Fairness for Robust Learning to Rank

Omid Memarrast[1], Ashkan Rezaei[1], Rizal Fathony[2], Brian Ziebart[1]

[1]Deprtment of Computer Science, University of Illinois at Chicago
[2]Bosch Center for Artificial Intelligence, Carnegie Mellon University

Summary

- Conventional ranking systems: Maximize the utility of the ranked items to users.
- Fairness-aware ranking systems: Maximize the utility + balance the exposure for different protected attributes such as gender or race.
- **FairRobust LTR** : Fairness-aware + Distributionally Robust
- We derive a new ranking system based on the first principles of distributional robustness.
- Provide better utility for highly fair rankings than existing baseline methods.

Background and Notation

❑ **Ranking Problem**

- Rank a candidate set of items $D = \{d_1, d_2, \ldots, d_N\}$
- Ranking/Permutation π
- Permutation π places item d_i at rank j $\pi_i = j$
- Relevance judgment for d_i rel_i

❑ **Utility of Ranking** [general form]:

$$U(\pi) = \sum_{d_i \in D} u(d_i)\, v(j)$$

- $u(d_i)$: utility of a single item d_i
- $v(j)$: attention d_i gets by being placed at rank j by permutation π

❑ **DCG**: common evaluation measure in ranking system

$$DCG(\pi) = \sum_{d_i \in D} \frac{2^{rel(d_i)} - 1}{\log(1+j)}$$

❑ Searching the space of all rankings is Exponential in N

- **Solution**: use doubly stochastic matrix $\mathbf{P} \in \mathbb{R}^{N\times N}$ instead of permutation π.
- $P_{j,k}$: Probability that $\pi_j = k$

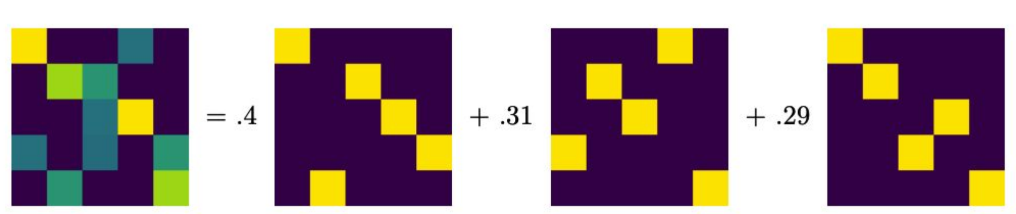

❑ Expected utility for a probabilistic ranking in a vectorized form: $U(\mathbf{P}) = \mathbf{u}^{\mathrm{T}}\mathbf{P}\mathbf{v}$

Fairness of Exposure in Ranking

- Fairness notions used in probabilistic ranking: $\mathbf{f}^T\mathbf{P}\mathbf{v} = h$ [1].
- $\mathbf{f}$: a vector to encode group membership and/or relevance of each document. h is scalar.

Distributionally Robust Learning to Rank

- The fair ranking optimization can be expressed as a linear programming problem (post-processing) [1]:

$$\max_{\mathbf{P}\in\Delta\cap\Gamma_{fair}} \mathbf{u}^{\mathrm{T}}\mathbf{P}\mathbf{v} \quad where: \Delta: \mathbf{P1} = \mathbf{P}^T\mathbf{1}, \quad \mathbf{P}_{j,k} \geq 0, \forall\, 1 \leq j,k \leq M$$

❑ **Extension:** Derive an LTR approach. [in-processing]

- Learn to optimize utility and fairness simultaneously.

Definition: The fair probabilistic ranking $\mathbf{P} \in \mathbb{R}^{M\times M}$ in adversarial LTR learns a fair ranking that maximizes the worst-case ranking utility approximated by an adversary $\mathbf{q}(\breve{\mathbf{u}})$, constrained to match the feature statistics of the training data:

$$\max_{\mathbf{P}\in\Delta\cap\Gamma_{fair}} \min_{\mathbf{q}} \mathbb{E}_{\mathbf{X}\sim\tilde{P}}\, [\mathrm{U}(\mathbf{X}, \mathbf{P}, \mathbf{q})]$$

$$s.t.\quad \mathbb{E}_{\mathbf{X}\sim\tilde{P}}\left[\sum_{j=1}^{M} \mathbb{E}_{\breve{u}_j|\mathbf{X}\sim\mathbf{q}}\left[\breve{u}_j \mathbf{X}_{j,:}\right]\right] = \mathbb{E}_{\mathbf{X},\mathbf{u}\sim\tilde{P}}\left[\sum_{j=1}^{M} u_j \mathbf{X}_{j,:}\right]$$

Optimization

- Solve the formulation in Lagrangian dual form:

$$\max_{\theta} \mathbb{E}_{\mathbf{X},\mathbf{u}\sim\tilde{P}}\left[\max_{\mathbf{P}\in\Delta} \min_{0\leq\mathbf{q}\leq 1} \mathbf{q}^{\mathrm{T}}\mathbf{P}\mathbf{v} + \left\langle \mathbf{q} - \mathbf{u}, \sum_{l} \theta_l\, \mathbf{X}_{:,l} \right\rangle + \lambda \mathbf{f}^{\mathbf{T}}\mathbf{P}\mathbf{v}\right]$$

- Optimize the dual parameters $\boldsymbol{\theta} \in \mathbb{R}^{L\times 1}$ for the feature matching constraint of L features by gradient decent.
- Use λ as a penalty parameter for fairness constraint.

- $\mathbf{u} \in \mathbb{R}^{M\times 1}$ is a vector of true labels.
- $\mathbf{q} \in \mathbb{R}^{M\times 1}$ is a probability vector of the adversary's estimation of each item being relevant.
- $\mathbf{X}_{:,l} \in \mathbb{R}^{M\times 1}$denotes the l_{th} feature of M samples.
- S is the set of protected attributes
- $\mathbf{v} \in \mathbb{R}^{M\times 1}$ is a vector containing the values of position bias function for each position.

Ranker player[P]: Choosing a distribution over rankings constrained to provide fairness while maximizing utility.

Adversary player[q]: Choosing a distribution of item relevancies that minimizes utility while being similar to training data properties.

Experiments

- Construct LTR dataset from fair classification datasets.
- **Goal:** provide low fairness violation and high NDCG
- **Fair_Robust** achieves preferable utility-fairness trade-off

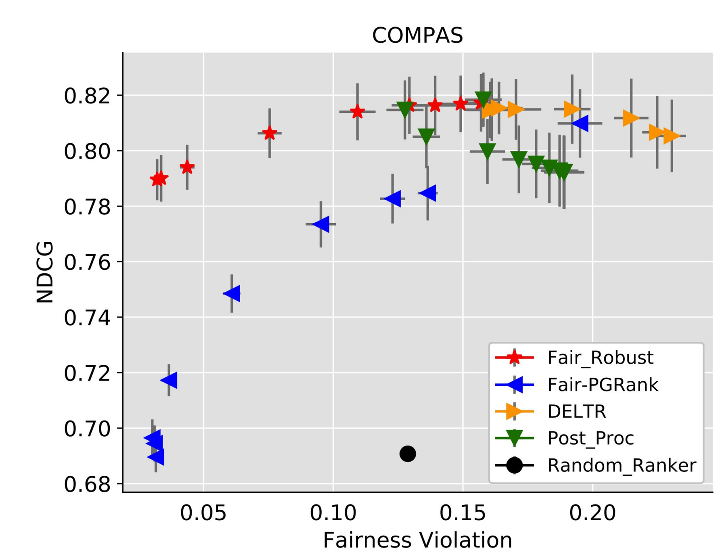

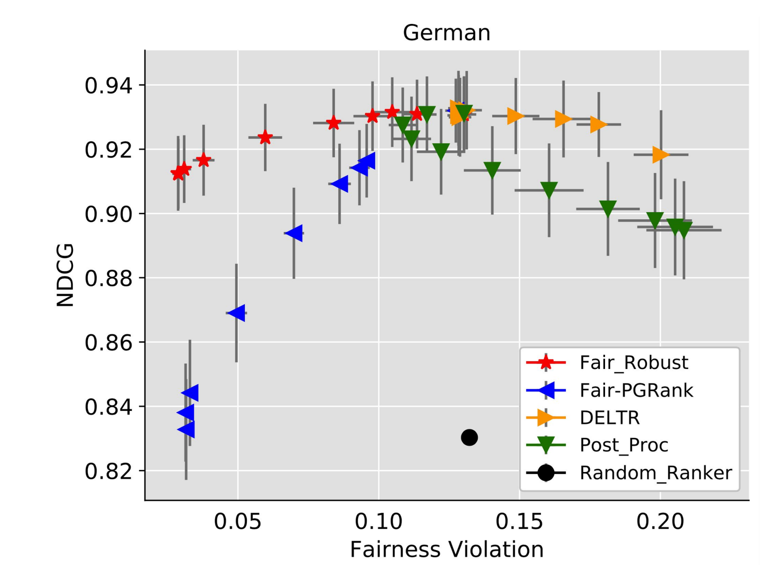

❑ Performance-Fairness Trade-off:

Y axis: Performance (top is better)
X axis: Fairness violation (left is better)

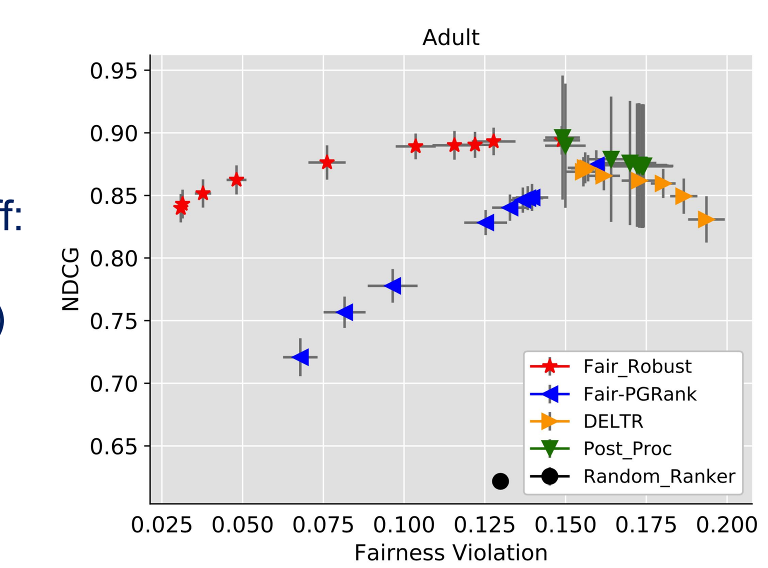

❑ Fairest results achieved by **Fair-Robust** and baseline methods.

Method \ Dataset	COMPAS NDCG	COMPAS $\hat{D}_{group}$	German NDCG	German $\hat{D}_{group}$
Fair_Robust	**0.789 ± 0.007**	0.032 ± 0.001	**0.912 ± 0.011**	**0.029 ± 0.002**
Fair_PGRank	0.696 ± 0.006	**0.030 ± 0.001**	0.838 ± 0.015	0.031 ± 0.001
DELTR	0.815 ± 0.010	0.160 ± 0.009	0.933 ± 0.011	0.128 ± 0.003
Post_Proc	0.818 ± 0.009	0.158 ± 0.007	0.927 ± 0.011	0.108 ± 0.005

References

[1] Singh, Ashudeep, and Thorsten Joachims. "Fairness of exposure in rankings." In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2219-2228. 2018.