Sampling Individually-Fair Rankings that are Always Group Fair

Sruthi Gorantla, Anay Mehrotra, Amit Deshpande, Anand Louis

IISc Microsoft Yale

Need of Fairness in Ranking Algorithms

Rankings are ubiquitous on online platforms and quintessential in finding relevant information quickly [Liu09; MRS10; ER15; Nob18; BW21]

Without fairness consideration, ranking algorithms output rankings that (1) skew representations of social groups–altering perception and polarizing opinions [ER15; KMM15] and (2) skew exposure at an individual level–leading to denial of economic opportunities to users [Han+17]

Group Fairness Constraints Alone Are Insufficient

A growing number algorithms try to output "fair" rankings by ensuring a desired representation of each protected group [PSK21; ZYS22]

While group fairness can mitigate under-representation across groups, they may not mitigate (and can even exacerbate) harms at individual level

E.g., specific individuals (whose utilities have high uncertainty) may be systematically assigned to poorer positions [Han+17; SKJ21]

Individual Fairness Requires Stochasticity

Recent work in fair ranking proposes to incorporate fairness, from the individuals' perspective, in rankings [SJ18; BGW18; SKJ21]. At a high level, they aim to equalize the exposure of individuals with similar "utilities"

Problem: Since exposure is linked to the position of individuals in the ranking, a single deterministic ranking cannot satisfy individual fairness

Example: For instance, suppose individuals i_1 and i_2 have similar utilities. Any deterministic ranking, must place one before the other, thereby imparting a systematically higher exposure to the earlier individual.

Idea: Introduce stochasticity in the output rankings [SJ18; BGW18; SKJ21]

However, satisfying individual fairness constraints alone, may introduce misrepresentation at the group-level

Goal: Output rankings satisfying both individual and group constraints

Model of Fair Ranking

Ranking [MRS10; Liu09]:

- There are m individuals or items and n positions
- Placing the i-th item at the j-th position gives $\rho_i \cdot v_j$ utility

Output: A ranking of n items with the highest utility

Fair Ranking [YS17; CSV18; SJ18…]:

Protected groups: A Laminar family of sets $G_1, G_2, \ldots, G_p \subseteq \{1, 2, \ldots, m\}$

Blocks: $B_1, B_2, \ldots, B_q \subseteq \{1, 2, \ldots, n\}$
Disjoint subsets of positions (e.g., different pages in web search)

Group fairness constraints: Lower/upper bounds specified by L/U
Require at least $L_{k\ell}$ and at most $U_{k\ell}$ items from G_ℓ in B_k
E.g., $L_{k\ell} = U_{k\ell} = |B_k|/2 \Rightarrow$ Equal representation constraints

Individual fairness constraints: Lower/upper bounds specified by C/A
Ensure item i appears in B_k with probability at least C_{ik} and at most A_{ik}

Output: Ranking maximizing expected utility subject to fairness constraints

Existing Algorithms Can Violate Group Fairness

Frameworks of some recent works can be adapted to incorporate both individual and group fairness constraints [SJ19; SJ18; SKJ21]

They solve a linear program (LP), whose solution is a distribution over rankings that (1) satisfies individual fairness constraints and (2) rankings sampled from it satisfy the group fairness constraints in expectation

Problem: Even though sampled rankings satisfy group fairness constraints in aggregate, each output ranking can violate group fairness constraints

Theoretical Guarantees on Sampling Group-Fair Rankings from Individually-Fair Distributions

An efficient algorithm that samples rankings from an individually-fair distribution while ensuring that every output ranking is group fair

Main Result. There is a polynomial time algorithm, that given L, U, A, C specifying fairness constraints, and ρ and v specifying utilities, outputs a ranking R sampled from distribution $\mathcal{D}$ such that

1. $\mathcal{D}$ satisfies (C, A)-individual-fairness constraints; and
2. R satisfies (L, U)-group-fairness constraints.

The expected utility of R is at least α-times the optimal, where, e.g.,

$$\alpha \geq \frac{v_1 + v_2 + \cdots + v_k}{k \cdot v_1}.$$

We present tighter bounds on α in the paper. α approaches 1 as either (1) the range of utilities or (2) the range of position-discounts shrinks.

E.g., with DCG position discounts [JK02] and $k = 1,2,3$, α is $\geq 0.8, 0.7, 0.6 \ldots$

Technical Approach and Challenges

Existing works [SJ19; SJ18; SKJ21] use the Birkhoff-von-Neumann algorithm [Bru82] to sample rankings from individually-fair distribution such that they satisfy group fairness constraints in aggregate

Idea: If the polytope of group-fair rankings is integral, one can use adaptations of the Birkhoff-von-Neumann algorithm

Problem: Group-fair rankings form a non-integral polytope (Section A)

Idea: Consider a family of "coarse rankings." A coarse ranking is the same as a ranking except that it does not order items within blocks

Fact: The polytope of coarse rankings is integral [PLN22] and, hence, one can use adaptations of the Birkhoff-von-Neumann algorithm

Algorithm: (1) Sample a group-fair coarse ranking from an individually-fair distribution and (2) Convert the sampled coarse ranking to a proper ranking (this step loses at most $(1-\alpha)$-fraction of the utility)

Empirical Results on Real-World Data

Data: Occupations data [CK20] which contains top 100 Google image search results for 96 occupation related queries. The data has crowd-sourced gender attributes–which we use as ground truth

Notion of Fairness: Equal representation across genders and individuals receive at least a certain exposure determined by their utility

Metrics: (1) Probability of violating equal representation, (2) deviation from the specified individual fairness constraints, and (3) DCG-based utility

Main Observation: Our algorithm outputs rankings which always satisfy the specified fairness constraints, while losing ≤6% utility with respect to baselines

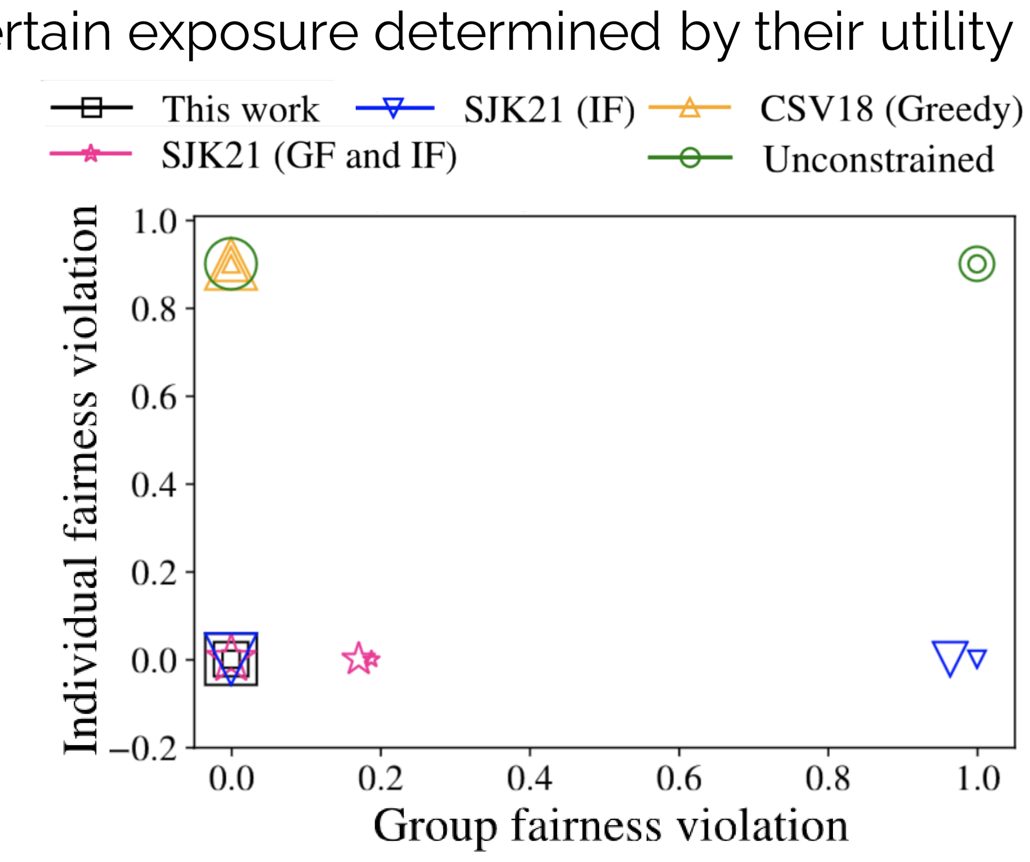

The paper also has empirical results with other datasets

Limitations and Future Work

We consider the utility of a ranking to be a linear function of the utilities items in the ranking. While this captures many applications [ZYS22a; ZYS22b; PSK21], in some applications, the utilities may be non-linear [Agr+09; Asa+22; KRT22].

Our algorithm works for Laminar family of groups, but not for an arbitrary family of protected groups