Spectral MLE: Top-K Rank Aggregation from Pairwise Comparisons

Yuxin Chen YXCHEN@STANFORD.EDU
Department of Statistics, Stanford University, Stanford, CA 94305, USA

Changho Suh CHSUH@KAIST.AC.KR
Department of Electrical Engineering, Korea Advanced Institute of Science and Technology, Korea

Abstract

This paper explores the preference-based top-K rank aggregation problem. Suppose that a collection of items is repeatedly compared in pairs, and one wishes to recover a consistent ordering that emphasizes the top-K ranked items, based on partially revealed preferences. We focus on the Bradley-Terry-Luce model that postulates a set of latent preference scores underlying all items, where the odds of paired comparisons depend only on the relative scores of the items involved.

We characterize the minimax limits on identifiability of top-K ranked items, in the presence of random and non-adaptive sampling. Our results highlight a separation measure that quantifies the gap of preference scores between the K^{th} and $(K+1)^{\text{th}}$ ranked items. The minimum sample complexity required for reliable top-K ranking scales inversely with the separation measure. To approach this minimax limit, we propose a nearly linear-time ranking scheme, called *Spectral MLE*, that returns the indices of the top-K items in accordance to a careful score estimate. In a nutshell, Spectral MLE starts with an initial score estimate with minimal squared loss (obtained via a spectral method), and then successively refines each component with the assistance of coordinate-wise MLEs. Encouragingly, Spectral MLE allows perfect top-K item identification under minimal sample complexity. The practical applicability of Spectral MLE is further corroborated by numerical experiments.

1. Introduction and Motivation

The task of rank aggregation is encountered in a wide spectrum of contexts like social choice (Caplin & Nalebuff, 1991; Soufiani et al., 2014b), web search and information retrieval (Dwork et al., 2001), crowd sourcing (Chen et al., 2013), recommendation systems (Baltrunas et al., 2010), to name just a few. Given partial preference information over a collection of items, the aim is to identify a consistent ordering that best respects the revealed preference. In the high-dimensional regime, one is often faced with two challenges: 1) the number of items to be ranked is ever growing, which makes it increasingly harder to recover a consistent total ordering over all items; 2) the observed data is highly incomplete and inconsistent: only a small number of noisy pairwise / listwise preferences can be acquired.

In an effort to address such challenges, this paper explores a popular pairwise preference-based model, which postulates the existence of a ground-truth ranking. Specifically, consider a parametric model involving n items, each assigned a preference score that determines its rank. Concrete examples of preference scores include the overall rating of an athlete, the academic performance and competitiveness of a university, the dining quality of a restaurant, etc. Each item is then repeatedly compared against a few others in pairs, yielding a set of *noisy* binary comparisons generated based on the relative preference scores. In many situations, the number of repeated comparisons essentially reflects the signal-to-noise ratio (SNR) / quality of the information revealed for each pair of items. The goal is then to develop a "denoising" procedure that recovers the ground-truth ranking with minimal sample complexity.

There has been a proliferation of ranking schemes that suggest partial solutions. While the ranking that we are seeking is better treated as a function of the preference parameters, most of the aforementioned schemes adopt the natural "plug-in" procedure, that is, start by inferring the preference scores, and then return a ranking in accordance to the parametric estimates. The most popular paradigm is arguably the maximum likelihood estimation (MLE) (Ford, 1957), where the main appeal of MLE is its inherent convexity under several comparison models, e.g. the Bradley-Terry-Luce (BTL) model (Bradley & Terry, 1952; Luce, 1959). Encouragingly, MLE often achieves low ℓ_2 estimation loss while retaining efficient finite-sample complexity. Another prominent alternative concerns a family of spectral

Proceedings of the 32^{nd} International Conference on Machine Learning, Lille, France, 2015. JMLR: W&CP volume 37.

ranking algorithms (e.g. *PageRank* (Brin & Page, 1998)). A provably efficient choice within this family is *Rank Centrality* (Negahban et al., 2012), which produces an estimate with nearly minimax mean squared error (MSE). While both MLE and Rank Centrality allow intriguing guarantees towards finding faithful parametric estimates, the squared loss metric considered therein does not necessarily imply optimality of the ranking accuracy. In fact, there is no shortage of high-dimensional situations that admit parametric estimates with low squared loss while precluding reliable ranking. Furthermore, many realistic scenarios emphasize only a few items that receive the highest ranks. Unfortunately, the above MSE results fall short of ensuring recovery of the top-ranked items.

In this work, we consider accurate identification of top-K ranked items under the popular BTL pairwise comparison model, assuming that the item pairs we can compare are selected in a random and non-adaptive fashion (termed *passive ranking*). In particular, we aim to explore the following questions: (a) what is the minimum number of repeated comparisons necessary for reliable ranking? (b) how is the ranking accuracy affected by the underlying preference score distributions? We will address these two questions from both statistical and algorithmic perspectives.

1.1. Main Contributions

This paper investigates minimax optimal procedures for top-K rank aggregation. Our contributions are two-fold.

To begin with, we characterize the fundamental three-way tradeoff between the number of repeated comparisons, the sparsity of the comparison graph, and the preference score distribution, from a minimax perspective. In particular, we emphasize a separation measure that quantifies the gap of preference scores between the K^{th} and $(K+1)^{\text{th}}$ ranked items. Our results demonstrate that the minimal sample complexity or quality of paired evaluation (reflected by the number of repeated comparisons per an observed pair) scales inversely with the separation measure.

Secondly, we propose a nearly linear-time two-stage algorithm, called *Spectral MLE*, which allows perfect top-K identification as soon as the sample complexity exceeds the minimax limits (modulo some constant). Specifically, Spectral MLE starts by obtaining careful initial scores that are faithful in the ℓ_2 sense (e.g. via a spectral method), and then iteratively sharpens the pointwise estimates by comparing the preceding estimates with coordinate-wise MLE. This algorithm is designed primarily in an attempt to seek a score estimate with minimal pointwise loss. Furthermore, numerical experiments demonstrate that Spectral MLE outperforms Rank Centrality by achieving higher ranking accuracy and lower ℓ_∞ estimation error.

1.2. Prior Art

There are two distinct families of observation models that receive considerable interest: (1) *value-based model*, where the observation on each item is drawn only from the distribution underlying this individual; (2) *preference-based model*, where one observes the relative order among a few items instead of revealing their individual values. Best-K identification in the value-based model with adaptive sampling (termed *active ranking*) is closely related to the multi-armed bandit problem, where the fundamental identification complexity has been characterized (Gabillon et al., 2011; Bubeck et al., 2013; Jamieson et al., 2014). The value-based and preference-based models have also been compared in terms of minimax error rates in estimating the latent quantities (Shah et al., 2014).

In the realm of pairwise preference settings, many *active ranking* schemes (Busa-Fekete & Hüllermeier, 2014) have been proposed in an attempt to optimize the exploration-exploitation tradeoff. For instance, in the noise-free case, Jamieson et al. (Jamieson & Nowak, 2011) considered perfect total ranking and characterized the query complexity gain of adaptive sampling relative to random queries, provided that the items under study admit a low-dimensional Euclidean embedding. Furthermore, the works (Ailon, 2012; Jamieson & Nowak, 2011; Braverman & Mossel, 2008; Wauthier et al., 2013) explored the query complexity in the presence of noise, but focused on "approximately correct" total rankings—a solution with loss at most a factor $(1+\epsilon)$ from optimal—rather than accurate ordering. Another path-based approach has been proposed to accommodate accurate top-K queries from noisy pairwise data (Eriksson, 2013), where the observation error is assumed to be i.i.d. instead of being item-dependent. Motivated by the success of value-based racing algorithms, (Busa-Fekete et al., 2013; Busa-Fekete & Hüllermeier, 2014) came up with a generalized racing algorithm that often led to efficient sample complexity. In contrast, our paper concentrates on top-K identification in a *passive* setting, assuming that partial preferences are collected in a noisy, random, and non-adaptive manner. This was previously out of reach.

Apart from Rank Centrality and MLE, the most relevant work is (Rajkumar & Agarwal, 2014). For a variety of rank aggregation methods, they developed intriguing sufficient statistical hypotheses that guarantee the convergence to an optimal ranking, which in turn leads to sample complexity bounds for Rank Centrality and MLE. Nevertheless, they focused on perfect total ordering instead of top-K selection, and their results fall short of a rigorous justification as to whether or not the derived sample complexity bounds are statistically optimal.

Finally, there are many related yet different problem settings considered in the prior literature. For instance, the work (Ammar & Shah, 2012) approached top-K ranking using a maximum entropy principle, assuming the existence of a distribution μ over all possible permutations.

Recent work (Soufiani et al., 2013; 2014a) investigated consistent *rank breaking* under more generalized models involving full rankings. A family of distance measures on rankings has been studied and justified based on an axiomatic approach (Farnoud & Milenkovic, 2014). Another line of works considered the popular distance-based Mallows model (Lu & Boutilier, 2011; Busa-Fekete et al., 2014; Awasthi et al., 2014). An online ranking setting has also been studied (Harrington, 2003; Farnoud et al., 2014). The minimax recovery limits under general pairwise measurements have been determined by (Chen et al., 2015a).

1.3. Notation

Let $[n]$ represent $\{1, 2, \cdots, n\}$. A graph $\mathcal{G}$ is said to be an Erdős–Rényi random graph, denoted by $\mathcal{G}_{n,p_{\rm obs}}$, if each pair (i, j) is connected by an edge independently with probability $p_{\rm obs}$. Besdies, we use $\deg(i)$ to represent the degree of vertex i in $\mathcal{G}$.

2. Problem Setup

Comparison Model and Assumptions. Suppose that we observe a few pairwise evaluations over n items. To pursue a statistical understanding towards the ranking limits, we assume that the pairwise comparison outcomes are generated according to the BTL model (Bradley & Terry, 1952; Luce, 1959), a long-standing model that has been applied in numerous applications (Agresti, 2014; Hunter, 2004).

- *Preference Scores*. The BTL model hypothesizes on the existence of some hidden preference vector $\boldsymbol{w} = [w_i]_{1\le i\le n}$, where w_i represents the underlying preference score / weight of item i. The outcome of each paired comparison depends only on the scores of the items involved. Unless otherwise specified, we will assume without loss of generality that

$$w_1 \ge w_2 \ge \cdots \ge w_n > 0. \tag{1}$$

- *Comparison Graph*. Denote by $\mathcal{G} = ([n], \mathcal{E})$ the comparison graph such that items i and j are compared if and only if (i, j) belongs to the edge set $\mathcal{E}$. We will mostly assume that $\mathcal{G}$ is drawn from the Erdős–Rényi model $\mathcal{G} \sim \mathcal{G}_{n,p_{\rm obs}}$ for some observation factor $p_{\rm obs}$.

- *(Repeated) Pairwise Comparisons*. For each $(i, j) \in \mathcal{E}$, we observe L paired comparisons between items i and j. The outcome of the $l^{\rm th}$ comparison between them, denoted by $y_{i,j}^{(l)}$, is generated as per the BTL model:

$$y_{i,j}^{(l)} = \begin{cases} 1, & \text{with probability } \frac{w_i}{w_i+w_j}, \\ 0, & \text{else}, \end{cases} \tag{2}$$

  where $y_{i,j}^{(l)} = 1$ indicates a win by i over j. We adopt the convention that $y_{j,i}^{(l)} = 1 - y_{i,j}^{(l)}$. It is assumed throughout that conditional on $\mathcal{G}$, $y_{i,j}^{(l)}$'s are jointly independent across all l and $i > j$. For ease of presentation, we introduce the collection of sufficient statistics as

$$\boldsymbol{y}_i := \{y_{i,j} \mid j : (i, j) \in \mathcal{E}\}; \quad y_{i,j} := \frac{1}{L}\sum_{l=1}^{L} y_{i,j}^{(l)}.$$

- *Signal to Noise Ratio (SNR) / Quality of Comparisons*. The overall faithfulness of the acquired evaluation between items i and j is captured by the sufficient statistic $y_{i,j}$. Its SNR can be captured by

$$\mathsf{SNR} := \mathbb{E}^2[y_{i,j}]/\mathbf{Var}[y_{i,j}] \asymp L. \tag{3}$$

  As a result, the number L of repeated comparisons measures the SNR or the *quality* of comparisons over any observed pair of items.

- *Dynamic Range of Preference Scores*. It is assumed throughout that the dynamic range of the preference scores is fixed irrespective of n, namely,

$$w_i \in [w_{\min}, w_{\max}], \qquad 1 \le i \le n \tag{4}$$

  for some positive constants $w_{\min}$ and $w_{\max}$ bounded away from 0, which amounts to the most challenging regime (Negahban et al., 2012). In fact, the case in which the range $\frac{w_{\max}}{w_{\min}}$ grows with n can be readily translated into the above fixed-range regime by first separating out those items with vanishing scores (e.g. via a simple voting method like Borda count (Ammar & Shah, 2011)).

Performance metric. Given these pairwise observations, one wishes to see whether or not the top-K ranked items are identifiable. To this end, we consider the probability of error $P_{\rm e}$ in isolating the *set* of top-K ranked items, i.e.

$$P_{\rm e}(\psi) := \mathbb{P}\Big\{\psi(\boldsymbol{y}) \ne [K]\Big\}, \tag{5}$$

where ψ is any ranking scheme that returns a set of K indices. Here, $[K]$ denotes the (unordered) set of the first K indices. We aim to characterize the fundamental *admissible region* of $(L, p_{\rm obs})$ where reliable top-K ranking is feasible, i.e. $P_{\rm e}$ can be vanishingly small as n grows.

3. Minimax Ranking Limits

We explore the fundamental ranking limits from a minimax perspective, which centers on the design of *robust* ranking schemes that guard against the worst case in probability of error. The most challenging component of top-K rank aggregation hinges upon distinguishing the two items near the decision boundary, i.e. the $K^{\rm th}$ and $(K+1)^{\rm th}$ ranked items. Due to the random nature of the acquired finite-bit comparisons, the information concerning their relative preference could be obliterated by noise, unless their latent preference scores are sufficiently separated. In light of this, we single out a preference separation measure as follows

$$\Delta_K := \frac{w_K - w_{K+1}}{w_{\max}}. \tag{6}$$

As will be seen, this measure plays a crucial role in determining information integrity for top-K identification.

To model non-adaptive sampling and partial observation, we employ the Erdős–Rényi model $\mathcal{G} \sim \mathcal{G}_{n,p_{\rm obs}}$. As already noted by (Ford, 1957), if the comparison graph $\mathcal{G}$ is not connected, then there is absolutely no basis to determine relative preferences between two disconnected components. Therefore, a reasonable necessary condition that one would expect is the connectivity of $\mathcal{G}$, which requires

$$p_{\rm obs} > \log n \,/\, n. \tag{7}$$

All results in this paper will operate under this assumption.

A main finding of this paper is an order-wise tight sufficient condition for top-K identifiability, as stated below.

Theorem 1 (**Identifiability**). *Suppose that $\mathcal{G} \sim \mathcal{G}_{n,p_{\rm obs}}$ with $p_{\rm obs} \geq c_0 \log n/n$. Assume that $L = O\left(\mathrm{poly}\left(n\right)\right)$ and $\frac{w_{\max}}{w_{\min}} = O\left(1\right)$. With probability exceeding $1 - c_1 n^{-2}$, the set of top-K ranked items can be identified exactly by an algorithm that runs in time $O\left(|\mathcal{E}| \log^2 n\right)$, provided that*

$$L \;\geq\; \frac{c_2 \log n}{n p_{\rm obs} \Delta_K^2}. \tag{8}$$

Here, $c_0, c_1, c_2 > 0$ are some universal constants.

Remark 1. *We assume throughout that the input fed to each ranking algorithm is the sufficient statistic $\{y_{i,j} \mid (i,j) \in \mathcal{E}\}$ rather than the entire collection of $y_{i,j}^{(l)}$, otherwise the complexity is at least $O\left(L \cdot |\mathcal{E}|\right)$.*

Theorem 1 characterizes an *identifiable region* within which exact identification of top-K items is plausible by nearly linear-time algorithms. The algorithm we propose, as detailed in Section 4, attempts recovery by computing a score estimate whose errors can be uniformly controlled across all entries. Afterwards, the algorithm reports the K items that receive the highest estimated scores.

Encouragingly, the above identifiable region is minimax optimal. Consider a given separation condition Δ_K, and suppose that nature behaves in an adversarial manner by choosing the worst-case scores $\boldsymbol{w}$ compatible with Δ_K. This imposes a minimax lower bound on the quality of comparisons necessary for reliable ranking, as given below.

Theorem 2 (**Minimax Lower Bounds**). *Fix $\epsilon \in \left(0, \frac{1}{2}\right)$, and let $\mathcal{G} \sim \mathcal{G}_{n,p_{\rm obs}}$. If*

$$L \;\leq\; c\frac{(1-\epsilon)\log n - 2}{n p_{\rm obs} \Delta_K^2} \tag{9}$$

holds for some constant[1] *$c > 0$, then for any ranking scheme ψ, there exists a preference vector $\boldsymbol{w}$ with separation Δ_K such that $P_{\rm e}\left(\psi\right) \geq \epsilon$.*

[1]More precisely, $c = w_{\min}^4/(2w_{\max}^4)$.

Theorem 2 taken collectively with Theorem 1 determines the scaling of the fundamental ranking boundary on L. Since the sample size sharply concentrates around $n^2 p_{\rm obs} L$ in our model, this implies that the required sample complexity for top-K ranking scales inversely with the preference separation at a quadratic rate. Put another way, Theorem 2 justifies the need for a minimum separation criterion that applies to any ranking scheme:

$$\Delta_K \;\gtrsim\; \sqrt{\log n \,/\, (n p_{\rm obs} L)}. \tag{10}$$

Somewhat unexpectedly, there is no computational barrier away from this statistical limit. Several other remarks of Theorems 1-2 are in order.

- ℓ_2 **Loss vs.** ℓ_∞ **Loss**. A dominant fraction of prior methods focus on the mean squared error in estimating the latent scores $\boldsymbol{w}$. It was established by (Negahban et al., 2012) that the minimax ℓ_2 regret is squeezed between

$$\frac{1}{\sqrt{n p_{\rm obs} L}} \;\lesssim\; \inf_{\hat{\boldsymbol{w}}} \sup_{\boldsymbol{w}} \frac{\mathbb{E}\left[\|\hat{\boldsymbol{w}} - \boldsymbol{w}\|\right]}{\|\boldsymbol{w}\|} \;\lesssim\; \sqrt{\frac{\log n}{n p_{\rm obs} L}},$$

  where the infimum is over all score estimators $\hat{\boldsymbol{w}}$. This limit is almost identical to the minimax separation criterion (10) we derive for top-K identification. In fact, if the pointwise error of $\hat{\boldsymbol{w}}$ is *uniformly* bounded by $\sqrt{\log n/(n p_{\rm obs} L)}$, then $\hat{\boldsymbol{w}}$ necessarily achieves the minimax ℓ_2 error. Moreover, the pointwise error bound presents a fundamental bottleneck for top-K ranking—it will be impossible to differentiate the $K^{\rm th}$ and $(K+1)^{\rm th}$ ranked items unless their score separation exceeds the aggregate error of the corresponding score estimates for these two items. Based on this observation, our algorithm is mainly designed to control the elementwise estimation error. As will be seen in Section 4, the resulting estimation error will be uniformly spread over all entries, which is optimal in both ℓ_2 and ℓ_∞ sense.

- **From Coarse to Detailed Ranking**. The identifiable region we present depends only on the preference separation between items K and $K+1$. This arises since we only intend to identify the group of top-K items without specifying fine details within this group (termed *coarse ranking*). In fact, our results readily uncover the minimax separation requirements for the case where one further expects fine ordering among these K items. Specifically, this task is feasible (in the minimax sense) iff

$$\Delta_i \;\gtrsim\; \sqrt{\log n/(n p_{\rm obs} L)}, \quad 1 \leq i \leq K. \tag{11}$$

- **High SNR Requirement for Total Ordering**. In many situations, the separation criterion (11) immediately demonstrates the hardness (or even impossibility) of recovering the ordering over all items. In fact, to figure out the total order, one expects sufficient score separation between all pairs of consecutive items, namely,

$$\Delta_i \;\gtrsim\; \sqrt{\log n \,/\, (n p_{\rm obs} L)}, \quad \forall i \; (1 \leq i < n).$$

Since the Δ_i's are defined in a normalized way (6), they need to satisfy

$$\sum_{i=1}^{n-1} \Delta_i = \frac{w_1 - w_n}{w_{\max}} \leq 1.$$

As can be easily verified, the preceding two conditions would be incompatible unless

$$L \gtrsim n \log n \,/\, p_{\mathrm{obs}},$$

which imposes a fairly stringent SNR requirement. For instance, under a sparse graph where $p_{\mathrm{obs}} \asymp \frac{\log n}{n}$, the number of repeated comparisons (and hence the SNR) needs to be at least $\Theta(n^2)$, regardless of the method employed. Such a high SNR requirement could be increasingly more difficult to guarantee as n grows.

- **Passive Ranking vs. Active Ranking**. In our passive ranking model, the sample complexity requirement $n^2 p_{\mathrm{obs}} L$ for reliable top-K identification is given by

$$n^2 p_{\mathrm{obs}} L \gtrsim n \log n \,/\, \Delta_K^2.$$

In comparison, when adaptive sampling is employed for the preference-based model, the most recent upper bound on the sample complexity (e.g. Theorem 1 of (Busa-Fekete et al., 2013)) is about the order of

$$\sum\nolimits_{i=1}^{n-1} \Delta_i^{-2} \log n.$$

In the challenging regime where a dominant fraction of consecutive pairs are minimally separated (e.g. $\Delta_1 = \cdots = \Delta_{n-1}$), the above results seem to suggest that active ranking may not outperform passive ranking. For the other extreme case where only a single pair is minimally separated (e.g. $\Delta_1 \ll \Delta_i$ $(i \geq 2)$), active ranking is more desirable as it will adaptively acquire more paired evaluation over the minimally separated items.

4. Ranking Scheme: Spectral Method Meets MLE

This section presents a nearly linear-time algorithm that attempts recovery of the top-K ranked items. The algorithm proceeds in two stages: (1) an appropriate initialization that concentrates around the ground truth in an ℓ_2 sense, which can be obtained via a spectral ranking method; (2) a sequence of iterative updates sharpening the estimates in a pointwise manner, which consists in computing coordinate-wise MLE solutions. The two stages operate upon different sets of samples, while *no further sample splitting* is needed within each stage. The combination of these two stages will be referred to as *Spectral MLE*.

Before continuing to describe the details of our algorithm, we introduce a few notations that will be used throughout.

- $\mathcal{L}(\boldsymbol{w}; \boldsymbol{y}_i)$: the likelihood function of a latent preference vector $\boldsymbol{w}$, given the part of comparisons $\boldsymbol{y}_i$ that have bearing on item i.
- $\boldsymbol{w}_{\setminus i}$: for any preference vector $\boldsymbol{w}$, let $\boldsymbol{w}_{\setminus i}$ represent $[w_1, \cdots, w_{i-1}, w_{i+1}, \cdots, w_n]$ excluding w_i.
- $\mathcal{L}\left(\tau, \boldsymbol{w}_{\setminus i}; \boldsymbol{y}_i\right)$: with a slight abuse of notation, denote by $\mathcal{L}\left(\tau, \boldsymbol{w}_{\setminus i}; \boldsymbol{y}_i\right)$ the likelihood of the preference vector $[w_1, \cdots, w_{i-1}, \tau, w_{i+1}, \cdots, w_n]$.

4.1. Algorithm: Spectral MLE

It has been established that the spectral ranking method, particularly *Rank Centrality*, is able to discover a preference vector $\hat{\boldsymbol{w}}$ that incurs minimal ℓ_2 loss. To enable reliable ranking, however, it is more desirable to obtain an estimate that is faithful in an elementwise sense. Fortunately, the solution returned by the spectral method will serve as an ideal initial guess to seed our algorithm. The two components of the proposed Spectral MLE are described below.

1. **Initialization via Spectral Ranking**. We generate an initialization $\boldsymbol{w}^{(0)}$ via Rank Centrality. In words, Rank Centrality proceeds by constructing a Markov chain based on the pairwise observations, and then returning its stationary distribution by computing the leading eigenvector of the associated probability transition matrix. Under the Erdős–Rényi model, the estimate $\boldsymbol{w}^{(0)}$ is reasonably faithful in terms of the mean squared loss (Negahban et al., 2012), that is, with high probability,

$$\|\boldsymbol{w}^{(0)} - \boldsymbol{w}\| / \|\boldsymbol{w}\| \lesssim \sqrt{\log n / (n p_{\mathrm{obs}} L)}.$$

2. **Successive Refinement via Coordinate-wise MLE.** Note that the state-of-the-art finite-sample analyses for MLE (e.g. (Negahban et al., 2012)) involve only the ℓ_2 accuracy of the global MLE when the locations of all samples are i.i.d. (rather than the graph-based model considered herein). Instead of seeking a global MLE solution, we propose to carefully utilize the coordinate-wise MLE. Specifically, we cyclically iterate through each component, one at a time, maximizing the log-likelihood function with respect to that component. In contrast to the coordinate-descent method for solving the global MLE, we replace the preceding estimate with the new coordinate-wise MLE only when they are far apart. Theorem 4 (to be stated in Section 4.2) guarantees the contraction of the pointwise error for each cycle, which leads to a geometric convergence rate.

The algorithm then returns the indices of top-K items in accordance to the score estimate. A formal and more detailed description of the procedure is summarized in Algorithm 1.

Remark 2. *We split $\mathcal{E}$ into $\mathcal{E}^{\mathrm{init}}$ and $\mathcal{E}^{\mathrm{iter}}$ for analytical convenience. Empirically, if we keep $\mathcal{E}^{\mathrm{init}} = \mathcal{E}^{\mathrm{iter}} = \mathcal{E}$*

Algorithm 1 Spectral MLE.

Input: The average comparison outcome $y_{i,j}$ for all $(i,j) \in \mathcal{E}$; the score range $[w_{\min}, w_{\max}]$.

Partition $\mathcal{E}$ randomly into two sets $\mathcal{E}^{\text{init}}$ and $\mathcal{E}^{\text{iter}}$ each containing $\frac{1}{2}|\mathcal{E}|$ edges. Denote by $\boldsymbol{y}_i^{\text{init}}$ (resp. $\boldsymbol{y}_i^{\text{iter}}$) the components of $\boldsymbol{y}_i$ obtained over $\mathcal{E}^{\text{init}}$ (resp. $\mathcal{E}^{\text{iter}}$).

Initialize $\boldsymbol{w}^{(0)}$ to be the estimate computed by *Rank Centrality* on $\boldsymbol{y}_i^{\text{init}}$ $(1 \le i \le n)$.

Successive Refinement: for $t = 0 : T$ **do**

1) Compute the coordinate-wise MLE

$$w_i^{\text{mle}} \leftarrow \arg \max_{\tau \in [w_{\min}, w_{\max}]} \mathcal{L}\left(\tau, \boldsymbol{w}_{\setminus i}^{(t)}; \boldsymbol{y}_i^{\text{iter}}\right). \tag{12}$$

2) For each $1 \le i \le n$, set

$$w_i^{(t+1)} \leftarrow \begin{cases} w_i^{\text{mle}}, & \text{if } \left|w_i^{\text{mle}} - w_i^{(t)}\right| > \xi_t, \\ w_i^{(t)}, & \text{else.} \end{cases} \tag{13}$$

Output the indices of the K largest components of $\boldsymbol{w}^{(T)}$.

and reuse all samples, then it will outperform the graph-splitting procedure (but not by much). Thus, we recommend the sample-reuse procedure for practical use.

Remark 3. *Spectral MLE is inspired by recent advances in solving non-convex programs by means of iterative methods (Keshavan et al., 2010; Candes et al., 2015; Chen & Candes, 2015b; Jain et al., 2013; Netrapalli et al., 2013; Balakrishnan et al., 2014). A key message conveyed from these works is: once we arrive at an appropriate initialization (often via a spectral method), the iterative estimates will often be rapidly attracted towards the global optimum.*

Remark 4. *While our analysis is restricted to $\mathcal{G}_{n,p_{\text{obs}}}$, Spectral MLE can be applied to general graphs. We caution, however, that spectral ranking is not guaranteed to achieve minimal ℓ_2 loss, particularly for those graphs with small spectral gaps. Therefore, Spectral MLE is not necessarily minimax optimal under general graph patterns.*

Notably, the successive refinement stage is based on the observation that we are able to characterize the confidence intervals of the coordinate-wise MLEs at each iteration. Such confidence intervals allow us to detect outlier components that incur large pointwise loss. Since the initial guess is optimal in an overall ℓ_2 sense, a large fraction of its entries are already faithful compared to the ground truth. As a result, it suffices to disentangle the set of "sparse" outliers.

One appealing feature of Spectral MLE is its low computational complexity. Recall that the initialization step by Rank Centrality can be solved for ϵ accuracy within $O\left(|\mathcal{E}| \log \frac{1}{\epsilon}\right)$ time instances by means of a power method. In addition, for each component i, the coordinate-wise likelihood function involves a sum of $\deg(i)$ terms. Since finding the coordinate-wise MLE (12) can be cast as an one-dimensional convex program, one can get ϵ accuracy via a bisection method within $O\left(\deg(i) \cdot \log \frac{1}{\epsilon}\right)$ time. Therefore, each iteration cycle of the successive refinement stage can be accomplished in time $O\left(|\mathcal{E}| \cdot \log \frac{1}{\epsilon}\right)$.

The following theorem establishes the ranking accuracy of Spectral MLE under the BTL model.

Theorem 3. *Let $c_0, c_1, c_2, c_3 > 0$ be some universal constants. Suppose that $L = O(\text{poly}(n))$, the comparison graph $\mathcal{G} \sim \mathcal{G}_{n,p_{\text{obs}}}$ with $p_{\text{obs}} > c_0 \log n/n$, and assume that the separation measure (6) satisfies*

$$\Delta_K \;>\; c_1 \sqrt{\log n \,/\, (np_{\text{obs}}L)}. \tag{14}$$

Then with probability exceeding $1 - 1/n^2$, Spectral MLE perfectly identifies the set of top-K ranked items, provided that the parameters obey $T \ge c_2 \log n$ and

$$\xi_t := c_3 \left\{ \xi_{\min} + \frac{1}{2^t}\left(\xi_{\max} - \xi_{\min}\right) \right\}, \tag{15}$$

where $\xi_{\min} := \sqrt{\frac{\log n}{np_{\text{obs}}L}}$ and $\xi_{\max} := \sqrt{\frac{\log n}{p_{\text{obs}}L}}$.

Theorem 3 basically implies that the proposed algorithm succeeds in separating out the high-ranking objects with high probability, as long as the preference score satisfies the separation condition

$$\Delta_K \;\gtrsim\; \sqrt{\log n \,/\, (np_{\text{obs}}L)}.$$

Additionally, Theorem 3 asserts that the number of iteration cycles required in the second stage scales at most logarithmically, revealing that Spectral MLE achieves the desired ranking precision with nearly linear time complexity.

4.2. Successive Refinement: Convergence and Contraction of ℓ_∞ Error

In the sequel, we would like to provide some interpretation as to why we expect the pointwise error of the score estimates to be controllable. The argument is heuristic in nature, since we will assume for simplicity that each iteration employs a fresh set of samples $\boldsymbol{y}$ independent from the present estimate $\boldsymbol{w}^{(t)}$.

Denote by $\ell^*(\tau)$ the true log-likelihood function

$$\ell^*(\tau) := \frac{1}{L}\log\mathcal{L}\left(\tau, \boldsymbol{w}_{\setminus i}; \boldsymbol{y}_i\right). \tag{16}$$

One can easily verify that its expectation around w_i can be controlled through a locally strongly-concave function, due to the existence of a second-order lower bound

$$\left|\mathbb{E}_{\boldsymbol{w}}\left[\ell^*(w_i) - \ell^*(\tau)\right]\right| \gtrsim |\tau - w_i|^2 np_{\mathrm{obs}}. \tag{17}$$

This measures the penalty when τ deviates from the ground truth. Note, however, that we don't have direct access to $\ell^*(\cdot)$ since it relies on the latent scores $\boldsymbol{w}$. To obtain a computable surrogate, we replace $\boldsymbol{w}$ with the present estimate $\boldsymbol{w}^{(t)}$, resulting in the plug-in likelihood function

$$\hat{\ell}_i(\tau) := \frac{1}{L}\log\mathcal{L}\left(\tau, \boldsymbol{w}_{\setminus i}^{(t)}; \boldsymbol{y}_i\right).$$

Fortunately, the surrogate loss incurred by employing $\hat{\ell}_i(\tau)$ is locally stable in the sense that

$$\begin{aligned}&\left|\mathbb{E}_{\boldsymbol{w}}\left[\hat{\ell}_i(\tau) - \hat{\ell}_i(w_i) - \left(\ell^*(\tau) - \ell^*(w_i)\right)\right]\right| \\ &\quad\lesssim np_{\mathrm{obs}}|\tau - w_i|\frac{\|\hat{\boldsymbol{w}} - \boldsymbol{w}\|}{\|\boldsymbol{w}\|}.\end{aligned} \tag{18}$$

As a result, any candidate $\tau \neq w_i$ will be viewed as *less likely* than and hence distinguishable from the ground truth w_i (i.e. $\hat{\ell}(w_i) > \hat{\ell}(\tau)$) in the mean sense, provided that its deviation penalty (17) dominates the surrogate loss (18), namely,

$$|\tau - w_i| \;\gtrsim\; \|\hat{\boldsymbol{w}} - \boldsymbol{w}\| \,/\, \|\boldsymbol{w}\|.$$

Thus, if the aforementioned likelihood functions concentrate around their means, then our procedure should be able to converge to a solution whose pointwise error is as low as the normalized ℓ_2 error of the initial guess.

Encouragingly, the ℓ_∞ estimation error not only converges, but converges at a geometric rate as well. This rapid convergence property does not rely on the "fresh-sample" assumption imposed in the above heuristic argument, as formally stated in the following theorem.

Theorem 4. *Suppose that $\mathcal{G} \sim \mathcal{G}_{n,p_{\mathrm{obs}}}$ with $p_{\mathrm{obs}} > c_0 \log n/n$ for some large constant c_0, and there exists a score vector $\hat{\boldsymbol{w}}^{\mathrm{ub}} \in [w_{\min}, w_{\max}]^n$ independent of $\mathcal{G}$ satisfying*

$$\left|\hat{w}_i^{\mathrm{ub}} - w_i\right| \;\le\; \xi w_{\max}, \quad 1 \le i \le n; \tag{19}$$

$$\|\hat{\boldsymbol{w}}^{\mathrm{ub}} - \boldsymbol{w}\| \le\; \delta\|\boldsymbol{w}\|. \tag{20}$$

Then with probability at least $1 - c_1 n^{-4}$ for some constant $c_1 > 0$, the coordinate-wise MLE

$$w_i^{\mathrm{mle}} := \arg\max_{\tau\in[w_{\min}, w_{\max}]} \mathcal{L}\left(\tau, \hat{\boldsymbol{w}}_{\setminus i}; \boldsymbol{y}_i\right) \tag{21}$$

satisfies

$$\begin{aligned}\left|w_i - w_i^{\mathrm{mle}}\right| < \frac{20\left(6 + \log L/\log n\right)w_{\max}^5}{w_{\min}^4}\cdot \\ \max\left\{\delta + \frac{\xi\log n}{np_{\mathrm{obs}}},\ \sqrt{\frac{\log n}{np_{\mathrm{obs}}L}}\right\}\end{aligned} \tag{22}$$

simultaneously for all scores $\hat{\boldsymbol{w}} \in [w_{\min}, w_{\max}]^n$ obeying $|\hat{w}_i - w_i| \le \left|\hat{w}_i^{\mathrm{ub}} - w_i\right|$, $1 \le i \le n$.

In the regime where $L = O(\mathrm{poly}(n))$ and $\delta \asymp \sqrt{\frac{\log n}{np_{\mathrm{obs}}L}}$, Theorem 4 asserts that given an appropriate initialization, the coordinate-wise MLE procedure is guaranteed to drag down the elementwise estimation error at a rate

$$\left\|\boldsymbol{w}^{(t+1)} - \boldsymbol{w}\right\|_\infty \lesssim \frac{\left\|\boldsymbol{w}^{(t)} - \boldsymbol{w}\right\|}{\|\boldsymbol{w}\|} + \frac{\log n}{np_{\mathrm{obs}}}\left\|\boldsymbol{w}^{(t)} - \boldsymbol{w}\right\|_\infty.$$

The same collection of samples can be *reused* across all iterations at the successive refinement stage, provided that we can identify in each cycle another slightly looser estimate that is independent from the samples. From Theorem 4, the pointwise estimation error will converge to

$$\left\|\boldsymbol{w} - \boldsymbol{w}^{(t+1)}\right\|_\infty \lesssim \sqrt{\log n/(np_{\mathrm{obs}}L)},$$

which is minimally apart from the ground truth.

4.3. Discussion

Choice of Initialization. Careful readers will remark that the success of Spectral MLE can be guaranteed by a broader selection of initialization procedures beyond Rank Centrality. Indeed, Theorem 4 and subsequent analyses lead to the following assertion: as long as the initialization method is able to produce an initial estimate $\boldsymbol{w}^{(0)}$ that is reasonably faithful in the ℓ_2 sense

$$\|\boldsymbol{w}^{(0)} - \boldsymbol{w}\|/\|\boldsymbol{w}\| \lesssim \sqrt{\log n/(np_{\mathrm{obs}}L)}, \tag{23}$$

then Spectral MLE will converge to a pointwise optimal preference $\boldsymbol{w}^{(T)}$ obeying

$$\|\boldsymbol{w}^{(T)} - \boldsymbol{w}\|_\infty \lesssim \sqrt{\log n \,/\, (np_{\mathrm{obs}}L)}.$$

Initialization via Global MLE. One would naturally wonder whether we can employ the global MLE (computed over $\boldsymbol{y}^{\mathrm{init}}$) to seed the iterative refinement stage (applied over $\boldsymbol{y}^{\mathrm{iter}}$). In fact, the state-of-the-art analysis (with a different but order-wise equivalent model) (Negahban et al., 2012) asserts that the global MLE satisfies the desired ℓ_2 property (23) for at least two cases: (a) complete graphs, i.e. $p_{\mathrm{obs}} = 1$, and (b) Erdős–Rnyi graphs with (almost) no repeated comparisons, i.e. $L = 1$. In these two cases, the proposed algorithm achieves minimal ℓ_∞ errors if we initialize it via the global MLE.

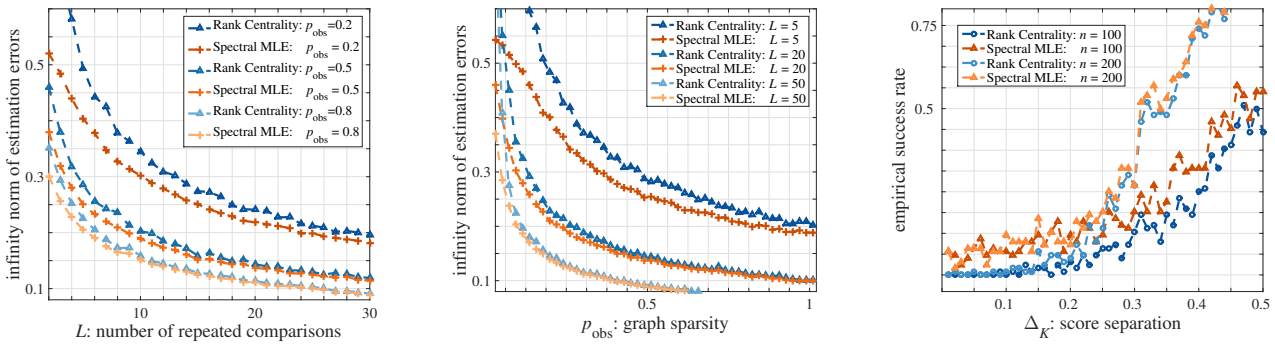

Figure 1. (Left) Empirical ℓ_∞ loss v.s. L; (Middle) ℓ_∞ loss v.s. $p_{\rm obs}$; (Right) Rate of success in top-K identification ($n = 100, 200$).

Nevertheless, whether the global MLE achieves minimal ℓ_2 loss for other configurations $(L, p_{\rm obs})$ has not been established. The analytical bottleneck seems to stem from an underlying bias-variance tradeoff when accounting for two successive randomness mechanisms: the random graph $\mathcal{G}$ and the repeated comparisons generated over $\mathcal{G}$. In general, $y_{i,j}^{(l)}$'s are not jointly independent unless we condition on $\mathcal{G}$. In contrast, the above two special cases amount to two extreme situations: (a) the randomness of $\mathcal{G}$ goes away when $p_{\rm obs} = 1$; (b) the condition $L = 1$ avoids repeated sampling. Nevertheless, these two cases alone (as well as the model in Theorem 4 of (Negahban et al., 2012)) are not sufficient in characterizing the complete tradeoff between graph sparsity and the quality of the acquired comparisons.

4.4. Numerical Experiments

A series of synthetic experiments is conducted to demonstrate the practical applicability of Spectral MLE. The important implementation parameters in our approach is the choice of c_2 and c_3 given in Theorem 3, which specify T and ξ_t. In all numerical simulations performed here, we pick $c_2 = 5$ and $c_3 = 1$, and do not split samples. We focus on the case where $n = 100$, where each reported result is calculated by averaging over 200 Monte Carlo trials.

We first examine the ℓ_∞ error of the score estimates. The latent scores are generated uniformly over $[0.5, 1]$. For each $(p_{\rm obs}, L)$, the paired comparisons are randomly generated as per the BTL model, and we perform score inference by means of both Rank Centrality and Spectral MLE. Fig. 1(a) (resp. Fig. 1(b)) illustrates the empirical tradeoff between the pointwise score estimation accuracy and the number L of repeated comparisons (resp. graph sparsity $p_{\rm obs}$). It can be seen from these plots that the proposed Spectral MLE outperforms Rank Centrality uniformly over all configurations, corroborating our theoretical results. Interestingly, the performance gain is the most significant under sparse graphs in the presence of low-resolution comparisons (i.e. when $p_{\rm obs}$ and L are small).

Next, we study the success rate of top-K identification as the number n of items varies. We generate the latent scores randomly over $[0.5, 1]$, except that a separation Δ_K is imposed between items K and $K + 1$. The results are shown in Fig. 1(c) for the case where $p_{\rm obs} = 0.2$, and $L = 5$. As can be seen, Spectral MLE achieves higher ranking accuracy compared to Rank Centrality for all these situations. Interestingly, the benefit of Spectral MLE relative to Rank Centrality is more apparent in the regime where the score separation is small. In addition, it seems that Rank Centrality is capable of achieving good ranking accuracy in the randomized model we simulate, and we leave the theoretical analysis for future work.

5. Conclusion

This paper investigates rank aggregation from pairwise data that emphasizes the top-K items. We developed a nearly linear-time algorithm that performs as well as the best model aware paradigm, from a minimax perspective.

This paper comes with some limitations in developing tight sample complexity bounds under general graphs. The performances of Spectral MLE under other sampling models are worth investigating (Osting et al., 2015). In addition, it remains to characterize both statistical and computational limits for other choice models (e.g. the Plackett-Luce model (Hajek et al., 2014)). It would also be interesting to consider the case where the paired comparisons are drawn from a mixture of BTL models (e.g. (Oh & Shah, 2014)), as well as the collaborative ranking setting where one aggregates the item preferences from a pool of different users in order to infer rankings for each individual user (e.g. (Lu & Negahban, 2014; Park et al., 2015))

Acknowledgements

C. Suh was partly supported by the ICT R&D program of MSIP/IITP [B0101-15-1272]. Y. Chen would like to thank Yuejie Chi for helpful discussion.

References

Agresti, A. *Categorical data analysis*. John Wiley & Sons, 2014.

Ailon, N. Active learning ranking from pairwise preferences with almost optimal query complexity. *Journal of Machine Learning Research*, 13:137–164, 2012.

Ammar, A. and Shah, D. Ranking: Compare, don't score. In *Allerton Conference*, pp. 776–783. IEEE, 2011.

Ammar, A. and Shah, D. Efficient rank aggregation using partial data. In *SIGMETRICS*, volume 40, pp. 355–366. ACM, 2012.

Awasthi, P., Blum, A., Sheffet, O., and Vijayaraghavan, A. Learning mixtures of ranking models. In *Neural Information Processing Systems*, pp. 2609–2617, 2014.

Balakrishnan, S., Wainwright, M. J., and Yu, B. Statistical guarantees for the EM algorithm: From population to sample-based analysis. *arXiv:1408.2156*, 2014.

Baltrunas, L., Makcinskas, T., and Ricci, F. Group recommendations with rank aggregation and collaborative filtering. In *ACM conference on Recommender systems*, pp. 119–126. ACM, 2010.

Bradley, R. A. and Terry, M. E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, pp. 324–345, 1952.

Braverman, M. and Mossel, E. Noisy sorting without resampling. In *ACM-SIAM SODA*, pp. 268–276, 2008.

Brin, S. and Page, L. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1):107–117, 1998.

Bubeck, S., Wang, T., and Viswanathan, N. Multiple identifications in multi-armed bandits. *ICML*, 2013.

Busa-Fekete, R. and Hüllermeier, E. A survey of preference-based online learning with bandit algorithms. In *Algorithmic Learning Theory*, pp. 18–39, 2014.

Busa-Fekete, R., Szörényi, B., Weng, P., Cheng, W., and Hüllermeier, E. Top-k selection based on adaptive sampling of noisy preferences. In *ICML*, 2013.

Busa-Fekete, R., Hüllermeier, E., and Szörényi, B. Preference-based rank elicitation using statistical models: The case of Mallows. In *ICML*, pp. 1071–1079, 2014.

Candes, Emmanuel, Li, Xiaodong, and Soltanolkotabi, Mahdi. Phase retrieval via Wirtinger flow: Theory and algorithms. *to appear, IEEE Trans. on Inf. Theory*, 2015.

Caplin, A. and Nalebuff, B. Aggregation and social choice: a mean voter theorem. *Econometrica*, pp. 1–23, 1991.

Chen, X., Bennett, P. N., Collins-Thompson, K., and Horvitz, E. Pairwise ranking aggregation in a crowd-sourced setting. In *WSDM*, pp. 193–202, 2013.

Chen, Y., Suh, C., and Goldsmith, A. J. Information recovery from pairwise measurements: A Shannon-theoretic approach. *arXiv preprint arXiv:1504.01369*, 2015a.

Chen, Yuxin and Candes, Emmanuel J. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *arXiv preprint arxiv:1505.05114*, 2015b.

Dwork, C., Kumar, R., Naor, M., and Sivakumar, D. Rank aggregation methods for the web. In *Proceedings of the Tenth International World Wide Web Conference*, 2001.

Eriksson, B. Learning to top-k search using pairwise comparisons. In *AISTATS*, pp. 265–273, 2013.

Farnoud, F. and Milenkovic, O. An axiomatic approach to constructing distances for rank comparison and aggregation. *IEEE Trans on Inf Theory*, 60(10):6417–6439, 2014.

Farnoud, F., Yaakobi, E., and Bruck, J. Approximate sorting of data streams with limited storage. In *Computing and Combinatorics*, pp. 465–476. Springer, 2014.

Ford, L. R. Solution of a ranking problem from binary comparisons. *American Mathematical Monthly*, 1957.

Gabillon, V., Ghavamzadeh, M., Lazaric, A., and Bubeck, S. Multi-bandit best arm identification. In *Neural Information Processing Systems*, pp. 2222–2230, 2011.

Hajek, B., Oh, S., and Xu, J. Minimax-optimal inference from partial rankings. In *NIPS*, pp. 1475–1483, 2014.

Harrington, E. F. Online ranking/collaborative filtering using the perceptron algorithm. In *ICML*, volume 20, pp. 250–257, 2003.

Hunter, D. R. MM algorithms for generalized Bradley-Terry models. *Annals of Statistics*, pp. 384–406, 2004.

Jain, P., Netrapalli, P., and Sanghavi, S. Low-rank matrix completion using alternating minimization. In *Symposium on Theory of computing*, pp. 665–674. ACM, 2013.

Jamieson, K., Malloy, M., Nowak, R., and Bubeck, S. lil'UCB: An optimal exploration algorithm for multi-armed bandits. *Conference on Learning Theory*, 2014.

Jamieson, K. G. and Nowak, R. Active ranking using pairwise comparisons. In *Advances in Neural Information Processing Systems*, pp. 2240–2248, 2011.

Keshavan, R. H., Montanari, A., and Oh, S. Matrix completion from a few entries. 56(6):2980–2998, June 2010.

Lu, T. and Boutilier, C. Learning Mallows models with pairwise preferences. In *ICML*, pp. 145–152, 2011.

Lu, Y. and Negahban, S. N. Individualized rank aggregation using nuclear norm regularization. *arXiv preprint arXiv:1410.0860*, 2014.

Luce, R. D. *Individual choice behavior: A theoretical analysis*. Wiley, 1959.

Negahban, S., Oh, S., and Shah, D. Rank centrality: Ranking from pair-wise comparisons. 2012. URL http://arxiv.org/abs/1209.1688.

Netrapalli, P., Jain, P., and Sanghavi, S. Phase retrieval using alternating minimization. In *Advances in Neural Information Processing Systems*, pp. 2796–2804, 2013.

Oh, S. and Shah, D. Learning mixed multinomial logit model from ordinal data. In *Neural Information Processing Systems*, pp. 595–603, 2014.

Osting, B., Xiong, J., Xu, Q., and Yao, Y. Analysis of crowdsourced sampling strategies for hodgerank with sparse random graphs. *arXiv:1503.00164*, 2015.

Park, D., Neeman, J., Zhang, J., Sanghavi, S., and Dhillon, I. S. Preference completion: Large-scale collaborative ranking from pairwise comparisons. 2015.

Rajkumar, A. and Agarwal, S. A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In *ICML*, pp. 118–126, 2014.

Shah, N. B., Balakrishnan, S., Bradley, J., Parekh, A., Ramchandran, K., and Wainwright, M. When is it better to compare than to score? *arXiv:1406.6618*, 2014.

Soufiani, H. A., Chen, W. Z., Parkes, D. C., and Xia, L. Generalized method-of-moments for rank aggregation. In *Neural Information Processing Systems*, 2013.

Soufiani, H. A., Parkes, D., and Xia, L. Computing parametric ranking models via rank-breaking. In *International Conference on Machine Learning (ICML)*, 2014a.

Soufiani, H. A., Parkes, D. C., and Xia, L. A statistical decision-theoretic framework for social choice. In *Neural Information Processing Systems*, 2014b.

Wauthier, F., Jordan, M., and Jojic, N. Efficient ranking from pairwise comparisons. In *ICML*, pp. 109–117, 2013.