L}_{L\hat{z}}^\dagger\| \le 32\kappa_1\kappa_2/(np)$. Using the perturbation bound of pseudo-inverse (see Theorem 4.1 in [Wed73]), we have

$$\begin{aligned}\|\boldsymbol{L}_{Lz}^\dagger - \boldsymbol{L}_{L\hat{z}}^\dagger\| &\le 3\cdot\|\boldsymbol{L}_{Lz}^\dagger\|\cdot\|\boldsymbol{L}_{L\hat{z}}^\dagger\|\cdot\|\boldsymbol{L}_{Lz} - \boldsymbol{L}_{L\hat{z}}\| \\ &\le \frac{C_9\kappa_1^3\kappa_2^{5/2}}{(np)^{3/2}}\end{aligned}$$

for some constant $C_9 > 0$.

E Proofs for inference

In this section, we present the analysis for the inference results we presented in Section 3.3.

E.1 Proof of Theorem 3

We condition the whole analysis on the high probability event when Lemmas 4 and 3 holds. Let $\widehat{\boldsymbol{\delta}} := \widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star$. Expanding (15), we have

$$\begin{aligned}\widehat{\boldsymbol{\delta}} &= -\boldsymbol{L}_{Lz}^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} + \boldsymbol{r} \\ &= -\boldsymbol{L}_{Lz}^\dagger \sum_{i>j:(i,j)\in\mathcal{E}_Y}\sum_{l=1}^{L_{ij}}\left[Y_{ji}^{(l)} - \sigma(\theta_i^\star - \theta_j^\star)\right](\boldsymbol{e}_i - \boldsymbol{e}_j) + \boldsymbol{r},\end{aligned}$$

where $\|\boldsymbol{r}\|_\infty \le C_1\kappa_1^6\log^2(n)/(np)$ for some constant C_1.

We start with the case with $\widetilde{\boldsymbol{z}} = \boldsymbol{z}$. Consider $\boldsymbol{L}_{Lz}\widehat{\boldsymbol{\delta}}$. As $\boldsymbol{L}_{Lz}\mathbf{1}_m = \mathbf{0}$, $\lambda_{m-1}(\boldsymbol{L}_{Lz}) > 0$, and $\widehat{\boldsymbol{\delta}}^\top\mathbf{1}_m = 0$,

$$\begin{aligned}\boldsymbol{L}_{Lz}^{1/2}\widehat{\boldsymbol{\delta}} &= -(\boldsymbol{L}_{Lz}^\dagger)^{1/2}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}} + \boldsymbol{L}_{Lz}^{1/2}\boldsymbol{r} \\ &= -\sum_{i>j:(i,j)\in\mathcal{E}_Y}\sum_{l=1}^{L_{ij}}\underbrace{\left[Y_{ji}^{(l)} - \sigma(\theta_i^\star - \theta_j^\star)\right](\boldsymbol{L}_{Lz}^\dagger)^{1/2}(\boldsymbol{e}_i - \boldsymbol{e}_j)}_{=:\boldsymbol{u}_{ij}^{(l)}} + \boldsymbol{L}_{Lz}^{1/2}\boldsymbol{r}.\end{aligned}$$

Conditional on $(\mathcal{E}_Y, \{L_{ij}\})$, $\boldsymbol{u}_{ij}^{(l)}$ are independent random variables. It is easy to see that $\boldsymbol{u}_{ij}^{(l)}$ is zero-mean and has covariance

$$\mathbb{E}\left[\boldsymbol{u}_{ij}^{(l)}\boldsymbol{u}_{ij}^{(l)\top}\right] = z_{ij}(\boldsymbol{L}_{Lz}^{\dagger})^{1/2}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top(\boldsymbol{L}_{Lz}^{\dagger})^{1/2}.$$

It is also bounded in third moment of its spectral norm by

$$\mathbb{E}\left[\|\boldsymbol{u}_{ij}^{(l)}\|^3\right] \le 2^{3/2}\|\boldsymbol{L}_{Lz}^{\dagger}\|^{3/2} \le \frac{2^{15/2}\kappa_1^{3/2}\kappa_2^{3/2}}{(np)^{3/2}},$$

where the last inequality uses Lemma 3. Summing up across i, j and l, we have

$$\begin{aligned}\sum_{i>j:(i,j)\in\mathcal{E}_Y}\sum_{l=1}^{L_{ij}}\mathbb{E}\left[\boldsymbol{u}_{ij}^{(l)}\boldsymbol{u}_{ij}^{(l)\top}\right] &= (\boldsymbol{L}_{Lz}^{\dagger})^{1/2}\sum_{i>j:(i,j)\in\mathcal{E}_Y} L_{ij}z_{ij}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top(\boldsymbol{L}_{Lz}^{\dagger})^{1/2} \\ &= (\boldsymbol{L}_{Lz}^{\dagger})^{1/2}\boldsymbol{L}_{Lz}(\boldsymbol{L}_{Lz}^{\dagger})^{1/2} \\ &= \boldsymbol{I}_m - \frac{1}{m}\boldsymbol{1}_m\boldsymbol{1}_m^\top.\end{aligned}$$

The last line holds since $\boldsymbol{L}_{Lz}\boldsymbol{1}_m = \boldsymbol{0}$ and $\lambda_{m-1}(\boldsymbol{L}_{Lz}) > 0$. Now using multivariate CLT (see, e.g., [Rai19]), we have that as $\sum_{i>j:(i,j)\in\mathcal{E}_Y} L_{ij} \to \infty$,

$$-(\boldsymbol{L}_{Lz}^{\dagger})^{1/2}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}} \xrightarrow{\mathrm{d}} \mathcal{N}\left(0, \boldsymbol{I}_m - \frac{1}{m}\boldsymbol{1}_m\boldsymbol{1}_m^\top\right).$$

As m is fixed, Lemma 2 implies that $\sum_{i>j:(i,j)\in\mathcal{E}_Y} L_{ij} \to \infty$ is equivalent to $np \to \infty$. For the residual term, Lemma 4 and Lemma 3 imply that with probability at least $1 - O(n^{-10})$,

$$\|\boldsymbol{L}_{Lz}^{1/2}\boldsymbol{r}\| \le \|\boldsymbol{L}_{Lz}^{1/2}\| \cdot \|\boldsymbol{r}\| \le \|\boldsymbol{L}_{Lz}^{1/2}\| \cdot \sqrt{m}\|\boldsymbol{r}\|_\infty \le \frac{C_1\sqrt{m}\kappa_1^6\log^2(n)}{\sqrt{np}}.$$

Then as $np/\log^4(n) \to \infty$,

$$\boldsymbol{L}_{Lz}^{1/2}\boldsymbol{r} \to \boldsymbol{0}$$

in probability. Combining the limiting distribution of both two terms, we have

$$\boldsymbol{L}_{Lz}^{1/2}\widehat{\boldsymbol{\delta}} = -(\boldsymbol{L}_{Lz}^{\dagger})^{1/2}\boldsymbol{B}\widehat{\boldsymbol{\epsilon}} + \boldsymbol{L}_{Lz}^{1/2}\boldsymbol{r} \xrightarrow{\mathrm{d}} \mathcal{N}\left(0, \boldsymbol{I}_m - \frac{1}{m}\boldsymbol{1}_m\boldsymbol{1}_m^\top\right).$$

In the case of $\widetilde{\boldsymbol{z}} = \widehat{\boldsymbol{z}}$, we introduce the following lemma. The proof is deferred to the end of this section.

Lemma 8. *Instate the assumptions of Theorem 1. Then with probability at least $1 - C_1 n^{-10}$ for some constant C_1, one has*

$$\|\boldsymbol{L}_{Lz}^{1/2} - \boldsymbol{L}_{L\hat{z}}^{1/2}\| \le C_2\kappa_1^2\kappa_2^{1/2}\sqrt{\log(n)}$$

for some constant C_2.

Now the above lemma and Theorem 1 imply

$$\left\|\boldsymbol{L}_{L\widehat{z}}^{1/2}\widehat{\boldsymbol{\delta}} - \boldsymbol{L}_{Lz}^{1/2}\widehat{\boldsymbol{\delta}}\right\| \le \|\boldsymbol{L}_{L\widetilde{z}}^{1/2} - \boldsymbol{L}_{Lz}^{1/2}\| \cdot \|\widehat{\boldsymbol{\delta}}\| \le C_3\kappa_1^3\kappa_2\sqrt{\frac{m\log^2(n)}{np}}$$

for some constant $C_3 > 0$. Therefore, as $np/\log^4(n) \to \infty$,

$$\boldsymbol{L}_{L\widehat{z}}^{1/2}\widehat{\boldsymbol{\delta}} - \boldsymbol{L}_{Lz}^{1/2}\widehat{\boldsymbol{\delta}} \to 0$$

in probability and hence

$$\boldsymbol{L}_{L\widehat{z}}^{1/2}\widehat{\boldsymbol{\delta}} = \left(\boldsymbol{L}_{L\widehat{z}}^{1/2}\widehat{\boldsymbol{\delta}} - \boldsymbol{L}_{Lz}^{1/2}\widehat{\boldsymbol{\delta}}\right) + \boldsymbol{L}_{Lz}^{1/2}\widehat{\boldsymbol{\delta}} \xrightarrow{\mathrm{d}} \mathcal{N}\left(0, \boldsymbol{I}_m - \frac{1}{m}\boldsymbol{1}_m\boldsymbol{1}_m^\top\right).$$

Proof of Lemma 8. Let $\boldsymbol{R} \in \mathbb{R}^{(n-1)\times n}$ be a partial isometry such that $\boldsymbol{R}\boldsymbol{R}^\top = \boldsymbol{I}_{m-1}$ and $\boldsymbol{R}\boldsymbol{1} = \boldsymbol{0}$. Then

$$\|\boldsymbol{L}_{L\hat{z}}^{1/2} - \boldsymbol{L}_{Lz}^{1/2}\| = \|\boldsymbol{R}(\boldsymbol{L}_{L\hat{z}}^{1/2} - \boldsymbol{L}_{Lz}^{1/2})\boldsymbol{R}^\top\|.$$

We also observe that $(\boldsymbol{R}\boldsymbol{L}_{L\widetilde{z}}\boldsymbol{R}^\top)^{1/2} = \boldsymbol{R}\boldsymbol{L}_{L\widetilde{z}}^{1/2}\boldsymbol{R}^\top$ for $\widetilde{z} \in \{z, \widehat{z}\}$. As shown in the proof of Lemma 7,

$$\|\boldsymbol{L}_{L\hat{z}} - \boldsymbol{L}_{Lz}\| \le C_1\kappa_1^{3/2}\sqrt{np\log(n)}$$

for some constant $C_1 > 0$.

Lemma 3 implies that

$$\lambda_{m-1}(\boldsymbol{R}\boldsymbol{L}_{Lz}^{1/2}\boldsymbol{R}^\top) = [\lambda_{m-1}(\boldsymbol{L}_{Lz})]^{1/2} \ge \frac{\sqrt{np}}{4\kappa_1^{1/2}\kappa_2^{1/2}}.$$

Then by Lemma 2.1 in [Sch92],

$$\|\boldsymbol{R}(\boldsymbol{L}_{L\widehat{z}}^{1/2} - \boldsymbol{L}_{Lz}^{1/2})\boldsymbol{R}^\top\| \le \frac{\|\boldsymbol{R}(\boldsymbol{L}_{L\widehat{z}} - \boldsymbol{L}_{Lz})\boldsymbol{R}^\top\|}{\lambda_{m-1}(\boldsymbol{R}\boldsymbol{L}_{Lz}^{1/2}\boldsymbol{R}^\top)} \le 4C_1\kappa_1^2\kappa_2^{1/2}\sqrt{\log(n)}.$$

E.2 Proof of Corollary 1

Observe that

$$\frac{\widehat{\theta}_i - \theta_i^\star}{\sqrt{[\boldsymbol{L}_{L\tilde{z}}^\dagger]_{ii}}} = \frac{\boldsymbol{e}_i^\top(\boldsymbol{L}_{L\tilde{z}}^\dagger)^{1/2}\boldsymbol{L}_{L\tilde{z}}^{1/2}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^\star)}{\sqrt{[\boldsymbol{L}_{L\tilde{z}}^\dagger]_{ii}}}.$$

Then Theorem 3 implies

$$\frac{\boldsymbol{e}_i^\top(\boldsymbol{L}_{L\tilde{z}}^\dagger)^{1/2}(\boldsymbol{I}_m - \frac{1}{m}\boldsymbol{1}_m\boldsymbol{1}_m^\top)(\boldsymbol{L}_{L\tilde{z}}^\dagger)^{1/2}\boldsymbol{e}_i}{[\boldsymbol{L}_{L\tilde{z}}^\dagger]_{ii}} = 1.$$

The last inequality follows from the fact that $\boldsymbol{1}_m^\top(\boldsymbol{L}_{L\tilde{z}}^\dagger)^{1/2} = 0$. As a side note, we compare this with a similar result for the BTL model. Proposition 4.1 in [GSZ23] studies the uncertainty quantification of MLE in the BTL model with uniform observation. It says that

$$\frac{\widehat{\theta}_i - \theta_i^\star}{\sqrt{[\boldsymbol{L}_{Lz}]_{ii}^{-1}}} \xrightarrow{\mathrm{d}} \mathcal{N}(0, 1).$$

This is close to our result as $[\boldsymbol{L}_{Lz}]_{ii}^{-1}$ can be viewed as an approximation of $[\boldsymbol{L}_{Lz}^\dagger]_{ii}$ in the context of the BTL model. To see why, consider the simplified case where $\boldsymbol{\theta}^\star = \boldsymbol{0}_n$ and $L = 1$. As $np \to \infty$, the Laplacian $\boldsymbol{L}_{Lz}$ concentrate to its expectation $\mathbb{E}[\boldsymbol{L}_{Lz}] = np\boldsymbol{I}_n - p\boldsymbol{1}_n\boldsymbol{1}_n^\top$. Then $\boldsymbol{L}^\dagger \approx (1/np)(\boldsymbol{I}_n - \frac{1}{n}\boldsymbol{1}_n\boldsymbol{1}_n^\top)$ and $[\boldsymbol{L}_{Lz}^\dagger]_{ii} \approx [\boldsymbol{L}_{Lz}]_{ii}^{-1}$.

F Auxiliary lemmas

In this section, we gather some auxiliary results that are useful throughout this paper.

Lemma 9 (Range of z_{ij})**.** *Recall*

$$z_{ij} = \frac{e^{\theta_i^\star}e^{\theta_j^\star}}{(e^{\theta_i^\star} + e^{\theta_j^\star})^2} = \frac{e^{\theta_i^\star - \theta_j^\star}}{(1 + e^{\theta_i^\star - \theta_j^\star})^2}.$$

For any (i, j),

$$\frac{1}{4\kappa_1} \le z_{ij} \le \frac{1}{4}.$$

Proof. Consider the function $f : [0, \infty) \to \mathbb{R}$ defined by $f(x) = x/(1+x)^2$. It has derivative $(1-x^2)/(1+x)^4$, so it is increasing at $x \in [0, 1)$ and decreasing at $x \in (1, \infty)$. By the definition of κ_1, $|\theta_i^\star - \theta_j^\star| \le \log(\kappa_1)$. Then

$$\frac{1}{4\kappa_1} \le f(e^{-\log(\kappa_1)}) \wedge f(e^{\log(\kappa_1)}) \le z_{ij} \le f(1) = \frac{1}{4}.$$

□

Lemma 10 (Maximum eigenvalue of Laplacian). *Let* $\boldsymbol{L} = \sum_{(i,j):i>j} w_{ij}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top$ *be a weighted graph Laplacian. Then* $\lambda_1(\boldsymbol{L}) \le 2\max_i \sum_j w_{ij}$.

Proof. Let $\boldsymbol{v} \in \mathbb{R}^m$, then

$$\begin{aligned}
\boldsymbol{v}^\top \boldsymbol{L}\boldsymbol{v} &= \boldsymbol{v}^\top \sum_{(i,j):i>j} w_{ij}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top \boldsymbol{v} \\
&= \sum_{(i,j):i>j} w_{ij}(v_i - v_j)^2 \\
&\le 2 \sum_{(i,j):i>j} w_{ij}(v_i^2 + v_j^2) \\
&\le 2 \sum_i \sum_{j \ne i} w_{ij} v_j^2 \\
&\le 2 \sum_i \max_i \sum_j w_{ij} \|\boldsymbol{v}\|^2
\end{aligned}$$

So $\lambda_1(\boldsymbol{L}) = \max_{\boldsymbol{v}\in\mathbb{R}^m, \|\boldsymbol{v}\|=1} \boldsymbol{v}^\top \boldsymbol{L}\boldsymbol{v} \le 2\max_i \sum_j w_{ij}$. □

Lemma 11 (A quantitative version of Sylvester's law of inertia, [Ost59]). *For any real symmetric matrix* $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ *and* $\boldsymbol{S} \in \mathbb{R}^{n\times n}$ *be a non-singular matrix. Then for any* $i \in [n]$, $\lambda_i(\boldsymbol{S}\boldsymbol{A}\boldsymbol{S}^\top)$ *lies between* $\lambda_i(\boldsymbol{A})\lambda_1(\boldsymbol{S}^\top\boldsymbol{S})$ *and* $\lambda_i(\boldsymbol{A})\lambda_n(\boldsymbol{S}^\top\boldsymbol{S})$.

Fact 2. *Let* $\mathcal{G}$ *be an arbitrary graph with* m *vertices and let* $\boldsymbol{L}_w$ *be a weighted graph Laplacian defined by*

$$\boldsymbol{L}_w \coloneqq \sum_{i>j:(i,j)\in\mathcal{G}} w_{ij}(\boldsymbol{e}_i - \boldsymbol{e}_j)(\boldsymbol{e}_i - \boldsymbol{e}_j)^\top.$$

If $w_{ij} > 0$ *for all* $(i,j) \in \mathcal{G}$ *and* $\mathcal{G}$ *is a connected graph, then* $\boldsymbol{L}_w$ *is rank* $m-1$, $\boldsymbol{L}_w \boldsymbol{1}_m = \boldsymbol{0}_m$ *and* $\boldsymbol{L}_w^\dagger \boldsymbol{1}_m = 0$. *Moreover for any* $i \in [n-1]$, $\lambda_i(\boldsymbol{L}_w^\dagger) = \lambda_{n-i}(\boldsymbol{L}_w)$.

Proof. The fact that $\boldsymbol{L}_w$ is rank $m-1$ when $\mathcal{G}$ is connected is well-known. See e.g. [Spi07] for reference. Since $\boldsymbol{L}_w$ is a real symmetric matrix, it has an eigendecomposition $\boldsymbol{L}_w = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^\top$ and then $\boldsymbol{L}_w^\dagger = \boldsymbol{U}\boldsymbol{\Sigma}^\dagger\boldsymbol{U}^\top$. The rest follows from this decomposition and the form of $\boldsymbol{L}_w$. □