
Aggregating Incomplete and Noisy Rankings

Dimitris Fotakis

Alkis Kalavasis

Konstantinos Stavropoulos

National Technical University of Athens

Abstract

We consider the problem of learning the true ordering of a set of alternatives from largely incomplete and noisy rankings. We introduce a natural generalization of both the classical Mallows model of ranking distributions and the extensively studied model of noisy pairwise comparisons. Our *selective Mallows model* outputs a noisy ranking on any given subset of alternatives, based on an underlying Mallows distribution. Assuming a sequence of subsets where each pair of alternatives appears frequently enough, we obtain strong asymptotically tight upper and lower bounds on the sample complexity of learning the underlying complete ranking and the (identities and the) ranking of the top- k alternatives from selective Mallows rankings. Moreover, building on the work of (Braverman and Mossel, 2009), we show how to efficiently compute the maximum likelihood complete ranking from selective Mallows rankings.

1 Introduction

Aggregating a collection of (possibly noisy and incomplete) ranked preferences into a complete ranking over a set of alternatives is a fundamental and extensively studied problem with numerous applications. Ranking aggregation has received considerable research attention in several fields, for decades and from virtually all possible aspects.

Most relevant, Statistics investigates the properties of *ranking distributions*, which provide principled ways to generate noisy rankings from structural information about the alternatives' relative order. Best known among them are the distance-based model of Mallows

[1957] and the parametric models of Thurstone [1927], Smith [1950], Bradley and Terry [1952], Plackett [1975] and Luce [2012]. Moreover, Machine Learning and Statistical Learning Theory aim to develop (statistically and computationally) efficient ways of retrieving the true ordering of the alternatives from noisy (and possibly incomplete) rankings (see e.g., [Xia, 2019] and the references therein).

Virtually all previous work in the latter research direction assumes that the input is a collection of either complete rankings (chosen adversarially, e.g., Ailon et al. [2008], Kenyon-Mathieu and Schudy [2007], or drawn from an unknown ranking distribution, e.g., [Caragiannis et al., 2013, Busa-Fekete et al., 2019]), or outcomes of noisy pairwise comparisons (see e.g., [Feige et al., 1994, Mao et al., 2018a]). Due to a significant volume of relatively recent research, the computational and statistical complexity of determining the best ranking based on either complete rankings or pairwise comparisons are well understood.

However, in most modern applications of ranking aggregation, the input consists of incomplete rankings of more than two alternatives. E.g., think of e-commerce or media streaming services, with a huge collection of alternatives, which generate personalized recommendations based on rankings aggregated by user ratings (see also Hajek et al. [2014]). Most users are able to rank (by rating or reviewing) several alternatives, definitely much more than two, but it is not even a remote possibility that a user is familiar with the entire inventory (see also [Moreno-Centeno and Escobedo, 2016] for applications of incomplete rankings to ranked preference aggregation, and [Yildiz et al., 2020] on why incomplete rankings are preferable in practice).

Motivated by the virtual impossibility of having access to complete rankings in modern applications, we introduce the *selective Mallows model*, generalizing both the classical Mallows model of ranking distributions and the extensively studied model of noisy pairwise comparisons. Under the selective Mallows model, we investigate the statistical complexity of learning the central ranking and the (identities and the) ranking of the top- k alternatives, and the computational com-

plexity of maximum likelihood estimation.

1.1 The Selective Mallows Model

The *Mallows model* [Mallows, 1957] is a fundamental and extensively studied family of ranking distributions over the symmetric group \mathfrak{S}_n . A *Mallows distribution* $\mathcal{M}_{\pi_0, \beta}$ on a set of n alternatives is parameterized by the *central ranking* $\pi_0 \in \mathfrak{S}_n$ and the *spread parameter* $\beta > 0$. The probability of observing a ranking $\pi \in \mathfrak{S}_n$ is proportional to $\exp(-\beta d(\pi_0, \pi))$, where d is a notion of ranking distance of π_0 to π . In this work, we consider the number of discordant pairs, a.k.a. the Kendall tau distance, defined as $d_{KT}(\pi_0, \pi) = \sum_{i < j} \mathbb{1}\{(\pi_0(i) - \pi_0(j))(\pi(i) - \pi(j)) < 0\}$.

The problem of aggregating a collection $\pi_1, \dots, \pi_r \in \mathfrak{S}_n$ of complete rankings asks for the *median ranking* $\pi^* = \arg \min_{\sigma \in \mathfrak{S}_n} \sum_{j=1}^r d_{KT}(\sigma, \pi_j)$. Computing the median is equivalent to a weighted version of Feedback Arc Set in tournaments, which is NP-hard [Ailon et al., 2008] and admits a polynomial-time approximation scheme [Kenyon-Mathieu and Schudy, 2007]. If the rankings are independent samples from a Mallows distribution, the median coincides with the *maximum likelihood ranking* and can be computed efficiently with high probability [Braverman and Mossel, 2009].

The *selective Mallows model* provides a principled way of generating noisy rankings over any given subset of alternatives, based on an underlying Mallows distribution $\mathcal{M}_{\pi_0, \beta}$. Given a central ranking $\pi_0 \in \mathfrak{S}_n$, a spread parameter $\beta > 0$ and a selection sequence $\mathcal{S} = (S_1, \dots, S_r)$, where each $S_i \subseteq [n]$, the *selective Mallows distribution* $\mathcal{M}_{\pi_0, \beta}^{\mathcal{S}}$ assigns a probability of

$$\Pr[(\pi_1, \dots, \pi_r) | \pi_0, \beta, \mathcal{S}] = \prod_{i \in [r]} \frac{1}{Z(S_i, \beta)} e^{-\beta d_{KT}(\pi_0, \pi_i)}$$

to each incomplete ranking profile (π_1, \dots, π_r) . In the probability above, each π_i is a permutation of S_i , $d_{KT}(\pi_0, \pi_i)$ is the number of pairs in S_i ranked reversely in π_0 and π_i (which naturally generalizes the Kendall tau distance to incomplete rankings), and the normalization constants $Z(S_i, \beta)$ correspond to a Mallows distribution on alternatives S_i and depend only on $|S_i|$ and β . The samples π_1, \dots, π_r are independent, conditioned on the selection sequence \mathcal{S} . We refer to $\Pi = (\pi_1, \dots, \pi_r)$ as a sample profile of length r .

In a selective Mallows sample π_i , the probability that two alternatives in S_i are not ranked as in π_0 depends on their distance in the restriction of π_0 to S_i (instead of their original distance in π_0). E.g., if π_0 is the identity permutation and $S_i = \{1, n\}$, the probability that $\pi_i = (1, n)$ (resp. $\pi_i = (n, 1)$) is the i -th sample in Π is $1/(1 + e^{-\beta})$ (resp. $e^{-\beta}/(1 + e^{-\beta})$). Hence,

the selective Mallows model generalizes both the standard Mallows model (if each $S_i = [n]$) and the model of noisy pairwise comparisons (if \mathcal{S} consists entirely of sets S_i with $|S_i| = 2$). Moreover, using sets S_i of different cardinality, one can smoothly interpolate between complete rankings and pairwise comparisons.

The amount of information provided by a selective Mallows model $\mathcal{M}_{\pi_0, \beta}^{\mathcal{S}}$ about π_0 is quantified by how frequently different pairs of alternatives compete against each other in Π . We say that a selective Mallows model $\mathcal{M}_{\pi_0, \beta}^{\mathcal{S}}$ is *p-frequent*, for some $p \in (0, 1]$, if every pair of alternatives appears in at least a p fraction of the sets in \mathcal{S} (we assume that each pair appears together at least once in \mathcal{S}). E.g., for $p = 1$, we recover the standard Mallows model, while $p \approx 2/n^2$ corresponds to pairwise comparisons. The definition of (p -frequent) selective Mallows model can be naturally generalized to unbounded selection sequences \mathcal{S} , which however is beyond the scope of this work.

In this work, we investigate the statistical complexity of retrieving either the central ranking π_0 or its top- k ranking from p -frequent selective Mallows samples, and the computational complexity of finding a maximum likelihood ranking from a fixed number r of p -frequent samples. In *learning from incomplete rankings*, for any given $p, \beta, \varepsilon > 0$, we aim to upper and lower bound the least number of samples $r^*(p, \beta, \varepsilon)$ (resp. $r_k^*(p, \beta, \varepsilon)$) from a selective Mallows distribution $\mathcal{M}_{\pi_0, \beta}^{\mathcal{S}}$ required to learn π_0 (resp. the top- k ranking of π_0) with probability at least $1 - \varepsilon$, where \mathcal{S} is any p -frequent selection sequence. In *maximum likelihood estimation*, for any given $p, \beta, \varepsilon > 0$, given a sample profile Π of length r from a p -frequent selective Mallows distribution $\mathcal{M}_{\pi_0, \beta}^{\mathcal{S}}$, we aim to efficiently compute either a ranking that is at least as likely as π_0 , or even a maximum likelihood ranking π^* . The interesting regime for maximum likelihood estimation is when r is significantly smaller than $r^*(p, \beta, \varepsilon)$.

We shall note here that the p -frequent assumption can be replaced by a milder one, where each selection set is drawn independently from a given distribution over the subsets of $[n]$, such that the probability that any specific pair of alternatives appears in a sampled set is at least p . Although we focus (for simplicity) on the deterministic p -frequent assumption, we expect similar results to hold for the randomized case. For an detailed discussion of the randomized p -frequent assumption, we refer the reader to the supplementary material.

1.2 Contribution

On the conceptual side, we introduce the selective Mallows model, which allows for a smooth interpolation

between learning from noisy complete rankings and sorting from noisy pairwise comparisons. On the technical side, we practically settle the statistical complexity of learning the central ranking and the top- k ranking of a p -frequent selective Mallows model. Moreover, we show how to efficiently compute a maximum likelihood ranking from r selective samples.

We believe that a significant advantage of our work lies in the simplicity and the uniformity of our approach. Specifically, all our upper bounds are based on the so-called *positional estimator* (Algorithm 1), which ranks an alternative i before any other alternative j ranked after i in the majority of the samples. The positional estimator belongs to the class of pairwise majority consistent rules [Caragiannis et al., 2013].

Generalizing the (result and the) approach of Caragiannis et al. [2013, Theorem 3.6], we show (Theorem 1) that if $pr = O(\frac{\log(n/\varepsilon)}{(1-e^{-\beta})^2})$, the central ranking of a p -frequent selective Mallows model can be recovered with probability at least $1 - \varepsilon$. Namely, observing a logarithmic number of noisy comparisons per pair of alternatives suffices for determining their true order. Theorem 1 generalizes (and essentially matches) the best known bounds on the number of (passively chosen¹) comparisons required for sorting [Mao et al., 2018a].

Interestingly, we show that the above upper bound is practically tight. Specifically, Theorem 2 shows that for any $p \in (0, 1/2]$, unless $rp = \Omega(\log(n/\varepsilon)/\beta)$ noisy comparisons per pair of alternatives are observed, any estimator of the central ranking from p -frequent selective samples fails to recover π_0 with probability larger than ε . Hence, observing incomplete rankings with (possibly much) more than two alternatives may help in terms of the number of samples, but it does not improve the number of noisy comparisons per pair required to recover the true ordering of the alternatives.

In Section 3, we generalize the proof of Theorem 3 and show that the positional estimator smoothly (and uniformly wrt different alternatives) converges to the central ranking π_0 of a p -frequent selective Mallows model $\mathcal{M}_{\pi_0, \beta}^S$, as the number of samples r (and the number of noisy comparisons pr per pair) increase. Specifically, Theorem 3 shows that the positional estimator has a remarkable property of the *average position estimator* [Braverman and Mossel, 2008, Lemma 18]: as r increases, the position of any alternative i in the estimated ranking converges fast to $\pi_0(i)$, with high probability. Since we cannot use the average posi-

tion estimator, due to our incomplete rankings, where the positions of the alternatives are necessarily relative, we need to extend [Braverman and Mossel, 2008, Lemma 18] to the positional estimator.

Combining Theorem 1 and Theorem 3, we show (Section 4.2) that for any $k = \Omega(1/(p\beta))$, we can recover the identities and the true ranking of the top- k alternatives in π_0 , with probability at least $1 - \varepsilon$, given $r = O(\frac{\log(k/\varepsilon)}{p(1-e^{-\beta})^2} + \frac{\log(n/\varepsilon)}{p^2\beta k})$ p -frequent selective samples. The second term accounts for learning the identities of the top $O(k)$ elements in π_0 (Theorem 3), while the first term accounts for learning the true ranking of these $O(k)$ elements (Theorem 1). For sufficiently large k , the first term becomes dominant. Applying the approach of Theorem 2, we show that such a sample complexity is practically best possible.

Moreover, building on the approach of Braverman and Mossel [2008, Lemma 18] and exploiting Theorem 3, we show how to compute a maximum likelihood ranking (resp. a ranking that is at least as likely as π_0), given r samples of a p -frequent selective Mallows distribution $\mathcal{M}_{\pi_0, \beta}^S$, in time roughly $n^{O(1/(r\beta p^4))}$ (resp. $n^{O(1/(r\beta p^2))}$), with high probability (see also Theorem 4 for the exact running time). The interesting regime for maximum likelihood estimation is when r is much smaller than the sample complexity of learning π_0 in Theorem 1. Our result compares favorably against the results of Braverman and Mossel [2008] if pr is small. E.g., consider the extreme case where $pr = 1$ (i.e., each pair is compared once in Π). Then, for small values of β , the running time of [Braverman and Mossel, 2008, Theorem 8] becomes $n^{O(1/\beta^4)}$, while the running time of maximum likelihood estimation from p -frequent selective samples becomes $n^{O(1/(p^3\beta))}$. Thus, large incomplete rankings mitigate the difficulty of maximum likelihood estimation (compared against noisy pairwise comparisons), if $rp = \Theta(1)$, β is small, and $1/\beta$ is much smaller than $1/p$.

In the following, we provide the intuition and proof sketches for our main results. Due to lack of space, the full proofs of all our technical claims are deferred to the supplementary material.

Notation. We conclude this section with some additional notation required in the technical part of the paper. For any ranking π of some $S \subseteq [n]$ and any pair of alternatives $i, j \in S$, we let $i \succ_{\pi} j$ denote that i precedes j in π , i.e., that $\pi(i) < \pi(j)$. For any object B , we use the notation $B = B[\Pi]$ to denote that B depends on a sample profile Π . Moreover, for simplicity and brevity, we use the asymptotic notation O_{β} (or Ω_{β}) to hide polynomial terms in $1/\beta$.

¹If the algorithm can actively select which pairs of alternatives to compare, $O(n \log n)$ noisy comparisons suffice for sorting, e.g., Braverman and Mossel [2008], Feige et al. [1994].

1.3 Related Work

There has been a huge volume of research work on statistical models over rankings (see e.g., Fligner and Verducci [1993], Marden [1996], Xia [2019] and the references therein). The Mallows [1957] model plays a central role in the aforementioned literature. A significant part of this work concerns either extensions and generalizations of the Mallows model (see e.g., [Fligner and Verducci, 1986, Murphy and Martin, 2003, Lebanon and Lafferty, 2003], and also [Lebanon and Mao, 2008, Lu and Boutilier, 2014, Busa-Fekete et al., 2014] more closely related to partial rankings) or statistically and computationally efficient methods for recovering the parameters of Mallows distributions (see e.g., [Adkins and Fligner, 1998, Caragiannis et al., 2013, Liu and Moitra, 2018, Busa-Fekete et al., 2019]).

From a conceptual viewpoint, the work of Hajek et al. [2014] is closest to ours. Hajek et al. [2014] introduce a model of selective incomplete Thurstone and Plackett-Luce rankings, where the selection sequence consists of sets of k alternatives selected uniformly at random. They provide upper and lower bounds on how fast optimizing the log-likelihood function from incomplete rankings (which in their case is concave in the parameters of the model and can be optimized via e.g., gradient descent) converges to the model's true parameters. In our case, however, computing a maximum likelihood ranking is not convex and cannot be tackled with convex optimization methods. From a technical viewpoint, our work builds on previous work by Braverman and Mossel [2008], Caragiannis et al. [2013] and Busa-Fekete et al. [2019].

For almost three decades, there has been a significant interest in ranked preference aggregation and sorting from noisy pairwise comparisons. One branch of this research direction assumes that the algorithm actively selects which pair of alternatives to compare in each step and aims to minimize the number of comparisons required for sorting (see e.g., [Feige et al., 1994, Braverman and Mossel, 2008], or [Ailon, 2012] for sorting with few errors, or [Braverman et al., 2016] for parallel algorithms). A second branch, closest to our work, studies how many passively (see e.g., [Mao et al., 2018b,a]) or randomly (see e.g., [Wauthier et al., 2013]) selected noisy comparisons are required for ranked preference aggregation and sorting. A more general problem concerns the design of efficient approximation algorithms (based on either sorting algorithms or common voting rules) for aggregating certain types of incomplete rankings, such as top- k rankings, into a complete ranking (see e.g., Ailon [2010], Mathieu and Mauras [2020]). Moreover, there has been recent work on assigning ranking scores to the alternatives based on the results of noisy pairwise comparisons, by likeli-

hood maximization through either gradient descent or majorize-maximization (MM) methods (see e.g., Vojnovic et al. [2020]). Such works on learning from pairwise comparisons are also closely related to our work from a graph-theoretic viewpoint, since they naturally correspond to weighted graph topologies, whose properties (e.g., Fiedler eigenvalue of the comparison matrix [Hajek et al., 2014, Shah et al., 2016, Khetan and Oh, 2016, Vojnovic and Yun, 2016, Negahban et al., 2017, Vojnovic et al., 2020] or degree sequence [Pananjady et al., 2020]) characterize the sample complexity and convergence rate of various learning approaches. The comparison graph of p -frequent Mallows is the clique.

Another related line of research (which goes back at least to Conitzer and Sandholm [2005]) investigates how well popular voting rules (e.g., Borda count, Kemeny ranking, approval voting) behave as maximum likelihood estimators for either the complete central ranking of the alternatives, or the identities of the top- k alternatives, or the top alternative (a.k.a. the winner). In this line of work, the input may consist of complete or incomplete noisy rankings [Xia and Conitzer, 2011, Procaccia et al., 2012], the results of noisy pairwise comparisons [Shah and Wainwright, 2017], or noisy k -approval votes [Caragiannis and Micha, 2017].

2 Retrieving the Central Ranking

In this section, we settle the sample complexity of learning the central ranking π_0 under the selective Mallows model. We show that a practically optimal strategy is to neglect the concentration of alternatives' positions around their initial positions in π_0 and act as if the samples are a set of pairwise comparisons with common probability of error only depending on β .

Positional Estimator. Given a sample profile $\Pi = (\pi_1, \dots, \pi_r)$ corresponding to a selection sequence $\mathcal{S} = (S_1, \dots, S_r)$, we denote with $\hat{\pi} = \hat{\pi}[\Pi]$ the permutation of $[n]$ output by Algorithm 1.

Algorithm 1 Positional Estimator of profile Π

```

1: procedure POSEST( $\Pi$ )
2:    $\hat{\pi} \leftarrow \mathbf{0}_n$ 
3:   for  $i \in [n]$  do
4:     for  $j \in [n]$  do
5:        $\Pi_{i,j} \leftarrow \{\pi \in \Pi : i, j \in \pi\}$ 
6:        $n_{j \succ i} \leftarrow |\{\pi \in \Pi_{i,j} : j \succ_{\pi} i\}|$ 
7:       if  $n_{j \succ i} \geq |\Pi_{i,j}|/2$  then
8:          $\hat{\pi}(i) \leftarrow \hat{\pi}(i) + 1$ 
9:   Break ties in  $\hat{\pi}$  uniformly at random
10: return  $\hat{\pi}$ 
    
```

Algorithm 1 first calculates the pairwise majority posi-

tion of each alternative, by comparing each $i \in [n]$ with any other $j \in [n]$ in the joint subspace of the sample profile where i and j appear together. Intuitively, $\hat{\pi}(i)$ equals $|\{j : j \text{ ranked before } i \text{ most times}\}|$. The algorithm, after breaking ties uniformly at random, outputs $\hat{\pi}$. We call $\hat{\pi}$ the *positional estimator* (of the sample profile Π), which belongs to the class of pairwise majority consistent rules, introduced by Caragiannis et al. [2013].

Sample Complexity. Motivated by the algorithm of Caragiannis et al. [2013] for retrieving the central ranking from complete rankings, we utilize the POS-EST (Algorithm 1) and establish an upper bound on the sample complexity of learning the central ranking in the selective Mallows case.

Theorem 1. *For any $\epsilon \in (0, 1)$, $\pi_0 \in \mathfrak{S}_n$, $\beta > 0$, $p \in (0, 1]$, there exists an algorithm such that, given a sample profile from $\mathcal{M}_{\pi_0, \beta}^S$, where \mathcal{S} is a p -frequent selection sequence of length $O(\frac{1}{p(1-e^{-\beta})^2} \log(n/\epsilon))$, Algorithm 1 retrieves π_0 with probability at least $1 - \epsilon$.*

Proof (Sketch). If we have enough samples so that every pair of alternatives is ranked correctly in the majority of its comparisons, with probability at least $1 - \epsilon/n^2$, then, by union bound, all pairs are ranked correctly in the majority of their comparisons with probability at least $1 - \epsilon$, which, in turn, would imply the theorem. If the number of samples is r , then each pair is compared at least pr times in the sample. The Hoeffding bound implies that the probability that a pair is swapped in the majority of its appearances decays exponentially with pr . \square

In fact, the positional estimator is an optimal strategy with respect to the sample complexity of retrieving the central ranking. This stems from the fact that if for some pair, the total number of its comparisons in the sample is small, then there exists a family of possible central rankings where different alternatives cannot be easily ranked, due to lack of information. We continue with an essentially matching lower bound:

Theorem 2. *For any $p \in (0, 1]$, $\epsilon \in (0, 1/2]$, $\beta > 0$ and $r = o(\frac{1}{\beta p} \log(n/\epsilon))$ there exists a p -frequent selection sequence \mathcal{S} with $|\mathcal{S}| = r$, such that for any central ranking estimator, there exists a $\pi_0 \in \mathfrak{S}_n$, such that the estimator, given a sample profile from $\mathcal{M}_{\pi_0, \beta}^S$, fails to retrieve π_0 with probability more than ϵ .*

Proof (Sketch). Let \mathcal{S} contain $p|\mathcal{S}|$ sets with all alternatives and $(1-p)|\mathcal{S}|$ sets of size at most $n\sqrt{p/(1-p)}$. For any $i, j \in [n]$, let $W_{ij}(\mathcal{S})$ be the number of sets of \mathcal{S} containing both i and j , that is, the number of the

appearances of pair (i, j) . Clearly, \mathcal{S} is p -frequent and:

$$\sum_{i < j} W_{ij}(\mathcal{S}) \leq pn^2 |\mathcal{S}|. \quad (1)$$

Assume that $|\mathcal{S}| < \frac{1}{8p\beta} \log(\frac{n(1-\epsilon)}{4\epsilon})$. We will show that there exists a set of $n/2$ disjoint pairs of alternatives which we observe only a few times in the samples. Assume that n is even. We design a family $\{P_t\}_{t \in [n/2]}$ of perfect matchings on the set of alternatives $[n]$. Specifically, we consider $n/2$ sets of $n/2$ disjoint pairs $P_1 = \{(1, 2), (3, 4), \dots, (n-1, n)\}$, $P_2 = \{(1, 4), (3, 6), \dots, (n-1, 2)\}$ and, in general, $P_t = \{(1, (2t) \bmod n), \dots, (n-1, (2t+n-2) \bmod n)\}$ for $t \in [n/2]$.

Observe that no pair of alternatives appears in more than one perfect matching of the above family. Therefore:

$$\sum_{t \in [n/2]} \sum_{(i,j) \in P_t} W_{ij}(\mathcal{S}) \leq \sum_{i < j} W_{ij}(\mathcal{S}) \quad (2)$$

Combining (1), the bound for $|\mathcal{S}|$ and (2), we get that:

$$\exists t \in [n/2] : \sum_{(i,j) \in P_t} W_{ij}(\mathcal{S}) < \frac{n}{4\beta} \log\left(\frac{n(1-\epsilon)}{4\epsilon}\right)$$

Hence, since $|P_t| = [n/2]$, there exist at least $n/4$ pairs $(i, j) \in P_t$ with $W_{ij}(\mathcal{S}) < \frac{1}{\beta} \log\left(\frac{n(1-\epsilon)}{4\epsilon}\right)$.

We conclude the proof with an information-theoretic argument based on the observation that if the pairs of P_t , $n/4$ of which are observed few times, are adjacent in the central ranking, then the probability of swap is maximized for each pair. Moreover, the knowledge of the relative order of the elements in some pairs in the matching does not provide any information about the relative order of the elements in any of the remaining pairs. Intuitively, since each of $n/4$ pairs is observed only a few times, no central ranking estimator can be confident enough about the relative order of the elements in all these pairs. \square

3 Approximating the Central Ranking with Few Samples

We show next that the positional estimator smoothly approximates the position of each alternative in the central ranking, within an additive term that diminishes as the number of samples r increases.

The average position estimation approximates the positions of the alternatives under the Mallows model, as shown by Braverman and Mossel [2009]. However, the average position is not meaningful under the selective Mallows model, because the lengths of the selective ranking may vary.

Also, although under the Mallows model, the probability of displacement of an alternative decays exponentially in the length of the displacement, under the selective Mallows model, distant elements might be easily swapped in a sample containing only a small number of the alternatives that are ranked between them in the central ranking. For example, denoting with id the identity permutation ($1 \succ 2 \succ \dots \succ n$), if $m \in [n]$, $m \gg 1$, $\pi_1 \sim \mathcal{M}_{\text{id},\beta}^{\{[n]\}}$ and $\pi_2 \sim \mathcal{M}_{\text{id},\beta}^{\{\{1,m\}\}}$, then:

$$\Pr[m \succ_{\pi_1} 1] \leq 2e^{-\beta m/2},$$

since in order for 1 and m to be swapped, either has to be displaced at least $m/2$ places, and as shown by Bhatnagar and Peled [2015], under Mallows model, the probability of displacement of an alternative by t places is bounded by $2e^{-\beta t}$. However:

$$\Pr[m \succ_{\pi_2} 1] \geq e^{-\beta}/(1 + e^{-\beta}),$$

using the bound for swap probability provided by Chierichetti et al. [2014].

Since $m \gg 1$: $\Pr[m \succ_{\pi_1} 1] \ll \Pr[m \succ_{\pi_2} 1]$.

However, even though some selection sets may weaken the concentration property of the positions of the alternatives, we show that the positional estimator works under the selective Mallows model. This happens due to the requirement that each pair of alternatives should appear frequently: the majority of distant (in π_0) alternatives remain distant in the majority of the resulting incomplete rankings obtained by restricting π_0 to the selection sets. This is summarized by the following:

Theorem 3. *Let $\Pi \sim \mathcal{M}_{\pi_0,\beta}^{\mathcal{S}}$, where $\pi_0 \in \mathfrak{S}_n$, $\beta > 0$, $|\mathcal{S}| = r$ and \mathcal{S} is p -frequent, for some $p \in (0, 1]$, and $\epsilon \in (0, 1)$. Then, for the positional estimator $\hat{\pi} = \hat{\pi}[\Pi]$, there exists some $N = O(\frac{\beta^2+1}{\beta^3 p^2 r} \log(n/\epsilon))$ such that:*

$$\Pr[\exists i \in [n] : |\hat{\pi}(i) - \pi_0(i)| > N] \leq \epsilon$$

Proof (Sketch). We show that with high probability, for any alternative $i \in [n]$, only $N = O(\frac{1}{p}(\frac{1}{\beta} + \frac{1}{\beta p r} \log \frac{n}{\epsilon}))$ other alternatives j are ranked incorrectly relatively to i in the majority of the samples of Π where both i and j appear. Therefore, by the definition of the POEST, even after tie breaking, each alternative is ranked by the output ranking within an $O(N)$ margin from its original position in π_0 .

If two alternatives i, j are ranked L positions away by the reduced central ranking corresponding to a sample, then the probability that they appear swapped in the sample is at most $2e^{-\beta L/2}$. However, even if i, j are distant in π_0 , they might be ranked close by a reduced central ranking.

For any $i \in [n]$, we define a neighborhood $\mathcal{N}_i(L, \lambda)$ containing the other alternatives j which appear less than L positions away from i in the corresponding reduced central rankings of at least r/λ samples. Intuitively, those alternatives j outside $\mathcal{N}_i(L, \lambda)$ are far from i in the corresponding reduced central ranking of many samples. Hence, in these samples where i, j are initially far (according to L), the probability of observing them swapped is small enough so that, with high probability, the number of samples where they are ranked correctly is dominant among all the appearances of the pair, since we have additionally forced the number of samples where i, j are initially close (in which swaps are easy) to be small (according to λ).

Additionally, we bound the size of the neighborhood by $|\mathcal{N}_i(L, \lambda)| \leq 2L\lambda$, because in each sample there is only a small number of candidate neighbors (according to L) and an element of $\mathcal{N}_i(L, \lambda)$ uses many of the total number places. We conclude the proof by setting $L = O(\frac{1}{\beta} + \frac{1}{\beta p r} \log \frac{n}{\epsilon})$ and $\lambda = O(\frac{1}{p})$. Intuitively, λ is chosen so that the number of samples where swaps are difficult is comparable to pr (the minimum number of samples where each pair appears). The margin of error is $N = O(L\lambda)$. \square

Remark 1. *In Theorem 3, the margin of the approximation accuracy N can be refined as follows:*

$$N = \begin{cases} O(\frac{1}{\beta p^2 r} \log(n/\epsilon)), & \text{when } r = O(\frac{1}{p} \log(n/\epsilon)) \\ O(\frac{1}{\beta p}), & \text{when } r = \omega(\frac{1}{p} \log(n/\epsilon)) \\ 0, & \text{when } r = \omega(\frac{1}{\beta^2 p} \log(n/\epsilon)) \end{cases}$$

According to Remark 1, the error margin of approximation for the POEST provably diminishes inversely proportionally to $\beta p^2 r$, when r is sufficiently small and eventually becomes zero when r exceeds the sample complexity of Theorem 1.

4 Applications of Approximation

Assume we are given a sample from $\mathcal{M}_{\pi_0,\beta}^{\mathcal{S}}$, where \mathcal{S} is p -frequent for some $p \in (0, 1]$. Unless $|\mathcal{S}|$ is sufficiently large, we cannot find the central ranking with high probability. However, due to Theorem 3, we know that the positional estimator approximates the true positions of the alternatives within a small margin. This implies two possibilities which will be analyzed shortly: First, in Section 4.1, we present an algorithm that finds the maximum likelihood estimation of the central ranking with high probability. The algorithm is quite efficient, assuming that the frequency p and the spread parameter β are not too small. In Section 4.2, we show how to retrieve the top- k ranking (assuming we have enough samples to sort $O(k)$ elements), when k is sufficiently large.

4.1 Maximum Likelihood Estimation of the Central Ranking

We work in the regime where r is (typically much) smaller than the sample complexity of Theorem 1. We start with some necessary notation. For any $\mathcal{A} \subseteq \mathfrak{S}_n$, let $\pi_{\mathcal{A}}^* = \pi_{\mathcal{A}}^*[\Pi]$ be a *maximal likelihood* estimation of π_0 among elements of \mathcal{A} , that is:

$$\pi_{\mathcal{A}}^* \in \arg \max_{\pi \in \mathcal{A}} \Pr[\Pi | \pi, \beta, \mathcal{S}] \quad (3)$$

If $\mathcal{A} = \mathfrak{S}_n$, $\pi_{\mathcal{A}}^* = \pi^*$, is a *maximum likelihood* estimation of π_0 , while if $\pi_0 \in \mathcal{A}$, $\pi_{\mathcal{A}}^* = \pi^\circ$ is a *likelier than nature* estimation of π_0 [Rubinstein and Vardi, 2017].

The following lemma states that computing the maximum likelihood ranking (MLR) is equivalent to maximizing the total number of pairwise agreements (MPA) between the selected ranking and the samples.

Lemma 1. *Let \mathcal{A} be a subset of \mathfrak{S}_n and assume that we draw a sample profile $\Pi \sim \mathcal{M}_{\pi_0, \beta}^{\mathcal{S}}$. Consider the following two problems:*

$$\text{(MLR)} \quad \arg \max_{\pi \in \mathcal{A}} \Pr[\Pi | \pi, \beta, \mathcal{S}] \quad \text{and} \quad (4)$$

$$\text{(MPA)} \quad \arg \max_{\pi \in \mathcal{A}} \sum_{i \succ_{\pi} j} |\{\pi' \in \Pi : i \succ_{\pi'} j\}| \quad (5)$$

Then, the solutions of (MLR) and (MPA) coincide.

Proof. If $\Pi = (\pi_1, \dots, \pi_r) \sim \mathcal{M}_{\pi_0, \beta}^{\mathcal{S}}$, then, starting from (4), we get that:

$$\arg \max_{\pi \in \mathcal{A}} \Pr[\Pi | \pi, \beta, \mathcal{S}] = \arg \min_{\pi \in \mathcal{A}} \sum_{\ell \in [r]} d_{KT}(\pi, \pi_{\ell})$$

That is, maximizing the likelihood function is equivalent to minimizing the total number of pairwise disagreements. Equivalently, we have to maximize the total number of pairwise agreements and, hence, retrieving (5). Note that the samples in Π are incomplete and therefore each pair of alternatives is compared only in some of the samples. \square

Let us consider a subset \mathcal{A} of \mathfrak{S}_n . According to Lemma 1, there exists a function $f = f[\Pi] : [n] \times [n] \rightarrow \mathbb{N}$ such that solving (MLR) is equivalent to maximizing a score function $s : \mathcal{A} \rightarrow \mathbb{N}$ of the form:

$$s(\pi) = \sum_{i \succ_{\pi} j} f(i, j) \quad (6)$$

Then, as shown by Braverman and Mossel [2009], there exists a dynamic programming, which given an initial approximation of the maximizer of s , computes a ranking that maximizes s in time linear in n , but exponential in the error of the initial approximation. More specifically, Braverman and Mossel [2009] show that:

Lemma 2 (Braverman and Mossel [2009]). *Consider a function $f : [n] \times [n] \rightarrow \mathbb{N}$. Suppose that there exists an optimal ordering $\pi \in \mathfrak{S}_n$ that maximizes the score (6) such that $|\pi(i) - i| \leq R, \forall i \in [n]$. Then, there exists an algorithm which computes π in time $O(n \cdot R^2 \cdot 2^{6R})$.*

Recall that the positional estimator finds such an approximation $\hat{\pi}$ of the central ranking. Also, a careful examination of the proof of Lemma 2 shows that given any initial permutation $\sigma \in \mathfrak{S}_n$, the dynamic programming algorithm finds, in time $O(n \cdot R^2 \cdot 2^{6R})$, a maximizer (of the score function s) in $\mathcal{A} \subseteq \mathfrak{S}_n$, where \mathcal{A} contains all the permutations that are R -pointwise close² to the initial permutation σ . Therefore, we immediately get an algorithm that computes a likelier than nature estimation π° by finding $\pi_{\mathcal{A}}^*$, for \mathcal{A} such that $\hat{\pi}, \pi_0 \in \mathcal{A}$.

Furthermore, if π^* is an approximation of π_0 , then $\hat{\pi}$ is an approximation of π^* . Hence, we get an algorithm for computing a maximum likelihood estimation π^* . It turns out that π^* approximates π_0 , but with a larger margin of error. Thus, we obtain the following:

Theorem 4. *Let Π be a sample profile from $\mathcal{M}_{\pi_0, \beta}^{\mathcal{S}}$, where \mathcal{S} is a p -frequent selection sequence, $p \in (0, 1]$, $|\mathcal{S}| = r$, $\pi_0 \in \mathfrak{S}_n$, $\beta > 0$ and let $\alpha > 0$. Then:*

1. *There exists an algorithm that finds a likelier than nature estimation of π_0 with input Π with probability at least $1 - n^{-\alpha}$ and in time:*

$$T = O(n^2 + n^{1+O(\frac{2+\alpha}{r\beta p^2})} 2^{O(\frac{1}{p\beta})} \log^2 n)$$

2. *There exists an algorithm that finds a maximum likelihood estimation of π_0 with input Π with probability at least $1 - n^{-\alpha}$ and in time:*

$$T = O(n^2 + n^{1+O(\frac{2+\alpha}{r\beta p^4})} 2^{O(\frac{1}{p^3\beta})} \log^2 n)$$

To summarize the algorithm of Theorem 4, we note that it consists of two main parts. First, using the fact that our samples are drawn from a selective Mallows distribution, in which the positions of the alternatives exhibit some notion of locality, we approximate the central ranking within some error margin for the positions of alternatives. Second, beginning from the approximation we obtained at the previous step, we explore (using dynamic programming instead of exhaustive search, see Lemma 2) a subset of \mathfrak{S}_n which is reasonably small and provably contains with high probability either π_0 (for finding a likelier than nature ranking) or π^* (for finding a maximum likelihood one).

²We say that $\pi, \sigma \in \mathfrak{S}_n$ are R -pointwise close, if it holds that: $|\pi(i) - \sigma(i)| \leq R$ for all $i \in [n]$.

4.2 Retrieving the Top- k Ranking

In this section, we deal with the problem of retrieving the top- k sequence and their ranking in π_0 . In this problem, we are given a sample profile from a selective Mallows model and a positive integer k . We aim to compute the (identities and the) order of the top- k sequence in the central ranking π_0 .

Based on the properties of the positional estimator, it suffices to show that (after sufficiently many selective samples) any of the alternatives of the top- k sequence i in π_0 is ranked correctly with respect to any other alternative j in the majority of samples where both i and j appear. Then, every other alternative will be ranked after the top- k by the POSEST.

The claim above follows from the approximation property of the POSEST. Theorem 3 ensures that for any alternative of the top- k sequence i , the correct majority order against most other alternatives (those that are far from i in most reduced central rankings). So, we can focus on only $O(k)$ pairs, which could appear swapped with noticeable probability.

To formalize the intuition above, we can restrict our attention to the regime where the number of alternatives n is sufficiently large and $k = \omega(1/(p\beta))$. By Remark 1, this ensures that the accuracy of the approximation of POSEST diminishes inversely proportional to $\beta p^2 r$, since we only aim to ensure that the accuracy is $O(k)$. Specifically, Theorem 5 provides an upper bound on the estimation in that regime and Corollary 1 gives a general lower bound for the case where $k = O(n)$.

Theorem 5. *Let $k = \omega(1/(p\beta))$ be a positive integer. For any $\epsilon \in (0, 1)$ and $p \in (0, 1]$, there exists an algorithm which given a sample profile from $\mathcal{M}_{\pi_0, \beta}^S$, where \mathcal{S} is a p -frequent selection sequence with:*

$$|\mathcal{S}| = O\left(\frac{\log(k/\epsilon)}{p(1 - e^{-\beta})^2} + \frac{\log(n/\epsilon)}{p^2 \beta k}\right),$$

retrieves the top- k ranking of the alternatives of π_0 , with probability at least $1 - \epsilon$.

Proof (Sketch). Let $\Pi \sim \mathcal{M}_{\pi_0, \beta}^S$ be our sample profile. We will make use of the POSEST $\hat{\pi} = \hat{\pi}[\Pi]$ and, without loss of generality, assume that π_0 is the identity permutation. We will bound the probability that there exists some $i \in [k]$ such that $\hat{\pi}(i) \neq i$.

For any $i \in [k]$, we can partition the remaining alternatives into $A_1(i) = \mathcal{N}_i(L, \lambda)$ and $A_2(i) = [n] \setminus (A_1(i) \cup \{i\})$. From the proof sketch of Theorem 3, we recall that $\mathcal{N}_i(L, \lambda)$ contains the alternatives that are ranked no more than L places away from i in the reduced central rankings corresponding to at least r/λ samples.

From an intermediate result occurring during the proof of Theorem 3, it holds that for some L, λ such that $|A_1(i)| = O(\frac{1}{p\beta} + \frac{1}{p^2\beta|\mathcal{S}|} \log(n/\epsilon))$, with probability at least $1 - \epsilon/2$, for every $i \in [n]$, for every alternative $j \in A_2(i)$ (distant from i in most samples), j is ranked in the correct order relatively to i in the majority of the samples where they both appear.

Picking L, λ so that the above result holds, there exists some $r_1 = O(\frac{1}{p^2\beta k} \log(n/\epsilon))$ such that, if $|\mathcal{S}| \geq r_1$, then $|A_1(i)| = O(k)$.

Furthermore, following the same technique used to prove Theorem 1, we get that for some $r_2 = O(\frac{1}{p(1 - e^{-\beta})^2} \log(k/\epsilon))$, if $|\mathcal{S}| \geq r_2$, then, with probability at least $1 - \epsilon/2$, every pair of alternatives (i, j) such that $i \in [k]$ and $j \in A_1(i)$ is ranked correctly by the majority of samples where both i and j appear, since the total number of such pairs is $O(k^2)$.

Therefore, with probability at least $1 - \epsilon$, both events hold and for any fixed $i \in [k]$, $\hat{\pi}(i) = i$, because i is ranked correctly relatively to every other alternative in the majority of their pairwise appearances and also because for every other alternative $j > k$: $\hat{\pi}(j) > k$, since each of the alternatives in $[k]$ appear before it in the majority of samples where they both appear. \square

From a macroscopic and simplistic perspective, the sample complexity of learning the top- k ranking can be interpreted as follows. The first term, i.e., $O_\beta(\frac{1}{p} \log(k/\epsilon))$, accounts for learning the ranking of the top- k sequence (as well as some $O(k)$ other alternatives), since each of them is close to the others in the central ranking (and in each reduced rankings where they appear). Hence, it is probable that their pairs appear swapped. The second term, i.e., $O_\beta(\frac{1}{p^2 k} \log(n/\epsilon))$, accounts for identifying the top- k sequence, by approximating their positions. Intuitively, the required accuracy of the approximation diminishes to the value of k , since we aim to keep $O(k)$ probable swaps for each of the alternatives of the top- k sequence. Combining the two parts, we conclude that given enough samples, POSEST outputs a ranking where the top- k ranking coincides with the top- k ranking of π_0 .

We conclude with the lower bound, followed by a discussion about the tightness of our results.

Corollary 1. *For any $k \leq n$, $p \in (0, 1]$, $\epsilon \in (0, 1/2]$, $\beta > 0$ and $r = o(\frac{1}{\beta p} \log(k/\epsilon))$, there exists a p -frequent selection sequence \mathcal{S} with $|\mathcal{S}| = r$, such that for any central ranking estimator, there exists a $\pi_0 \in \mathfrak{S}_n$ such that the estimator, given a sample profile from $\mathcal{M}_{\pi_0, \beta}^S$, fails to retrieve the top- k ranking of π_0 with probability at least ϵ .*

Corollary 1 is an immediate consequence of Theorem 2. The bounds we provided in Theorem 5 and Corollary 1 become essentially tight if $k = \Omega(\frac{1}{p} \log(n/\epsilon))$, since the term $O_\beta(\frac{1}{p} \log(k/\epsilon))$ becomes dominant in the upper bound, which then essentially coincides with the lower bound. In the intuitive interpretation we provided for the two terms of the sample complexity in Theorem 5, this observation suggests that when k is sufficiently large, the sample complexity of identifying the top- k ranking under the selective Mallows model is dominated by the sample complexity of sorting them.

We conclude with an informative example, where we compare the sample complexity of retrieving the complete central ranking and the top- k ranking in an interesting special case. Let us assume that $p = 1/\log \log n$ and that $k = \Theta_\beta(\log(n/\epsilon))$. Then we only need $O_\beta(\log \log n \cdot \log \frac{\log n}{\epsilon})$ samples to retrieve the top- k ranking, while learning the complete central ranking requires $\Omega_\beta(\log(n/\epsilon))$ samples. Namely, we have an almost exponential improvement in the sample complexity, for values of k that suffice for most practical applications.

5 Experiments

In this section, we present some experimental evaluation of our main results, using synthetic data. First, we empirically verify that the sample complexity of learning the central ranking from p -frequent selective Mallows samples using POSEST is $\Theta(1/p)$, assuming every other parameter to be fixed. Furthermore, we illustrate empirically that POSEST is a smooth estimator of the central ranking, in the sense that POSEST outputs rankings that are, on average, closer in Kendall Tau distance to the central ranking as the size of the sample profile grows.

5.1 Empirical sample complexity

We estimate the sample complexity of retrieving the central ranking from selective Mallows samples where $n = 20$ and $\beta = 2$, with probability at least 0.95, using POSEST by performing binary search over the size of the sample profile. During a binary search, for every value, say r , of the sample profile size we examine, we estimate the probability that POSEST outputs the central ranking by drawing 100 independent p -frequent selective Mallows profiles of size r , computing POSEST for each one of them and counting successes. We then compare the empirical success rate to 0.95 and proceed with our binary search accordingly. For a specific value of p , we estimate the corresponding sample complexity, by performing 100 independent binary searches and computing the average value. The results, which are




Figure 1: Estimated sample complexity of retrieving, with probability at least 0.95 and using POSEST, the central ranking from selective Mallows samples, with $n = 20$, $\beta = 2$, over the frequency parameter’s inverse.




Figure 2: Average Kendall Tau distance between the output of POSEST and the central ranking with respect to the size of the sample profile, for different values of the frequency parameter p , when $n = 20$, $\beta = 0.3$.

shown in Figure 1, indicate that the dependence of sample complexity on the frequency parameter p is indeed $\Theta(1/p)$.

5.2 Smoothness of PosEst

We plot, for different values of the frequency parameter p , the average Kendall Tau distance between the central ranking and the output of POSEST with respect to the size of the sample profile. For each value, say r , of the sample profile size, considering $\beta = 0.3$ and $n = 20$, we draw 100 independent selective Mallows sample profiles, each of size r , we compute the distance between the output of POSEST for each sample profile and the central ranking and take the average of these distances. The results are presented in Figure 2.

Acknowledgements

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions.

Dimitris Fotakis and Alkis Kalavasis are partially supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the “First Call for H.F.R.I. Research Projects to support Faculty members and Researchers and the procurement of high-cost research equipment grant”, project BALSAM, HFRI-FM17-1424.

References

- Colin L Mallows. Non-Null Ranking Models. I. *Biometrika*, 44(1/2):114–130, 1957.
- Louis L Thurstone. A Law of Comparative Judgment. *Psychological review*, 34(4):273, 1927.
- B Babington Smith. Discussion of Professor Ross’s Paper. *Journal of the Royal Statistical Society B*, 12(1):41–59, 1950.
- 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.
- Robin L Plackett. The Analysis of Permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 24(2):193–202, 1975.
- R Duncan Luce. *Individual Choice Behavior: A Theoretical Analysis*. Courier Corporation, 2012.
- Lirong Xia. Learning and Decision-Making from Rank Data. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 13(1):1–159, 2019.
- Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating Inconsistent Information: Ranking and Clustering. *Journal of the ACM (JACM)*, 55(5):1–27, 2008. URL http://dimacs.rutgers.edu/~alantha/papers2/aggregating_journal.pdf.
- Claire Kenyon-Mathieu and Warren Schudy. How to Rank with Few Errors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 95–103. ACM, 2007.
- Ioannis Caragiannis, Ariel D Procaccia, and Nisarg Shah. When Do Noisy Votes Reveal the Truth? In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 143–160, 2013. URL <https://dl.acm.org/doi/10.1145/2892565>.
- Robert Busa-Fekete, Dimitris Fotakis, Balázs Szörényi, and Manolis Zampetakis. Optimal Learning of Mallows Block Model. In *Conference on Learning Theory*, pages 529–532, 2019. URL <https://arxiv.org/pdf/1906.01009.pdf>.
- Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with Noisy Information. *SIAM Journal on Computing*, 23(5):1001–1018, 1994.
- Cheng Mao, Jonathan Weed, and Philippe Rigollet. Minimax Rates and Efficient Algorithms for Noisy Sorting. In *Algorithmic Learning Theory*, pages 821–847. PMLR, 2018a.
- Bruce Hajek, Sewoong Oh, and Jiaming Xu. Minimax-optimal Inference from Partial Rankings. In *Advances in Neural Information Processing Systems*, pages 1475–1483, 2014.
- Erick Moreno-Centeno and Adolfo R. Escobedo. Axiomatic aggregation of incomplete rankings. *IIE Transactions*, 48(6):475–488, 2016.
- Ilkay Yildiz, Jennifer G. Dy, Deniz Erdogmus, Jayashree Kalpathy-Cramer, Susan Ostmo, J. Peter Campbell, Michael F. Chiang, and Stratis Ioannidis. Fast and Accurate Ranking Regression. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020*, volume 108 of *Proceedings of Machine Learning Research*, pages 77–88. PMLR, 2020. URL <http://proceedings.mlr.press/v108/yildiz20a.html>.
- Mark Braverman and Elchanan Mossel. Sorting from Noisy Information. *arXiv preprint arXiv:0910.1191*, 2009. URL <https://arxiv.org/pdf/0910.1191.pdf>.
- Mark Braverman and Elchanan Mossel. Noisy Sorting Without Resampling. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 268–276. Society for Industrial and Applied Mathematics, 2008.
- Michael A Fligner and Joseph S Verducci. *Probability Models and Statistical Analyses for Ranking Data*, volume 80. Springer, 1993.
- John I Marden. *Analyzing and Modeling Rank Data*. CRC Press, 1996.
- Michael A Fligner and Joseph S Verducci. Distance Based Ranking Models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 48(3):359–369, 1986.
- Thomas Brendan Murphy and Donal Martin. Mixtures of distance-based models for ranking data. *Computational statistics & data analysis*, 41(3-4):645–655, 2003.
- Guy Lebanon and John D Lafferty. Conditional Models on the Ranking Poset. In *Advances in Neural Information Processing Systems*, pages 431–438, 2003.
- Guy Lebanon and Yi Mao. Non-Parametric Modeling of Partially Ranked Data. *Journal of Machine Learning Research*, 9(Oct):2401–2429, 2008.

- Tyler Lu and Craig Boutilier. Effective Sampling and Learning for Mallows Models with Pairwise-Preference Data. *The Journal of Machine Learning Research*, 15(1):3783–3829, 2014.
- Róbert Busa-Fekete, Eyke Hüllermeier, and Balázs Szörényi. Preference-Based Rank Elicitation using Statistical Models: The Case of Mallows. 2014.
- Laura Adkins and Michael Fligner. A non-iterative procedure for maximum likelihood estimation of the parameters of Mallows’ model based on partial rankings. *Communications in Statistics-Theory and Methods*, 27(9):2199–2220, 1998.
- Allen Liu and Ankur Moitra. Efficiently Learning Mixtures of Mallows Models. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 627–638. IEEE, 2018.
- Nir Ailon. An Active Learning Algorithm for Ranking from Pairwise Preferences with an Almost Optimal Query Complexity. *Journal of Machine Learning Research*, 13(Jan):137–164, 2012.
- Mark Braverman, Jieming Mao, and S Matthew Weinberg. Parallel Algorithms for Select and Partition with Noisy Comparisons. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 851–862, 2016.
- Cheng Mao, Ashwin Pananjady, and Martin J Wainwright. Breaking the $1/\sqrt{n}$ Barrier: Faster Rates for Permutation-based Models in Polynomial Time. *arXiv preprint arXiv:1802.09963*, 2018b.
- Fabian Wauthier, Michael Jordan, and Nebojsa Jojic. Efficient Ranking from Pairwise Comparisons. In *International Conference on Machine Learning*, pages 109–117, 2013.
- Nir Ailon. Aggregation of Partial Rankings, p -Ratings and Top- m Lists. *Algorithmica*, 57(2):284–300, 2010.
- Claire Mathieu and Simon Mauras. How to aggregate Top-lists: Approximation algorithms via scores and average ranks. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms*, pages 2810–2822. SIAM, 2020.
- Milan Vojnovic, Se-Young Yun, and Kaifang Zhou. Convergence Rates of Gradient Descent and MM Algorithms for Bradley-Terry Models. In *The 23rd International Conference on Artificial Intelligence and Statistics, (AISTATS 2020)*, volume 108 of *Proceedings of Machine Learning Research*, pages 1254–1264. PMLR, 2020.
- Nihar B Shah, Sivaraman Balakrishnan, Joseph Bradley, Abhay Parekh, Kannan Ramchandran, and Martin J Wainwright. Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. *The Journal of Machine Learning Research*, 17(1):2049–2095, 2016. URL <https://jmlr.org/papers/volume17/15-189/15-189.pdf>.
- Ashish Khetan and Sewoong Oh. Data-driven rank breaking for efficient rank aggregation. *The Journal of Machine Learning Research*, 17(1):6668–6721, 2016.
- Milan Vojnovic and Seyoung Yun. Parameter estimation for generalized thurstone choice models. In *International Conference on Machine Learning*, pages 498–506, 2016.
- Sahand Negahban, Sewoong Oh, and Devavrat Shah. Rank centrality: Ranking from pairwise comparisons. *Operations Research*, 65(1):266–287, 2017.
- Ashwin Pananjady, Cheng Mao, Vidya Muthukumar, Martin J Wainwright, Thomas A Courtade, et al. Worst-case versus average-case design for estimation from partial pairwise comparisons. *Annals of Statistics*, 48(2):1072–1097, 2020.
- Vincent Conitzer and Tuomas Sandholm. Common Voting Rules as Maximum Likelihood Estimators. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI ’05)*, pages 145–152. AUAI Press, 2005.
- Lirong Xia and Vincent Conitzer. Determining Possible and Necessary Winners Given Partial Orders. *Journal of Artificial Intelligence Research*, 41:25–67, 2011.
- Ariel D. Procaccia, Sashank Jakkam Reddi, and Nisarg Shah. A Maximum Likelihood Approach For Selecting Sets of Alternatives. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI ’12)*, pages 695–704. AUAI Press, 2012.
- Nihar B Shah and Martin J Wainwright. Simple, robust and optimal ranking from pairwise comparisons. *The Journal of Machine Learning Research*, 18(1):7246–7283, 2017.
- Ioannis Caragiannis and Evi Micha. Learning a Ground Truth Ranking Using Noisy Approval Votes. In *IJCAI*, pages 149–155, 2017.
- Nayantara Bhatnagar and Ron Peled. Lengths of Monotone Subsequences in a Mallows Permutation. *Probability Theory and Related Fields*, 161(3-4):719–780, 2015.
- Flavio Chierichetti, Anirban Dasgupta, Ravi Kumar, and Silvio Lattanzi. On Reconstructing a Hidden Permutation. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2014.
- Aviad Rubinfeld and Shai Vardi. Sorting from Noisier Samples. In *Proceedings of the Twenty-*

Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 960–972. SIAM, 2017.
URL <http://shaivardi.com/research/Sorting%20Noisier.pdf>.