

Top- K ranking with a monotone adversary

Yuepeng Yang¹, Antares Chen², Lorenzo Orecchia², and Cong Ma¹

¹Department of Statistics, University of Chicago

²Department of Computer Science, University of Chicago

February 12, 2024

Abstract

In this paper, we address the top- K ranking problem with a monotone adversary. We consider the scenario where a comparison graph is randomly generated and the adversary is allowed to add arbitrary edges. The statistician’s goal is then to accurately identify the top- K preferred items based on pairwise comparisons derived from this semi-random comparison graph. The main contribution of this paper is to develop a weighted maximum likelihood estimator (MLE) that achieves near-optimal sample complexity, up to a $\log^2(n)$ factor, where n denotes the number of items under comparison. This is made possible through a combination of analytical and algorithmic innovations. On the analytical front, we provide a refined ℓ_∞ error analysis of the weighted MLE that is more explicit and tighter than existing analyses. It relates the ℓ_∞ error with the spectral properties of the weighted comparison graph. Motivated by this, our algorithmic innovation involves the development of an SDP-based approach to reweight the semi-random graph and meet specified spectral properties. Additionally, we propose a first-order method based on the Matrix Multiplicative Weight Update (MMWU) framework. This method efficiently solves the resulting SDP in nearly-linear time relative to the size of the semi-random comparison graph.

1 Introduction

In this paper we consider the problem of ranking n items given pairwise comparisons among them. This problem finds numerous applications in recommendation system [WYH⁺18], rating players [Elo67], web search [DKNS01], etc. One widely adopted model for pairwise comparison data is the Bradley-Terry-Luce (BTL) model [BT52, Luc05]. In this model, one assumes a latent score vector $\theta^* \in \mathbb{R}^n$, and that the Bernoulli outcome of the comparison between items i and j follows

$$\mathbb{P}[\text{item } i \text{ is preferred over item } j] = \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_j^*}}.$$

It is intuitive that under the BTL model, a higher score indicates a higher chance of winning a comparison.

In practice, the comparisons are often made for a subset of all possible pairs. A popular model to accommodate this situation is the uniform sampling model [CS15, CFMW19], i.e., each pair is compared independently with probability p . Uniform sampling is quite convenient for theory. An example, under this sampling mechanism, [CFMW19] shows that with high probability, the (regularized) maximum likelihood estimator (MLE) [FJ57] exactly identifies the top K preferred items with an optimal sample complexity:

$$n^2 p \gtrsim \frac{n \log(n)}{\Delta_K^2}, \tag{1}$$

where Δ_K measures the latent score difference between the K -th and the $(K + 1)$ -th preferred items, and $n^2 p$ is the expected number of comparisons.

Uniform sampling, while convenient for theoretical purposes, is often too ideal to match practice. This motivates a line of work [SBB⁺16, LSR22, Che23] that goes beyond uniform sampling and focuses on general




Figure 1: Adjacency matrix of semi-random graph. Each non-white square corresponds to a non-zero entry in the adjacency matrix.




Figure 2: Accuracy of top- K recovery for MLE under uniform sampling and weighted MLE under semi-random sampling. See Section F for the experiment setup.

Figure 3: Comparison graph and top- K recovery for weighted MLE under semi-random sampling. See Section F for the experiment setup.

sampling mechanisms. However, the theoretical guarantees are far from satisfactory. Take the recent work [LSR22] on the general sampling case as an example. The (regularized) MLE requires a sample complexity of $n^2 p \gtrsim \Delta_K^{-2} p^{-1} n \log(n)$ when applied to the special case of uniform sampling. For a sparse random graph, i.e., when $p \gtrsim \log(n)/n$, this sample complexity could be n times larger than the optimal one (1).

In this paper we aim to find a middle ground between uniform and general sampling mechanisms. Inspired by a line of work [BS95, FK01, MMV12, MPW16, AV18, CG18, KLL+23, GC23], we consider the so-called *semi-random sampling* where some benign adversary is allowed to make more comparisons in addition to uniform sampling; see Figure 1. Given its monotone nature, this is sometimes also referred to as a *monotone adversary*.¹ Intuitively, the monotone adversary should bring no harm to the ranking problem as it only reveals more information about the underlying score vector θ^* . However, it is well documented in the literature that the monotone adversary poses serious algorithmic and analytical challenges for a variety of problems. In community detection, [MPW16] shows that the information-theoretic detection limit could shift given a monotone adversary. In problems including sparse recovery [KLL+23] and matrix completion [CG18], methods and analyses that work well for uniform sampling can fail dramatically with a semi-random adversary. In this paper, we investigate top- K ranking under semi-random sampling. Our goal is to address the following question:

Can we identify the top- K items with minimal sample complexity, even under a monotone adversary?

1.1 Technical challenges

Top- K ranking under semi-random sampling brings some unique challenges. To begin with, it is worth noting that bounding the ℓ_2 estimation error for the score vector θ^* is not sufficient to guarantee exact recovery of the top- K items with an optimal sample complexity; see [CS15, JKSO16, CFMW19]. Instead, one would need a more fine-grained ℓ_∞ error bound.

To further complicate matters, even under uniform sampling, obtaining optimal control of the ℓ_∞ error—thus ensuring optimal sample complexity for top- K ranking—poses a significant challenge. [CFMW19] and its follow-up work [CGZ22, GSZ23] successfully characterize the optimal ℓ_∞ error of the MLE, leveraging a powerful leave-one-out argument [EK18, ZB18, AFWZ20, MWCC18, CCFM19]; see [CCF+21] for more references. However, a successful application of this argument relies crucially on the independence of the edges and certain homogeneity (e.g., degree homogeneity) in the Erdős–Rényi random graphs—a model for uniform sampling. These properties are easily violated for general comparison graphs, let alone the semi-random model we consider herein. In fact, as a manifestation, [LSR22] recently applies the leave-one-out technique to the BTL model in a general (deterministic) sampling mechanism, and obtains a somewhat loose

¹In this paper, we use the terms semi-random and monotone interchangeably.

control on the ℓ_∞ error of the MLE. Notably, even in the special case of uniform sampling, the required sample complexity for MLE can be n times larger than the optimal one. As a result, to address the problem of top- K ranking with a monotone adversary, one needs to develop a novel analysis that goes beyond uniform sampling and the leave-one-out technique.

1.2 Main contributions

The key result of this paper is to answer the main question affirmatively: we show that the weighted maximum likelihood estimator (MLE) with proper choices of weights is able to recover the top- K items with near-optimal sample complexity, albeit under a semi-random adversary. Moreover, the weights can be computed efficiently, in nearly-linear time in the size of the comparison graph. We achieve this through a combination of analytical and algorithmic innovations, which we detail below.

Analytical contributions. We provide a novel ℓ_∞ error analysis of the weighted MLE with explicit dependence on the spectral properties of the weighted comparison graph (e.g., the maximum degree, and the spectral gap of the weighted graph Laplacian); see Theorem 3. While the dependence on spectral properties has been characterized for the ℓ_2 error of the MLE [SBB⁺16, HOX14], we remark again that the ℓ_2 error alone cannot guarantee top- K recovery with optimal sample complexity.

Inspired by the recent work [Che23], we analyze the weighted MLE via a preconditioned gradient descent method that iteratively approximates the weighted MLE. This analysis bypasses the use of the leave-one-out argument. As opposed to two mysterious B, Q factors appearing in the performance bound (cf. Theorem 1 in [Che23]), our characterization of the ℓ_∞ error of the weighted MLE depends explicitly on the spectral properties of the weighted comparison graph. In particular, it is tight when applied to uniform sampling, in stark contrast with the previously mentioned result in [LSR22]. We expect this novel ℓ_∞ error analysis to be broadly applicable to more general sampling mechanisms beyond the semi-random case.

Algorithmic contributions. Motivated by the ℓ_∞ analysis of the weighted MLE, our goal boils down to finding a reweighting of the semi-random comparison graph such that it satisfies the required spectral properties. These amount to a constant lower bound on the spectral gap and upper bounds on the maximum degree and maximum weight in the reweighting. Taking a convex optimization approach, we show that the problem of finding such a reweighting can be cast exactly as a semi-definite program (SDP). We then develop a fast first-order method—based on the Matrix Multiplicative Weight Update (MMWU) framework [AK16] to approximately solve the resulting SDP; see Algorithm 2. We further show that such an approximate solution yields a desired set of weights, and the solution can be found efficiently, in nearly-linear time in the size of the semi-random graph. We believe that our SDP approach may find broader applications in learning problems over semi-random graphs where we need to restore spectral properties that have been disrupted by adversarial perturbations.

1.3 Related work

Ranking with the BTL model. The BTL model is a classic model for the ranking problem and has been extensively studied in the literature. Various methods have been proposed to tackle this problem, including Borda counting [Bor81], the maximum likelihood estimator [FJ57], and the spectral method [NOS12], among others. Since the conventional ℓ_2 analysis [NOS12] fails to capture the accuracy of top- K recovery, recent advances [CS15, JKSO16, CFMW19, CGZ22] focus on establishing the ℓ_∞ estimation error of the score vector θ^* . The story is mostly successful under the uniform sampling model. For instance, [CFMW19] first shows that both the spectral method and the (regularized) MLE achieve minimax optimal ℓ_∞ estimation error, and recover exactly the top- K items under the minimal sample complexity. [CGZ22] further proves that the vanilla MLE without regularization is optimal, and is superior to the spectral method in terms of the leading constant in the sample complexity. As uniform sampling is often too ideal, several attempts have been made to go beyond it. Most recently, [LSR22] and [Che23] investigate the ℓ_∞ guarantee of the MLE for general comparison graphs. As we mentioned, their analyses are loose, even when applied to the special case of uniform sampling.

Semi-random adversary. Semi-random adversary has been examined in a number of settings. Early work in this line studies problems related to semi-random graphs, including graph partitioning [MMV12], coloring [BS95, FK01], and finding independent sets [FK01]. In recent years, researchers have started to consider semi-random adversary in non-graphical data structures such as sparse recovery [KLL⁺23], Gaussian mixture model [AV18], matrix sensing [GC23], and dueling optimization [BGL⁺23]. While the exact definition of semi-random varies, it usually involves some seemingly benign manipulation on top of random sampling. For instance, [LM22, MPW16, MMV16, FC20] study stochastic block model, where the data is corrupted by monotone adversary that can arbitrarily add edges within the clusters and remove edges between the clusters. [CG18] considers the low-rank matrix completion problem where the adversary can only provide additional observed entries.

Notation. For a positive integer k , we use the shorthand $[k] := \{1, 2, \dots, k\}$. For any symmetric matrix $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$, $\mathbf{A} \preceq \mathbf{B}$ means $\mathbf{B} - \mathbf{A}$ is positive semidefinite, i.e., $\mathbf{v}^\top (\mathbf{B} - \mathbf{A}) \mathbf{v} \geq 0$ for any $\mathbf{v} \in \mathbb{R}^n$. For any real symmetric matrix \mathbf{A} , we use $\lambda_1(\mathbf{A}) \geq \lambda_2(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$ to denote its eigenvalues and \mathbf{A}^\dagger its Moore-Penrose pseudo-inverse. We use \mathbf{e}_i to denote the standard unit vector with 1 at i -th coordinate and 0 elsewhere. For any two real number a and b , $a \vee b$ denotes the maximum of a, b and $a \wedge b$ denotes the minimum of a, b . Additionally, the standard notation $f(n) = O(g(n))$ or $f(n) \lesssim g(n)$ means that there exists a constant $c > 0$ such that $|f(n)| \leq c|g(n)|$; $f(n) \gtrsim g(n)$ means that there exists a constant $c > 0$ such that $|f(n)| \geq c|g(n)|$. Also, $f(n) \gg g(n)$ means that there exists some large enough constant $c > 0$ such that $|f(n)| \geq c|g(n)|$; $f(n) \ll g(n)$ means that there exists some sufficiently small constant $c > 0$ such that $|f(n)| \leq c|g(n)|$.

2 Main results

We begin with formally introducing the problem setup for top- K ranking with a monotone adversary.

2.1 Problem formulation

Semi-random comparison graph. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a comparison graph over the n items of interest. In other words, items i and j are compared if and only if $(i, j) \in \mathcal{E}$. Prior work often assumes a homogeneous random comparison graph, e.g., \mathcal{G} is an Erdős–Rényi random graph where each pair (i, j) is an edge with probability p independently. Our focus in this work is to investigate the ranking problem with a monotone adversary described as follows. Let $\mathcal{G}_{\text{ER}} = (\mathcal{V}, \mathcal{E}_{\text{ER}})$ be the initial Erdős–Rényi random graph. An adversary observes \mathcal{G}_{ER} and is allowed to add edges to \mathcal{E}_{ER} *arbitrarily*. We denote the semi-random comparison graph with added edges $\mathcal{G}_{\text{SR}} = (\mathcal{V}, \mathcal{E}_{\text{SR}})$. From now on, \mathcal{G}_{SR} will be the comparison graph on which pairwise comparisons are made.

Latent scores and pairwise comparisons. In the Bradley-Terry-Luce (BTL) model, each item $i \in [n]$ is associated with a latent score θ_i^* that represents the skill level of item i . Without loss of generality, we assume that the scores are ordered, i.e., $\theta_1^* \geq \theta_2^* \geq \dots \geq \theta_n^*$.

For each pair $(i, j) \in \mathcal{E}_{\text{SR}}$ with $i > j$, we observe L outcomes $\{y_{ij}^{(l)}\}_{l \in [L]}$, which are independent Bernoulli random variables obeying

$$\mathbb{P}(y_{ij}^{(l)} = 1) = \frac{e^{\theta_j^*}}{e^{\theta_i^*} + e^{\theta_j^*}}.$$

Correspondingly, we define the average winning rate $y_{ij} := (1/L) \sum_{l=1}^L y_{ij}^{(l)}$. A simple observation is that the BTL model is shift invariant in $\boldsymbol{\theta}^*$ so we assume $\mathbf{1}_n^\top \boldsymbol{\theta}^* = 0$ without loss of generality. Finally, we define a sort of condition number $\kappa := e^{\theta_1^*} / e^{\theta_n^*}$ to characterize the range of $\boldsymbol{\theta}^*$.

Top- K recovery. Our goal is to recover the top- K items. Clearly, the hardness of the problem relies on the score difference between the K -th and the $(K+1)$ -th preferred items, which we denote by

$$\Delta_K := \theta_K^* - \theta_{K+1}^*.$$

Algorithm 1 Weighted MLE for top- K recovery under the semi-random sampling

1. Observe $\mathcal{G}_{\text{SR}} = (\mathcal{V}, \mathcal{E}_{\text{SR}})$ and $\{y_{ij}\}_{(i,j) \in \mathcal{E}_{\text{SR}}, i > j}$.
 2. Use Algorithm 2 with appropriate input to compute $\{w_{ij}\}$. The input does not include $\{y_{ij}\}_{(i,j) \in \mathcal{E}_{\text{SR}}, i > j}$.
 3. Output $\hat{\boldsymbol{\theta}} := \arg \min_{\boldsymbol{\theta}: \mathbf{1}^\top \boldsymbol{\theta} = 0} \mathcal{L}_w(\boldsymbol{\theta})$.
-

Throughout the paper, we assume $\Delta_K > 0$ so that the top- K items are well defined, and are given by $[K]$.

Now we turn to the main message of this paper: the weighted MLE, with proper weights, achieves exact top- K recovery.

2.2 Weighted MLE achieves exact recovery

Under uniform sampling, it has been shown in the work [CFMW19] that the MLE achieves exact recovery of the top- K items with optimal sample complexity. This motivates us to consider a weighted MLE for the semi-random graph \mathcal{G}_{SR} that can approximate the vanilla MLE under the purely random sampling case (i.e., with the comparison graph being \mathcal{G}_{ER}). More formally, let $\{w_{ij}\}_{i > j}$ be a set of non-negative weights supported on \mathcal{E}_{SR} , that is, $w_{ij} = 0$ if $(i, j) \notin \mathcal{E}_{\text{SR}}$. We define the weighted negative log-likelihood function $\mathcal{L}_w(\cdot)$:

$$\begin{aligned} \mathcal{L}_w(\boldsymbol{\theta}) &:= -\frac{1}{L} \sum_{i,j:i>j} \sum_{l=1}^L w_{ij} \log \left(y_{ji}^{(l)} \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} + (1 - y_{ji}^{(l)}) \frac{e^{\theta_j}}{e^{\theta_i} + e^{\theta_j}} \right) \\ &= \sum_{i,j:i>j} w_{ij} (-y_{ji}(\theta_i - \theta_j) + \log(1 + e^{\theta_i - \theta_j})), \end{aligned} \quad (2)$$

where we recall $y_{ji} = (1/L) \sum_{l=1}^L y_{ji}^{(l)}$. We then define the weighted MLE to be

$$\hat{\boldsymbol{\theta}} := \arg \min_{\boldsymbol{\theta}: \mathbf{1}^\top \boldsymbol{\theta} = 0} \mathcal{L}_w(\boldsymbol{\theta}). \quad (3)$$

The top- K items are identified by selecting the top- K entries in the estimate $\hat{\boldsymbol{\theta}}$.

The key to the success of the weighted MLE lies in a proper construction of the set of weights $\{w_{ij}\}$ that can mimic the vanilla MLE under an Erdős–Rényi random graph. It turns out that such a goal can be achieved, and we have the following guarantees for the estimation error as well as the top- K recovery performance.

Theorem 1. *Suppose that $np \geq C_1 \log(n)$ and $npL \geq C_2 \kappa^4 \log^3(n)$ for some large enough constants $C_1, C_2 > 0$. With probability at least $1 - n^{-10}$, Algorithm 1 returns the weighted MLE $\hat{\boldsymbol{\theta}}$ that obeys*

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq C_3 \kappa \sqrt{\frac{\log(n)}{npL}} \quad (4)$$

for some constant $C_3 > 0$. On this event, the top- K items are recovered exactly as long as

$$n^2 p L \geq C_4 \frac{\kappa^2 n \log(n)}{\Delta_K^2}$$

for some large enough constant $C_4 > 0$. In addition, the reweighting algorithm (i.e., Algorithm 2) runs in nearly-linear time in the size of \mathcal{G}_{SR} .

We defer the details on the construction of the weights to Section 3.2 and focus on interpreting the performance of the weighted MLE now. Similar to the literature [CS15, CFMW19, CGZ22], we assume $\kappa = O(1)$ when interpreting the results.




Figure 4: Sample complexity required to exactly recover top- K items v.s. score gap Δ_K . The solid line represent the required sample complexity of the weighted MLE as given in Theorem 1. The dashed line represents the minimax lower bound given in Theorem 2.

Near-optimal sample complexity under monotone adversary. Under the uniform sampling, the minimax sample complexity for top- K recovery has been identified in [CS15].

Theorem 2 (Informal). *Assume that $\kappa = O(1)$. If $n^2 pL \leq cn \log(n) / \Delta_K^2$ for some small constant $c > 0$, then no method whatsoever can achieve exact recovery with constant probability.*

Since uniform sampling is a special case of semi-random sampling (i.e., the adversary adds no edges at all), Theorem 2 is also a valid lower bound for the semi-random case. Comparing the performance of the weighted MLE (Theorem 1) and the lower bound (Theorem 2), we see that the weighted MLE achieves near-optimal sample complexity for top- K recovery with a monotone adversary; see Figure 4. More precisely, when $\Delta_K \lesssim 1/\log(n)$, the weighted MLE requires $n^2 pL \asymp n \log(n) / \Delta_K^2$ number of observations, which is exactly the minimax limit. On the other hand, when $1/\log(n) \lesssim \Delta_K \lesssim 1$, the weighted MLE needs $n^2 pL \asymp n \log^3(n)$ comparisons, which exceeds the minimax lower bound by at most a $\log^2(n)$ factor.

Computational complexity. Compared to a vanilla MLE, Algorithm 1 requires an additional reweighting step, which is performed by Algorithm 2. Crucially, we show that this step does not fundamentally alter the complexity of the whole procedure, as Algorithm 2 runs in nearly-linear time in the size of the graph \mathcal{G}_{SR} . To ensure this fast running time, we combine the Matrix Multiplicative Weight Update (MMWU) framework [AK16] with a number of known approximation schemes, including fast computations of the action of the matrix exponential [OSV12], randomized dimensional reduction [Ach03] and fast solvers for packing linear programs [AZO18]. We also describe a simpler algorithm which does not require solving a packing linear program as a subroutine but relies on an easy-to-compute greedy approximation. This is the algorithm we implement in Section F.

Improved dependence on κ . We remark in passing that compared to the ℓ_∞ bound $\|\hat{\theta} - \theta^*\|_\infty \lesssim \kappa^2 \sqrt{\frac{\log(n)}{npL}}$ provided in [CFMW19], which is based on a probabilistic leave-one-out argument, our ℓ_∞ error bound (4) is tighter by a κ factor.

3 Two algorithmic components

In this section we present the high-level analyses of the two algorithmic components in Algorithm 1, namely the weighted MLE, and the SDP-based reweighting method.

3.1 Weighted MLE

As the selected weight can be dependent on all edges of the comparison graph, we cannot use the popular and powerful leave-one-out technique in recent papers [CFMW19, CGZ22] to achieve entrywise control. Instead, we rely on and refine a new analysis for MLE [Che23] that is geared towards general graphs.

Recall that our estimator is a weighted MLE with weights $\{w_{ij}\}$ supported on the edges of the semi-random graph \mathcal{G}_{SR} . Let $w_{\max} := \max_{i,j} w_{ij}$ be the maximum weight, $d_{\max} := \max_{i \in [m]} \sum_{j:j \neq i} w_{ij}$ be the maximum (weighted) degree, and $d_{\min} := \min_{i \in [m]} \sum_{j:j \neq i} w_{ij}$ be the minimum (weighted) degree. In addition, we define the weighted graph Laplacian of $\mathcal{G}_{\text{SR}} = (\mathcal{V}, \mathcal{E}_{\text{SR}}, \{w_{ij}\})$ to be

$$\mathbf{L}_w := \sum_{(i,j):i>j} w_{ij}(\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top. \quad (5)$$

We have the following performance bounds for the weighted MLE whenever the weights are independent with the observed comparisons $\{y_{ij}\}$.

Theorem 3. *Suppose that the weighted graph $\mathcal{G}_{\text{SR}} = (\mathcal{V}, \mathcal{E}_{\text{SR}}, \{w_{ij}\})$ is connected. Assume that $w_{\max} \leq n^2$ and $d_{\min} \geq 1$. Further suppose that*

$$L \geq C_1 \frac{\kappa^4 w_{\max} (d_{\max})^4 \log^2(n)}{(\lambda_{n-1}(\mathbf{L}_w))^5}$$

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

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq C_2 \kappa \sqrt{\frac{w_{\max} \log(n)}{\lambda_{n-1}(\mathbf{L}_w) L}},$$

where $C_2 > 0$ is a constant. On this event, the top- K items are recovered exactly as long as

$$L \geq C_3 \frac{\kappa^2 w_{\max} \log(n)}{\lambda_{n-1}(\mathbf{L}_w) \Delta_K^2}$$

for some large enough constant $C_3 > 0$.

See Section 4 for the proof of this theorem. Note that the assumptions on w_{\max} and d_{\min} are mild and stated only to simplify the log factor.

Theorem 3 provides an ℓ_∞ error bound for the weighted MLE under general comparison graphs. The bounds depend explicitly on the spectral properties of the weighted graph, including the maximum degree, and the spectral gap of the graph Laplacian.

An important by-product of our novel analysis of the weighted MLE is to demonstrate the optimal ℓ_∞ estimation error of the vanilla MLE in the uniform sampling case.

Corollary 1. *Consider the uniform sampling case, that is the comparison graph is an Erdős–Rényi graph. Assume that $\kappa = O(1)$, and that $p \gtrsim \log(n)/n$. The vanilla MLE with high probability achieves*

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \lesssim \sqrt{\frac{\log(n)}{npL}},$$

as long as $npL \gtrsim \log^2(n)$.

Proof. Observe that the vanilla MLE is equivalent to the weighted MLE with a uniform weight 1 on all the edges. With this choice, it is easy to show (see Lemma 10) that with high probability,

$$w_{\max} \leq 1, \tag{6a}$$

$$d_{\max} \leq 2np, \tag{6b}$$

$$\lambda_{n-1}(\mathbf{L}_w) \geq np/2. \tag{6c}$$

Moreover by Lemma 14, $\lambda_{n-1}(\mathbf{L}_w) \geq np/2$ together with $p \gtrsim \log(n)/n$ implies $d_{\min} \geq 1$. Apply Theorem 3 verbatim to arrive at the desired conclusion. \square

3.2 An SDP-based reweighting

In view of the proof of Corollary 1, to mimic the vanilla MLE under uniform sampling, it is sufficient to construct the weights $\{w_{ij}\}$ such that the weighted graph $\mathcal{G}_{\text{SR}} = (\mathcal{V}, \mathcal{E}_{\text{SR}}, \{w_{ij}\})$ satisfies the spectral properties (6). In this section, we describe how to formulate this problem as a saddle-point semi-definite program (SDP) and approximately solve it in nearly-linear time in the size of \mathcal{G}_{SR} by designing a fast first-order method.

We formulate our task in terms of a convex optimization problem with variables $\{w_{ij} \in \mathbb{R}\}_{(i,j) \in \mathcal{E}_{\text{SR}}}$ representing our desired reweighting. The convex feasible set \mathcal{F} for such weights is given by the rescaled Equations (6b) and (6c):

$$\mathcal{F} := \{\mathbf{w} \in \mathbb{R}_{\geq 0}^{\mathcal{E}_{\text{SR}}} : \forall i \in \mathcal{V}, \sum_{j:(i,j) \in \mathcal{E}_{\text{SR}}} w_{ij} \leq 2np \wedge \forall (i,j) \in \mathcal{E}_{\text{SR}}, w_{ij} \leq 1\}. \tag{7}$$

For a choice of weights \mathbf{w} , we denote by \mathbf{L}_w the corresponding weighted Laplacian as defined in (5). With this notation, we consider the problem of maximizing the spectral gap $\lambda_{n-1}(\mathbf{L}_w)$ over $\mathbf{w} \in \mathcal{F}$. It is a well-known fact that this is a convex optimization problem in the variables \mathbf{w} [BDX04]. Indeed, we can write $\lambda_{n-1}(\mathbf{L}_w)$ as the minimum of the matrix inner product $\langle \mathbf{L}_w, \mathbf{X} \rangle$ over \mathbf{X} in the set

$$\Delta := \{\mathbf{X} \in \mathbb{S}^{\mathcal{V}} : \mathbf{X} \succeq 0 \wedge \langle \Pi_{\perp \mathbf{1}}, \mathbf{X} \rangle = 1\},$$

where $\mathbb{S}^{\mathcal{V}}$ is the set of symmetric linear operators over \mathcal{V} and $\Pi_{\perp \mathbf{1}}$ is the orthogonal projector over the orthogonal complement of the vector $\mathbf{1}$, which is the eigenvector of \mathbf{L}_w with the smallest eigenvalue. Therefore, our desired convex optimization problem can be recast as the following saddle point problem between an SDP variable \mathbf{X} and the weight \mathbf{w} . This formulation will be crucial for the solvers designed in the later section:

$$\text{OPT} := \max_{\mathbf{w} \in \mathcal{F}} \min_{\mathbf{X} \in \Delta} \langle \mathbf{L}_w, \mathbf{X} \rangle. \tag{Saddle-Point SDP}$$

By considering the weighting \mathbf{w} corresponding to the original graph in \mathcal{G}_{ER} as a feasible reweighting for **Saddle-Point SDP**, the proof of Corollary 1 immediately implies a lower bound on OPT: with probability at least $1 - n^{-10}$,

$$\text{OPT} \geq \frac{np}{2}, \tag{8}$$

as long as $p \geq C \log(n)/n$ for some sufficiently large constant $C > 0$.

While it is not always possible to recover the underlying graph \mathcal{G}_{ER} , the lower bound (8) ensures that, by approximately solving **Saddle-Point SDP**, we can find a reweighting of \mathcal{G}_{SR} that satisfies the required spectral properties (6). The next lemma, proved in Section 5, shows that this approximate solution can be computed in nearly-linear time in the size of \mathcal{G}_{SR} . The corresponding algorithm, Algorithm 2 to be detailed in Section 5, is based on the Matrix Multiplicative Weight Update (MMWU) framework [AK16], a first-order method for non-smooth convex SDP optimization.

Lemma 1. *Given the observed comparison graph \mathcal{G}_{SR} , there is an algorithm that computes a set of non-negative weights $\{w_{ij}\}_{i>j}$ supported on \mathcal{E}_{SR} that satisfy the properties (6) with high probability. In addition, the algorithm runs in nearly-linear time in the size of \mathcal{G}_{SR} .*

We finish this section by noting that once the spectral properties are met, repeating the proof of Corollary 1 yields the desired results in Theorem 1 in the semi-random case.

4 Analysis of weighted MLE

In this section, we present detailed analysis of the weighted MLE with the aim of proving Theorem 3.

Given a set of weights $\{w_{ij}\}$, the Hessian of the weighted loss at ground truth is given by

$$\nabla^2 \mathcal{L}(\boldsymbol{\theta}^*) = \sum_{(i,j):i>j} w_{ij} \underbrace{\frac{e^{\theta_i^*} e^{\theta_j^*}}{(e^{\theta_i^*} + e^{\theta_j^*})^2}}_{=:z_{ij}} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top.$$

This is exactly the graph Laplacian of the weighted graph \mathcal{G}_{SR} with weights $\{w_{ij}z_{ij}\}$. Hence we denote this by \mathbf{L}_{wz} . We also define the effective resistance to be

$$\Omega_{kl}(\mathbf{L}_{wz}) := (\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l),$$

where \mathbf{L}_{wz}^\dagger is the pseudo-inverse of \mathbf{L}_{wz} .

Inspired by Theorem 1 in the paper [Che23], the first step of the proof relates the performance of the weighted MLE with two crucial quantities $\{B_{kl}\}$ and $\{Q_{kl}\}$.

Lemma 2. *Suppose that the weighted graph $(\mathcal{V}, \mathcal{E}_{\text{SR}}, \{w_{ij}\})$ is connected by edges of non-zero weight and the weights $\{w_{ij}\}$ are independent with the observations $\{y_{ij}\}$. For any $(k, l) \in [n]^2, k \neq l$, let B_{kl} and Q_{kl} be some quantities obeying*

$$B_{kl} \geq C \sqrt{\frac{\kappa}{L} \left(\max_{i,j} w_{ij} \right) \Omega_{kl}(\mathbf{L}_{wz}) \log(n)}; \quad (9a)$$

$$Q_{kl} \geq \sum_{(i,j):i>j} w_{ij} B_{ij}^2 |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_i - \mathbf{e}_j)|. \quad (9b)$$

Here $C > 0$ is some large enough constant. Suppose that $Q_{kl} \leq 4B_{kl}$ for any (k, l) . Then with probability at least $1 - n^{-10}$, we have that for any $(k, l) \in [n]^2$,

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq B_{kl}.$$

On this event, the top- K items are recovered exactly as long as

$$\max_{k,l} B_{kl} \leq \Delta_K/2.$$

Admittedly, the two quantities $\{B_{kl}\}$ and $\{Q_{kl}\}$ appearing in the performance bound of MLE is quite mysterious. A key contribution of this paper is to further relate these two quantities to basic spectral properties of the weighted graph $(\mathcal{V}, \mathcal{E}_{\text{SR}}, \{w_{ij}\})$. We start with the characterization of B_{kl} . Recall that $w_{\max} = \max_{i,j} w_{ij}$ is the maximum weight.

Lemma 3. *For any $(k, l) \in [n]^2$, the effective resistance $\Omega_{kl}(\mathbf{L}_{wz})$ satisfies*

$$\Omega_{kl}(\mathbf{L}_{wz}) \leq \frac{8\kappa}{\lambda_{n-1}(\mathbf{L}_w)}.$$

As a result, it is sufficient to take

$$B_{kl} = C\kappa \sqrt{\frac{w_{\max} \log(n)}{L\lambda_{n-1}(\mathbf{L}_w)}},$$

where $C > 0$ is some large enough constant.

Now we move on to controlling the Q factors via spectral properties of the weighted graph. Recall $d_{\max} = \max_{i \in [m]} \sum_{j:j \neq i} w_{ij}$ is the maximum (weighted) degree and $d_{\min} = \min_{i \in [m]} \sum_{j:j \neq i} w_{ij}$ is the minimum (weighted) degree. We have the following bound.

Lemma 4. Suppose that $w_{\max} \leq n^2$ and that $d_{\min} \geq 1$. Then for any $(k, l) \in [n]^2$, we have

$$\sum_{(i,j):i>j} w_{ij} |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_i - \mathbf{e}_j)| \leq C\kappa \cdot \frac{(d_{\max})^2 \log(n)}{(\lambda_{n-1}(\mathbf{L}_w))^2},$$

where $C_1 > 0$ is some constant. As a result, it is sufficient to take

$$Q_{kl} = C_2 \kappa^3 \cdot \frac{w_{\max} (d_{\max})^2 \log^2(n)}{L(\lambda_{n-1}(\mathbf{L}_w))^3}$$

for some large enough constant $C_2 > 0$.

Combining Lemmas 3-4, we see that $Q_{kl} \leq 4B_{kl}$ holds as long as

$$L \geq C_3 \frac{\kappa^4 w_{\max} (d_{\max})^4 \log^2(n)}{(\lambda_{n-1}(\mathbf{L}_w))^5}$$

for some large enough constant $C_3 > 0$. This together with Lemma 2 completes the proof of Theorem 3. In what follows, we present the proofs of Lemmas 3-4, and defer the proof of Lemma 2 to Appendix B.

4.1 Proof of Lemma 3

Recall that $\Omega_{kl}(\mathbf{L}_{wz}) = (\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l)$, which implies

$$\Omega_{kl}(\mathbf{L}_{wz}) \leq 2 \|\mathbf{L}_{wz}^\dagger\| = \frac{2}{\lambda_{n-1}(\mathbf{L}_{wz})}.$$

Regarding $\lambda_{n-1}(\mathbf{L}_{wz})$, by definition, one has

$$\begin{aligned} \lambda_{n-1}(\mathbf{L}_{wz}) &= \inf_{\mathbf{v} \in \mathbb{R}, \|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{1}_n} \mathbf{v}^\top \mathbf{L}_{wz} \mathbf{v} \\ &= \inf_{\mathbf{v} \in \mathbb{R}, \|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{1}_n} \mathbf{v}^\top \left[\sum_{(i,j):i>j} w_{ij} z_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top \right] \mathbf{v} \\ &\geq \frac{1}{4\kappa} \inf_{\mathbf{v} \in \mathbb{R}, \|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{1}_n} \mathbf{v}^\top \left[\sum_{(i,j):i>j} w_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top \right] \mathbf{v} \\ &= \frac{1}{4\kappa} \lambda_{n-1}(\mathbf{L}_w), \end{aligned} \tag{10}$$

where the inequality follows from Lemma 9.

4.2 Proof of Lemma 4

For readers' convenience, we copy the key quantity appearing in Lemma 4 below:

$$\sum_{(i,j):i>j} w_{ij} |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_i - \mathbf{e}_j)|. \tag{11}$$

It turns out that this quantity is closely related to the so-called conductance of the weighted graph defined as follows.

Definition 1 (Conductance). For a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$, we define its conductance to be

$$\Phi(\mathcal{G}) := \min_{\mathcal{S} \subset \mathcal{V}} \frac{\sum_{i \in \mathcal{S}, j \in \mathcal{V} \setminus \mathcal{S}} w_{ij}}{\min\{\text{vol}(\mathcal{S}), \text{vol}(\mathcal{V} \setminus \mathcal{S})\}},$$

where the volume of a vertex set $\mathcal{S} \subset \mathcal{V}$ is defined by

$$\text{vol}(\mathcal{S}) := \sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{V}: j \neq i} w_{ij}.$$

Algorithm 2 MMWU algorithm for **Saddle-Point SDP**

Input: Graph $\mathcal{G}_{\text{SR}} = (V, \mathcal{E}_{\text{SR}})$, $|V| = n$, $|\mathcal{E}_{\text{SR}}| = m$, error parameter $\epsilon \in (0, 1/2]$.

- 1: $\mathbf{w}^{(0)} = \mathbf{0}, \eta = \epsilon/4pn, k = O(\log(n)/\epsilon^2), T = 8\log(n)/\epsilon^2$
 - 2: **for** $t = 0, \dots, T - 1$ **do**
 - 3: Sample $\mathbf{R} = \{\pm 1/\sqrt{k}\}^{n \times k}$ uniformly at random. // $\{\pm 1\}$ JL matrix with $1/\sqrt{k}$ scaling [Ach03]
 - 4: Compute $\mathbf{U}^{(t)} = \exp\{-\eta \sum_{ij \in \mathcal{E}_{\text{SR}}} w_{ij}^{(t)} \mathbf{L}_{ij}\} \mathbf{R}$. // Use matrix exponentiation in [AK16]
 - 5: Compute $\mathbf{V}^{(t)} = \mathbf{U}^{(t)} / \sqrt{\langle \Pi_{\perp 1}, \mathbf{U}^{(t)} (\mathbf{U}^{(t)})^T \rangle}$. // MMWU Normalization
 - 6: For each $ij \in \mathcal{E}_{\text{SR}}$, let $c_{ij} = \langle \mathbf{L}_{ij}, \mathbf{V}^{(t)} (\mathbf{V}^{(t)})^T \rangle$. // Compute edge gains
 - 7: Compute $\hat{\mathbf{w}}^{(t)} \in \mathcal{F}$, an ϵ - or $1/2$ -approximate solution to $\max_{\mathbf{w} \in \mathcal{F}} \langle \mathbf{c}, \mathbf{w} \rangle$. // Use Theorem 4 or 5
 - 8: Update $\mathbf{w}^{t+1} = \mathbf{w}^{(t)} + \hat{\mathbf{w}}^{(t)}$.
 - 9: **end for**
- Output:** Edge weighting $\mathbf{w}_{\text{out}} = 1/T \cdot \sum_{t=0}^{T-1} \mathbf{w}^{(t)}$.
-

The following lemma links the quantity in (11) with graph conductance. This lemma is modified from Lemma 28 in [KLOS14]. The proof is deferred to Section C.

Lemma 5. *Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ equipped with two set of weights $\{w_{ij}\}$ and $\{\tilde{w}_{ij}\}$ both supported on \mathcal{E} . Suppose that the minimum $\{w_{ij}\}$ -weighted degree is at least 1. Then one has*

$$\sum_{(i,j):i>j} w_{ij} \left| (\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{L}_{\mathcal{G}_w}^\dagger (\mathbf{e}_k - \mathbf{e}_l) \right| \leq \max_{(i,j) \in \mathcal{E}} \left\{ \frac{w_{ij}}{\tilde{w}_{ij}} \right\} \cdot \frac{8 \log \left(\sum_{(i,j):i>j} w_{ij} \right)}{\Phi(\mathcal{G}_w)^2}. \quad (12)$$

Set $\tilde{w}_{ij} = w_{ij} z_{ij}$. We can apply Lemma 5 to obtain

$$\begin{aligned} \sum_{(i,j):i>j} w_{ij} \left| (\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l) \right| &\leq \max_{(i,j):w_{ij}>0} \left\{ \frac{w_{ij}}{w_{ij} z_{ij}} \right\} \cdot \frac{8 \log \left(\sum_{(i,j):i>j} w_{ij} \right)}{\Phi(\mathcal{G}_w)^2} \\ &\leq 4\kappa \cdot \frac{8 \log \left(\sum_{(i,j):i>j} w_{ij} \right)}{\Phi(\mathcal{G}_w)^2}, \end{aligned}$$

where the last inequality again follows from Lemma 9.

For the numerator, since $w_{\max} \leq n^2$, we have $\log(\sum_{(i,j):i>j} w_{ij}) \leq 3\log(n)$. Now we focus on the denominator, i.e., the graph conductance. It is well known that the graph conductance is controlled by the eigenvalue of the normalized Laplacian (see Lemma 11), that is

$$\Phi(\mathcal{G}_w) \geq \frac{1}{2} \lambda_{n-1} \left(\mathbf{D}_w^{-1/2} \mathbf{L}_w \mathbf{D}_w^{-1/2} \right),$$

where \mathbf{D}_w is a diagonal matrix composed of the weighted degrees. By Sylvester's law of inertia (Lemma 13), we further have

$$\lambda_{n-1} \left(\mathbf{D}_w^{-1/2} \mathbf{L}_w \mathbf{D}_w^{-1/2} \right) \geq \lambda_n \left(\mathbf{D}_w^{-1} \right) \lambda_{n-1} \left(\mathbf{L}_w \right) = (d_{\max})^{-1} \cdot \lambda_{n-1} \left(\mathbf{L}_w \right).$$

Taking the above bounds collectively yields the desired claim in Lemma 4.

5 Analysis of SDP-based reweighting

In this section, we describe and analyze the MMWU algorithm for solving the reweighting SDP problem **Saddle-Point SDP** in Section 3.2. We conclude by proving Lemma 1.

5.1 MMWU algorithm for Saddle-Point SDP

We present the pseudocode for the the MMWU Algorithm in Algorithm 2. Our algorithm instantiates the MMWU framework of Arora and Kale [AK16], where we avoid maintaining full $|\mathcal{V}| \times |\mathcal{V}|$ matrices by relying on the Johnson-Lindestrauss Lemma [Ach03] (see lines 3 and 4).

At every iteration, the MMWU algorithm produces a candidate solution $\mathbf{X}^{(t)} \in \Delta$ to which we respond with a loss matrix $\mathbf{L}^{(t)} := \mathbf{L}_{\mathbf{w}^{(t)}} = \sum_{ij \in \mathcal{E}_{\text{SR}}} w_{ij}^{(t)} \mathbf{L}_{ij}$ with $\mathbf{w}^{(t)} \in \mathcal{F}$. In this way, the loss $\langle \mathbf{L}^{(t)}, \mathbf{X}^{(t)} \rangle$ incurred by the MMWU algorithm at iteration t equals the value of Saddle-Point SDP on the pair of solutions $(\mathbf{X}^{(t)}, \mathbf{w}^{(t)})$. At every iteration t , given $\mathbf{X}^{(t)}$, our goal is then to choose $\mathbf{w}^{(t)} \in \mathcal{F}$ to maximize the loss $\langle \mathbf{L}^{(t)}, \mathbf{X}^{(t)} \rangle$. The regret minimization property of MMWU then allows us to turn this per-iteration guarantee into a global guarantee on $\lambda_{n-1}(\sum_{t=0}^{T-1} \mathbf{L}^{(t)})$.

To maximize the loss of the MMWU algorithm, we choose $\mathbf{w}^{(t)} \in \mathcal{F}$ to approximate the best response $\max_{\mathbf{w} \in \mathcal{F}} \langle \mathbf{L}_{\mathbf{w}}, \mathbf{X}^{(t)} \rangle$. We describe two algorithms (oracles in the language of [AK16]) for approximately solving this task. The first one is based on directly applying a nearly-linear-time packing LP solver. It is described in Appendix D and yields the following theorem.

Theorem 4. *Given $\mathbf{X} \in \Delta$, one can $(1 - \epsilon)$ -approximate $\max_{\mathbf{w} \in \mathcal{F}} \langle \mathbf{L}_{\mathbf{w}}, \mathbf{X} \rangle$ multiplicatively in time $\tilde{O}(|\mathcal{E}_{\text{SR}}|/\epsilon)$.*

Our second oracle exploits the fact that $\max_{\mathbf{w} \in \mathcal{F}} \langle \mathbf{L}_{\mathbf{w}}, \mathbf{X} \rangle = \max_{\mathbf{w} \in \mathcal{F}} \langle \mathbf{c}, \mathbf{w} \rangle$ is a LP relaxation of the maximum weight b -matching problem over \mathcal{G}_{SR} with edge weights \mathbf{c} . We can approximate the maximum weight b -matching by a greedy procedure: iterate through the edges of \mathcal{E}_{SR} once in decreasing order of edge gains $c_{ij} = \langle \mathbf{L}_{ij}, \mathbf{X} \rangle$, and add an edge so long as both adjacent vertices possess available demand. Though this only achieves a $1/2$ -approximation [Mes06], the simplicity of the algorithm makes it extremely suitable for implementation. We prove the following theorem in Appendix D.

Theorem 5. *Given $\mathbf{X} \in \Delta$, one can $1/2$ -approximate $\max_{\mathbf{w} \in \mathcal{F}} \langle \mathbf{L}_{\mathbf{w}}, \mathbf{X} \rangle$ multiplicatively in time $\tilde{O}(|\mathcal{E}_{\text{SR}}|)$.*

5.2 Analysis of Algorithm 2

We are now ready to prove Lemma 1. To analyze the correctness and running time of Algorithm 2, we first recall the regret bound achieved by applying MMWU to the vector space $\mathbb{R}^V \perp \mathbf{1}$.

Theorem 6 (Theorem 3.1 [AK16]). *Consider a sequence of loss matrices $\{\mathbf{L}^{(t)} \in \mathbb{S}^V\}_{t \in [T]}$ with $\mathbf{L}^{(t)} \mathbf{1} = 0$ for all t . Let*

$$\mathbf{W}^{(t)} = \exp \left\{ -\eta \sum_{s=0}^{t-1} \mathbf{L}^{(s)} \right\}, \quad \text{and} \quad \mathbf{X}^{(t)} = \frac{\mathbf{W}^{(t)}}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{W}^{(t)} \rangle}.$$

Then, we have the regret bound:

$$\min_{\mathbf{X} \in \Delta} \left\langle \sum_{t=0}^{T-1} \mathbf{L}^{(t)}, \mathbf{X} \right\rangle \geq \sum_{t=0}^{T-1} \langle \mathbf{L}^{(t)}, \mathbf{X}^{(t)} \rangle - \eta \sum_{s=0}^{T-1} \langle (\mathbf{L}^{(s)})^2, \mathbf{X}^{(s)} \rangle - \frac{\log(n)}{\eta}. \quad (13)$$

Next, we use Lemma 6 to show that $\mathbf{V}^{(t)}$, produced by Algorithm 2, approximate the MMWU updates $\mathbf{X}^{(t)}$ in Theorem 6 when computing the squared distances $\langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle$. This is proven in Appendix D.2.

Lemma 6. *Let $\tilde{\mathbf{X}}^{(t)} = \mathbf{V}^{(t)}(\mathbf{V}^{(t)})^T$ as defined in Algorithm 2. Then, $\tilde{\mathbf{X}}^{(t)} \in \Delta$. Moreover, for large enough $k \in O(\log(n)/\epsilon^2)$, for all pairs $(i, j) \in \mathcal{V} \times \mathcal{V}$, the squared distance $\langle \mathbf{L}_{ij}, \tilde{\mathbf{X}}^{(t)} \rangle$ is an ϵ -multiplicative approximation to the squared distance $\langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle$ with high probability.*

We can use the above to establish the algorithm's correctness; we prove the following lemma in Section 5.3.

Lemma 7. *With high probability, Algorithm 2 outputs a reweighting $\mathbf{w}_{\text{out}} \in \mathcal{F}$ of \mathcal{G}_{SR} such that*

$$\min_{\mathbf{X} \in \Delta} \langle \mathbf{L}_{\mathbf{w}_{\text{out}}}, \mathbf{X} \rangle \geq \begin{cases} (1 - O(\epsilon)) \cdot \text{OPT}, & \text{if Theorem 4 is used in line 7,} \\ (1/2 - O(\epsilon)) \cdot \text{OPT}, & \text{if Theorem 5 is used in line 7.} \end{cases}$$

Finally, we can bound the running time of Algorithm 2 using known algorithms for computing the action of the matrix exponential [AK16, OSV12]. The proof of the next lemma is also given in Appendix D.2.

Lemma 8. *For a constant $\epsilon > 0$, Algorithm 2 runs in nearly-linear time in the size of \mathcal{G}_{SR} .*

Combining Lemmata 7 and 8 yields our main result on the reweighting of \mathcal{G}_{SR} , i.e., Lemma 1.

Proof of Lemma 1. We claim that the output \mathbf{w}_{out} of Algorithm 2 satisfies the required properties. The conditions of (6b) and (6a) are immediately satisfied by the fact that $\mathbf{w}_{\text{out}} \in \Delta$, which is proved in Lemma 7. The spectral condition in (6c) follows from the approximation guarantee of Lemma 7 and the lower bound (8):

$$\lambda_{n-1}(\mathbf{L}_{\mathbf{w}_{\text{out}}}) = \cdot \min_{\mathbf{X} \in \Delta} \langle \mathbf{L}_{\mathbf{w}_{\text{out}}}, \mathbf{X} \rangle \geq \Omega(\text{OPT}) \geq \Omega(np).$$

The nearly-linear running time is proved in Lemma 8. \square

5.3 Proof of Lemma 7

Notice that $\mathbf{w}_{\text{out}} \in \mathcal{F}$ as it is the average of elements of \mathcal{F} and \mathcal{F} is convex. Moreover, we have $\mathbf{L}_{\mathbf{w}_{\text{out}}} = 1/T \cdot \sum_{t=0}^{T-1} \mathbf{L}_{\mathbf{w}^{(t)}} = 1/T \cdot \sum_{t=0}^{T-1} \mathbf{L}^{(t)}$. By the regret bound in Equation (13), we then obtain:

$$\min_{\mathbf{X} \in \Delta} \langle \mathbf{L}_{\mathbf{w}_{\text{out}}}, \mathbf{X} \rangle \geq \frac{1}{T} \left(\sum_{t=0}^{T-1} \langle \mathbf{L}_{\mathbf{w}^{(t)}}, \mathbf{X}^{(t)} \rangle \right) - \frac{\eta}{T} \left(\sum_{s=0}^{T-1} \langle (\mathbf{L}_{\mathbf{w}^{(t)}})^2, \mathbf{X}^{(t)} \rangle \right) - \frac{\log(n)}{\eta T}. \quad (14)$$

We can exploit the positive semi-definiteness of $\mathbf{L}_{\mathbf{w}^{(t)}}$ to bound the second term as a function of the first. By the definition of \mathcal{F} , the reweighting of \mathcal{G} by $\mathbf{w}^{(t)} \in \mathcal{F}$ has maximum degree at most $2pn$. Hence:

$$\mathbf{L}_{\mathbf{w}^{(t)}} \preceq 2 \cdot 2pn \cdot \Pi_{\perp \mathbf{1}} \text{ and } \langle (\mathbf{L}_{\mathbf{w}^{(t)}})^2, \mathbf{X}^{(t)} \rangle \leq 4pn \cdot \langle \mathbf{L}_{\mathbf{w}^{(t)}}, \mathbf{X}^{(t)} \rangle.$$

We can now rewrite Equation (14) as:

$$\min_{\mathbf{X} \in \Delta} \langle \mathbf{L}_{\mathbf{w}_{\text{out}}}, \mathbf{X} \rangle \geq \left(\frac{1}{T} - \frac{4pn \cdot \eta}{T} \right) \left(\sum_{t=0}^{T-1} \langle \mathbf{L}_{\mathbf{w}^{(t)}}, \mathbf{X}^{(t)} \rangle \right) - \frac{\log(n)}{\eta T}.$$

For all iterations t , we have:

$$\langle \mathbf{L}_{\mathbf{w}^{(t)}}, \mathbf{X}^{(t)} \rangle \geq (1 - \epsilon) \langle \mathbf{L}_{\mathbf{w}^{(t)}}, \tilde{\mathbf{X}}^{(t)} \rangle = (1 - \epsilon) \sum_{ij \in \mathcal{E}_{\text{SR}}} w_{ij}^{(t)} \langle \mathbf{L}_{ij}, \mathbf{V}^{(t)} (\mathbf{V}^{(t)})^T \rangle = (1 - \epsilon) \max_{\mathbf{w} \in \mathcal{F}} \langle \mathbf{L}_{\mathbf{w}}, \tilde{\mathbf{X}}^{(t)} \rangle,$$

where the first inequality follows from Lemma 6 and the last equality follows from Lines 6 and 7 in Algorithm 2. As $\tilde{\mathbf{X}}^{(t)} \in \Delta$, the maximum in the last expression has value at least OPT. Hence, for the two oracles of Theorems 4 and 5, we have:

$$\text{Theorem 4: } \langle \mathbf{L}_{\mathbf{w}^{(t)}}, \mathbf{X}^{(t)} \rangle \geq (1 - \epsilon)^2 \cdot \text{OPT} \text{ and Theorem 5: } \langle \mathbf{L}_{\mathbf{w}^{(t)}}, \mathbf{X}^{(t)} \rangle \geq \frac{(1 - \epsilon)}{2} \cdot \text{OPT}.$$

Therefore:

$$\min_{\mathbf{X} \in \Delta} \langle \mathbf{L}_{\mathbf{w}_{\text{out}}}, \mathbf{X} \rangle \geq \begin{cases} (1 - 4pn \cdot \eta) \cdot (1 - \epsilon)^2 \cdot \text{OPT} - \frac{\log(n)}{\eta T} \\ (1 - 4pn \cdot \eta) \cdot \frac{(1 - \epsilon)}{2} \cdot \text{OPT} - \frac{\log(n)}{\eta T} \end{cases}$$

Substituting the definitions of $\eta = \epsilon/4pn$ yields:

$$\min_{\mathbf{X} \in \Delta} \langle \mathbf{L}_{\mathbf{w}_{\text{out}}}, \mathbf{X} \rangle \geq \begin{cases} (1 - \epsilon)^3 \cdot \text{OPT} - \frac{4pn \log(n)}{\epsilon T} \\ \frac{(1 - \epsilon)^2}{2} \cdot \text{OPT} - \frac{4pn \log(n)}{\epsilon T} \end{cases}$$

By the lower bound (8), the last term in both expressions can be upper bound by $8\text{OPT} \log(n)/\epsilon T$. The statement of the lemma follows from the definition $T = 8 \log(n)/\epsilon^2$.

6 Discussion

In this paper we consider top- K ranking with a monotone adversary. We carefully construct a weighted MLE that achieves the optimal ℓ_∞ estimation error and top- K recovery sample complexity. This leaves open quite a few interesting directions. We single out several of them below.

- *Shaving the $\log^2(n)$ factor.* Compared to the results for uniform sampling, ours requires an extra assumption that $npL \gtrsim \log^3(n)$. While this is optimal for a wide range of Δ_K , it does incur an extra $\log^2(n)$ factor when $1/\log(n) \lesssim \Delta_K \lesssim 1$. We do not expect this to be the fundamental gap between uniform sampling and semi-random sampling. Successfully shaving this extra log factor will potentially bring more insights to understanding the geometry of the comparison graph on the ranking problem.
- *Is weighted MLE necessary?* While our approach relies on the weighted MLE, it is not clear whether the unweighted vanilla MLE succeeds or not with a monotone adversary. In fact, in Appendix E, we present a specific example of semi-random sampling where the spectral properties of the semi-random graph are drastically different from those of random graph, yet the vanilla MLE still succeeds.
- *Extension to general graphs.* Theorem 1 in [Che23] and Theorem 2 in our paper both consider general comparison graphs. Our analysis can recover a good rate with small sample complexity when the spectrum of the comparison graph has small dynamic range, i.e. λ_1/λ_{n-1} is small. Our reweighting procedure indicates that any comparison graph which can be reweighted to match this spectral condition will yield small sample complexity. It is an interesting question to explore the connection between the reweighting SDP and similar SDPs used in the approximation of general partitioning objectives [LTW23] to find novel combinatorial or geometric properties enabling the success of the weighted MLE.
- *Extensions to other ranking models.* In this paper we focus on one of the most popular model in ranking, namely the BTL model. It is certainly interesting to see whether our algorithm and analysis for semi-random sampling can be extended to other models, for instance the Thurstone model [Thu27], the Plackett-Luce model [Luc05], and other models for multi-way comparisons [FLWY22, FLWY23].

Acknowledgements

AC is supported by NSF DGE 2140001. LO is supported by NSF CAREER 1943510. CM was partially supported by the National Science Foundation via grant DMS-2311127.

References

- [Ach03] Dimitris Achlioptas. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *Journal of Computer and System Sciences*, 66(4):671–687, June 2003.
- [AFWZ20] Emmanuel Abbe, Jianqing Fan, Kaizheng Wang, and Yiqiao Zhong. Entrywise eigenvector analysis of random matrices with low expected rank. *Annals of statistics*, 48(3):1452, 2020.
- [AK16] Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. *Journal of the ACM*, 63(2):12:1–12:35, May 2016.
- [AV18] Pranjali Awasthi and Aravindan Vijayaraghavan. Clustering semi-random mixtures of Gaussians. In *International Conference on Machine Learning*, pages 294–303, 2018.
- [AZO18] Zeyuan Allen-Zhu and Lorenzo Orecchia. Nearly linear-time packing and covering lp solvers: Achieving width-independence and -convergence. *Mathematical Programming*, 175, 02 2018.
- [BDX04] Stephen Boyd, Persi Diaconis, and Lin Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, January 2004.

- [BGL⁺23] Avrim Blum, Meghal Gupta, Gene Li, Naren Sarayu Manoj, Aadirupa Saha, and Yuanyuan Yang. Dueling optimization with a monotone adversary. *arXiv preprint arXiv:2311.11185*, 2023.
- [Bol98] Béla Bollobás. *Modern graph theory*, volume 184. Springer Science & Business Media, 1998.
- [Bor81] de Borda. Mémoire sur les élections au scrutin: Histoire de l’Académie royale des sciences. *Paris, France*, 12, 1781.
- [BS95] Avrim Blum and Joel Spencer. Coloring random and semi-random k -colorable graphs. *Journal of Algorithms*, 19(2):204–234, 1995.
- [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.
- [CCF⁺21] Yuxin Chen, Yuejie Chi, Jianqing Fan, Cong Ma, et al. Spectral methods for data science: A statistical perspective. *Foundations and Trends® in Machine Learning*, 14(5):566–806, 2021.
- [CCFM19] Yuxin Chen, Yuejie Chi, Jianqing Fan, and Cong Ma. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, 176(1-2):5–37, 2019.
- [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.
- [CG18] Yu Cheng and Rong Ge. Non-convex matrix completion against a semi-random adversary. In *Conference On Learning Theory*, pages 1362–1394. PMLR, 2018.
- [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.
- [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.
- [DKNS01] Cynthia Dwork, Ravi Kumar, Moni Naor, and Dandapani Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th international conference on World Wide Web*, pages 613–622, 2001.
- [EK18] Nouredine El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, 170(1):95–175, 2018.
- [Elo67] Arpad E Elo. The proposed USCF rating system, its development, theory, and applications. *Chess Life*, 22(8):242–247, 1967.
- [FC20] Yingjie Fei and Yudong Chen. Achieving the Bayes error rate in synchronization and block models by SDP, robustly. *IEEE Transactions on Information Theory*, 66(6):3929–3953, 2020.
- [FJ57] Lester R Ford Jr. Solution of a ranking problem from binary comparisons. *The American Mathematical Monthly*, 64(8P2):28–33, 1957.
- [FK01] Uriel Feige and Joe Kilian. Heuristics for semirandom graph problems. *Journal of Computer and System Sciences*, 63(4):639–671, 2001.
- [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.
- [GC23] Xing Gao and Yu Cheng. Robust matrix sensing in the semi-random model. In *Proceedings of the 37th Conference on Neural Information Processing Systems*. Curran Associates, 2023.
- [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.
- [HOX14] Bruce Hajek, Sewoong Oh, and Jiaming Xu. Minimax-optimal inference from partial rankings. *Advances in Neural Information Processing Systems*, 27, 2014.
- [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.
- [KLL⁺23] Jonathan Kelner, Jerry Li, Allen X Liu, Aaron Sidford, and Kevin Tian. Semi-random sparse recovery in nearly-linear time. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2352–2398. PMLR, 2023.
- [KLOS14] Jonathan A Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 217–226. SIAM, 2014.
- [LM22] Allen Liu and Ankur Moitra. Minimax rates for robust community detection. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 823–831. IEEE, 2022.
- [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.
- [LTW23] Lap Chi Lau, Kam Chuen Tung, and Robert Wang. Cheeger Inequalities for Directed Graphs and Hypergraphs using Reweighted Eigenvalues. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023*, pages 1834–1847, New York, NY, USA, June 2023. Association for Computing Machinery.
- [Luc05] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2005.
- [Mes06] Julián Mestre. Greedy in Approximation Algorithms. In Yossi Azar and Thomas Erlebach, editors, *Algorithms – ESA 2006*, Lecture Notes in Computer Science, pages 528–539, Berlin, Heidelberg, 2006. Springer.
- [MMV12] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. Approximation algorithms for semi-random partitioning problems. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 367–384, 2012.
- [MMV16] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. Learning communities in the presence of errors. In *Conference on learning theory*, pages 1258–1291. PMLR, 2016.
- [MPW16] Ankur Moitra, William Perry, and Alexander S Wein. How robust are reconstruction thresholds for community detection? In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 828–841, 2016.
- [MWCC18] Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval and matrix completion. In *International Conference on Machine Learning*, pages 3345–3354. PMLR, 2018.

- [NOS12] Sahand N. Negahban, Sewoong Oh, and Devavrat Shah. Rank centrality: Ranking from pairwise comparisons. *Oper. Res.*, 65:266–287, 2012.
- [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.
- [OSV12] Lorenzo Orecchia, Sushant Sachdeva, and Nisheeth K. Vishnoi. Approximating the exponential, the Lanczos method and an $\tilde{O}(m)$ -time spectral algorithm for balanced separator. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC ’12, pages 1141–1160, New York, NY, USA, May 2012. Association for Computing Machinery.
- [SBB⁺16] Nihar B Shah, Sivaraman Balakrishnan, Joseph Bradley, Abhay Parekh, Kannan Ramch, Martin J Wainwright, et al. Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. *Journal of Machine Learning Research*, 17(58):1–47, 2016.
- [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.
- [Thu27] L. L. Thurstone. A law of comparative judgment. *Psychological Review*, 34:273–286, 1927.
- [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.
- [WYH⁺18] Weiqing Wang, Hongzhi Yin, Zi Huang, Qinyong Wang, Xingzhong Du, and Quoc Viet Hung Nguyen. Streaming ranking based recommender systems. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 525–534, 2018.
- [ZB18] Yiqiao Zhong and Nicolas Boumal. Near-optimal bounds for phase synchronization. *SIAM Journal on Optimization*, 28(2):989–1016, 2018.

A Auxiliary lemmas

This section collects several auxiliary lemmas we use in the proofs of our main results.

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

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

We have for any i, j ,

$$\frac{1}{4\kappa} \leq z_{ij} \leq \frac{1}{4}.$$

Proof. The function $f(x) = x/(1+x)^2$ has derivative $(1-x^2)/(1+x)^4$ so it is increasing in $(0, 1)$ and decreasing in $(1, \infty)$. Furthermore from the definition of κ , $|\theta_i^* - \theta_j^*| \leq \log(\kappa)$. Then we have

$$\frac{1}{4\kappa} \leq \min \{f(\kappa), f(1/\kappa)\} \leq z_{ij} \leq f(1) = \frac{1}{4}.$$

This completes the proof. □

Lemma 10 (spectral gap of a Erdős-Rényi graph). *Let $\mathcal{G} = ([n], \mathcal{E}) \sim G(n, p)$ be an Erdős-Rényi graph with n vertices and edge probability p . Let \mathbf{L} be its corresponding graph Laplacian matrix. Let d_{\max} be the maximum degree of the vertices. Suppose that $np \geq C \log(n)$ for some large enough constant $C > 0$, then with probability at least $1 - n^{-10}$,*

$$\lambda_{n-1}(\mathbf{L}) \geq \frac{np}{2}$$

$$d_{\max} \leq 2np.$$

Proof. See, for instance, Section 5.3.3 of [Tro15]. \square

Lemma 11 ([Spi07]). *Let \mathcal{G} be a weighted graph and $\Phi(\mathcal{G})$ be its conductance (see Definition 1). Let \mathbf{L} be its graph Laplacian matrix. Let \mathbf{D} be the diagonal matrix with the weighted vertex degrees as the entries. Then*

$$\Phi(\mathcal{G}) \geq \frac{1}{2} \lambda_{n-1}(\mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}).$$

Lemma 12 (Rayleigh's monotonicity law, [Bol98] Corollary 7 in Ch.IX). *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph with weights $\{w_{ij}\}$ on \mathcal{E} , and $\tilde{\mathcal{G}} = (\mathcal{V}, \tilde{\mathcal{E}})$ be a graph with weights $\{\tilde{w}_{ij}\}$ on $\tilde{\mathcal{E}}$. Suppose that $\mathcal{E} \subset \tilde{\mathcal{E}}$ and $w_{ij} \leq \tilde{w}_{ij}$ for any $(i, j) \in \mathcal{E}$, then for any $(k, l) \in \mathcal{V}^2$, the effective resistance satisfies*

$$\Omega_{kl}(\mathbf{L}_{\mathcal{G}_w}) \leq \Omega_{kl}(\mathbf{L}_{\tilde{\mathcal{G}}_{\tilde{w}}}).$$

Lemma 13 (A quantitative version of Sylvester's law of inertia, [Ost59]). *For any real symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{S} \in \mathbb{R}^{n \times n}$ be a non-singular matrix. Then for any $i \in [n]$, $\lambda_i(\mathbf{S} \mathbf{A} \mathbf{S}^\top)$ lies between $\lambda_i(\mathbf{A}) \lambda_1(\mathbf{S}^\top \mathbf{S})$ and $\lambda_i(\mathbf{A}) \lambda_n(\mathbf{S}^\top \mathbf{S})$.*

Lemma 14. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \{w_{ij}\}_{(i,j) \in \mathcal{E}})$ be a weighted graph. Recall $d_{\min} = \min_{i \in [n]} \sum_{j: (i,j) \in \mathcal{E}} w_{ij}$ is the minimum weighted degree and $\mathbf{L}_w = \sum_{i>j} w_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top$ is the weighted graph Laplacian. Then*

$$d_{\min} \geq \frac{1}{2} \lambda_{n-1}(\mathbf{L}_w).$$

Proof. Let $k \in [n]$ be an arbitrary vertex. Let $\mathbf{v} \in \mathbb{R}^n$ be a vertex defined by $v_k = 1$ and $v_j = -1/(n-1)$ for all $j \neq k$. It is easy to see that $\mathbf{v}^\top \mathbf{1}_n = 0$. Moreover, for any $i, j \neq k$, $(\mathbf{e}_k - \mathbf{e}_j)^\top \mathbf{v} = n/(n-1)$ and $(\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{v} = 0$. By the definition of $\lambda_{n-1}(\mathbf{L}_w)$, we have that

$$\begin{aligned} \lambda_{n-1}(\mathbf{L}_w) &= \min_{\mathbf{u} \in \mathbb{R}^n: \mathbf{u}^\top \mathbf{1}_n = 0} \frac{\mathbf{u}^\top \mathbf{L}_w \mathbf{u}}{\|\mathbf{u}\|^2} \\ &\leq \frac{\mathbf{v}^\top \mathbf{L}_w \mathbf{v}}{\|\mathbf{v}\|^2} \\ &= \frac{\sum_{j: j \neq k} w_{kj} \left(\frac{n}{n-1}\right)^2}{n/(n-1)} \\ &\leq 2 \sum_{j: j \neq k} w_{kj} \end{aligned}$$

as long as $n \geq 2$. This holds for all $k \in [n]$ so $d_{\min} \geq \frac{1}{2} \lambda_{n-1}(\mathbf{L}_w)$. \square

B Proof of Lemma 2

We start by defining some useful notations. Let $\epsilon_{ij}^{(l)} = y_{ji}^{(l)} - \sigma(\theta_i^* - \theta_j^*)$ where $\sigma(x) = e^x / (1 + e^x)$ is the sigmoid function. Let $\mathbf{B} := [\dots, \sqrt{w_{ij} z_{ij}} (\mathbf{e}_i - \mathbf{e}_j), \dots] \in \mathbb{R}^{n \times n(n-1)L/2}$ (each entry repeats for L times), $\hat{\boldsymbol{\epsilon}} = [\dots, \sqrt{w_{ij} / z_{ij}} \epsilon_{ji}^{(l)}, \dots] \in \mathbb{R}^{n(n-1)L/2}$ and $\boldsymbol{\delta}^t := \boldsymbol{\theta}^t - \boldsymbol{\theta}^*$. Observe that

$$\mathbf{B} \mathbf{B}^\top = L \sum_{(i,j): i>j} w_{ij} z_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top = L \cdot \mathbf{L}_{wz}. \quad (15)$$

Inspired by [Che23], we analyze the weighted MLE by studying the preconditioned gradient descent starting from the ground truth. Let $\boldsymbol{\theta}_0 = \boldsymbol{\theta}^*$ be the starting point and $\eta > 0$ be the stepsize.

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta \mathbf{L}_{wz}^\dagger \nabla \mathcal{L}_w(\boldsymbol{\theta}^t). \quad (16)$$

It is worth noting that because the preconditioned gradient descent starts at ground truth and the gradient is of form

$$\nabla \mathcal{L}_w(\boldsymbol{\theta}) = \frac{1}{L} \sum_{i,j:i>j} \sum_{l=1}^L w_{ij} \left\{ -y_{ji}^{(l)} + \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} \right\} (\mathbf{e}_i - \mathbf{e}_j),$$

we have $\mathbf{1}_n^\top \boldsymbol{\theta}^t = 0$ for all $t \geq 0$. We break the proof of Lemma 2 into the following lemmas. The first states that $\boldsymbol{\theta}^t$ stays close to the ground truth in ℓ_∞ distance, and the second states that it converges to the weighted MLE.

Lemma 15. *Instate the assumptions of Lemma 2. With probability at least $1 - n^{-10}$, (20) is satisfied and for any k, l and iteration $t \geq 0$,*

$$|(\theta_k^t - \theta_l^t) - (\theta_k^* - \theta_l^*)| \leq B_{kl}. \quad (17)$$

For any $\mathbf{x} \in \mathbb{R}$ and $\mathbf{Z} \in \mathbb{R}^{n \times n}$, let $\|\mathbf{x}\|_{\mathbf{Z}} := (\mathbf{x}^\top \mathbf{Z} \mathbf{x})^{1/2}$.

Lemma 16. *Instate the assumptions of Lemma 2. On the event that (20) is satisfied, the following holds:*

1. *There exists a unique minimizer $\hat{\boldsymbol{\theta}}$ of (2).*
2. *There exists some $\alpha_1, \alpha_2 \in \mathbb{R}$ such that $0 < \alpha_1 \leq \alpha_2$ and for any $t \in \mathbb{N}$,*

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

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

We will prove these two lemmas in Section B.1 and Section B.2. Their proofs are similar to the proof of Theorem 1 in [Che23]. As we are studying weighted MLE instead of the unweighted MLE used in [Che23], we redo the proofs and add necessary modifications.

Proof of Lemma 2. It is easy to see that Lemma 16 shows the preconditioned gradient descent iterates $\boldsymbol{\theta}^t$ converge to a unique solution $\hat{\boldsymbol{\theta}}$. Combining this with Lemma 15, we have

$$\left| (\hat{\theta}_k - \hat{\theta}_l) - (\theta_k^* - \theta_l^*) \right| \leq B_{kl}.$$

As $\mathbf{1}_n^\top \boldsymbol{\theta}^* = \mathbf{1}_n^\top \hat{\boldsymbol{\theta}} = 0$, for any $i \in [n]$,

$$\begin{aligned} \left| \hat{\theta}_i - \theta_i^* \right| &= \left| \frac{1}{n} \sum_{j=1}^n (\hat{\theta}_i - \hat{\theta}_j) - (\theta_i^* - \theta_j^*) \right| \\ &\leq \frac{1}{n} \sum_{j=1}^n \left| (\hat{\theta}_i - \hat{\theta}_j) - (\theta_i^* - \theta_j^*) \right| \\ &\leq \max_{k,l} B_{kl}. \end{aligned}$$

It remains to show the exact recovery of the top- K items. Recall $\theta_1^* \geq \dots \geq \theta_K^* > \theta_{K+1}^* \geq \dots \geq \theta_n^*$ by assumption. It suffices to show $\hat{\theta}_i - \hat{\theta}_j > 0$ for any $i \leq K, j > K$. Let $i \leq K, j > K$,

$$\hat{\theta}_i - \hat{\theta}_j \geq (\theta_i^* - \theta_j^*) - \left| (\hat{\theta}_i - \hat{\theta}_j) - (\theta_i^* - \theta_j^*) \right| \geq \Delta_K - B_{ij}.$$

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

$$\max_{k,l} B_{kl} \leq \Delta_K/2.$$

Now the proof of Lemma 2 is completed.

B.1 Proof of Lemma 15

We will prove this lemma by induction on t . Since preconditioned gradient descent starts at ground truth, the base case of $t = 0$ is trivial. We will now study the dynamics of (16) to prove the induction step. Using the definition (2) we can compute the gradient and Hessian of the loss function:

$$\begin{aligned}\nabla\mathcal{L}(\boldsymbol{\theta}) &= \frac{1}{L} \sum_{i,j:i>j} \sum_{l=1}^L w_{ij} \left\{ -y_{ji}^{(l)} + \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} \right\} (\mathbf{e}_i - \mathbf{e}_j); \\ \nabla^2\mathcal{L}(\boldsymbol{\theta}) &= \sum_{i,j:i>j} w_{ij} \frac{e^{\theta_i} e^{\theta_j}}{(e^{\theta_i} + e^{\theta_j})^2} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top.\end{aligned}$$

Applying Taylor's expansion on $\nabla\mathcal{L}(\boldsymbol{\theta}^t)$, we have

$$\begin{aligned}\nabla\mathcal{L}(\boldsymbol{\theta}^t) &= \frac{1}{L} \sum_{i,j:i>j} \sum_{l=1}^L w_{ij} \left((-y_{ji}^{(l)} + \sigma(\theta_i^t - \theta_j^t)) (\mathbf{e}_i - \mathbf{e}_j) \right) \\ &= \frac{1}{L} \sum_{i,j:i>j} \sum_{l=1}^L w_{ij} \left((-\epsilon_{ji}^{(l)} - \sigma(\theta_i^* - \theta_j^*) + \sigma(\theta_i^t - \theta_j^t)) (\mathbf{e}_i - \mathbf{e}_j) \right) \\ &= \frac{1}{L} \sum_{i,j:i>j} \sum_{l=1}^L w_{ij} \left(\left(-\epsilon_{ji}^{(l)} + \sigma'(\theta_i^* - \theta_j^*)(\delta_i^t - \delta_j^t) + \frac{1}{2}\sigma''(\zeta_{ij}^t)(\delta_i^t - \delta_j^t)^2 \right) (\mathbf{e}_i - \mathbf{e}_j) \right).\end{aligned}$$

Here for all $i, j \in [n]$, $\delta_i^t := \theta_i^t - \theta_i^*$ and $\zeta_{ij}^t \in \mathbb{R}$ is some number that lies between $\theta_i^* - \theta_j^*$ and $\theta_i^t - \theta_j^t$. As $\sigma'(\theta_i^* - \theta_j^*) = z_{ij}$ and $\delta_i^t - \delta_j^t = (\mathbf{e}_i - \mathbf{e}_j)^\top \boldsymbol{\delta}$, we can rewrite the above formula as

$$\nabla\mathcal{L}(\boldsymbol{\theta}^t) = \frac{1}{L} (L \cdot \mathbf{L}_{wz} \boldsymbol{\delta}^t - \mathbf{B} \hat{\boldsymbol{\epsilon}} + L \cdot \mathbf{r}^t),$$

where $\mathbf{r}^t := \sum_{i,j:i>j} w_{ij} [\frac{1}{2}\sigma''(\zeta_{ij}^t)(\delta_i^t - \delta_j^t)^2 (\mathbf{e}_i - \mathbf{e}_j)]$. Putting this in (16), we have

$$\boldsymbol{\delta}^{t+1} = (1 - \eta) \boldsymbol{\delta}^t - \frac{\eta}{L} (\mathbf{L}_{wz}^\dagger \mathbf{B} \hat{\boldsymbol{\epsilon}} - L \cdot \mathbf{L}_{wz}^\dagger \mathbf{r}^t). \quad (18)$$

Now consider $\delta_k^t - \delta_l^t = (\mathbf{e}_k - \mathbf{e}_l)^\top \boldsymbol{\delta}^t$, we have

$$\delta_k^{t+1} - \delta_l^{t+1} = (1 - \eta) (\delta_k^t - \delta_l^t) - \frac{\eta}{L} (\mathbf{e}_k - \mathbf{e}_l)^\top (\mathbf{L}_{wz}^\dagger \mathbf{B} \hat{\boldsymbol{\epsilon}} - L \cdot \mathbf{L}_{wz}^\dagger \mathbf{r}^t). \quad (19)$$

We now control the size of $(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{B} \hat{\boldsymbol{\epsilon}}$ and $(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{r}^t$ separately.

Controlling $(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{B} \hat{\boldsymbol{\epsilon}}$. We rewrite this term with

$$(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{B} \hat{\boldsymbol{\epsilon}} = \langle \mathbf{B}^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l), \hat{\boldsymbol{\epsilon}} \rangle.$$

By Lemma 9, $z_{ij} \geq 1/(4\kappa)$ so each entry of $\hat{\boldsymbol{\epsilon}}$ is of sub-Gaussian norm $\sqrt{w_{ij}/z_{ij}} \leq \sqrt{4\kappa \max_{ij} w_{ij}}$. Then $\langle \mathbf{B}^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l)^\top, \hat{\boldsymbol{\epsilon}} \rangle$ has sub-Gaussian norm

$$\|\mathbf{B}^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l)\|_2 \sqrt{4\kappa \max_{ij} w_{ij}} \leq \sqrt{4\kappa L \max_{ij} w_{ij} \Omega_{kl}(\mathbf{L}_{wz})},$$

where the inequality comes from

$$\begin{aligned}\|\mathbf{B}^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l)^\top\|_2^2 &= (\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{B} \mathbf{B}^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l) \\ &= L \cdot (\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_k - \mathbf{e}_l) \\ &= L \cdot \Omega_{kl}(\mathbf{L}_{wz}).\end{aligned}$$

The second line follows from (15). Taking a union bound on concentration of sub-Gaussian random variables (see for instance Section 2.5 in [Ver18]), we have that with probability at least $1 - n^{-10}$, for all $k, l \in [n]$,

$$|(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{B} \hat{\boldsymbol{\epsilon}}| \leq C \sqrt{\kappa L \max_{i,j} w_{ij} \cdot \boldsymbol{\Omega}_{kl}(\mathbf{L}_{wz}) \log(n)} \leq \frac{1}{2} L \cdot B_{kl} \quad (20)$$

for some constant C .

Controlling $(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{r}^t$. We expand the term to get

$$\begin{aligned} |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{r}^t| &= \left| \sum_{i,j:i>j} w_{ij} \left[\frac{1}{2} \sigma''(\zeta_{ij}^t) (\delta_i^t - \delta_j^t)^2 (\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_i - \mathbf{e}_j) \right] \right| \\ &\leq \frac{1}{8} \sum_{i,j:i>j} w_{ij} (\delta_i^t - \delta_j^t)^2 |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_i - \mathbf{e}_j)|. \end{aligned}$$

For the inequality here we use the fact that $\sigma''(\zeta) \leq 1/4$ for any $\zeta \in \mathbb{R}$.

Now show the induction step. Assume (17) holds for iteration t , then

$$\begin{aligned} |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{r}^t| &\leq \frac{1}{8} \sum_{i,j:i>j} w_{ij} (\delta_i^t - \delta_j^t)^2 |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_i - \mathbf{e}_j)| \\ &\leq \frac{1}{8} \sum_{i,j:i>j} w_{ij} B_{ij}^2 |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger (\mathbf{e}_i - \mathbf{e}_j)| \\ &= \frac{1}{8} Q_{kl}. \end{aligned}$$

Substituting this and (20) into (19),

$$\begin{aligned} |\delta_k^{t+1} - \delta_l^{t+1}| &\leq (1 - \eta) |\delta_k^t - \delta_l^t| + \frac{\eta}{L} (|(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{r}^t| + L \cdot |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_{wz}^\dagger \mathbf{B} \hat{\boldsymbol{\epsilon}}|) \\ &\leq (1 - \eta) B_{kl} + \frac{\eta}{L} \left(\frac{1}{2} L \cdot B_{kl} + \frac{1}{8} L \cdot Q_{kl} \right) \\ &\leq B_{kl} \end{aligned}$$

as long as $Q_{kl} \leq 4B_{kl}$.

B.2 Proof of Lemma 16

The proof of this follows a similar strategy to Lemma 1 in [Che23]. We will only prove the first part here, i.e. the existence and uniqueness of a solution for (2), as the proof for the second part is the same as [Che23]. We first make the following claim that we will prove at the end of this subsection. This claim is analogous to a classical result in [FJ57], which states the same thing but with unweighted MLE.

Claim 1. Instate the assumptions of Lemma 16. Suppose that for any disjoint partition $\mathcal{S}_1 \cup \mathcal{S}_2 = [n]$, there exists some $i \in \mathcal{S}_1, j \in \mathcal{S}_2$ such that $w_{ij} > 0$ and $y_{ji}^{(l)} = 1$ for some $1 \leq l \leq L$. Then there exists a minimizer $\hat{\boldsymbol{\theta}}$ of (2) that satisfies $\mathbf{1}_n^\top \hat{\boldsymbol{\theta}} = 0$. Furthermore it is unique in the sense that no other minimizer $\boldsymbol{\theta}$ satisfies $\mathbf{1}_n^\top \boldsymbol{\theta} = 0$.

Now it suffices to show that the condition of this claim is true. Suppose for the sake of contradiction that it is false. That is there exists some disjoint partition $\mathcal{S}_1 \cup \mathcal{S}_2 = [n]$, such that for all $i \in \mathcal{S}_1, j \in \mathcal{S}_2$ with $w_{ij} > 0$, $y_{ji} = (1/L) \sum_{l=1}^L y_{ji}^{(l)} = 0$. By Lemma 15 and the fact that $\mathbf{1}_n^\top \boldsymbol{\theta}^t = 0$ for any $t \geq 0$, $\boldsymbol{\theta}^t \in B_R := \{\boldsymbol{\theta} : \|\boldsymbol{\theta}\|_\infty \leq R\}$ for some scalar R that depends on the observations. Now we consider the minimum size of the gradient in this close ball. Recall

$$\nabla \mathcal{L}(\boldsymbol{\theta}) = \sum_{i,j:i>j} w_{ij} \left\{ -y_{ji} + \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} \right\} (\mathbf{e}_i - \mathbf{e}_j).$$

Consider a vector $\mathbf{v} = \mathbf{1}_{\mathcal{S}_1}$ that is 1 on all entries for \mathcal{S}_1 and 0 otherwise. It is easy to see that $(\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{v} = 0$ if $i, j \in \mathcal{S}_1$ or $i, j \in \mathcal{S}_2$; $(\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{v} = 1$ if $i \in \mathcal{S}_1, j \in \mathcal{S}_2$. Then for any $\boldsymbol{\theta} \in B_R$, one can rearrange the summation to reach

$$\begin{aligned} \nabla \mathcal{L}(\boldsymbol{\theta})^\top \mathbf{v} &= \sum_{i \in \mathcal{S}_1, j \in \mathcal{S}_2} w_{ij} \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} (\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{v} \\ &= \sum_{i \in \mathcal{S}_1, j \in \mathcal{S}_2} w_{ij} \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j}} \\ &\stackrel{(i)}{\geq} \sum_{i \in \mathcal{S}_1, j \in \mathcal{S}_2} w_{ij} \frac{e^{-R}}{e^{-R} + e^R} \\ &\stackrel{(ii)}{\geq} \frac{1}{2} e^{-2R} \min\{w_{ij} : w_{ij} > 0\}. \end{aligned}$$

Here (i) holds since $\boldsymbol{\theta} \in B_R$ and the function $(x, y) \mapsto e^x / (e^x + e^y)$ is increasing in x and decreasing in y ; (ii) holds since we assume the weighted graph to be connected by edges with non-zero weights. Taking infimum over $\boldsymbol{\theta}$ in B_R , we have $\inf_{\boldsymbol{\theta} \in B_R} \|\nabla \mathcal{L}(\boldsymbol{\theta})\| > 0$.

It now suffices to show $\|\nabla \mathcal{L}(\boldsymbol{\theta}^t)\| \rightarrow 0$ as $t \rightarrow \infty$ which contradicts the claim that $\inf_{\boldsymbol{\theta} \in B_R} \|\nabla \mathcal{L}(\boldsymbol{\theta})\| > 0$. Recall \mathbf{L}_w is the Laplacian matrix of the weighted graph defined as $\sum_{(i,j):i>j} w_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top$. Since the graph is connected by edges with non-zero weights, $\lambda_{n-1}(\mathbf{L}_w) > 0$. Consider the Hessian

$$\nabla^2 \mathcal{L}(\boldsymbol{\theta}) = \sum_{(i,j):i>j} w_{ij} \frac{e^{\theta_i} e^{\theta_j}}{(e^{\theta_i} + e^{\theta_j})^2} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top.$$

For any unit vector $\mathbf{v} \in \mathbb{R}^n$ such that $\mathbf{v}^\top \mathbf{1}_n = 0$,

$$\begin{aligned} \mathbf{v}^\top \nabla^2 \mathcal{L}(\boldsymbol{\theta}) \mathbf{v} &= \sum_{(i,j):i>j} w_{ij} \frac{e^{\theta_i} e^{\theta_j}}{(e^{\theta_i} + e^{\theta_j})^2} \mathbf{v} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{v} \\ &\geq \sum_{(i,j):i>j} w_{ij} \frac{e^{-R} e^R}{(e^{-R} + e^R)^2} \mathbf{v} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{v} \\ &\geq \frac{1}{4e^{2R}} \lambda_{n-1}(\mathbf{L}_w). \end{aligned}$$

Here the last inequality follows from the definition of $\lambda_{n-1}(\mathbf{L}_w)$. We also have that

$$\lambda_{n-1}(\mathbf{L}_{wz}^\dagger) = \|\mathbf{L}_{wz}\|^{-1} > 0.$$

Now consider the precondition gradient descent update

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

By Taylor expansion we have

$$\begin{aligned} \nabla \mathcal{L}(\boldsymbol{\theta}^{t+1}) &= \nabla \mathcal{L}(\boldsymbol{\theta}^t) + \nabla^2 \mathcal{L}(\boldsymbol{\xi})(\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t) \\ &= \nabla \mathcal{L}(\boldsymbol{\theta}^t) - \eta \nabla^2 \mathcal{L}(\boldsymbol{\xi}) \mathbf{L}_{wz}^\dagger \nabla \mathcal{L}(\boldsymbol{\theta}^t) \\ &= (\mathbf{I} - \eta \nabla^2 \mathcal{L}(\boldsymbol{\xi}) \mathbf{L}_{wz}^\dagger) \nabla \mathcal{L}(\boldsymbol{\theta}^t) \end{aligned}$$

for some $\boldsymbol{\xi} \in B_R$. Since $\mathbf{1}_n^\top \nabla \mathcal{L}(\boldsymbol{\theta}^t) = 0$ and $\mathbf{1}_n^\top \mathbf{L}_{wz}^\dagger \nabla \mathcal{L}(\boldsymbol{\theta}^t) = 0$,

$$\|\nabla \mathcal{L}(\boldsymbol{\theta}^{t+1})\| \leq \left[1 - \frac{\eta}{4e^{2R}} \lambda_{n-1}(\mathbf{L}_w) \lambda_{n-1}(\mathbf{L}_{wz}^\dagger)\right] \|\nabla \mathcal{L}(\boldsymbol{\theta}^t)\|.$$

Then for sufficiently small $\eta > 0$, $\|\nabla \mathcal{L}(\boldsymbol{\theta}^t)\|$ goes to 0 as $t \rightarrow \infty$.

Proof of Claim 1. Recall

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{L} \sum_{i,j:i>j} \sum_{l=1}^L w_{ij} \left(-y_{ji}^{(l)} (\theta_i - \theta_j) + \log(1 + e^{\theta_i - \theta_j}) \right).$$

Observe that since $0 \leq y_{ji}^{(l)} \leq 1$,

$$-y_{ji}^{(l)} (\theta_i - \theta_j) + \log(1 + e^{\theta_i - \theta_j}) \geq 0 \quad (21)$$

for each i, j . Since $|\theta_i^* - \theta_j^*| \leq \log \kappa$ for any $(i, j) \in [n]^2$,

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}^*) &= \frac{1}{L} \sum_{i,j:i>j} \sum_{l=1}^L w_{ij} \left(-y_{ji}^{(l)} (\theta_i^* - \theta_j^*) + \log(1 + e^{\theta_i^* - \theta_j^*}) \right) \\ &\leq \sum_{i,j:i>j} w_{ij} (\log(\kappa) + \log(1 + \kappa)) \\ &\leq \sum_{i,j:i>j} w_{ij} 2 \log(2\kappa). \end{aligned} \quad (22)$$

Let R be a large enough scalar such that

$$\frac{R}{nL} \min_{\{i,j:w_{ij}>0\}} w_{ij} > \sum_{\{i,j:i>j\}} w_{ij} 2 \log(2\kappa). \quad (23)$$

Now consider any $\boldsymbol{\theta}$ with $\|\boldsymbol{\theta}\|_\infty > R$. Assume without loss of generality $\theta_1 \geq \theta_2 \geq \dots \geq \theta_n$. Since we only consider the case when $\boldsymbol{\theta}^\top \mathbf{1}_n = 0$, $\theta_1 - \theta_n > R - 0$ and there exists some $k \leq n - 1$ such that $\theta_k - \theta_{k+1} > R/n$. We then have a natural partition by two non-empty vertex set $\mathcal{S}_1 = \{i : i > k\}, \mathcal{S}_2 = \{j : j \leq k\}$. By assumption for some $i > k, j \leq k$, $w_{ij} > 0$ and there is some l such that $y_{ji}^{(l)} = 1$. Then combine this with (21),

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}) &\geq \frac{1}{L} w_{ij} \left(-y_{ji}^{(l)} (\theta_i - \theta_j) + \log(1 + e^{\theta_i - \theta_j}) \right) \\ &\geq \frac{1}{L} w_{ij} (R/n + \log(1)) \\ &\geq \frac{R}{nL} \min_{i,j:w_{ij}>0} w_{ij}. \end{aligned} \quad (24)$$

Putting (22), (23), and (24) together, we see that $\mathcal{L}(\boldsymbol{\theta}^*) \leq \mathcal{L}(\boldsymbol{\theta})$ for any $\boldsymbol{\theta}$ such that $\boldsymbol{\theta}^\top \mathbf{1}_n = 0$ and $\|\boldsymbol{\theta}\|_\infty > R$. Then by the continuity of \mathcal{L} there must exist a minimizer in the closed and bounded set $\{\boldsymbol{\theta} : \|\boldsymbol{\theta}\|_\infty \leq R : \boldsymbol{\theta}^\top \mathbf{1}_n = 0\}$. The restricted uniqueness is guaranteed by the restricted strong convexity shown earlier, that is for any unit vector $\mathbf{v} \in \mathbb{R}^n$ such that $\mathbf{v}^\top \mathbf{1}_n = 0$ and $\|\boldsymbol{\theta}\|_\infty \leq R$,

$$\mathbf{v}^\top \nabla^2 \mathcal{L}(\boldsymbol{\theta}) \mathbf{v} \geq \frac{1}{4e^{2R}} \lambda_{n-1}(\mathbf{L}_w) > 0.$$

C Proof of Lemma 5

This proof is mostly based on the proof of Lemma 28 in [KLOS14], which considers the case with integer weights. Let $\mathbf{v} = \mathbf{L}_{\mathcal{G}_{\tilde{w}}}^\dagger (\mathbf{e}_k - \mathbf{e}_l)$, where $\mathbf{L}_{\mathcal{G}_{\tilde{w}}}$ is the \tilde{w}_{ij} -weighted graph Laplacian. Then we can rewrite the LHS of (12) with

$$\sum_{(i,j):i>j} w_{ij} \left| (\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{L}_{\mathcal{G}_{\tilde{w}}}^\dagger (\mathbf{e}_k - \mathbf{e}_l) \right| = \sum_{(i,j):i>j} w_{ij} |v_i - v_j|.$$

For the rest of the proof we write (i, j) as edges so (i, j) and (j, i) are the same element. For any $c \in \mathbb{R}$ we define two vertex sets:

$$\mathcal{S}_c^> := \{i \in \mathcal{V} : v_i > c\} \quad \text{and} \quad \mathcal{S}_c^< := \{i \in \mathcal{V} : v_i < c\}.$$

Recall that the volume of a vertex set \mathcal{S} is defined as

$$\text{vol}(\mathcal{S}) := \sum_{(i,j) \in \mathcal{E}: i \in \mathcal{S}} w_{ij}.$$

Let

$$\bar{c} := \sup_c \left\{ c : \text{vol}(\mathcal{S}_c^<) \leq \frac{1}{2} \text{vol}(\mathcal{V}) \right\}$$

and

$$\underline{c} := \inf_c \left\{ c : \text{vol}(\mathcal{S}_c^>) \leq \frac{1}{2} \text{vol}(\mathcal{V}) \right\}.$$

For any $\epsilon > 0$,

$$\text{vol}(\mathcal{S}_{\bar{c}+\epsilon}^>) \leq \text{vol}(\mathcal{V}) - \text{vol}(\mathcal{S}_{\bar{c}+\epsilon}^<) \leq \text{vol}(\mathcal{V})/2.$$

By the definition of \underline{c} , $\bar{c} + \epsilon \geq \underline{c}$. This holds for every $\epsilon > 0$ so $\bar{c} \geq \underline{c}$. Fix some $c^* \in [\underline{c}, \bar{c}]$, we have that

$$\text{vol}(\mathcal{S}_{c^*}^>) \leq 1/2 \quad \text{and} \quad \text{vol}(\mathcal{S}_{c^*}^<) \leq 1/2. \quad (25)$$

For any vertex set $\mathcal{S} \subset \mathcal{V}$, we denote $\partial\mathcal{S} := \{(i, j) : i \in \mathcal{S}, v \in \mathcal{V} \setminus \mathcal{S}\}$. We also define the flow and weight of the graph cut corresponding to \mathcal{S} as

$$f(\mathcal{S}) := \sum_{(i,j) \in \partial\mathcal{S}} w_{ij} |v_i - v_j| \quad \text{and} \quad w(\mathcal{S}) := \sum_{(i,j) \in \partial\mathcal{S}} w_{ij}.$$

Abusing the notation, for any set of edges \mathcal{F} we write $w(\mathcal{F}) := \sum_{(i,j) \in \mathcal{F}} w_{ij}$. Now consider a sequence of real numbers $\{c_t\}$ defined recursively by $c_0 := c^*$ and for any $t \geq 0$,

$$c_{t+1} := c_t + \Delta_t \quad \text{for} \quad \Delta_t := \frac{2f(\mathcal{C}_t)}{w(\mathcal{C}_t)},$$

where to simplify the notation we also write $\mathcal{C}_t := \mathcal{S}_{c_t}^>$. At this point we will use the following fact.

Fact 1. For any vertex subset $\mathcal{S} \subset \mathcal{V}$,

$$\sum_{(i,j) \in \partial\mathcal{S}} \tilde{w}_{ij} (v_i - v_j) \leq 1$$

In a graph theoretical language, this fact says that the total flow going across any cut is at most the amount of flow going from source to sink. The exact proof is omitted as it requires the introduction of a number of notions that are irrelevant to the rest of this paper. Please see [KLOS14] for the details. For any $t \geq 0$, from the definition of \mathcal{C}_t we know that for any $i \in \mathcal{C}_t$ and $j \in \mathcal{V} \setminus \mathcal{C}_t$, $v_i \geq v_j$. Then using the above fact we have

$$\begin{aligned} f(\mathcal{C}_t) &= \sum_{(i,j): i \in \mathcal{C}_t, v \in \mathcal{V} \setminus \mathcal{C}_t} w_{ij} |v_i - v_j| \\ &= \sum_{(i,j): i \in \mathcal{C}_t, v \in \mathcal{V} \setminus \mathcal{C}_t} w_{ij} (v_i - v_j) \\ &\leq Z \sum_{(i,j): i \in \mathcal{C}_t, v \in \mathcal{V} \setminus \mathcal{C}_t} \tilde{w}_{ij} (v_i - v_j) \leq Z. \end{aligned}$$

Here Z is a normalized factor defined as $Z := \max_{ij} \{w_{ij}/\tilde{w}_{ij}\}$. We now show that $\text{vol}(\mathcal{C}_t)$ exponentially decreases with t . Since $\mathcal{C}_{t+1} \subset \mathcal{C}_t$ for any $t \geq 0$,

$$\begin{aligned} \text{vol}(\mathcal{C}_{t+1}) &\leq \text{vol}(\mathcal{C}_t) - w(\partial\mathcal{C}_t \setminus \partial\mathcal{C}_{t+1}) \\ &= \text{vol}(\mathcal{C}_t) - w(\partial\mathcal{C}_t) + w(\partial\mathcal{C}_t \cap \partial\mathcal{C}_{t+1}). \end{aligned}$$

For any $(i, j) \in \partial\mathcal{C}_t \cap \partial\mathcal{C}_{t+1}$, $i \in \mathcal{C}_{t+1}$ and $j \in \mathcal{V} \setminus \mathcal{C}_t$ (or the other way around). Then by the choice of Δ_t ,

$$\begin{aligned}
f(\mathcal{C}_t) &= \sum_{(i,j) \in \partial\mathcal{C}_t} w_{ij} |v_i - v_j| \\
&\geq \sum_{(i,j) \in \partial\mathcal{C}_t \cap \partial\mathcal{C}_{t+1}} w_{ij} |v_i - v_j| \\
&\geq \sum_{(i,j) \in \partial\mathcal{C}_t \cap \partial\mathcal{C}_{t+1}} w_{ij} |c_{t+1} - c_t| \\
&= w(\partial\mathcal{C}_t \cap \partial\mathcal{C}_{t+1}) \cdot \Delta_t \\
&= w(\partial\mathcal{C}_t \cap \partial\mathcal{C}_{t+1}) \cdot \frac{2f(\mathcal{C}_t)}{w(\mathcal{C}_t)}.
\end{aligned}$$

Therefore $w(\partial\mathcal{C}_t \cap \partial\mathcal{C}_{t+1}) \leq w(\mathcal{C}_t)/2$ and

$$\begin{aligned}
\text{vol}(\mathcal{C}_{t+1}) &\leq \text{vol}(\mathcal{C}_t) - w(\partial\mathcal{C}_t) + w(\partial\mathcal{C}_t \cap \partial\mathcal{C}_{t+1}) \\
&\leq \text{vol}(\mathcal{C}_t) - \frac{w(\mathcal{C}_t)}{2} \\
&\leq \text{vol}(\mathcal{C}_t) - \frac{1}{2} \text{vol}(\mathcal{C}_t) \Phi(\mathcal{G}_w),
\end{aligned}$$

where the last inequality follow from the definition of the conductance $\Phi(\mathcal{G}_w)$. Now applying this recursively and use (25),

$$\text{vol}(\mathcal{C}_t) \leq \left(1 - \frac{1}{2} \Phi(\mathcal{G}_w)\right)^t \text{vol}(\mathcal{C}_0) \leq \frac{1}{2} \left(1 - \frac{1}{2} \Phi(\mathcal{G}_w)\right)^t \text{vol}(\mathcal{V}).$$

As $\sum_{i \in \mathcal{S}, j \in \mathcal{V} \setminus \mathcal{S}} w_{ij} \leq \text{vol}(\mathcal{S})$ for any $\mathcal{S} \subset \mathcal{V}$,

$$\Phi(\mathcal{G}_w) := \min_{\mathcal{S} \subset \mathcal{V}} \frac{\sum_{i \in \mathcal{S}, j \in \mathcal{V} \setminus \mathcal{S}} w_{ij}}{\min\{\text{vol}(\mathcal{S}), \text{vol}(\mathcal{V} \setminus \mathcal{S})\}} \leq 1.$$

So $\text{vol}(\mathcal{C}_t)$ is decreasing in t and for any $t > 2 \log(\text{vol}(\mathcal{V}))/\Phi(\mathcal{G}_w)$,

$$\begin{aligned}
\text{vol}(\mathcal{C}_t) &\leq \frac{1}{2} \left(1 - \frac{1}{2} \Phi(\mathcal{G}_w)\right)^t \text{vol}(\mathcal{V}) \\
&\leq \frac{1}{2} e^{-\frac{1}{2} t \Phi(\mathcal{G}_w)} \text{vol}(\mathcal{V}) \\
&< \frac{1}{2}.
\end{aligned}$$

By assumption, for any vertex $i \in \mathcal{V}$, $\text{vol}(\{i\}) \geq d_{\min} \geq 1$, so $\text{vol}(\mathcal{C}_t) < 1/2$ implies $\text{vol}(\mathcal{C}_t) = \emptyset$. Now let r be the smallest integer such that $\text{vol}(\mathcal{C}_{r+1}) = \emptyset$, we have $r \leq 2 \log(\text{vol}(\mathcal{V}))/\Phi(\mathcal{G}_w)$. Let $d_i := \sum_{j:j \neq i} w_{ij}$. Then

$$\begin{aligned}
\sum_{i \in \mathcal{S}_{c^*}^>} d_i |v_i - c^*| &= \sum_{t=0}^r \sum_{i \in \mathcal{C}_t \setminus \mathcal{C}_{t+1}} d_i (v_i - c^*) \\
&\stackrel{(i)}{\leq} \sum_{t=0}^r \sum_{i \in \mathcal{C}_t \setminus \mathcal{C}_{t+1}} d_i (c_{t+1} - c^*) \\
&\stackrel{(ii)}{\leq} \sum_{t=0}^r (\text{vol}(\mathcal{C}_t) - \text{vol}(\mathcal{C}_{t+1})) \sum_{s=0}^t \Delta_s \\
&\stackrel{(iii)}{\leq} \sum_{t=0}^r \text{vol}(\mathcal{C}_t) \Delta_t.
\end{aligned}$$

Here (i) comes from the definition of \mathcal{C}_{t+1} , (ii) comes from the definition of d_i and choice of c_{t+1} , and (iii) comes from rearrangement. Now we control the term $\text{vol}(\mathcal{C}_t)\Delta_t$ for any $t \geq 0$. From the definition of $\Phi(\mathcal{G}_w)$ we have that

$$\Phi(\mathcal{G}_w) \geq \frac{w(\mathcal{C}_t)}{\min\{\text{vol}(\mathcal{C}_t), \text{vol}(\mathcal{V}) - \text{vol}(\mathcal{C}_t)\}} \geq \frac{w(\mathcal{C}_t)}{\text{vol}(\mathcal{C}_t)}$$

where the last inequality follows from (25). Moreover we established earlier in this proof that $f(\mathcal{C}_t) \leq Z$, then

$$\text{vol}(\mathcal{C}_t)\Delta_t \leq \frac{w(\mathcal{C}_t)}{\Phi(\mathcal{G}_w)} \cdot \frac{2Z}{w(\mathcal{C}_t)} = \frac{2Z}{\Phi(\mathcal{G}_w)}.$$

Therefore

$$\sum_{i \in \mathcal{S}_{c^*}^>} d_i |v_i - c^*| \leq \frac{2Zr}{\Phi(\mathcal{G}_w)}.$$

Similarly we can achieve

$$\sum_{i \in \mathcal{S}_{c^*}^<} d_i |v_i - c^*| \leq \frac{2Zr}{\Phi(\mathcal{G}_w)}.$$

Finally,

$$\begin{aligned} \sum_{(i,j):i>j} w_{ij} |v_i - v_j| &\leq \sum_{(i,j):i>j} w_{ij} (|v_i - c^*| + |v_j - c^*|) \\ &\leq \sum_{i \in \mathcal{V}} d_i |v_i - c^*| \\ &\leq \sum_{i \in \mathcal{S}_{c^*}^>} d_i |v_i - c^*| + \sum_{i \in \mathcal{S}_{c^*}^<} d_i |v_i - c^*| + \sum_{i:v_i=c^*} d_i |v_i - c^*| \\ &\leq \frac{4Zr}{\Phi(\mathcal{G}_w)} \leq \frac{8Z \log(\text{vol}(\mathcal{V}))}{\Phi(\mathcal{G}_w)^2}. \end{aligned}$$

This finishes the proof.

D Proofs for SDP-based reweighting

D.1 Proofs for oracles in Section 5

Proof of Theorem 4. Let $\mathbf{c} \in \mathbb{R}^{\mathcal{E}_{\text{SR}}}$ be defined by $c_{ij} = \langle \mathbf{L}_{ij}, \mathbf{X} \rangle$. Then, we are interested in approximately solving the linear program (LP) $\min_{\mathbf{w} \in \mathcal{F}} \langle \mathbf{c}, \mathbf{w} \rangle$. Because $\mathbf{c} \geq 0$ and the set \mathcal{F} consists of the intersection of the positive quadrant with entrywise non-negative linear upper bounds, this LP is an instance of a packing LP. This class of programs can be approximately solved by specialized first-order algorithms. In particular, the solver of [AZO18] yields a $(1 - \epsilon)$ multiplicative approximation in time $\tilde{O}(\text{nnz} + n/\epsilon)$, where nnz is the number of non-zero entries in the matrix defining the constraints. By the construction of \mathcal{F} , the sparsity of the corresponding constraints is simply $|\mathcal{E}_{\text{SR}}|$. This complete the proof of the theorem. \square

Proof of Theorem 5. Let $\mathbf{c} \in \mathbb{R}^{\mathcal{E}_{\text{SR}}}$ be defined by $c_{ij} := \langle \mathbf{L}_{ij}, \mathbf{X} \rangle$. Let $d := 2pn$. Without loss of generality, we can assume that d is an integer. Then, we are interested in approximately solving the following linear program (LP):

$$\begin{aligned} &\max \langle \mathbf{c}, \mathbf{w} \rangle \\ &\forall i \in V, \quad \sum_{ij \in \mathcal{E}_{\text{SR}}} w_{ij} \leq d, \\ &\forall ij \in \mathcal{E}_{\text{SR}}, \quad \mathbf{0} \leq w_{ij} \leq 1. \end{aligned}$$

Notice that this is a relaxation of the maximum weight d -matching problem with weight \mathbf{c} over the graph \mathcal{G}_{SR} . We are going to exploit this connection by showing that a greedy maximal-weight d -matching yields

a $1/2$ -approximation to the optimum of this LP. To bound the value of this optimum, we will rely on the following dual LP:

$$\begin{aligned} \min & d \cdot \sum_{i \in V} s_i + \sum_{ij \in \mathcal{E}_{\text{SR}}} \ell_{ij} \\ \forall ij \in \mathcal{E}_{\text{SR}}, & s_i + s_j + \ell_{ij} \geq c_{ij}, \\ \forall ij \in \mathcal{E}_{\text{SR}}, & \ell_{ij} \geq 0, \\ \forall i \in V, & s_i \geq 0 \end{aligned}$$

The greedy d -matching is constructed in the following natural way. First all edges are sorted in decreasing order of weight \mathbf{c} . Then, edges are added to the matching in this order as long as their addition does not cause a degree constraint to become violated. Let M be the resulting matching. We are now going to construct a feasible dual solution based on M . For each vertex $i \in V$, let s_i be the weight of the last matching edge incident to i that was added to M . Notice that some s_i may be 0 if all edges were exhausted before the degree constraint was violated. For any edge $ij \in M$, let

$$\ell_{ij} = \max\{c_{ij} - s_i - s_j, 0\}. \quad (26)$$

If $ij \notin M$, simply let $\ell_{ij} = 0$.

We will now argue that the resulting dual solution is feasible. It suffices to show that all constraints $s_i + s_j + \ell_{ij} \geq c_{ij}$ are satisfied. If $ij \in M$, this is an immediate consequence of Equation 26. If $ij \notin M$, the addition of ij must cause either i or j to violate the degree constraint. Suppose wlog that this was the case for vertex i . Then, it must be the case that $c_{ij} \leq s_i$, as edge ij was considered after the last matching edge was added to i . This proves that our dual solution is feasible.

We now compare the value of the dual solution to that of the greedy matching. By construction, for each vertex i , we have:

$$\begin{aligned} \sum_{j:ij \in M} c_{ij} &= \sum_{j:ij \in M \wedge \ell_{ij}=0} c_{ij} + \sum_{j:ij \in M \wedge \ell_{ij}>0} c_{ij} \geq \\ \sum_{j:ij \in M \wedge \ell_{ij}=0} s_i &+ \sum_{j:ij \in M \wedge \ell_{ij}>0} s_i + \ell_{ij}/2 = d \cdot s_i + \frac{1}{2} \sum_{j:ij \in M} \ell_{ij}, \end{aligned}$$

where the last equality uses the fact that $s_i = 0$ if the degree of i in M is less than d . Summing over all vertices yields:

$$2 \cdot \sum_{ij \in M} c_{ij} = \sum_{i \in V} \sum_{j:ij \in M} c_{ij} = d \cdot \sum_{i \in V} s_i + \sum_{ij \in M} \ell_{ij}.$$

As the right hand side is the dual value of our solution, we have shown that the greedy matching achieves a primal value that is within a factor of two of the optimal, as required. \square

D.2 Proofs for the analysis of Algorithms 2

In this section, we verify that our use of Johnson–Lindenstrauss preserves the ℓ_2^2 -distances up to an ϵ -multiplicative factor by proving Lemma 6. We also demonstrate Lemma 8 thereby justifying that Algorithm 2 runs in time nearly-linear with respect to the size of \mathcal{G}_{SR} .

Beginning with Lemma 6, we first recall the statement of Johnson–Lindenstrauss provided by [Ach03]. The following can be recovered by setting $\beta = 8$, and observing that $\epsilon^2/2 - \epsilon^2/3 \geq \epsilon^2/6$ for all $\epsilon \in [0, 1]$.

Theorem 7 (Theorem 1.1 in [Ach03]). *Suppose $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\epsilon \in [0, 1]$ are given. For any $k \geq 120 \cdot \log n / \epsilon^2$, let $\mathbf{R} = \{\pm 1/\sqrt{k}\}^{n \times k}$ be a random matrix with entries sampled independently and uniformly at random, and set $\mathbf{V} = \mathbf{AR}$. If $\mathbf{a}_i \in \mathbb{R}^d$ and $\mathbf{v}_i \in \mathbb{R}^k$ denote the i -th row of \mathbf{A} and \mathbf{V} respectively, then*

$$(1 - \epsilon) \cdot \|\mathbf{a}_i - \mathbf{a}_j\|^2 \leq \|\mathbf{v}_i - \mathbf{v}_j\|^2 \leq (1 + \epsilon) \cdot \|\mathbf{v}_i - \mathbf{v}_j\|^2$$

We will also require a result which approximately computes the action of a matrix exponential where the matrix is Symmetric and Diagonally Dominant (SDD). Recall that a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is SDD if it is both symmetric, and for all $i \in [n]$, the entries satisfy $A_{ii} \geq \sum_{j \neq i} |A_{ij}|$. In [OSV12], they apply the Lanczos method to obtain the following guarantee.

Theorem 8 (Theorem 1.2 in [OSV12]). *Given an SDD matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, a vector $\mathbf{v} \in \mathbb{R}^n$, and $\delta > 0$, there exists an algorithm \mathcal{A} which outputs a vector $\mathbf{u} \in \mathbb{R}^n$ satisfying*

$$\|\exp\{-\mathbf{A}\} \cdot \mathbf{v} - \mathbf{u}\|_2 \leq \delta \cdot \|\mathbf{v}\|_2$$

in time $\tilde{O}((\text{nnz}(\mathbf{A}) + n) \cdot \log(2 + \|\mathbf{A}\|))$. Here, $\text{nnz}(\mathbf{A})$ denotes the number of non-zero entries of \mathbf{A} , $\|\mathbf{A}\|$ denotes its spectral norm, and $\tilde{O}(\cdot)$ hides factors of $\text{poly}(\log n)$ and $\text{poly}(\log 1/\delta)$.

We can now prove Lemma 6.

Proof of Lemma 6. To show that $\tilde{\mathbf{X}}^{(t)} \in \Delta$, note that $\tilde{\mathbf{X}}^{(t)} \succeq \mathbf{0}$ follows immediately from $\tilde{\mathbf{X}}^{(t)}$ being the Gram matrix of $\mathbf{V}^{(t)}$. To check that its trace with $\Pi_{\perp \mathbf{1}}$ is unit, we can compute

$$\langle \Pi_{\perp \mathbf{1}}, \tilde{\mathbf{X}}^{(t)} \rangle = \langle \Pi_{\perp \mathbf{1}}, \mathbf{V}^{(t)}(\mathbf{V}^{(t)})^\top \rangle = \frac{\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle} = 1$$

We then show that $\langle \mathbf{L}_{ij}, \tilde{\mathbf{X}}^{(t)} \rangle$ is an ϵ -multiplicative factor of $\langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle$. Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be the Cholesky factorization of the matrix exponential $\mathbf{W}^{(t)} = \exp\left\{-\eta \sum_{ij \in \mathcal{E}_{\text{SR}}} w_{ij}^{(t)} \mathbf{L}_{ij}\right\}$, i.e. $\mathbf{W}^{(t)} = \mathbf{A}\mathbf{A}^\top$. Denote $\mathbf{a}_i \in \mathbb{R}^d$ and $\mathbf{v}_i \in \mathbb{R}^k$ by the i -th row of \mathbf{A} and \mathbf{V} respectively. Because Algorithm 2 sets

$$\mathbf{U}^{(t)} = \exp\left\{-\eta \sum_{ij \in \mathcal{E}_{\text{SR}}} w_{ij}^{(t)} \mathbf{L}_{ij}\right\} \mathbf{R},$$

and k is chosen so that $k = O(\log n/\epsilon^2)$, Johnson–Lindenstrauss in Theorem 7 implies

$$(1 - \epsilon) \cdot \|\mathbf{a}_i - \mathbf{a}_j\|^2 \leq \|\mathbf{v}_i - \mathbf{v}_j\|^2 \leq (1 + \epsilon) \cdot \|\mathbf{a}_i - \mathbf{a}_j\|^2$$

for all $(i, j) \in \mathcal{V} \times \mathcal{V}$. We use this to claim the following bound

$$(1 - \epsilon) \cdot \langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle \leq \langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle \leq (1 + \epsilon) \cdot \langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle. \quad (27)$$

One can see the RHS by applying Johnson–Lindenstrauss in this manner:

$$\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle = \frac{1}{n} \cdot \sum_{i < j} \|\mathbf{u}_i - \mathbf{u}_j\|^2 \leq (1 + \epsilon) \cdot \frac{1}{n} \cdot \sum_{i < j} \|\mathbf{a}_i - \mathbf{a}_j\|^2 = (1 + \epsilon) \cdot \langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle.$$

The LHS of inequality (27) then follows by applying the Johnson–Lindenstrauss lower bound.

Now, to finally show that $\langle \mathbf{L}_{ij}, \tilde{\mathbf{X}}^{(t)} \rangle$ is an ϵ -multiplicative factor of $\langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle$, it suffices to demonstrate

$$(1 - O(\epsilon)) \cdot \langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle \leq \langle \mathbf{L}_{ij}, \tilde{\mathbf{X}}^{(t)} \rangle \leq (1 + O(\epsilon)) \cdot \langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle. \quad (28)$$

Because Algorithm 2 sets $\mathbf{V}^{(t)} = \mathbf{U}^{(t)} / \sqrt{\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle}$, we have

$$\langle \mathbf{L}_{ij}, \tilde{\mathbf{X}}^{(t)} \rangle = \langle \mathbf{L}_{ij}, \mathbf{V}^{(t)}(\mathbf{V}^{(t)})^\top \rangle = \frac{\langle \mathbf{L}_{ij}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle} = \frac{\|\mathbf{u}_i - \mathbf{u}_j\|^2}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle}.$$

On the other hand, setting \mathbf{A} to be the Cholesky factorization of $\mathbf{W}^{(t)}$ determines

$$\langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle = \frac{\langle \mathbf{L}_{ij}, \mathbf{W}^{(t)} \rangle}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{W}^{(t)} \rangle} = \frac{\langle \mathbf{L}_{ij}, \mathbf{A}\mathbf{A}^\top \rangle}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle} = \frac{\|\mathbf{a}_i - \mathbf{a}_j\|^2}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle}$$

The RHS of inequality (28) then follows by applying the upper bound in Johnson–Lindenstrauss, and the lower bound in inequality (27) to get the following.

$$\langle \mathbf{L}_{ij}, \tilde{\mathbf{X}}^{(t)} \rangle = \frac{\|\mathbf{u}_i - \mathbf{u}_j\|^2}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle} \leq \frac{1 + \epsilon}{1 - \epsilon} \cdot \frac{\|\mathbf{a}_i - \mathbf{a}_j\|^2}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle} \leq (1 + 2\epsilon) \cdot \frac{\|\mathbf{a}_i - \mathbf{a}_j\|^2}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle} = (1 + O(\epsilon)) \cdot \langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle$$

The LHS of inequality (28) can then be derived similarly

$$\langle \mathbf{L}_{ij}, \tilde{\mathbf{X}}^{(t)} \rangle = \frac{\|\mathbf{u}_i - \mathbf{u}_j\|^2}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)} \rangle} \geq \frac{1 - \epsilon}{1 + \epsilon} \cdot \frac{\|\mathbf{a}_i - \mathbf{a}_j\|^2}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle} \geq (1 - 2\epsilon) \cdot \frac{\|\mathbf{a}_i - \mathbf{a}_j\|^2}{\langle \Pi_{\perp \mathbf{1}}, \mathbf{A}\mathbf{A}^\top \rangle} = (1 - O(\epsilon)) \cdot \langle \mathbf{L}_{ij}, \mathbf{X}^{(t)} \rangle$$

as required. \square

Using Theorem 8, we can also prove Lemma 8.

Proof of Lemma 8. We bound the running time by considering each line of Algorithm 2. As $k = O(\log/\epsilon^2)$, Line 3 only requires nearly-linear time. By Theorem 3.2 in [OSV12], for any graph $\mathcal{G} = (V, E)$, the action of the heat kernel $\exp\{-t\mathbf{L}_{\mathcal{G}}\}$ for $t \geq 0$ can be computed in nearly-linear time in the graph size. Line 4 requires k such computations and hence also runs in nearly-linear time. The normalization of the embedding $\mathbf{U}^{(t)}$ requires computing the inner product $\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle$. This can be achieved easily by noticing that:

$$\langle \Pi_{\perp \mathbf{1}}, \mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top \rangle = \text{Tr}(\mathbf{U}^{(t)}(\mathbf{U}^{(t)})^\top) - \frac{1}{n}\|\mathbf{U}^{(t)}\mathbf{1}\|^2.$$

As the dimension of $\mathbf{U}^{(t)}$ is $n \times k$, both of these terms can be computed in $O(nk)$ arithmetic operations, which is nearly-linear in n . Similarly, the computation of each edge gain \mathbf{c}_{ij} only requires computing the squared distance between the i and j column of $\mathbf{V}^{(t)}$, which can be achieved in time $O(k)$. Hence, all the gains can be computed in nearly-linear time. Finally, by Theorems 4 and 5, Line 7 also runs in nearly-linear time. Hence, every iteration of the main loop runs in nearly-linear time. As there are only $O(\log n/\epsilon^2)$ iterations, Algorithm 2 runs in nearly-linear time. \square

E Sampling in clusters

In this section we consider ranking in a sampling regime structured with underlying clusters. Suppose that there is an underlying cluster labeling $\varphi : [n] \rightarrow [k]$ such that item i belongs to cluster $\varphi(i)$. We assume each edge (i, j) in the comparison graph \mathcal{E} is drawn independently with probability p_t if $\varphi(i) = \varphi(j) = t$ for some $t \in [k]$ and probability q if otherwise. Without loss of generality, we let the first n_1 items to be cluster 1, the next n_2 items to be cluster 2, etc. To distinguish from the rest of the paper, we denote the sampled comparison graph as $\mathcal{G}_{\text{CL}} = (\mathcal{V} = [n], \mathcal{E}_{\text{CL}})$ where CL stands for clusters.

We consider the case of $p_t > q$ for all t . It is easy to see that this is a special case of semi-random sampling. The spectral profile of this setting can be vastly different from a Erdős–Rényi comparison graph. Thus a naive application of Theorem 3 would be unsatisfactory.

Proposition 1. *Suppose that there is an underlying cluster labeling $\varphi : [n] \rightarrow [k]$ and $(i, j) \in \mathcal{E}$ is drawn independently with probability p_t if $\varphi(i) = \varphi(j) = t$ for some $t \in [k]$ and probability q if otherwise. Let $\mathcal{G} = (\mathcal{V} = [n], \mathcal{E})$ be the comparison graph, then there exists a subgraph of \mathcal{G} that is an Erdős–Rényi graph.*

Proof. An equivalent way to write \mathcal{E} is the following: let $\delta_{ij} \sim \text{Unif}(0, 1)$ be i.i.d. random variables and

$$\mathcal{E} = \{(i, j) : (\delta_{ij} \leq p_t \text{ and } \varphi(i) = \varphi(j) = t \text{ for some } t \in [k]) \text{ or } \delta_{ij} \leq q\}.$$

We also consider another graph $\mathcal{G}_1 = (\mathcal{V}, \mathcal{E}_1)$ where $\mathcal{E}_1 = \{(i, j) : \delta_{ij} \leq q\}$. It is easy to see \mathcal{G}_1 is an Erdos-Renyi graph with uniform sampling probability q and also a subgraph of \mathcal{G} . \square

In addition to the assumptions in the semi-random setting, we assume $n_t p_t \geq C \log(n)$ for some absolute constant $C > 0$. We are interested in the vanilla, unweighted MLE, i.e. the solution of (3) with $w_{ij} = \mathbf{1}_{(i,j) \in \mathcal{E}_{\text{CL}}}$. The following result shows $\hat{\boldsymbol{\theta}}$ achieves $\tilde{O}(1/\sqrt{nqL})$ rate in ℓ_∞ error.

Theorem 9. *Suppose that $nq \geq C_1 \log(n)$ and $n_t p_t \geq C_1 \log(n)$. for some large enough constants $C_1 > 0$. Suppose $nqL \geq C_2 \kappa^4 \log^3(n)$ and $n^2 q^2 L \geq C_2 \kappa^4 n_t p_t \log^3(n)$ for some large enough constant $C_2 > 0$. for any $t \in [n]$. Then with probability at least $1 - O(n^{-10})$, for any (k, l) , the unweighted MLE $\hat{\boldsymbol{\theta}}$ satisfies*

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq C_3 \kappa \sqrt{\frac{\log(n)}{nqL}}$$

for some constants $C_3 > 0$. As a result, the top- K items are recovered exactly as long as

$$n^2 qL \geq C_4 \frac{\kappa^2 n \log(n)}{\Delta_K^2}$$

for some large enough constant C_5 .

The result of this theorem matches what we can get in Theorem 1 using reweighting, with a slight difference in the sample complexity requirement. This suggests that at least in some non-uniform sampling models, reweighting is unnecessary for MLE.

E.1 Proof of Theorem 9

We start by applying Lemma 2 with $w_{ij} = \mathbb{1}_{(i,j) \in \mathcal{E}_{\text{CL}}}$. For any $k \neq l$, let B_{kl}, Q_{kl} be some real number that we will specify later such that

$$B_{kl} \geq C_1 \sqrt{\frac{\kappa}{L} \boldsymbol{\Omega}_{kl}(\mathbf{L}_z) \log(n)}; \quad (29)$$

$$Q_{kl} \geq \sum_{(i,j) \in \mathcal{E}_{\text{CL}}: i > j} B_{ij}^2 |(\mathbf{e}_k - \mathbf{e}_l)^\top \mathbf{L}_z^\dagger (\mathbf{e}_i - \mathbf{e}_j)|. \quad (30)$$

Here $C_1 > 0$ is some large enough constant and

$$\mathbf{L}_z = \sum_{(i,j) \in \mathcal{E}_{\text{CL}}: i > j} z_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top.$$

Suppose that $Q_{kl} \leq 4B_{kl}$ for any (k, l) . By Lemma 2, with probability at least $1 - n^{-10}$, we have for any (k, l) ,

$$\left| (\hat{\theta}_k - \hat{\theta}_l) - (\theta_k^* - \theta_l^*) \right| \leq B_{kl}.$$

We now let

$$B_{kl} := C_2 \sqrt{\frac{\kappa^2 \log(n)}{(n_t p_t \vee nq)L}}$$

if k, l are both in some cluster t and

$$B_{kl} := C_2 \sqrt{\frac{\kappa^2 \log(n)}{nqL}}$$

if k, l are not in the same cluster. We also let

$$Q_{kl} := C_3 \frac{\kappa^2 \log^2(n)}{nqL}$$

Here $C_2, C_3 > 0$ are some large enough constants. constant. We denote \mathcal{G}_{B^2} as the graph \mathcal{G}_{CL} equipped with weights $\{B_{ij}^2\}$.

We now show (29) and (30) hold using the same strategy in the proof of Theorem 3. First, we control effective resistance and graph conductance with the following lemmas. The proofs are deferred to Section E.2 and E.3.

Lemma 17. *With probability at least $1 - O(n^{-10})$,*

$$\Omega_{kl}(\mathbf{L}_z) \leq \frac{\kappa}{n_t p_t \vee nq}$$

if k, l are both in some cluster t and

$$\Omega_{kl}(\mathbf{L}_z) \leq \frac{\kappa}{nq}$$

if k, l are not in the same cluster.

Lemma 18. *Let \mathcal{G}_{B^2} be \mathcal{G}_{CL} with weights $\{B_{ij}^2\}_{(i,j) \in \mathcal{E}_{\text{CL}}}$. The graph conductance of \mathcal{G}_{B^2} satisfies*

$$\Phi(\mathcal{G}_{B^2}) \geq \frac{1}{8}.$$

From Lemma 17, we can see that (29) is satisfied as long as C_2 is large enough. By Lemma 18 we have that

$$\sum_{(i,j) \in \mathcal{E}_{\text{CL}}: i > j} B_{ij}^2 |(e_i - e_j)^\top \mathbf{L}_z^\dagger (e_k - e_l)| \leq \max_{(i,j) \in \mathcal{E}_{\text{CL}}} \{B_{ij}^2\} \cdot \frac{8 \log \left(\sum_{(i,j) \in \mathcal{E}_{\text{CL}}: i > j} B_{ij}^2 \right)}{\Phi(\mathcal{G}_{B^2})^2}. \quad (31)$$

For all $(k, l) \in [n]^2$,

$$B_{kl}^2 \leq C_2^2 \frac{\kappa^2 \log(n)}{nqL} \leq 1$$

as long as $nqL \geq C_4 \kappa^4 \log^3(n)$ for some large enough constant $C_4 > 0$. Then provided that C_3 is large enough, (31) becomes

$$\begin{aligned} \sum_{(i,j) \in \mathcal{E}_{\text{CL}}: i > j} B_{ij}^2 |(e_i - e_j)^\top \mathbf{L}_z^\dagger (e_k - e_l)| &\leq C_2^2 \frac{\kappa^2 \log(n)}{nqL} \cdot \frac{8 \log(n^2)}{(1/8)^2} \\ &\leq C_3 \frac{\kappa^2 \log^2(n)}{nqL} = Q_{kl}. \end{aligned}$$

This satisfies (30). Using the assumption that $nqL \geq C_4 \kappa^4 \log^3(n)$ and $n^2 q^2 L \geq C_5 \kappa^4 n_t p_t \log^3(n)$ for any $t \in [n]$ with some large enough constants $C_4, C_5 > 0$, we have that $Q_{kl} \leq 4B_{kl}$ for any $(k, l) \in [n]^2$.

Now that we have checked all conditions for Lemma 2, we conclude that

$$\left| (\hat{\theta}_k - \hat{\theta}_l) - (\theta_k^* - \theta_l^*) \right| \leq B_{kl} = C_2 \kappa \sqrt{\frac{\log(n)}{nqL}}.$$

E.2 Proof of Lemma 17

We consider two subgraphs $\mathcal{G}_1 = (\mathcal{V}, \mathcal{E}_1)$ and $\mathcal{G}_2 = (\mathcal{V}, \mathcal{E}_2)$ of \mathcal{G} , where \mathcal{E}_1 consists of all edges within each cluster, and \mathcal{E}_2 consists of the edges of an underlying Erdős-Rényi subgraph as in Proposition 1. Let their unweighted graph Laplacian be \mathbf{L}_1 and \mathbf{L}_2 . Let their z -weighted graph Laplacian be $\mathbf{L}_{z,1}$ and $\mathbf{L}_{z,2}$. Combining Lemma 3 and 10, for k, l both in some cluster t ,

$$\Omega_{kl}(\mathbf{L}_{z,1}) \leq \frac{8\kappa}{\lambda_{n_t-1}(\mathbf{L}_1)} \leq \frac{16\kappa}{n_t p_t};$$

while for any k, l (including those in cluster t), we look at \mathcal{G}_2 to see

$$\Omega_{kl}(\mathbf{L}_{z,2}) \leq \frac{8\kappa}{\lambda_{n-1}(\mathbf{L}_2)} \leq \frac{16\kappa}{nq}.$$

Combining this with Rayleigh's law of monotonicity (see Lemma 12), we reach the claimed result.

E.3 Proof of Lemma 18

From Lemma 11 we have

$$\Phi(\mathcal{G}_{B^2}) \geq \frac{1}{2} \lambda_{n-1}(\mathbf{D}_{B^2}^{-1/2} \mathbf{L}_{B^2} \mathbf{D}_{B^2}^{-1/2}). \quad (32)$$

Here \mathbf{D}_{B^2} is the diagonal matrix with $[\mathbf{D}_{B^2}]_{ii} = \sum_{j:(i,j) \in \mathcal{E}} B_{ij}^2$ and

$$\begin{aligned} \mathbf{L}_{B^2} &= \sum_{(i,j) \in \mathcal{E}_{\text{CL}}, i>j} B_{ij}^2 (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top \\ &= \sum_{(i,j), i>j} \underbrace{B_{ij}^2 \delta_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top}_{=: \mathbf{L}_{ij}}, \end{aligned}$$

where $\delta_{ij} = \mathbb{1}_{\{(i,j) \in \mathcal{E}_{\text{CL}}\}}$. To control the right hand side of (32), we give a lower bound of $\lambda_{n-1}(\mathbf{L}_{B^2})$ and an upper bound of the diagonal entries of \mathbf{D}_{B^2} in the following lemmas. The proofs are deferred to the end of this section.

Lemma 19. *Instate the assumption of Lemma 18. With probability at least $1 - O(n^{-10})$,*

$$\lambda_{n-1}(\mathbf{L}_{B^2}) \geq \frac{C\kappa^2 \log(n)}{2L},$$

where C is some constant.

Lemma 20. *Instate the assumption of Lemma 18. With probability at least $1 - O(n^{-10})$,*

$$\sum_{j:(i,j) \in \mathcal{E}} B_{ij}^2 \leq \frac{4C\kappa^2 \log(n)}{L}$$

for all $i \in [n]$. Here C is a constant.

With these two lemmas, by Sylvester's law of inertia (Lemma 13),

$$\begin{aligned} \Phi(\mathcal{G}_{B^2}) &\geq \frac{1}{2} \lambda_{n-1}(\mathbf{D}_{B^2}^{-1/2} \mathbf{L}_{B^2} \mathbf{D}_{B^2}^{-1/2}) \\ &\geq \frac{1}{2} \lambda_{n-1}(\mathbf{L}_{B^2}) \lambda_n(\mathbf{D}_{B^2}^{-1}) \\ &\geq \frac{1}{8}. \end{aligned}$$

Proof of Lemma 19. The proof follows similar strategy to Section 5.3.3 in [Tro15]. Let $\mathbf{R} \in \mathbb{R}^{(n-1) \times n}$ be a partial isometry such that $\mathbf{R}\mathbf{R}^\top = \mathbf{I}_{n-1}$ and $\mathbf{R}\mathbf{1}_n = \mathbf{0}$. Then $\lambda_{n-1}(\mathbf{L}_{B^2}) = \lambda_{n-1}(\mathbf{R}\mathbf{L}_{B^2}\mathbf{R}^\top)$ and we will use matrix Chernoff to control the latter term.

Recall that

$$\mathbf{L}_{ij} = B_{ij}^2 \delta_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top.$$

For all (i, j) such that $i > j$,

$$0 \leq \lambda_{n-1}(\mathbf{R}\mathbf{L}_{ij}\mathbf{R}^\top) \leq \lambda_1(\mathbf{R}\mathbf{L}_{ij}\mathbf{R}^\top) = \lambda_1(\mathbf{L}_{ij}) \leq 2C_2 \frac{\kappa^2 \log(n)}{nqL}.$$

Because $\mathbb{P}[L_{ij}^t = 1 \mid R_{ij}^t = 1] \geq \frac{2\kappa}{(1+\kappa)^2} \geq 1/(2\kappa)$,

$$\begin{aligned} \lambda_{n-1}(\mathbb{E}\mathbf{R}\mathbf{L}_{B^2}\mathbf{R}^\top) &= \lambda_{n-1} \left(\mathbf{R} \sum_{t=1}^n \sum_{i>j} \mathbb{E}\mathbf{L}_{ij}\mathbf{R}^\top \right) \\ &= \lambda_{n-1} \left(\mathbf{R} \sum_{t=1}^n \sum_{i>j} B_{ij}^2 \mathbb{E}\delta_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{R}^\top \right) \end{aligned}$$

When i, j are in the same cluster t ,

$$B_{ij}^2 \mathbb{E} \delta_{ij} \geq C_2 \frac{p_t \kappa^2 \log(n)}{(n_t p_t \vee nq)L} \geq C_2 \frac{\kappa^2 \log(n)}{nL};$$

when i, j are not in the same cluster,

$$B_{ij}^2 \mathbb{E} \delta_{ij} \geq C_2 \frac{q \kappa^2 \log(n)}{nqL} = C_2 \frac{\kappa^2 \log(n)}{nL}.$$

Then

$$\begin{aligned} \lambda_{n-1}(\mathbb{E} \mathbf{R} \mathbf{L}_{B^2} \mathbf{R}^\top) &\geq C_2 \frac{\kappa^2 \log(n)}{nL} \cdot \lambda_{n-1} \left(\mathbf{R} \sum_{t=1}^n \sum_{i>j} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top \mathbf{R}^\top \right) \\ &= C_2 \frac{\kappa^2 \log(n)}{nL} \cdot \lambda_{n-1} [\mathbf{R} (n\mathbf{I}_n - \mathbf{1}_n \mathbf{1}_n^\top) \mathbf{R}^\top] \\ &= C_2 \frac{\kappa^2 \log(n)}{nL} \cdot \lambda_{n-1} [n\mathbf{I}_{n-1}] \\ &= C_2 \frac{\kappa^2 \log(n)}{L}. \end{aligned}$$

The second to last line holds since $\mathbf{R} \mathbf{R}^\top = \mathbf{I}_{n-1}$ and $\mathbf{R} \mathbf{1}_n = \mathbf{0}$. Then by matrix Chernoff (see e.g. Theorem 5.1.1 in [Tro15]),

$$\mathbb{P} \left\{ [\lambda_{n-1}(\mathbf{R} \mathbf{L}_{B^2} \mathbf{R}^\top)] \leq \frac{C_2 \kappa^2 \log(n)}{2L} \right\} \leq n \cdot \left[\frac{e^{-1/2}}{(1/2)^{1/2}} \right]^{nq/2} \leq n^{-10}$$

as long as $nq \geq C_1 \log(n)$ for some large enough constant C_1 . The $nq/2$ term in the exponent comes from the fact that $\lambda_{n-1}(\mathbb{E} \mathbf{R} \mathbf{L}_{B^2} \mathbf{R}^\top) / \lambda_1(\mathbf{L}_{ij}) \geq nq/2$.

Proof of Lemma 20. Let $\varphi(i) = t$. We split the summation with

$$\begin{aligned} \sum_{j:(i,j) \in \mathcal{E}} B_{ij}^2 &= \sum_{j:(i,j) \in \mathcal{E}, \varphi(j)=t} B_{ij}^2 + \sum_{j:(i,j) \in \mathcal{E}, \varphi(j) \neq t} B_{ij}^2 \\ &= \frac{C_1 \kappa^2 \log(n)}{L} \left[\sum_{j:(i,j) \in \mathcal{E}, \varphi(j)=t} \frac{1}{n_t p_t \vee nq} + \sum_{j:(i,j) \in \mathcal{E}, \varphi(j) \neq t} \frac{1}{nq} \right], \end{aligned}$$

where $C_1 > 0$ is a constant. By assumption for any $t \in [k]$, $nq \geq C_2 \log(n)$ and $n_t p_t \geq C_3 \log(n)$, where $C_2, C_3 > 0$ are some constants. Applying standard Chernoff bound, we have that with probability at least $1 - 4n^{-10}$, for all $i \in [n]$,

$$\frac{1}{2} n_t p_t \leq |\{j : (i, j) \in \mathcal{E}, \varphi(j) = t\}| \leq 2n_t p_t$$

and

$$\frac{1}{2} nq \leq |\{j : (i, j) \in \mathcal{E}\}| \leq 2nq$$

as long as C_2, C_3 are large enough. Then

$$\begin{aligned} \sum_{j:(i,j) \in \mathcal{E}} B_{ij}^2 &\leq C_1 \frac{\kappa^2 \log(n)}{L} \left[2n_t p_t \cdot \frac{1}{n_t p_t \vee nq} + 2nq \cdot \frac{1}{nq} \right] \\ &\leq 4C_1 \frac{\kappa^2 \log(n)}{L}. \end{aligned}$$

F Experiment setup for Figure 2

To check our result, we implement the Algorithm 1 and apply it to some simulated data generated with semi-random sampling. Our semi-random graph is generated as follows:

1. Generate a Erdos-Renyi random graph \mathcal{G}_{ER} with n vertices and edge probability p .
2. Randomly select a subset A of $n/3$ vertices in the first $n/2$ vertices, and a subset B of $n/3$ vertices in the last $n/2$ vertices.
3. Form \mathcal{G}_{SR} by adding all edges between vertices in A and all edges between vertices in B to \mathcal{G}_{ER} .

We then generate the comparison data following the BTL model with $\theta^* = \mathbf{0}_n$. Similar to the example in Figure 1, this quickly ruins the nice spectral properties observed in uniform sampling. In each trial We compute two MLE estimates of θ^* and the corresponding top- K items. The first one is the vanilla MLE $\hat{\theta}$ using all data on \mathcal{G}_{ER} and the second one is the weighted MLE $\hat{\theta}_w$ given by Algorithm 1. We compare the top- K recovery accuracy, i.e. the proportion of top- K items that are successfully recovered, under varying latent score gap Δ_K .