

Equality of Opportunity in Ranking: a Fair-Distributive Model

Original

Equality of Opportunity in Ranking: a Fair-Distributive Model / Beretta, Elena; Vetrò, Antonio; Lepri, Bruno; De Martin, Juan Carlos. - ELETTRONICO. - (2021), pp. 51-63. (Intervento presentato al convegno Second International Workshop on Algorithmic Bias in Search and Recommendation (Bias 2021) tenutosi a Online nel April 1, 2021) [10.1007/978-3-030-78818-6_6].

Availability:

This version is available at: 11583/2898812 since: 2021-08-06T09:30:33Z

Publisher:

Springer

Published

DOI:10.1007/978-3-030-78818-6_6

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Springer postprint/Author's Accepted Manuscript

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/10.1007/978-3-030-78818-6_6

(Article begins on next page)

Equality of Opportunity in Ranking: a Fair-Distributive Model

Elena Beretta^{1,3}[0000–0002–6090–2765], Antonio Vetrò¹[0000–0003–2027–3308],
Bruno Lepri³[0000–0003–1275–2333], and Juan Carlos De
Martin¹[0000–0002–7867–1926]

¹ Nexa Center for Internet & Society, DAUIN, Politecnico di Torino, Italy
{elena.beretta,antonio.vetro,demartin}@polito.it

² Fondazione Bruno Kessler, Italy
lepri@fbk.eu

Abstract. In this work, we define a Fair-Distributive ranking system based on Equality of Opportunity theory and fair division models. The aim is to determine the ranking order of a set of candidates maximizing utility bound to a fairness constraint. Our model extends the notion of protected attributes to a pool of individual's *circumstances*, which determine the membership to a specific *type*. The contribution of this paper are i) a Fair-Distributive Ranking System based on criteria derived from distributive justice theory and its applications in both economic and social sciences; ii) a class of fairness metrics for ranking systems based on the Equality of Opportunity theory. We test our approach on an hypothetical scenario of a selection university process. A follow up analysis shows that the Fair-Distributive Ranking System preserves an equal exposure level for both minority and majority groups, providing a minimal system utility cost.

Keywords: fairness in rankings · algorithmic fairness · position bias · algorithmic bias · distributive justice · equality of opportunity

1 Introduction

Ranking systems have rapidly spread in nowadays economies: despite such tools have been widely employed since decades in Information Retrieval field [21], they have recently come back at the cutting edge thanks to the explosive growth of computational power and data availability [15]. Ranking is one of the predominant forms by which both online and offline software systems present results in a wide variety of domains ranging from web search engines [19] to recommendation systems [21]. The main task of ranking systems is to find an allocation of elements to each of the n positions so that the total value obtained is maximized. The key technical principle that for decades has driven this optimization is the *Probability Ranking Principle* [23], according to which elements are ranked in descending order depending on their probability of relevance for a certain query q . Consequently, each element will have a probability of exposure given by its

relevance for the query [27]. It is widely recognized that the position of an element in the ranking has a crucial influence on its exposure and its success; hence, systems with ranking algorithms, whose only task is to maximize utility, do not necessarily lead to fair or desirable scenarios [13]. In fact, a number of researches [33],[12],[32] have demonstrated that rankings produced in this way can lead to the over-representation of one element to the detriment of another, causing forms of algorithmic biases that in some cases can lead to serious social implications. Web search engine results that inadvertently promote stereotypes through over-representation of sensitive attributes such as gender, ethnicity and age are valid examples [30] [1] [16] [37]. In order to mitigate and overcome biased results, researchers have proposed a number of fairness metrics [3]. However, the majority of these studies formalizes the notion of equity only in supervised machine learning systems, keeping equity in ranking systems a poorly explored ground despite the increasing influence of rankings on our society and economy. The lower attention devoted to this field is probably due to the complexity of ranking and recommendation systems, which are characterized by dynamics difficult to predict, multiple models and antithetic goals, and are difficult to evaluate due to the great sparsity (e.g., see [6], [7] and [20]). In addition, a leading path to exploring the trade-off between the expected utility of a ranking system and its fairness has not yet been mapped out. We address these challenges through developing a multi-objective ranking system that optimizes the utility of the system and simultaneously satisfies some ethical constraints. Our model is inspired by *fair division* models [34] dealing on how to divide a set of resources among a series of individuals. The main contributions of this article are the following: first, we introduce a Fair-Distributive Ranking System combining methods of supervised machine learning and criteria derived from economics and social sciences; secondly, we define a class of fairness metrics for ranking systems based on the Equality of Opportunity theory [25]. Finally, we conduct an empirical analysis to study the trade-off between fairness and utility in ranking systems.

2 Related Work

Several recent works have addressed the issue of fairness in ranking systems. Some studies minimize the difference in representation between groups in a ranking through the concept of demographic parity, that requires members of protected groups are treated similarly to the advantaged ones or to the entire population [2], [36], [35], [26]. In particular, Yang and Stoyanovich [35] have dealt with this issue as a multi-objective programming problem, while Celis *et al.* [8] have approached it from the perspective of the ranking results' diversification, as in [36]. More recently, Asudeh *et al.* [2] have proposed a class of fair scoring functions that associates non-negative weights to item attributes in order to compute an item score. Singh and Joachims [28] have proposed a Learning-to-Rank algorithm that optimizes the system's utility according to the merit of achieving a certain level of exposure. Lastly, some recent studies have investigated the notion of fairness through equalizing exposure; in this specific strand studies differ

in the way exposure is allocated: while Biega *et al.* [4] have investigated individual fairness alongside utility, Singh and Joachims [27] have proposed an optimal probabilistic ranking to equalize exposure among groups. It is worth noting that the majority of the previous works have established equity constraints reflecting demographic parity constraints by narrowing the elements fraction for each attribute in the ranking or balancing utility with exposure merit. The methodology we propose goes beyond these parity and merit constraints for several reasons: a) the protected attributes are not a-priori established but are updated on the basis of the sample features, and b) the exposure is defined on the basis of the *effort* variable; this variable represents the real effort that the elements have made to be included in the Top-N-rank according to Roemer’s Equality of Opportunity theory [25].

3 Problem statement

In Information Retrieval the notion of utility is commonly stated as the expected ranking that maximizes the system’s utility by exploiting the n th ranking r and a query q , such that $\text{argmax}U(\text{ranking}_r|q)$, where $r = 1..R$ (R is the rankings set); it is generally achieved through a series of utility measures in ranking system domain that leverage a mapping function β to detect the relevance of an item to each user given a certain query q , $\beta(\text{Rel}(\text{item}_i|\text{user}_u, q))$, where $i = 1..I$ and $u = 1..U$ (I and U are the items set and the users set). Several recent works establish a certain exposure degree to each individual or group of individuals as a fairness constraint. The exposure indicates the probability of attention that the item gets based on the query and its ranking position, and is generically calculated as $\frac{1}{\log(1+j)}$, where j is the position of the item_i in the ranking_r . We adapt the example proposed by Singh and Joachims[27] to our scenario: suppose a group of students has to enroll at university, the decision-maker then sorts the students according to their relevance for the expressed query and draws up a certain number of rankings to evaluate the system response accuracy. Relevance is thus derived from the probability that the candidate is relevant for the query.

In this example, 8 individuals are divided into 3 groups based on ethnicity attribute. Individuals belonging to the white group have relevance 1, 0.98, 0.95, the Asians have 0.93, 0.91, 0.88, the African-Americans 0.86, 0.84. Students are sorted in ranking according to relevance. Since exposure is a measure exploiting relevance and ranking position, it is computed after sorting. As shown in Figure 1a, Asian and African-American students, despite being placed a few positions below white ones, get a very low exposure; this means their average exposure is significantly lower compared to the white group, despite a minimal deviation in relevance. Efforts to enforce a fairness constraint on exposure, even if important, are missing the real point that is instead tied to relevance. As a matter of fact, exposure is calculated on the basis of the candidate’s position, regardless of the student’s traits. Consider the new ranking in Figure 1b. In this case, a fairness constraint is applied to proportionally allocate exposure among ethnic groups; despite the constraint, the African-American minority remains in




Fig. 1: Both pictures 1a and 1b revise the example of Singh and Joachims [27]. Blue: white group; green: Asian group; red: African-American group.

lower positions compared to the other two groups. This problem is even more serious in case of binary relevance: assuming the decision-maker would admit to the university the first 3 students, no African-American individuals would be included in the top-3-ranking. To address the problem of fairness in rankings we suggest to marginally consider exposure and to focus on analyzing how relevance is computed and which features correlate with the query q . This means that a ranking is considered unfair if the students' relevance, hence their position, is systematically established on the basis of irrelevant features such as protected attributes.

4 A Fair-Distributive Ranking Model

Preliminary Egalitarian theories [24] such as EOp arise from the notion of distributive justice, which recognizes that all goods should be equally distributed across society. The key principle of Roemer's Equality of Opportunity (EOp) theory is based on the assumption that the resources obtained by individuals depend on two factors: individual choices, which lie within the sphere of personal responsibility, and circumstances, which are exogenous to individual control. He claims that if inequalities in a set of individuals are caused by birth circumstances, which include variables such as gender, race, or familiar socio-economic status and so forth, then these are morally unacceptable and must be compensated by society. The theory is therefore based on four key principles: circumstances, effort, responsibility and reward. Responsibility is a theoretical notion reflecting the effort degree that individuals invest in achieving the acts they perform. The reward is the fraction of resources that individuals belonging to a disadvantaged group get in case an inequality of opportunity occurs, and it is established by a certain policy [9],[22]. According to Roemer, policies should

be oriented to equalize the opportunities that different *types*, or groups of individuals, categorized in accordance with diverse circumstances, must be able to have in order to achieve a given goal. A *type* is a set of individuals sharing the same circumstances, while the set of individuals characterized by the same degree of effort is called a *tranche*. Below We provide our notion of fair ranking according to Roemer’s EOp theory:

A ranking is said to be fair according to Roemer’s EOp theory if all individuals belonging to different types have the same opportunity to reach high ranking position, and if relevance and exposure are given by personal responsibility and not by birth circumstances.

In order to offer a better understanding of our notation, a summary of some basic notions of EOp’s theory is briefly provided:

- i **Effort**: proxy indicating individual’s personal responsibility. According to Roemer, effort is derived by computing the Cumulative Distribution Function of each type, and then by extracting the quantile’s distribution.
- ii **Types**: groups of individuals, categorized in accordance with diverse circumstances;
- iii **Tranche**: the set of individuals characterized by the same degree of effort;
- iv **Policy**: actions oriented to equalize type’s opportunity;
- v **Utility**: a measure of the system’s ability to provide relevant results. In the case of ranking systems, the utility is defined as the ability to provide the best ranking according to a certain query q .

Notation According to Roemer, individuals are fully described by a two-dimensional list: the first one is the set of circumstances beyond individuals’ control denoting the population partition into a set of non-overlapping types T characterized by certain attributes; the second one expresses the traits for which individuals are fully responsible and is described by the scalar variable effort π . Given a certain query q , the Fair-Distributive Ranking Γ is the one in which the utility maximization (e.g., the best candidate selection) is a function of circumstances, degrees of effort and a policy:

$$\Gamma = \operatorname{argmax}_{\theta \in \Theta} u^t(q|e_i(\pi), \theta) \quad (1)$$

As a corollary, exposure is allocated equally among individuals: ranking Γ is therefore the one that maximizes the area below the lowest function \exp^t , i.e. the type-specific-exposure:

$$\max_{\theta \in \Theta} \int_0^1 \min_t \exp^t(\pi, \theta) d\pi \quad (2)$$

Equation 2 denotes the opportunity-equalizing policy in ranking according to EOp theory.

Fair-distributive ranking criteria Ranking Γ has to show the following properties: (i) *Ranked type fairness*. Γ is exploited by overcoming the a-priori assessment of the sensitive attributes through partitioning population in types;

(ii) *Circumstances-effort based ranking*. Γ is based on a counterfactual score obtained from effort and circumstances' variables; the counterfactual score indicates which score individuals would have got if they were not belonging to their type T ; and (iii) *Counterfactual-ordered ranking utility*. Γ is ordered by decreasing counterfactual score.

4.1 Ranked type fairness

To split the population in types we perform the conditional inference trees method proposed by Hothorn [14]; to the best of our knowledge, the only application of this algorithm in exploring inequalities has been carried out by Brunori [5] to study socio-economic differences on panel data. The algorithm exploits the permutation test theory developed by Strasser [29] to generate recurring binary partitions and recursively split the Euclidean space of the individuals' variables in convex sets of hyperplanes, overcoming the problem of overfitting and variable selection [18]. Let $P(Y|X)$ be a conditional probability distribution of the response Y given the vector of n covariates $X = X_1, \dots, X_n$; the recursion measures correlation among Y and the X vector and performs multiple hypothesis tests to assess correlation significance. If a statistically significant correlation doesn't exist, recursion stops. Formally, types are permuted in this way:

$$T_k = \begin{cases} S_i^X & \text{if } H_0^i : P(Y|X_i) = P(Y), \\ \text{recursion stops} & \text{otherwise.} \end{cases} \quad (3)$$

Where S_i^X is the set of x^i possible realizations.

4.2 Circumstances-effort based ranking and counterfactual-ordered ranking utility

Since effort is not directly observable, we need a proxy in order to measure it. Roemer argues that it exists an effort distribution function that characterizes the entire subgroup within which the location of the individual is set and what is needed is a measure of effort that is comparable between different types. The basic assumption is that two individuals belonging to a different type t who occupy the same position in their respective distribution functions have exerted the same level of effort - and therefore of responsibility. Since, under the same circumstances, individuals who make different choices exercise different degrees of effort - and thus achieve a different outcome -, the differences in outcome within the same type are by definition determined by different degrees of effort, and therefore are not considered in the computation of the EOp. In general, Roemer states that to estimate effort it is necessary to:

- I aggregate individuals according to their circumstances;
- II compare outcome distributions;
- III measure the effort that an individual has exerted through the quantile occupied in his or her type distribution.

Consequently, all the individuals positioned at the same quantile in the distribution of the respective type are by assumption characterized by the same level of effort. Hence, the counterfactual score distribution \tilde{y} is computed following these steps:

1. approximate each type-score cumulative distribution $ecdf(y)$ through ten-fold cross-validation on Bernstein polynomials log-likelihood: $LL_B(p_m = \sum_{i=1}^n \log f_B(x_j, p_m))^3$;
2. identify the degree of effort through the type-score distribution $y^k(\pi)$;
3. apply a smoothing transformation on $y^k(\pi)$ by scaling each tranche distribution (i.e., individuals have exerted same degree of effort) until individuals have the same tranche-score and then multiplying the tranche-score by each individual's score. The result of the smoothing process is a standardized distribution \tilde{y} where all the unexplained inequalities are removed, i.e., only inequalities due to circumstances or degrees of effort are observed.

Inequality of opportunity is captured by applying an inequality index on the standardized distribution \tilde{y} . The counterfactual score $u_\pi^t(\theta)$ is then computed by assigning to individuals the fraction of inequality of their respective quantile-type-score distribution $y^k(\pi)$. Finally, individuals are ranked according to their counterfactual score in decreasing order. The ranking Γ is therefore given by:

$$\Gamma = u_\pi^t(\tilde{y}, \text{IneqIndex}(\tilde{y})) \quad (4)$$

5 Experiment

5.1 Data and Settings

We design an hypothetical scenario where high school students compete to be admitted at university (i.e. a typical scenario in several countries). Firstly, the whole students' population is simply ordered by decreasing counterfactual score (relevance is treated as not binary); then, the Fair-Distributive ranking Γ for fair-top-N-rankings is studied (as the available positions are assumed as finite, relevance is treated as binary). The conducted experiments compare three types of ranking: the first one is based on real students' test scores serving as a benchmark; the second one is based on the standardized distribution \tilde{y} ; the last one is our ranking Γ based on counterfactual score $u_\pi^t(\theta)$. For the analysis, we employ the Student Performance Dataset [10] reporting students test scores of two Portuguese schools. Dataset contains 649 instances and 33 attributes, especially demographic ones (i.e. gender, parents' education, extra educational support).

5.2 Metrics

We apply three types of metrics in order to fulfill both ranking and fairness constraints: i) ranking domain metrics. ii) inequality domain metrics (i.e., Gini

³ For a detailed explanation of Bernstein Polynomials Log Likelihood, see [17],[38],[5]

index and Theil index), and iii) a set of metrics we propose to study our fairness constraints (i.e., Opportunity Loss/Gain Profile, Opportunity Loss/Gain Set, Unexplained Inequality Rate, Reward Profile, Reward Rate). Regarding the inequality metrics, Gini index is a statistical concentration index ranging from 0 to 1 that measures the inequality degree distribution [11]; a low or equal to zero Gini index indicates the tendency to the equidistribution and expresses perfect equality, while a high or equal to 1 value indicates the highest concentration and expresses maximum inequality. Theil index [31] is a measure of entropy to study segregation; a zero Theil value means perfect equality. Finally, we have proposed a new fairness metrics set: the Opportunity-Loss/Gain Profile and the Opportunity-Loss/Gain Set are computed to study inequality in the original distribution. They indicate which levels of score could be reached by each type with different effort degrees. The Unexplained Inequality Rate calculates the amount of fair removed inequality due to individuals' responsibility. The Reward Profile calculates the type that obtained the highest gain/loss from the re-allocation of scores - i.e. after applying fairness constraints; while the Reward Rate calculates the average re-allocation score rate for each type. All formulas are summarized in Table 1.

Metric	Formula	Input	Metrics domain
Gini Index	$1 - \frac{1}{\mu} \int_0^{\infty} (1 - F(y))^2 dy$	All Distr.	Inequality
Theil Index	$\frac{1}{N} \sum_{i=1}^N \ln(\frac{\mu}{y_i})$	All Distr.	Inequality
Opportunity-L/G Profile	$\min/\max(y_{\pi}^t - \mu(y_{\pi}))$	Sc.Distr.	New set of fairness measures
Opportunity-L/G Set	$y_{\pi}^t - \mu(y_{\pi})$	Sc.Distr.	New set of fairness measures
Unexplained Inequality Rate	$\frac{1}{N} \sum y_i - \tilde{y}_i$	Sc.Distr., Stnd.Sc.Distr.	New set of fairness measures
Reward Profile	$\min/\max(y_{\pi}^t - adj(\tilde{y}_{\pi}^t))$	Sc.Distr., Adj Sc.Distr.	New set of fairness measures
Reward Rate	$y_{\pi}^t - adj(\tilde{y}_{\pi}^t)$	Sc.Distr., Adj Sc.Distr.	New set of fairness measures

Table 1: Summary of inequality domain metrics and of a set of novel metrics proposed to study fairness constraints. *Notation:* $F(y)$ = cumulative distribution function of the score, μ = mean score; R = number of types, p_i = frequency of types; y_{π}^t = score distribution aggregated by type and quantile; \tilde{y}_i = standardized score; $adj(\tilde{y}_{\pi}^t)$ = adjusted mean-type score at each effort degree (after policy).

5.3 Results and Discussion

Table 2 summarizes our main results: first, we observe that the Γ ranking shows less inequality compared to benchmark ranking, as emerges from Gini's index trend. Both ranking exhibit same value of Theil index, revealing entropy is similarly balanced. By observing the Outcome Set, we notice that D and F types get the lower average outcomes for all degrees of effort; this doesn't necessarily mean they are in a disadvantaged position. There are indeed multiple reasons, which do not necessarily indicate inequality, why some types systematically show a lower average outcome. We compute Gini Index on the standardized distribution to observe if there are types that systematically receive a lower outcome due

Table 2: Summary of Results

Initial outcome					Policy outcome						
Overall results											
Gini	0.144					0.112					
Theil	0.025					0.025					
Unexp.Ineq.	-					0.126					
Tranche results											
	1	2	3	4	5		1	2	3	4	5
Outcome.Set A	0.02	0.01	-0.04	0.04	0.04	Reward A	0.11	-0.11	-0.04	-0.09	-0.08
B	0.09	0.09	0.08	0.16	0.17	B	0.06	-0.15	-0.12	-0.18	-0.21
C	0.07	0.12	0.02		0.12	C	0.07	-0.21	-0.08	-	-0.14
D	-0.08	-0.22	-0.07	-0.16	-0.31	D	-0.07	0.13	-0.01	0.09	0.28
E	0.03	0.05	0	0.03	0	E	0.14	-0.21	-0.09	-0.12	-0.01
F	-0.11	-0.06	-0.05	-0.08	-0.07	F	0.22	-0.05	-0.03	0.03	0.04
G	-0.01	0.01	0.06	0.01	0.04	G	0.16	-0.07	-0.17	-0.09	-0.09


to their type membership. In this way, only and exclusively inequalities caused by circumstances and effort degrees are obtained. This explains why types showing a lower average outcome are not directly compensated with higher rewards. Reward Rate is expressed as a dispersion around the quantile mean for each type, thus showing that it doesn't produce a significantly change in expected ranking utility. The aggregated exposure is computed recursively on all rankings with $n+1$ individuals. Analysis shows extremely disproportionate exposure values for the initial score ranking for all top-N-ranking. The Γ ranking keeps a proportionate aggregate exposure level among types for large subsets, while for smaller subsets it tends to favor individuals in groups which displays high levels of inequality (Figure 2a). Overall, these results indicate that our approach produces a fair ranking Γ with minimal cost for utility, that we compute in terms of relevance for the query "best candidates selection" (Figure 2b).

6 Conclusions

The method we have proposed generates a ranking with a guaranteed fair division score, with minimal cost for utility. Our ranking is based on a counterfactual score indicating which score students would have gotten if they had not belonged to different type. In this sense, the ranking is drawn up on the basis of the effort (aka, individual responsibility) that individuals have exerted to reach the initial score. As a result, our ranking presents equal opportunities for each group of individuals exhibiting the same circumstances (*types*) to achieve high ranking positions (high relevance) and good exposure rates. Moreover, the paper provides a set of new metrics to measure fairness in Fair-Distributive ranking. Finally, we study the trade-off between the aggregated type exposure and the system's utility: our analyses show that the counterfactual score doesn't affect significantly the expected ranking utility and preserves level of exposure proportionally among



(a) Barplot of aggregate exposure for the top-50, top-150 and top-300, initial score ranking, standardized ranking and T ranking.



(b) Utility results on the Initial and T ranking in terms of relevance for a query q

Fig. 2: Exposure and Relevance Results

groups. The method presented has some limitations, including for example: the need to have a dataset containing several demographic attributes, and the need to have one or more target variables to calculate conditional inference trees (alternatively, the method is subordinate to the construction of score indices). As far as next steps are concerned, it is important to i) verify the robustness of the model (internal validity) with larger datasets (also synthetic ones) and ii) verify the external validity of the approach by applying the model on different fields of application. In the long term, our intention is to implement a ranking simulator that tests the results of different distributive justice theories.

References

1. ALRossais, N.A., Kudenko, D.: Evaluating stereotype and non-stereotype recommender systems. In: Proceedings of the 12th ACM Conference on Recommender Systems. RecSys '18, ACM, Vancouver, Canada (2018)

2. Asudeh, A., Jagadish, H.V., Stoyanovich, J., Das, G.: Designing fair ranking schemes. In: Proceedings of the 2019 International Conference on Management of Data. p. 1259–1276. SIGMOD '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3299869.3300079>
3. Barocas, S., Hardt, M., Narayanan, A.: Fairness and machine learning (2018), <http://www.fairmlbook.org>
4. Biega, A.J., Gummadi, K.P., Weikum, G.: Equity of attention: Amortizing individual fairness in rankings. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. p. 405–414. SIGIR '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3209978.3210063>
5. Brunori, P., Neidhöfer, G.: The evolution of inequality of opportunity in germany: A machine learning approach. ZEW - Centre for European Economic Research Discussion **20** (2020). <https://doi.org/10.2139/ssrn.3570385>
6. Burke, R.: Multisided fairness for recommendation (jul 2017), <http://arxiv.org/abs/1707.00093>
7. Burke, R., Sonboli, N., Ordóñez-Gauger, A.: Balanced neighborhoods for multisided fairness in recommendation. In: Conference on Fairness, Accountability and Transparency in Proceedings of Machine Learning Research. vol. 81, pp. 202–214. Proceedings of Machine Learning Research (PMLR), New York, NY, USA (23–24 Feb 2018), <http://proceedings.mlr.press/v81/burke18a.html>
8. Celis, L.E., Straszak, D., Vishnoi, N.K.: Ranking with fairness constraints (2018)
9. Checchi, D., Peragine, V.: Inequality of opportunity in italy. Journal of Economic Inequality **8**(4), 429–450 (2010). <https://doi.org/10.1007/s10888-009-9118-3>
10. Cortez, P., Silva, A.: Using data mining to predict secondary school student performance. In: Proceedings of 5th FUTURE BUSINESS TECHNOLOGY CONFERENCE (FUBUTEC 2008). pp. 5–12. A. Brito and J. Teixeira Eds. (April 2008), <https://archive.ics.uci.edu/ml/datasets/Student+Performance>
11. Gastwirth, J.L.: The estimation of the lorenz curve and gini index. The Review of Economics and Statistics **54**, 306–316 (1972)
12. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. p. 3323–3331. NIPS'16, Curran Associates Inc., Red Hook, NY, USA (2016)
13. Helberger, N., Karppinen, K., D'Acunzio, L.: Exposure diversity as a design principle for recommender systems. Information, Communication & Society **21**(2), 191–207 (2016)
14. Hothorn, T., Hornik, K., Zeileis, A.: Unbiased recursive partitioning: A conditional inference framework. Journal of Computational and Graphical Statistics **15**(3), 651–674 (2006)
15. Irfan, S., Babu, B.V.: Information retrieval in big data using evolutionary computation: A survey. In: 2016 International Conference on Computing, Communication and Automation (ICCCA). pp. 208–213. IEEE, New York, NY, USA (2016)
16. Karako, C., Mangala, P.: Using image fairness representations in diversity-based re-ranking for recommendations. In: Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization. pp. 23–28. UMAP '18, ACM, Singapore, Singapore (2018). <https://doi.org/10.1145/3213586.3226206>
17. Leblanc, A.: On estimating distribution functions using bernstein polynomials. Annals of the Institute of Statistical Mathematics **64**, 919–943 (2012). <https://doi.org/10.1007/s10463-011-0339-4>

18. Li Donni, P., Rodríguez, J.G., Rosa Dias, P.: Empirical definition of social types in the analysis of inequality of opportunity: a latent classes approach. *Social Choice and Welfare* **44**(3), 673–701 (2015). <https://doi.org/10.1007/s00355-014-0851-6>
19. Makhijani, R., Chakrabarti, S., Struble, D., Liu, Y.: Lore: A large-scale offer recommendation engine with eligibility and capacity constraints. In: Proceedings of the 13th ACM Conference on Recommender Systems. p. 160–168. RecSys '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3298689.3347027>
20. Patro, G.K., Biswas, A., Ganguly, N., Gummadi, K.P., Chakraborty, A.: Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In: Huang, Y., King, I., Liu, T., van Steen, M. (eds.) WWW '20: The Web Conference 2020. pp. 1194–1204. ACM / IW3C2, Taipei, Taiwan (2020). <https://doi.org/10.1145/3366423.3380196>
21. Pei, C., Zhang, Y., Zhang, Y., Sun, F., Lin, X., Sun, H., Wu, J., Jiang, P., Ge, J., Ou, W., Pei, D.: Personalized re-ranking for recommendation. In: Proceedings of the 13th ACM Conference on Recommender Systems. p. 3–11. RecSys '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3298689.3347000>
22. Ramos, X., Van de Gaer, D.: Approaches to inequality of opportunity: Principles, measures and evidence. *Journal of Economic Surveys* **30**, 855–883 (2016). <https://doi.org/10.1111/joes.12121>
23. Robertson, S.E.: The Probability Ranking Principle in IR, p. 281–286. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1997)
24. Roemer, J.E.: A pragmatic theory of responsibility for the egalitarian planner. *Philosophy and Public Affairs* **22**(2), 146–166 (1993)
25. Roemer, J.E., Trannoy, A.: Equality of opportunity. *Handbook of Income Distribution* **2**(2), 217–300 (2015)
26. Singh, A., Joachims, T.: Equality of opportunity in rankings (2017), https://www.k4all.org/wp-content/uploads/2017/09/WPOC2017_paper9.pdf
27. Singh, A., Joachims, T.: Fairness of exposure in rankings. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. p. 2219–2228. KDD '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3219819.3220088>
28. Singh, A., Joachims, T.: Policy learning for fairness in ranking. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32, pp. 5426–5436. Curran Associates, Inc., New York, US (2019)
29. Strasser, Helmut and Weber, C.: On the asymptotic theory of permutation statistics. *Mathematical Methods of Statistics* **2**, 220–250 (1999)
30. Sweeney, L.: Discrimination in online ad delivery. *Queue - Storage* **11**(3), 10:10–10:29 (2013). <https://doi.org/10.1145/2460276.2460278>
31. Theil, H.: *Statistical Decomposition Analysis: With Applications in the Social and Administrative Sciences*, vol. 14. North-Holland Publishing Company, Amsterdam, Holland (1972)
32. Tsintzou, V., Pitoura, E., Tsaparas, P.: Bias disparity in recommendation systems (nov 2018)
33. Weinsberg, U., Bhagat, S., Ioannidis, S., Taft, N.: Blurme: Inferring and obfuscating user gender based on ratings. In: Proceedings of the Sixth ACM Conference on Recommender Systems. p. 195–202. RecSys '12, Association for Computing Machinery, New York, NY, USA (2012). <https://doi.org/10.1145/2365952.2365989>

34. Yaari, M.E., Bar-Hillel, M.: On dividing justly. *Social Choice and Welfare* **1**(1), 1–24 (1984)
35. Yang, K., Stoyanovich, J.: Measuring fairness in ranked outputs. In: Proceedings of the 29th International Conference on Scientific and Statistical Database Management. SSDBM '17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3085504.3085526>
36. Zehlike, M., Bonchi, F., Castillo, C., Hajian, S., Megahed, M., Baeza-Yates, R.: Fa*ir: A fair top-k ranking algorithm. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management. p. 1569–1578. CIKM '17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3132847.3132938>
37. Zehlike, M., Sühr, T., Castillo, C., Kitanovski, I.: Fairsearch: A tool for fairness in ranked search results. In: Companion Proceedings of the Web Conference 2020. p. 172–175. WWW '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3366424.3383534>
38. Zhong, G.: Efficient and robust density estimation using bernstein type polynomials. *Journal of Nonparametric Statistics* **28**(2), 250–271 (2016). <https://doi.org/10.1080/10485252.2016.1163349>