Fair Rank Aggregation

Anonymous Author(s)
Affiliation
Address
email

Abstract

Ranking algorithms find extensive usage in diverse areas such as web search, employment, college admission, voting, etc. The related rank aggregation problem deals with combining multiple rankings into a single aggregate ranking. However, algorithms for both these problems might be biased against some individuals or groups due to implicit prejudice or marginalization in the historical data. We study ranking and rank aggregation problems from a fairness or diversity perspective, where the candidates (to be ranked) may belong to different groups and each group should have a fair representation in the final ranking. We allow the designer to set the parameters that define fair representation. These parameters specify the allowed range of the number of candidates from a particular group in the top-k positions of the ranking. Given any ranking, we provide a linear time exact algorithm for finding the closest fair ranking for the Kendall tau metric under *strong fairness*, i.e., when the final ranking is fair for all values of k. We also provide an exact algorithm for finding the closest fair ranking for the Ulam metric under strong fairness when there are only $\mathcal{O}(1)$ number of groups. Our algorithms are simple, fast, and might be extendable to other relevant metrics. We also give a novel meta-algorithm for the general rank aggregation problem under the fairness framework. Surprisingly, this meta-algorithm works for any generalized mean objective (including center and median problems) and any fairness criteria. As a byproduct, we obtain 3-approximation algorithms for both center and median problems, under both Kendall tau and Ulam metrics. Furthermore, using sophisticated techniques we obtain a $(3-\varepsilon)$-approximation algorithm, for a constant $\varepsilon > 0$, for the Ulam metric under strong fairness.

1 Introduction

Ranking a set of candidates or items is a ubiquitous problem arising in diverse areas ranging from social choice theory [BCE$^+$16] to information retrieval [Har92]. Given a set of d candidates and a set of n different, potentially conflicting, rankings of these candidates, one fundamental task is to determine a single ranking that best summarizes the preference orders in the individual rankings. This summarizing task, popularly termed *rank aggregation*, has been widely studied from a computational viewpoint over the last two decades [DKNS01, FKS03, GL11, ASCPX13]. Most well-studied rank aggregation paradigms are *median rank aggregation* (or simply *rank aggregation*) [Kem59, You88, YL78, DKNS01] and *maximum rank aggregation* [BBGH15, BBD09, Pop07], which are based on finding the *median* and *center* of the given set of rankings, respectively.

Recently, fairness and diversity have become a natural prerequisite for ranking algorithms where individuals are rated for access to goods and services or ranked for seeking facilities in education (e.g., obtaining scholarship or admission), employment (e.g., hiring or promotion in a job), medical (e.g., triage during a pandemic), or economic opportunities (e.g., loan lending). Some concrete examples include university admissions through affirmative action in the USA [Des05] or the reservation system

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

in jobs in India [Bor10], where we want rankings to be fair to mitigate the prevalent disparities due to historical marginalization. Rankings not being fair may risk promoting extreme ideology [CHRG16] or certain stereotypes about dominating/marginalized communities based on sensitive attributes like gender or race [KMM15, BCZ$^+$16]. There has been a series of works on fair ranking algorithms, see [ZBC$^+$17, AJSD19, Cas19, GSB21, GDL21, PPM$^+$22, ZYS21] and the references therein.

A substantial literature on algorithmic fairness focuses on *group fairness* to facilitate *demographic parity* [DHP$^+$12] or *equal opportunity* [HPS16]: typically this is done by imposing fairness constraints which require that top-k positions in the ranking contain *enough* candidates from *protected* groups that are typically underrepresented due to prevalent discrimination (e.g., due to gender, caste, age, race, sex, etc.). In many countries, group fairness constraints are being enforced by legal norms [Eur, USD]. For example, in Spain 40% of candidates for elections in large voting districts must be women [Ver10], in India 10% of the total recruitment for civil posts and services in government are reserved for people from Economically Weaker Society (EWS) [SIN19], etc.

In this paper, we study group fairness, more specifically proportional fairness (sometimes also referred to as p-fairness [BCPV96]). Inspired by the Disparate Impact doctrine, this notion of fairness mandates that the output of an algorithm must contain a fair representation of each of the 'protected classes' in the population. In the context of ranking, the set of candidates is considered to be partitioned into g groups $G_1, G_2, \ldots, G_g$. For each group $G_i, i = 1, 2, \ldots, g$, we have two parameters $\alpha_i \in (0, 1], \beta_i \in (0, 1]$. A ranking π of the set of items is called *proportionally fair* if for every position $k \in \{1, 2, \ldots, n\}$ and for every group G_i, the following two properties are satisfied: (a) *Minority Protection*: The number of items from group G_i, which are in the top-k positions $\pi(1), \pi(2), \ldots, \pi(k)$, is at least $\lfloor \alpha_i \cdot k \rfloor$, and (b) *Restricted Dominance*: The number of items from group C_i, which are in the top-k positions $\pi(1), \pi(2), \ldots, \pi(k)$, is at most $\lceil \beta_i \cdot k \rceil$.

To compare different rankings several distance functions have been considered defined on the set of permutations/rankings, such as Kendall tau distance [Ken38, DG77, Kem59, You88, YL78, DKNS01, ACN08, KMS07, KV10] (also called *Kemeny distance* in case of rank aggregation), Ulam distance [AD99, CMS01, CK06, AK10, AN10, NSS17, BS19, CDK21, CGJ21], Spearman footrule distance [Spe04, Spe06, DG77, DKNS01, KV10, BBGH15], etc. Among these, Kendall tau distance is perhaps the most common measure used in ranking as it is the only known measure to simultaneously satisfy several required properties such as neutrality, consistency, and the extended Condorcet property [Kem59, You88]. The Ulam metric is another widely-used measure in practice as it is also a simpler variant of the general edit distance metric which finds numerous applications in computational biology, DNA storage system, speech recognition, classification, etc. (e.g., see [CMS01, CDK21, CGJ21]).

One natural computational question related to fairness in ranking is, given a ranking, how to find its closest fair ranking under p-fairness. Celis et al. [CSV18] considered this problem and gave exact and approximation algorithms under several ranking metrics such as discounted cumulative gain (DCG), Spearman footrule, and Bradley-Terry. However, their algorithms do not extend to Kendall tau and Ulam metric, two of the most commonly used ranking metrics.

Fair rank aggregation is relatively less studied. Recently, Kulman et al. [KR20] initiated the study of fair rank aggregation under Kendall tau metric. However, their fairness notion is based on *top-k statistical parity* and *pairwise statistical parity*. These notions are quite restricted. For example, their results only hold for binary protected attributes (i.e., $g = 2$) and and does not satisfy p-fairness. Informally, pairwise statistical parity considers pairs of items from different groups in an aggregated manner and does not take into account the actual rank of the items in the final ranking. See [WISR22] for an example on why the fairness notion in [KR20] does not satisfy p-fairness. In fact, as p-fairness satisfies statistical parity for all top k-positions in the ranking, it is a much stronger notion compared to statistical parity. Thus achieving p-fairness is a significantly more challenging problem.

1.1 Our Contributions.

Our first main contribution is exact algorithms for the closest fair ranking (CFR) problem under proportional fairness (see Definition 2.4) for Kendall tau and Ulam metrics. For the Kendall tau metric, we give the *first exact algorithm* for the closest fair ranking problem (Theorem 3.4). Our algorithm is simple and based on a greedy strategy; however, the analysis is delicate. It exploits the following interesting and perhaps surprising fact. Under the Kendall tau metric, given a fixed (possibly unfair) ranking π, there exists a closest fair ranking π' to π such that for every group $G_i, i = 1, 2, \ldots, g$, the relative ordering of elements in G_i remains unaltered in π' compared to π (Theorem 3.3). Then, for the *Ulam metric*, we give a polynomial time dynamic programming algorithm for the closest fair ranking problem when the number of groups g is a constant (Theorem 3.10). In practice, the number of protected classes is relatively few, and hence our result gives an efficient algorithm for such cases.

Our second significant contribution is the study of *rank aggregation problem under a very general notion of proportional fairness*. Our main contribution is to develop a novel algorithmic toolbox for the fair rank aggregation that solves a wide variety of rank aggregation objectives satisfying such generic fairness constraints. An essential takeaway of our work is that a set of potentially biased rankings can be aggregated into a fair ranking with only a small loss in the 'quality' of the ranking. We study q-mean Fair Rank Aggregation (FRA), where given a set of rankings $\pi_1, \pi_2, \ldots, \pi_n$, a (dis)similarity measure (or distance) ρ between two rankings, and any $q \geq 1$, the task is to determine a *fair* ranking σ that minimizes the generalized mean objective: $\sum_{i=1}^{n} \left(\rho(\pi_i, \sigma)^q\right)^{1/q}$. We would like to emphasize that in general, q-mean objective captures two classical data aggregation tasks: One is *median* which asks to minimize the sum of distances (i.e., $q = 1$) and another is *center* which asks to minimize the maximum distance to the input points (i.e., $q = \infty$). Without the fairness requirement, if ρ is defined as the *Kendall tau* distance between two rankings and $q = 1$, then our objective boils down to the classical *Kemeny optimization* [Kem59]. The rank aggregation problems (without any fairness constraint) are already hard to solve under certain metrics. E.g., it is known to be NP-hard under the Kendall tau [BTT89] (even for 4 input rankings [DKNS01]). Thus the best we could hope for is to provide approximate algorithms for the fair rank aggregation problem.[1]

We show generic reductions of the q-mean Fair Rank Aggregation (FRA) to the problem of determining the *closest fair ranking* (CFR) to a given ranking. More specifically, we show that any c-approximation algorithm for the closest fair ranking problem can be utilized as a blackbox to give a $(c + 2)$-approximation to the FRA for any $q \geq 1$ (Theorem 4.3). This result is oblivious to the specifics of the (dis)similarity measure and only requires the measure to be a metric. Using the exact algorithms for the CFR for the Kendall tau, Spearman footrule, and Ulam metrics (for constantly many groups), we thus obtain 3-approximation algorithms for the FRA problem under these three (dis)similarity measures, respectively. Further, we provide yet another simple algorithm that even breaks below 3-factor for the Ulam metric. For $q = 1$, by combining the above-stated 3-approximation algorithm with an additional procedure, we achieve a $(3 - \varepsilon)$-approximation factor (for some $\varepsilon > 0$) for the FRA under the Ulam, for constantly many groups (Theorem 4.11). We also provide another reduction from FRA to one rank aggregation computation (without fairness) and a CRF computation (Theorem 4.8), and as a corollary get an $\mathcal{O}(d^3 \log d + n^2 d)$-time algorithm for Spearman footrule when $q = 1$ (Corollary 4.10). We summarize our main results in Table 1.1.

Problem	Metric	#Groups	Approx Ratio	Runtime	Reference
CFR	Kendall tau	Arbitrary	Exact	$\mathcal{O}(d)$	Theorem 3.4
	Ulam	Constant	Exact	$\mathcal{O}(d^{g+2})$	Theorem 3.10
	Spearman footrule	Arbitrary	Exact	$\mathcal{O}(d^3 \log d)$	[CSV18]
FRA	Kendall tau	Arbitrary	3	$\mathcal{O}(n^2 d \log d)$	Corollary 4.5
	Ulam	Constant	3	$\mathcal{O}(nd^{g+2} + n^2 d \log d)$	Corollary 4.7
	Ulam ($q = 1$)	Constant	$3 - \varepsilon$	$\mathrm{poly}(n, d)$	Theorem 4.11
	Spearman footrule	Arbitrary	3	$\mathcal{O}(nd^3 \log d + n^2 d)$	Corollary 4.6

Comparison with concurrent work. Independently and concurrently to our work, Wei et al. [WISR22] considers the fair ranking problem under a setting that is closely related to ours. However, the fairness criteria in their work are much more restrictive compared to ours as follows. Their algorithms for CFR are only designed for a special case of our formulation where for each group G_i and any position k in the output ranking, $\alpha_i = \beta_i = p(i)$, where $p(i)$ denotes the proportion of group G_i in the entire population. Further, under the Kendall tau metric, they give a polynomial time exact algorithm for CFR only for the special case of binary groups ($g = 2$). They also give additional algorithms for multiple groups - an exact algorithm that works in time exponential in the number

[1] For a minimization problem, an algorithm $\mathcal{A}$ is an α-approximation ($\alpha \geq 1$) algorithm if for all input instance I, $\mathcal{A}(I) \leq \alpha\mathrm{OPT}(I)$. Here $\mathcal{A}(I)$ and $\mathrm{OPT}(I)$ are the cost of the solution returned by $\mathcal{A}$ and the optimal algorithm, respectively, on input I.

of groups and a polynomial time 2-approximation. In contrast, we fully resolve the CFR problem under the Kendall tau metric by giving a linear time algorithm for the case of multiple groups and any arbitrary bounds on α_i and β_i for each group G_i. Further, we give the first results for CFR and FRA under the Ulam metric as well.

2 Preliminaries

Notations. For any $n \in \mathbb{N}$, let $[n]$ denote the set $\{1, 2, \cdots, n\}$. We refer to the set of all permutations/rankings over $[d]$ by $\mathcal{S}_d$. Throughout this paper we consider any permutation $\pi \in \mathcal{S}_d$ as a sequence of numbers $a_1, a_2, \cdots, a_d$ such that $\pi(i) = a_i$, and we say that the *rank* of a_i is i. For any two $x, y \in [d]$ and a permutation $\pi \in \mathcal{S}_d$, we use the notation $x <_\pi y$ to denote that the rank of x is less than that of y in π. For any subset $I = \{i_1 < i_2 < \cdots < i_r\} \subseteq [d]$, let $\pi(I)$ be the sequence $\pi(i_1), \pi(i_2), \cdots, \pi(i_r)$ (which is essentially a subsequence of the sequence represented by π). When clear from the context, we use $\pi(I)$ also to denote the set of elements in the sequence $\pi(i_1), \pi(i_2), \cdots, \pi(i_r)$. For any $k \in [d]$ and a permutation $\pi \in \mathcal{S}_d$, we refer to $\pi([k])$ as the k-length *prefix* of π. For any prefix P, let $|P|$ denote the length of that prefix. For any two prefixes P_1, P_2, we use $P_1 \subseteq P_2$ to denote $|P_1| \leq |P_2|$.

Distance measures on rankings. There are different distance functions being considered to measure the dissimilarity between any two rankings/permutations. Among them, perhaps the most commonly used one is the *Kendall tau distance*.

Definition 2.1 (Kendall tau distance)**.** Given two permutations $\pi_1, \pi_2 \in \mathcal{S}_d$, the *Kendall tau distance* between them, denoted by $\mathcal{K}(\pi_1, \pi_2)$, is the number of pairwise disagreements between x and y, i.e.,

$$\mathcal{K}(\pi_1, \pi_2) := |\{(a, b) \in [d] \times [d] \mid a <_{\pi_1} b \text{ but } b <_{\pi_2} a\}|.$$

Another important distance measure is the *Spearman footrule* (aka. *Spearman's rho*) which is essentially the ℓ_1-norm between two permutations.

Definition 2.2 (Spearman footrule distance)**.** Given two permutations $\pi_1, \pi_2 \in \mathcal{S}_d$, the *Spearman footrule distance* between them is defined as $\mathcal{F}(\pi_1, \pi_2) := \sum_{i \in [d]} |\pi_1(i) - \pi_2(i)|$.

Another interesting distance measure is the *Ulam distance* which counts the minimum number of character move operations between two permutations [AD99]. This definition is motivated by the classical *edit distance* that is used to measure the dissimilarity between two strings. A character move operation in a permutation can be thought of as "picking up" a character from its position and then "inserting" that character in a different position[2].

Definition 2.3 (Ulam distance)**.** Given two permutations $\pi_1, \pi_2 \in \mathcal{S}_d$, the *Ulam distance* between them, denoted by $\mathcal{U}(\pi_1, \pi_2)$, is the minimum number of character move operations that is needed to transform π_1 into π_2.

Alternately, the Ulam distance between π_1, π_2 can be defined as $d - \mathsf{LCS}(\pi_1, \pi_2)$, where $\mathsf{LCS}(\pi_1, \pi_2)$ denotes the *longest common subsequence* between the sequences π_1 and π_2.

Fair rankings. We are given a set C of d candidates, which are partitioned into g groups. We call a ranking (of these d candidates) *fair* if all sufficiently large prefixes of it have certain proportion of representatives from each group. Formally,

Definition 2.4 ($(\bar{\alpha}, \bar{\beta})$-$k$-fair ranking)**.** Consider a set C of d candidates partitioned into g groups $G_1, \cdots, G_g$, and $\bar{\alpha} = (\alpha_1, \cdots, \alpha_g) \in [0, 1]^g$, $\bar{\beta} = (\beta_1, \cdots, \beta_g) \in [0, 1]^g$, $k \in [d]$. A ranking $\pi \in \mathcal{S}_d$ is said to be $(\bar{\alpha}, \bar{\beta})$-$k$-*fair* if for any prefix P of size at least k, of π and each group $i \in [g]$, there are at least $\lfloor \alpha_i \cdot |P| \rfloor$ and at most $\lceil \beta_i \cdot |P| \rceil$ elements from the group G_i in P, i.e.,

$$\forall_{\text{prefix } P : |P| \geq k},\ \forall_{i \in [g]},\ \lfloor \alpha_i \cdot |P| \rfloor \leq |P \cap G_i| \leq \lceil \beta_i \cdot |P| \rceil.$$

We also define a weak fairness notion that preserves the proportionate representation only for a fixed k-length prefix.

[2] One may also consider one deletion and one insertion operation instead of a character move, and define the Ulam distance accordingly as in [CMS01].

Definition 2.5 ($(\bar{\alpha}, \bar{\beta})$-weak k-fair ranking)**.** Consider a set C of d candidates partitioned into g groups $G_1, \cdots, G_g$, and $\bar{\alpha} = (\alpha_1, \cdots, \alpha_g) \in [0,1]^g$, $\bar{\beta} = (\beta_1, \cdots, \beta_g) \in [0,1]^g$, $k \in [d]$. A ranking $\pi \in \mathcal{S}_d$ is said to be $(\bar{\alpha}, \bar{\beta})$-*weak k-fair* if for the k-length prefix P of π and each group $i \in [g]$, there are at least $\lfloor \alpha_i \cdot k \rfloor$ and at most $\lceil \beta_i \cdot k \rceil$ elements from the group G_i in P, i.e.,

$$\forall_{i \in [g]},\ \lfloor \alpha_i \cdot k \rfloor \leq |P \cap G_i| \leq \lceil \beta_i \cdot k \rceil.$$

Note, an $(\bar{\alpha}, \bar{\beta})$-$k$-fair ranking is also $(\bar{\alpha}, \bar{\beta})$-weak k-fair, but the converse need not be true. We would like to emphasize that all the results presented in this paper hold for both $(\bar{\alpha}, \bar{\beta})$-$k$-fairness and $(\bar{\alpha}, \bar{\beta})$-weak k-fairness.

3 Closest Fair Ranking

In this section, we consider the problem of computing the closest fair ranking of a given input ranking. Below we formally define the problem.

Definition 3.1 (Closest fair ranking problem)**.** Consider a metric space $(\mathcal{S}_d, \rho)$ for a $d \in \mathbb{N}$. Given a ranking $\pi \in \mathcal{S}_d$ and $\bar{\alpha}, \bar{\beta} \in [0,1]^g$ for some $g \in \mathbb{N}$, $k \in [d]$, the objective of the *closest fair ranking problem* (resp. closest weak fair ranking problem) is to find a $(\bar{\alpha}, \bar{\beta})$-$k$-fair ranking (resp. $(\bar{\alpha}, \bar{\beta})$-(weak) k-fair ranking) $\pi^* \in \mathcal{S}_d$ that minimizes the distance $\rho(\pi, \pi^*)$.

Unless stated explicitly, we consider the notion of $(\bar{\alpha}, \bar{\beta})$-$k$-fairness (not the weak one) in all the results presented in this section.

3.1 Closest fair ranking under Kendall tau

Closest weak fair ranking. We first show that we can compute a closest weak fair ranking under the Kendall tau, exactly in linear time.

Theorem 3.2. *There exists a linear time algorithm that, given a ranking $\pi \in \mathcal{S}_d$, a partition of $[d]$ into g groups $G_1, \cdots, G_g$ for some $g \in \mathbb{N}$, and $\bar{\alpha} = (\alpha_1, \cdots, \alpha_g) \in [0,1]^g$, $\bar{\beta} = (\beta_1, \cdots, \beta_g) \in [0,1]^g$, $k \in [d]$, outputs a closest $(\bar{\alpha}, \bar{\beta})$-weak k-fair ranking under the Kendall tau distance.*

Let us first describe the algorithm. Our algorithm follows a simple greedy strategy. For each group G_i, it picks the top $\lfloor \alpha_i k \rfloor$ elements according to the input ranking π, and add them in a set P. If P contains k elements, then we are done. Otherwise, we iterate over the remaining elements and add them in P as long as for each group G_i, $|P \cap G_i| \leq \lceil \beta_i k \rceil$ (each group has at most $\lceil \beta_i k \rceil$ elements in P) until the size of P becomes exactly k. Then we use the relative ordering of the elements in P as in the input ranking π and make it the k-length prefix of the output ranking σ. Fill the last $d-k$ positions of σ by the remaining elements ($[d] \setminus P$) by following their relative ordering as in the input π. See Algorithm 1 in appendix for the pseudocode of the algorithm.

By the construction of the set P, at the end, for each group G_i, $\lfloor \alpha_i k \rfloor \leq |P \cap G_i| \leq \lceil \beta_i k \rceil$. Since we use the elements of P in the k-length prefix of the output ranking σ, σ is an $(\bar{\alpha}, \bar{\beta})$-weak k-fair ranking. For the running time, a straightforward implementation our algorithms takes $\mathcal{O}(d)$ time. It only remains to argue that σ is a closest $(\bar{\alpha}, \bar{\beta})$-weak k-fair ranking to the input π. To show that, we use the following key observation.

Claim 3.3. *Under the Kendall tau distance, there always exists a closest $(\bar{\alpha}, \bar{\beta})$-weak k-fair ranking π^* such that, for each group G_i ($i \in [g]$), for any two elements $a \neq b \in G_i$, $a <_{\pi^*} b$ if and only if $a <_{\pi} b$.*

We defer the proof of the above claim and how we use it to conclude the proof of Theorem 3.2, to the appendix (provided in the supplementary material).

Extension to general fairness notion. Previously, we provide an algorithm that outputs a weak fair ranking (see Definition 2.5 for the definition of weak fairness) closest to the input. Now, we present an algorithm that outputs a closest fair (according to Definition 2.4) ranking.

Theorem 3.4. *There exists an $\mathcal{O}(d)$ time algorithm that, given a ranking $\pi \in \mathcal{S}_d$, a partition of $[d]$ into g groups $G_1, \cdots, G_g$ for some $g \in \mathbb{N}$, and $\bar{\alpha} = (\alpha_1, \cdots, \alpha_g) \in [0,1]^g$, $\bar{\beta} = (\beta_1, \cdots, \beta_g) \in [0,1]^g$, $k \in [d]$, outputs a closest $(\bar{\alpha}, \bar{\beta})$-$k$-fair ranking under the Kendall tau distance.*

The main challenge with this stronger fairness notion is that now we need to satisfy the fairness criteria for all the prefixes not just the k-length prefix as in case of weak fairness. Surprisingly, we show that under the Kendall tau metric, by iteratively applying the algorithm for the closest weak fair ranking (Algorithm 1) as a black-box, over the prefixes of decreasing length, we can construct a closest fair (not just the weak one) ranking. It is worth noting that here at any iteration the input to Algorithm 1 is a prefix of π which is not a permutation. However, Algorithm 1 only treats the input as a sequence of numbers (not really as a permutation). See Algorithm 2 in appendix for a formal description of the algorithm.

It is worth noting that, since we iteratively apply Algorithm 1 on a prefix of π (not the whole sequence represented by π), it is not even clear whether the algorithm finally outputs a fair ranking (assuming it exists). Below we first argue that if there exists a fair ranking then the output σ must be a fair ranking. Next, we establish that σ is indeed a closest fair ranking to π.

Let π^* be a closest fair ranking to π that preserves relative orderings (of elements in $[d]$) maximally. We show that the output $\sigma = \pi^*$. We start the argument by considering any two prefixes of length k_1 and k_2, where $k_2 < k_1$. We argue that k_1 and k_2-length prefixes of σ and π^* are the same. Since this hold for any k_1 and k_2 (with $k_1, k_2 \geq k$), where $k_2 < k_1$, by using induction we can show that $\sigma = \pi^*$. We defer the induction argument to the appendix, and below provide the argument for the k_1 and k_2-length prefixes (which is a key to prove the correctness of Theorem 3.4).

For the sake of analysis, let us consider the following three permutations. Let π_1 be the $(\bar{\alpha}, \bar{\beta})$-weak k_1-fair ranking closest to π, output by Algorithm 1. Let π_2 be the ranking output by Algorithm 1 when given the k_1-length prefix of π_1 (i.e., the sequence $\pi_1([k_1])$) as input and is asked to output an $(\bar{\alpha}, \bar{\beta})$-weak k_2-fair ranking closest to π_1. Further, let π_2' be the $(\bar{\alpha}, \bar{\beta})$-weak k_2-fair ranking closest to π, output by Algorithm 1. In other words, π_2 be the ranking produced by first applying Algorithm 1 on π to make its k_1-length prefix fair and then apply Algorithm 1 again on that output to make its k_2-length prefix fair. Whereas, π_2' be the ranking produced by directly applying Algorithm 1 on π to make its k_2-length prefix fair.

Then the next claim argues about the existence of fair ranking π_2. We defer the proof of this claim to the appendix.

Claim 3.5. *If there is a ranking π' such that its k_1-length prefix P_1 and k_2-length prefix P_2 satisfies that for each group G_i ($i \in [g]$), $\lfloor \alpha_i k_1 \rfloor \leq |P_1 \cap G_i| \leq \lceil \beta_i k_1 \rceil$ and $\lfloor \alpha_i k_2 \rfloor \leq |P_2 \cap G_i| \leq \lceil \beta_i k_2 \rceil$, then π_2 exists.*

It follows from the construction that,

Claim 3.6. *The set of elements in $\pi_2([k_1])$ is the same as that in $\pi_1([k_1])$.*

Claim 3.7. $\pi_2([k_2]) = \pi_2'([k_2])$.

Proof. Consider an element $a \in \pi_2'([k_2]) \cap G_i$ for some $i \in [g]$. If a is among the top $\lfloor \alpha_i k_2 \rfloor$ elements (according to π) inside the group G_i, then by Algorithm 1, it would also be selected in $\pi_1([k_1])$ (since $k_1 \geq k_2$) and also in $\pi_2([k_2])$.

Now consider the case where a is among the top $\lceil \beta_i k_2 \rceil$ elements of G_i, but not among the top $\lfloor \alpha_i k_2 \rfloor$ elements. This means that a is also among the top $\lceil \beta_i k_1 \rceil (\geq \lceil \beta_i k_2 \rceil)$ elements of its group G_i. This means that if it is encountered during the execution of Algorithm 1 on π to get an $(\bar{\alpha}, \bar{\beta})$-weak k_1-fair ranking, then it will be selected in $\pi_1([k_1])$. However, we also know that it is selected in $\pi_2'([k_2])$ which is a shorter prefix. Since the upper bound constraints were not violated for G_i during the selection of the elements in $\pi_2'([k_2])$, the upper bound constraints cannot be violated during the selection of the elements of $\pi_1([k_1])$ as well. Hence, a will be selected in $\pi_1([k_1])$.

By a similar argument, when executing Algorithm 1 on π_1 (in the later iteration) to output an $(\bar{\alpha}, \bar{\beta})$-weak k_2-fair ranking, a will again be encountered and be selected in $\pi_2([k_2])$. Therefore, every element in $\pi_2'([k_2])$ is also in $\pi_2([k_2])$. Since the sizes of both the sets are equal, the two sets are infact the same, and so are the rankings (by Algorithm 1). □

Claim 3.8. *The set of elements in $\pi^*([k_1])$ is the same as that in $\pi_2([k_1])$.*

Proof. Assume towards contradiction that $\exists a \in \pi^*([k_1]) \setminus \pi_2([k_1])$ and $\exists b \in \pi_2([k_1]) \setminus \pi^*([k_1])$. If a, b were in the same group, then by Algorithm 1, we know that $b <_\pi a$, and hence by swapping the elements in π^*, the distance from π can only be reduced. Hence we can obtain a different solution $\bar{\pi}$ in which $b \in \bar{\pi}([k_1])$ and $a \notin \bar{\pi}([k_1])$, and thus is also fair. This contradicts that π^* is a closest fair ranking to π that preserves relative orderings (of elements in $[d]$) maximally. So, we can assume, $a \in G_i$ and $b \in G_j$ for some $i \neq j$.

Now we note that a cannot be among the top $\lfloor \alpha_i k_1 \rfloor$ elements, but is in the top $\lceil \beta_i k_1 \rceil$ elements in G_i. Similarly, b cannot be among the top $\lfloor \alpha_j k_1 \rfloor$ elements, but is in the top $\lceil \beta_j k_1 \rceil$ elements in G_j. Again, it follows from Algorithm 1, $b <_\pi a$. So by swapping these two elements in π^* we can only reduce the distance from π, while obtaining another fair ranking (because it has b in the k_1-length prefix instead of a). This again contradicts that π^* is a closest fair ranking to π that preserves relative orderings (of elements in $[d]$) maximally. The claim now follows. □

Claim 3.9. *The set of elements in $\pi^*([k_2])$ is the same as that in $\pi_2([k_2])$.*

The proof of the above is in the appendix. We apply Claim 3.8 and Claim 3.9 iteratively to complete the correctness of Algorithm 2 which we defer to the appendix.

It only remains to argue that the algorithm runs in time $\mathcal{O}(d)$. It is easy to see that a straightforward implementation takes $\mathcal{O}(d^2)$ time (since it invokes at most d calls to the subroutine Algorithm 1). However, with a slightly more intricate implementation (by maintaining an "active" prefix which in the beginning contains all the elements, and then at each iteration we remove lowest rank elements from specific groups, and thus avoiding re-processing of the elements again and again over the iterations), we can show that the algorithm runs in only $\mathcal{O}(d)$ time, which we defer to the appendix.

3.2 Closest fair ranking under Ulam Metric

Theorem 3.10. *There exists a polynomial time dynamic programming based algorithm that finds a $(\bar{\alpha}, \bar{\beta})$-$k$-fair ranking when there are constant number of groups.*

The proof of the lemma uses an intricate dynamic program exploiting the connection between the Ulam distance with the Longest Common Subsequence problem. We defer the proof to the appendix.

4 Fair Rank Aggregation

We start this section by formally defining the *fair rank aggregation* problem. Then we will provide two meta-algorithms that approximate the fair aggregated ranking.

Definition 4.1 (q-mean Rank Aggregation)**.** Consider a metric space $(\mathcal{S}_d, \rho)$ for a $d \in \mathbb{N}$. Given a set $S \subseteq \mathcal{S}_d$ of n input rankings, the *q-mean rank aggregation* problem asks to find a ranking $\sigma \in \mathcal{S}_d$ (not necessarily from S) that minimizes the objective function $\texttt{Obj}_q(S, \sigma) := \left(\sum_{\pi \in S} \rho(\pi, \sigma)^q\right)^{1/q}$.

Generalized mean or q-mean objective functions are well-studied in the context of clustering [CMV22], and division of goods [BKM22]. We study it for the first time in the context of rank aggregation. For $q = 1$, the above problem is also referred to as the *median ranking* problem or simply *rank aggregation* problem [Kem59, You88, YL78, DKNS01]. On the other hand, for $q = \infty$, the problem is also referred to as the *center ranking* problem or *maximum rank aggregation* problem [BBGH15, BBD09, Pop07]. Both these special cases are studied extensively in the literature with different distance measures, e.g., Kendall tau distance [DKNS01, ACN08, KMS07, Sch12, BBD09], Ulam distance [CDK21, BBGH15, CGJ21], Spearman footrule distance [DKNS01, BBGH15].

In the fair rank aggregation problem, we want the output aggregated rank to satisfy certain fairness constraint. Throughout this section, for brevity, we use the term (weak) fair ranking instead of $(\bar{\alpha}, \bar{\beta})$-(weak) k-fair ranking.

Definition 4.2 (q-mean Fair Rank Aggregation)**.** Consider a metric space $(\mathcal{S}_d, \rho)$ for a $d \in \mathbb{N}$. Given a set $S \subseteq \mathcal{S}_d$ of n input rankings/permutations, the *q-mean (weak) fair rank aggregation* problem asks to find a (weak) fair ranking $\sigma \in \mathcal{S}_d$ (not necessarily from S) that minimizes the objective function $\texttt{Obj}_q(S, \sigma) := \left(\sum_{\pi \in S} \rho(\pi, \sigma)^q\right)^{1/q}$.

It is worth noting that in the above definition, the minimization is over the set of all the (weak) fair rankings in $\mathcal{S}_d$. When clear from the context, we drop weak and refer it as the q-mean fair rank aggregation problem. Let σ^* be a (weak) fair ranking that minimizes $\texttt{Obj}_q(S, \sigma)$, i.e., $\sigma^* =$

$\arg\min_{\text{fair } \sigma \in \mathcal{S}_d} \texttt{Obj}_q(S, \sigma)$. Then we call σ^* a *q-mean fair aggregated rank* of S. We refer to $\texttt{Obj}_q(S, \sigma^*)$ as $\text{OPT}_q(S)$.

When $q = 1$, we refer the problem as the *fair median ranking* problem or simply *fair rank aggregation* problem. When $q = \infty$, the objective function becomes $\texttt{Obj}_\infty(S, \sigma) = \max_{\pi \in S} \rho(\pi, \sigma)$, and we refer the problem as the *fair center ranking* problem.

Next, we present two meta algorithms that work for any values of q and irrespective of strong or weak fairness constraint.

4.1 First Meta Algorithm

Theorem 4.3. *Consider any $q \geq 1$. Suppose there is a $t(d)$-time c-approximation algorithm $\mathcal{A}$, for some $c \geq 1$, for the closest fair ranking problem over the metric space $(\mathcal{S}_d, \rho)$. Then there exists a $(c + 2)$-approximation algorithm for the q-mean fair rank aggregation problem, that runs in $\mathcal{O}(n \cdot t(d) + n^2 \cdot f(d))$ time where $f(d)$ is the time to compute $\rho(\pi_1, \pi_2)$ for any $\pi_1, \pi_2 \in \mathcal{S}_d$.*

We devote this subsection in proving the above theorem. Let us start with describing the algorithm. It works as follows: Given a set $S \subseteq \mathcal{S}_d$ of rankings, it first computes c-approximate closest fair ranking σ (for some $c \geq 1$) for each $\pi \in S$. Next, output a σ that minimizes $\texttt{Obj}_q(S, \sigma)$. Let us denote the output ranking by $\bar{\sigma}$. See Algorithm 4 in appendix for a more formal description.

It is straightforward to verify that the running time of the above algorithm is $\mathcal{O}(n \cdot t(d) + n^2 \cdot f(d))$, where $f(d)$ is the time to compute $\rho(\pi_1, \pi_2)$ for any $\pi_1, \pi_2 \in \mathcal{S}_d$ and $t(d)$ denotes the running time of the algorithm $\mathcal{A}$. So it only remains to argue about the approximation factor of Algorithm 4. The following simple observation plays a pivotal role in establishing the approximation factor of Algorithm 4.

Lemma 4.4. *Given a set $S \subseteq \mathcal{S}_d$ of n rankings, let σ^* be a q-mean fair aggregated rank of S under a distance function ρ. Further, let $\bar{\pi}$ be a nearest neighbor (closest ranking) of σ^* in S, and $\bar{\sigma}$ be a c-approximate closest fair ranking to $\bar{\pi}$, for some $c \geq 1$. Then $\forall \pi \in S, \rho(\pi, \bar{\sigma}) \leq (c + 2) \cdot \rho(\pi, \sigma^*)$.*

We defer the proof of the above claim to the appendix. Now, we use the above lemma to show that the approximation factor of Algorithm 4 is $c + 2$. Let σ^* be an (arbitrary) optimal fair aggregate rank of S and $\bar{\sigma}$ be the output of Algorithm 4. The optimal value of the objective function is $\text{OPT} = \texttt{Obj}_q(S, \sigma^*) = \left(\sum_{\pi \in S} \rho(\pi, \sigma^*)^q\right)^{1/q}$. Next, we show that $\texttt{Obj}_q(S, \bar{\sigma}) \leq (c + 2) \cdot \text{OPT}$.

$$\texttt{Obj}_q(S, \bar{\sigma}) \leq \left(\sum_{\pi \in S} \left((c + 2) \cdot \rho(\pi, \sigma^*)\right)^q\right)^{1/q} = (c + 2) \cdot \left(\sum_{\pi \in S} \rho(\pi, \sigma^*)^q\right)^{1/q} = (c + 2) \cdot \text{OPT}.$$

where the second inequality follows from Theorem 4.4 This concludes the proof of Theorem 4.3.

Applications of Theorem 4.3. We have shown in Theorem 3.4 that the closest fair ranking problem for Kendall tau can be solved exactly in $\mathcal{O}(d)$ time, i.e., the approximation ratio is $c = 1$. We also know from [Kni66], that the Kendall tau distance between two permutations can be computed in $\mathcal{O}(d \log d)$ time. This gives us that,

Corollary 4.5. *For any $q \geq 1$, there exists an $\mathcal{O}(n^2 d \log d)$ time meta-algorithm, that finds a 3-approximate solution to the q-mean fair rank aggregation problem , under the Kendall tau metric.*

It is shown in [CSV18] that the closest fair ranking problem for Spearman Footrule can be solved exactly in $\mathcal{O}(d^3 \log d)$ time, i.e., the approximation ratio is $c = 1$. Since distance under Spearman Footrule can be trivially computed in $\mathcal{O}(d)$ we have that,

Corollary 4.6. *For any $q \geq 1$, there exists an $\mathcal{O}(nd^3 \log d + n^2 d)$ time meta-algorithm, that finds a 3-approximate solution to the q-mean fair rank aggregation problem , under the Spearman footrule metric.*

We have shown in Theorem 3.10 that for constant number of groups, the closest fair ranking problem for Ulam metric can be solved exactly in $\mathcal{O}(d^{g+2})$ time, i.e., the approximation ratio is $c = 1$. From [AD99] we know that Ulam distance between two permutations can be computed in $\mathcal{O}(d \log d)$ time. This gives us that,

Corollary 4.7. *For any $q \geq 1$, there exists an $\mathcal{O}(nd^{g+2} + n^2 d \log d)$ time meta-algorithm, that finds a 3-approximate solution to the q-mean fair rank aggregation problem , under the Ulam metric.*

We would like to emphasize that all the above results hold for any values of $q \geq 1$. Hence, they are also true for the special case of the fair median problem (i.e., for $q = 1$) and the fair center problem (i.e., for $q = \infty$).

4.2 Second Meta Algorithm

Theorem 4.8. *Consider any $q \geq 1$. Suppose there is a $t_1(n)$ time c_1-approximation algorithm $\mathcal{A}_1$ for some $c_1 \geq 1$ for q-mean rank aggregation problem; and a $t_2(d)$-time c_2-approximation algorithm $\mathcal{A}_{\in}$, for some $c_2 \geq 1$, for the closest fair ranking problem over the metric space $(\mathcal{S}_d, \rho)$. Then there exists a $(c_1c_2 + c_1 + c_2)$-approximation algorithm for the q-mean fair rank aggregation problem, that runs in $\mathcal{O}(t_1(n) + t_2(d) + n^2 \cdot f(d))$ time where $f(d)$ is the time to compute $\rho(\pi_1, \pi_2)$ for any $\pi_1, \pi_2 \in \mathcal{S}_d$.*

The algorithm works as follows: Given a set $S \subseteq \mathcal{S}_d$ of rankings, it first computes c_1-approximate aggregate rank π^*. Next, output a c_2-approximate closest fair ranking $\bar{\sigma}$, to π^*. See Algorithm 5 in appendix for a more formal description.

It is easy to see that the running time of the algorithm is $\mathcal{O}(t_1(n) + t_2(d) + n^2 \cdot f(d))$, where $f(d)$ is the time to compute $\rho(\pi, \sigma)$ for any $\pi, \sigma \in \mathcal{S}_d$, $t_1(n)$ denotes the running time of the algorithm $\mathcal{A}_1$, and $t_2(d)$ denotes the running time of the algorithm $\mathcal{A}_2$. It now remains to argue about the approximation ratio of the above algorithm. We again make a simple but crucial observation towards establishing the approximation ratio for Algorithm 5.

Lemma 4.9. *Given a set $S \subseteq \mathcal{S}_d$ of n rankings, let σ^* be a q-mean fair aggregated rank of S under a distance function ρ. Further, let π^* be the c_1-approximate aggregate rank of S and $\bar{\sigma}$ be a c_2-approximate closest fair ranking to π^*, for some $c_1, c_1 \geq 1$. Then*

$$\forall \pi \in S, \rho(\pi, \bar{\sigma}) \leq (c_1c_1 + c_1 + c_2) \cdot \rho(\pi, \sigma^*).$$

We defer the proof of this lemma to the appendix. Once we have this key lemma in place, the remaining proof of Theorem 4.8, follows exactly as the proof of Theorem 4.3.

The above algorithm can give similar approximation guarantees as Algorithm 4, but with potentially better running times depending on whether the rank aggregation problem is solved in a faster way for the particular problem in consideration. For instance consider the case for Spearman footrule. It is known that the rank aggregation problem for Spearman footrule can be solved in $\tilde{\mathcal{O}}(d^{2.5})$ [vdBLN+20]. So, using this in conjunction with Algorithm 5 we obtain the following result.

Corollary 4.10. *For $q = 1$, there exists an $\mathcal{O}(d^3 \log d + n^2 d)$ time meta-algorithm, that finds a 3-approximate solution to the q-mean fair rank aggregation problem (i.e., the fair median problem), under Spearman footrule metric.*

4.3 Breaking below 3-factor for Ulam

Theorem 4.11. *For $q = 1$, there exists a constant $\varepsilon > 0$ and a polynomial time algorithm, that finds a $(3 - \varepsilon)$-approximate solution to the q-mean fair rank aggregation problem (i.e., the fair median problem), under the Ulam metric for constantly many groups.*

We show the above result by designing a new algorithm based on the relative ordering of the elements (as in in majority of the input rankings). Then the final output is the best of that output by this new algorithm and that produced by our first meta-algorithm. We argue that when the whole optimal objective value is distributed among only a few elements, then the first meta-algorithm already achieves $(3 - \epsilon)$-approximation. Otherwise, this new relative ordering based approach will provide a $(3 - \epsilon)$-approximation. Although our new algorithm is also very simple, the whole analysis is quite delicate and involves a few important observations on the Ulam metric. We provide the description of our new algorithm along with the whole analysis in the appendix.

References

[ACN08] Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5):23:1–23:27, 2008.

[AD99] David Aldous and Persi Diaconis. Longest increasing subsequences: from patience sorting to the Baik-Deift-Johansson theorem. *Bulletin of the American Mathematical Society*, 36(4):413–432, 1999.

[AJSD19] Abolfazl Asudeh, HV Jagadish, Julia Stoyanovich, and Gautam Das. Designing fair ranking schemes. In *Proceedings of the 2019 International Conference on Management of Data*, pages 1259–1276, 2019.

[AK10] Alexandr Andoni and Robert Krauthgamer. The computational hardness of estimating edit distance. *SIAM J. Comput.*, 39(6):2398–2429, 2010.

[AN10] Alexandr Andoni and Huy L. Nguyen. Near-optimal sublinear time algorithms for Ulam distance. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 76–86, 2010.

[ASCPX13] Hossein Azari Soufiani, William Chen, David C Parkes, and Lirong Xia. Generalized method-of-moments for rank aggregation. *Advances in Neural Information Processing Systems*, 26, 2013.

[BBD09] Therese Biedl, Franz J Brandenburg, and Xiaotie Deng. On the complexity of crossings in permutations. *Discrete Mathematics*, 309(7):1813–1823, 2009.

[BBGH15] Christian Bachmaier, Franz J. Brandenburg, Andreas Gleißner, and Andreas Hofmeier. On the hardness of maximum rank aggregation problems. *Journal of Discrete Algorithms*, 31:2–13, 2015. 24th International Workshop on Combinatorial Algorithms (IWOCA 2013).

[BCE+16] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

[BCPV96] Sanjoy K Baruah, Neil K Cohen, C Greg Plaxton, and Donald A Varvel. Proportionate progress: A notion of fairness in resource allocation. *Algorithmica*, 15(6):600–625, 1996.

[BCZ+16] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29, 2016.

[BKM22] Siddharth Barman, Arindam Khan, and Arnab Maiti. Universal and tight online algorithms for generalized-mean welfare. In *AAAI*, 2022.

[Bor10] Vani K Borooah. Social exclusion and jobs reservation in india. *Economic and Political Weekly*, pages 31–35, 2010.

[BS19] Mahdi Boroujeni and Saeed Seddighin. Improved MPC algorithms for edit distance and Ulam distance. In *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2019, Phoenix, AZ, USA, June 22-24, 2019.*, pages 31–40, 2019.

[BTT89] John Bartholdi, Craig A Tovey, and Michael A Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and welfare*, 6(2):157–165, 1989.

[Cas19] Carlos Castillo. Fairness and transparency in ranking. In *ACM SIGIR Forum*, volume 52, pages 64–71. ACM New York, NY, USA, 2019.

[CDK21] Diptarka Chakraborty, Debarati Das, and Robert Krauthgamer. Approximating the median under the ulam metric. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 761–775. SIAM, 2021.

[CGJ21] Diptarka Chakraborty, Kshitij Gajjar, and Agastya Vibhuti Jha. Approximating the center ranking under ulam. In *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.

[CHRG16] Matthew Costello, James Hawdon, Thomas Ratliff, and Tyler Grantham. Who views online extremism? individual attributes leading to exposure. *Computers in Human Behavior*, 63:311–320, 2016.

[CK06] Moses Charikar and Robert Krauthgamer. Embedding the Ulam metric into l_1. *Theory of Computing*, 2(11):207–224, 2006.

[CMS01] Graham Cormode, Shan Muthukrishnan, and Süleyman Cenk Sahinalp. Permutation editing and matching via embeddings. In *International Colloquium on Automata, Languages, and Programming*, pages 481–492. Springer, 2001.

[CMV22] Eden Chlamtác, Yury Makarychev, and Ali Vakilian. Approximating fair clustering with cascaded norm objectives. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 2664–2683. SIAM, 2022.

[CSV18] L. Elisa Celis, Damian Straszak, and Nisheeth K. Vishnoi. Ranking with fairness constraints. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, volume 107 of *LIPIcs*, pages 28:1–28:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[Des05] Ashwini Deshpande. Affirmative action in india and the united states. 2005.

[DG77] Persi Diaconis and R. L. Graham. Spearman's footrule as a measure of disarray. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(2):262–268, 1977.

[DHP+12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.

[DKNS01] Cynthia Dwork, Ravi Kumar, Moni Naor, and Dandapani Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th international conference on World Wide Web*, pages 613–622, 2001.

[Eur] Eu platform of diversity charters. `https://ec.europa.eu/info/policies/justice-and-fundamental-rights/combatting-discrimination/tackling-discrimination/diversity-and-inclusion-initiatives/eu-platform-diversity-charters_en`. Accessed: May, 2022.

[FKS03] Ronald Fagin, Ravi Kumar, and Dandapani Sivakumar. Efficient similarity search and classification via rank aggregation. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*, pages 301–312, 2003.

[GDL21] Sruthi Gorantla, Amit Deshpande, and Anand Louis. On the problem of underranking in group-fair ranking. In *International Conference on Machine Learning*, pages 3777–3787. PMLR, 2021.

[GL11] David F Gleich and Lek-heng Lim. Rank aggregation via nuclear norm minimization. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 60–68, 2011.

[GSB21] David García-Soriano and Francesco Bonchi. Maxmin-fair ranking: individual fairness under group-fairness constraints. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 436–446, 2021.

[Har92] Donna Harman. Ranking algorithms. In William B. Frakes and Ricardo A. Baeza-Yates, editors, *Information Retrieval: Data Structures & Algorithms*, pages 363–392. Prentice-Hall, 1992.

[HPS16] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.

[Kem59] John G Kemeny. Mathematics without numbers. *Daedalus*, 88(4):577–591, 1959.

[Ken38] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

[KMM15] Matthew Kay, Cynthia Matuszek, and Sean A Munson. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd annual acm conference on human factors in computing systems*, pages 3819–3828, 2015.

[KMS07] Claire Kenyon-Mathieu and Warren Schudy. How to rank with few errors. In *Proceedings of the thirty-ninth annual ACM Symposium on Theory of Computing*, pages 95–103, 2007.

[Kni66] William R. Knight. A computer method for calculating kendall's tau with ungrouped data. *Journal of the American Statistical Association*, 61(314):436–439, 1966.

[KR20] Caitlin Kuhlman and Elke Rundensteiner. Rank aggregation algorithms for fair consensus. *Proceedings of the VLDB Endowment*, 13(12), 2020.

[KV10] Ravi Kumar and Sergei Vassilvitskii. Generalized distances between rankings. In *Proceedings of the 19th international conference on World wide web*, pages 571–580, 2010.

[NSS17] Timothy Naumovitz, Michael E. Saks, and C. Seshadhri. Accurate and nearly optimal sublinear approximations to Ulam distance. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2012–2031, 2017.

[Pop07] V Yu Popov. Multiple genome rearrangement by swaps and by element duplications. *Theoretical computer science*, 385(1-3):115–126, 2007.

[PPM+22] Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. Fair ranking: a critical review, challenges, and future directions. *arXiv preprint arXiv:2201.12662*, 2022.

[Sch12] Warren Schudy. Approximation schemes for inferring rankings and clusterings from pairwise data. *Ph.D. Thesis*, 2012.

[SIN19] SIMRANDEEP SINGH. An analysis of the constitutionality of the economically weaker sections (ews) reservation. *Available at SSRN 3543312*, 2019.

[Spe04] C Spearman. The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1):72–101, 1904.

[Spe06] Charles Spearman. Footrule for measuring correlation. *British Journal of Psychology*, 2(1):89, 1906.

[USD] Us equal employment opportunity commission. `https://www.eeoc.gov`. Accessed: May, 2022.

[vdBLN+20] Jan van den Brand, Yin Tat Lee, Danupon Nanongkai, Richard Peng, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Bipartite matching in nearly-linear time on moderately dense graphs. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 919–930. IEEE, 2020.

[Ver10] Tània Verge. Gendering representation in spain: Opportunities and limits of gender quotas. *Journal of Women, Politics & Policy*, 31(2):166–190, 2010.

[WISR22] David Wei, Md Moinul Islam, Baruch Schieber, and Senjuti Basu Roy. Rank aggregation with proportional fairness. In *SIGMOD*, page To appear, 2022.

[YL78] H Peyton Young and Arthur Levenglick. A consistent extension of condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.

[You88] H Peyton Young. Condorcet's theory of voting. *American Political science review*, 82(4):1231–1244, 1988.

[ZBC+17] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. Fa* ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1569–1578, 2017.

[ZYS21] Meike Zehlike, Ke Yang, and Julia Stoyanovich. Fairness in ranking: A survey. *arXiv preprint arXiv:2103.14000*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [No] The current work is theoretical. The setting and models under study are clearly described in the paper.
   (c) Did you discuss any potential negative societal impacts of your work? [No] Given the theoretical nature of the work, we do not envision any potential negative societal impact.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] complete proofs are given in the appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]