

# Facets of Fairness in Search and Recommendation

Sahil Verma<sup>1</sup>, Ruoyuan Gao<sup>2</sup><sup>[0000-0002-8784-4171]</sup>, and Chirag Shah<sup>1</sup><sup>[0000-0002-3797-4293]</sup>

<sup>1</sup> University of Washington, Seattle WA 98195, USA  
{vsahil,chirags}@uw.edu

<sup>2</sup> Rutgers University, New Brunswick NJ 08901, USA  
ruoyuan.gao@rutgers.edu

**Abstract.** Several recent works have highlighted how search and recommender systems exhibit bias along different dimensions. Counteracting this bias and bringing a certain amount of fairness in search is crucial to not only creating a more balanced environment that considers relevance and diversity but also providing a more sustainable way forward for both content consumers and content producers. This short paper examines some of the recent works to define relevance, diversity, fairness, and related concepts. Then, it shows how various metrics to measure such concepts are materialized on search results. In doing so, this paper presents comparisons and highlights contracts among various measures, and gaps in our conceptual and evaluative frameworks.

**Keywords:** Search bias · Fairness · Evaluation metrics · Fairness in recommendation · Fair ranking.

## 1 Introduction

Recommendations or ranking candidates for any purpose is an integral part of the technologies we use each day. Each potential candidate is scored with relevance which is used to rank them in a recommendation list. The algorithms used in the underlying software are not only complicated, but they also take clues from the previous actions of the users on the platform. This feedback loop potentially leads to discrimination against future users, for example, women less likely to be shown advertisements for high-paying jobs [10, 3].

Left unchecked, such implicit biases can amount to increased stereotypes and polarized opinions. Mitigation of bias in automated decisions is an emerging area in machine learning and related domains. Classification and ranking/filtering are the two important categorizations of automated decisions using machine learning. Fairness in automated decisions has gained significant traction, and many papers published in the recent years have attempted to 1) devise metrics to measure the fairness, 2) design frameworks using which fair models can be produced (according to some metric of fairness) or 3) synthesize data to fight bias in the historical data. There have been many works of the kinds mentioned beforehand

in both classification and recommendation setting. A recent work summarized and explained various fairness metrics used in the classification tasks [23]. Unlike classification, recommendations have widely different facets and application scenarios. One of the significant differences lies in the output space, which is very restricted in the case of classification. In contrast, the output space for ranking or recommendation could be the entire list of items. Owing to the sheer abundance of the fairness metrics and their applicability in specific scenarios, understanding their differences and similarities is complicated.

We review papers from major conferences, which received submissions related to fairness in ranking and recommendation, including KDD, WSDM, WWW, SIGIR, ECIR, RecSys, IP&M, and FAT\*, from 2015 to 2020. We found twenty-two relevant papers that devise fairness metrics and provide frameworks to optimize models using them. In this paper, we collect and intuitively explain the fairness metrics used in five recommendation settings: non-personalized recommendation setting, crowd-sourced recommendation setting, personalized recommendation setting, online advertisements and marketplace. Since literature has proposed several metrics for each of the settings above, we present all the metrics but we do not attempt to develop arguments in favor of any particular metric, we rather explain the underlying similarities and differences between them, and show how these metrics affect other dimensions of ranking. We also develop a clear distinction between fairness in various recommendation settings from often related terms such as diversity, novelty, and relevance. We have categorized fairness metrics according to the setting they are applied. The main contribution of this work is intuitive categorization and explanation of *twenty-five* fairness definitions in various recommendation settings and identification of relationships between them.

The remainder of this paper is organized as follows. Section 2 presents the definitions of commonly used terms in the recommendation literature. Sections 3 to 7 delineates the fairness metrics in various settings. Section 8 outlines the conclusions.

## 2 Dimensions of Search and Recommendation results evaluation

We formally define three dimensions for evaluating search and recommendation results – relevance, diversity, and novelty. These dimensions help to gauge the quality of ranking for a search query or recommendation.

- **Relevance** [12]: Search results are relevant if they accurately answer or describe various aspects of the query or recommendation. It focuses on whether and to what extent a search result relates to the given query. For example, it can be considered as the documents containing some keywords in the query, answering the query, or providing information related to the topic of the query. Relevance only considers the match between the query and the results, often disregarding the user’s intent.

- **Diversity** [12]: Search results for each query might have several topics. For example, given the query “Lisbon”, the topics include geographic and historical facts, tourism information, and the weather. Diversity refers to the constitution of the search results from its various topics. Various metrics have been defined to measure diversity, such as those found in [6, 20].
- **Novelty** [12]: Given several relevant documents pertaining to a query, novelty requires the presentation of results that deliver considerably different information content than the results already shown. It thus encourages uniqueness in the results.

### 3 Fairness Metrics in Non-Personalized Recommendation Settings

We collected the following metrics in literature to capture fairness in search and recommendations. We consider the setting where a ranker wants to rank a set of candidates that are relevant to a query or for a recommendation. The ranker does not account for individual preferences of the candidates. The fairness in this setting addresses how the candidates are ranked. We assume the existence of protected and unprotected groups (binary setting), which are defined by law [1].

#### 3.1 Accuracy-based fairness metrics

Most fairness metrics for recommendations either state the condition for user satisfaction or provide a measure of deviation from the ideal ranking. These metrics require a certain proportion of candidates from the protected group, or they calculate the divergence from that proportion in a ranked list. Collectively they are called *accuracy-based fairness metrics*.

1. **Statistical/demographic Parity:** A ranker is said to be satisfying statistical parity if an equal number of candidates from the protected and unprotected groups are ranked in top-k search results. For example, for the image search query, “CEO”, if displayed results show an equal proportion of male and female CEOs, the results are said to satisfy statistical parity. Singh et al. [22] define *exposure* as the resource allocated to a candidate that is computed as a measure of its relevance and position in the ranking. If the average exposure of the candidates belonging to different groups is not equal, the ranker is said to violate statistical parity.
2. **Disparate Treatment:** Disparate treatment refers to the unequal treatment of a candidate due to its membership in a protected class [15]. The candidate would have been treated differently if it had belonged to another group. If the average exposure of different groups is not proportional to their average relevance, then the ranker is said to exhibit disparate treatment [22].
3. **Disparate Impact:** Disparate impact is the practice of not allocating favorable outcomes to protected groups. For searches and recommendations, a click on a candidate is a favorable outcome, which is called as candidate’s

click-through rate. If the expected click-through rate for members of protected and unprotected groups is not proportional to their average relevance, then the ranker is said to cause disparate impact [22]. Click-through rates can be estimated with the help of several techniques [19].

4. **Search Neutrality:** Search neutrality refers to the search engines' lack of editorial power. It means that the search engine should return results to a query impartially and based solely on relevance. It should not promote/demote or differentiate based on websites [14]. Google, Facebook, and other tech companies have been accused of violating search neutrality to promote websites that pay them for higher rankings [2].
5. **Top-k fairness:** Zehlike et al. [27] describe a ranking as a top-k fair ranking if the top-k candidates in the ranking contain the required proportion of members from the protected group. Given a required proportion, the algorithm they propose ranks candidates in a manner that fairly represents the protected group. Within the protected and unprotected groups, the candidates are ranked by relevance. If the required proportion is 50%, top-k fairness and statistical parity are equivalent.
6. **Skew@k:** Geyik et al. [13] define skew@k as the logarithm of the ratio of proportions of the candidates belonging to the protected group in the top-k ranked candidates to the desired proportion of the protected group in top-k ranks. A negative skew implies lower than desired representation in the top-k ranks. Zero skew implies a top-k fair ranking (Definition 5). They also define minskew@k and maxskew@k to extend the metric beyond binary groups.
7. **Normalized Discounted Difference (rND)** Yang et al. [26] point out that establishing statistical parity at higher ranks (e.g., top-10) is more important than establishing it at lower ranks (e.g., top-100). To account for this, they measure set-based fairness metrics at discrete points (e.g., top-10, top-20) and use logarithmic discounts for lower ranks. Normalized discounted difference (rND) computes the difference between the proportion of the protected group members in top-k results and the overall population proportion. If the protected group is proportionally represented, a lower rND value is achieved, which is preferable. If the desired proportion equals to the actual population proportion, then the rND score is correlated with top-k fairness and Skew@k.
8. **Normalized Discounted KL-divergence (rKL)** Normalized discounted KL-divergence(rKL) [26] measures the expected difference between the proportion of the protected group members in the top-k and the overall population proportion. The metric rKL resembles rND, only it is a smoother measure (and therefore optimizable in gradient-based optimization setting) and can be applied to multiple group settings.
9. **Normalized Discounted Ratio (rRD)** Normalized discounted ratio(rRD) [26] takes the difference between the ratio of the protected to unprotected group members among the top-k ranking and the ratio of sizes of the protected and unprotected groups. The metric rRD is considered to be useful when the protected group is a minority, in which case it resembles rND and rKL values; otherwise, rRD is meaningless.

### 3.2 Error based fairness metrics

Kuhlman et al. [17] point out that fairness in classification setting has several metrics that are error-based, i.e., they require the classifier to have similar error-rates across the protected and unprotected groups. They claim that those metrics carry value and should be used for measuring fairness in rankings. Unlike a classification task where the error is readily computable, there exists no such error in case of ranking. Therefore, Kuhlman et al. propose to use pair-inversions to measure ranking errors. They assume the existence of a ground-truth rank for each candidate. If a ranker ranks a candidate higher than its ground-truth rank, they call it a false positive case. Similarly, a candidate that is ranked lower is called a false negative case. Each definition that follows has roots in the counterparts described in fairness in classification literature [23].

1. **Rank Equality:** Rank equality has its origins in equalized odds, which requires equal error rates (false positive and false negative error) across the protected and unprotected groups. Rank equality error captures the number of times a candidate from a group has been falsely given a higher rank than a candidate of another group; the score is calculated for each such inverted pair. This metric does not penalize the ranking where a candidate from the same group has been ranked higher.
2. **Rank Calibration:** Rank calibration [17] has roots in calibration which enforces equal precision of classifiers across the protected and unprotected groups. It checks how correctly the ranker predicts candidates in each group. Rank calibration error is calculated as the number of times a candidate from one group is falsely ranked higher than candidates of all groups; the score is calculated for each such inverted pair.
3. **Rank Parity:** Rank parity criterion [17] has roots in statistical parity. It requires equal representation of members from the protected and unprotected groups in the ranking. The rank parity error is computed as the number of candidates belonging to one group that were ranked higher than candidates from another group; the score is calculated for each such inverted pair.

### 3.3 Causal approach for mitigating discrimination

Wu et al. [25] use a causal graph to reduce discrimination contained in historical data. They use a score variable (instead of rank) to account for individual qualifications and the path-specific effect technique to capture direct and indirect discrimination based on one's membership in the protected group. Having detected discrimination, each individual's score is modified to remove the bias, keeping the distribution of new scores close to the original distribution. The modified scores are applied to create a fairer ranking.

## 4 Fairness Metrics in Crowd-Sourced Non-Personalized Recommendation Settings

Chakraborty et al. [9] consider the setting of top-k trending recommendations on platforms like Twitter or Yelp, in non-personalized recommendation settings. Generally, recommendations are the top-voted candidates in an election procedure. Each person on the platform can vote (e.g., via a click) for multiple candidates and that too multiple times.

1. **Equality of Voice:** In the setting described above, trends are subject to manipulations by hyper-active group or campaigners of all kinds. This can lead to a veneer of popularity for a particular candidate. To avoid this situation, Chakraborty et al. propose a “one person, one vote” election procedure in which everyone has an equal say. Each person is asked to specify their preferences across a set of candidates. The first position is assigned to the candidate, which is the first preference of the majority.
2. **Proportionality for Solid Coalitions:** Chakraborty et al. [9] point out that due to the abundance of options, user’s votes might get split across irrelevant or redundant alternatives, e.g. if there are three candidates out of which the first two are very close. Assume that 60% of the people are interested in the first two candidates. Due to their similarity, votes would split among them. Thus, even though the sum of the votes across these two candidates is more, a less popular third candidate would emerge as the winner. To avert this, the diversity of opinions in the overall population should be proportionally represented in the top-k recommendations.
3. **Anti-Plurality:** Chakraborty et al. also propose that if a majority of users dislike a candidate, it should not be in the top-k recommendations. In the previous example, the third candidate, disliked by 60% of the population, would not be recommended.

## 5 Fairness Metrics in Personalized Recommendation Settings

Beutel et al. [7] consider fairness metrics for personalized recommendation settings. Consider  $M$  total candidates, out of which  $M'$  are relevant to a query, but only  $K$  candidates are useful as part of the personalization. Since they are dealing with personalized ranking along with clicking, Beutel et al. also consider the engagement of the user with a recommended candidate. Engagement can be estimated between a given user and the recommended candidate. They compare ranked candidates pairwise and define **pairwise accuracy** as the probability that a clicked candidate is ranked above another relevant unclicked candidate.

1. **Pairwise Fairness:** A ranker is said to satisfy pairwise fairness if the probability of a clicked candidate being ranked higher than another relevant unclicked candidate is the same across groups, conditioned on the candidates

that have the same predicted engagement score. This does not eliminate systematic preference between groups. For example, one can rank all candidate belonging to a favored group that are not relevant to the query and give a lower rank to candidates from the other group that are relevant to the query.

2. **Inter-Group Pairwise Fairness:** A ranker satisfies inter-group pairwise fairness [7] if the probability of a clicked candidate being ranked higher than a relevant but unclicked candidate in the other group is the same across pairs of groups, conditioned on the candidates that have same engagement score.
3. **Intra-Group Pairwise Fairness:** A ranker is said to satisfy intra-group pairwise fairness [7] if the probability of a clicked candidate being ranked higher than another relevant but unclicked candidate in the same group is the same across groups, conditioned on the candidates that have the same predicted engagement score.

A combination of Intra-Group and Inter-Group Pairwise Fairness can reduce systematic bias against a group.

## 6 Fairness Metrics in Advertisement Settings

Chawla et al. [16] present fairness concerns from an entirely different perspective. Thus far, we have described metrics that view fairness from the perspective of a candidate to be ranked or recommended. Chawla et al. describe fairness from the perspective of individuals who are being served advertisements. Individual fairness [23] states that the advertisements shown to two similar individuals should be similar. Qualifications of an individual can characterize the similarity. Nevertheless, the solution to the problem does not lie in showing an equal proportion of all advertisements to similar people, as they have needs and preferences.

1. **Envy-Freeness:** Envy-freeness [16, 21] is a complementary notion to individual fairness, in which only a user’s preference is considered for advertisements and a user’s qualifications (therefore, similarities) are not reasoned. It requires every user to be content with his/her share of ads.
2. **Inter-Category Envy-Freeness:** Enforcing individual fairness for all categories of advertisements is problematic since it does not recognize one’s preferences. Inter-category envy-freeness [16] lets each user specify individual preferences, and the criterion requires that all users interested in a category should be served the same amount of ads from it. For example, two individuals interested in jobs should be shown the same number of job-related ads.
3. **Total Variation Fairness:** Inter-category envy-freeness does not ensure a fair distribution of ads within a category. Two equally qualified individuals belonging to different groups can be unfairly shown high-paying and low-paying jobs, respectively, while satisfying that metric. Total variation fairness [16] overcomes this limitation by requiring that all subsets of the ads from a category shown to two similar individuals must be the same. Consequently, it evades the problem of unfairly showing high-paying job ads to one user.

4. **Compositional Fairness:** Compositional fairness [16] combines inter-category envy-freeness and total variation fairness. Compositional fairness has two requirements: 1) a user’s preferences must be recognized, and ads served to them belong to their preferred categories only (*envy-freeness*) and 2) within each category, the proportions of ads should be the same across all users interested in that category. This lets an advertiser serve ads from different categories with varying probabilities to a user (based on their preference). However, the mix of ads from that category should be the same across all interested users.

## 7 Fairness Metrics in Marketplace Settings

Advertisement setting brings us to a discussion related to marketplaces, in which consumers are shown advertisements about products which the suppliers want to publicize. Marketplaces are ubiquitous. Almost all online platforms we interact with serve as a marketplace for consumers and service providers. These multi-sided recommendation platforms have complicated fairness constraints. Historically, most marketplaces have optimized for consumer satisfaction, but given the rising competition among different platforms, the satisfaction of service providers has also gathered attention. For example, Spotify would like to recommend tracks that a particular consumer would find relevant and is likely to listen. But it would be problematic if Spotify only recommends songs from a few popular artists to the consumer because: 1) it gives low exposure to less popular artists, and 2) the consumer may not find the recommendations interesting. There has been a few recent works in this area [18, 5, 24, 11, 4, 8].

There are several classes of multi-sided recommendation: 1) Multi-receiver recommendation – when a target audience is a group of people, rather than an individual, e.g., students on an education platform; 2) Multi-provider recommendation – when several suppliers provide the recommendation content, and the platform needs to choose between them, e.g., Airbnb and Spotify, and 3) Side stakeholder recommendation – when there are parties other than suppliers and consumers involved in the marketplace, the recommendations need to consider their preferences as well, e.g., drivers in the Uber Eats platform.

1. **Consumer Fairness:** A recommendation satisfies consumer fairness [5] if it does not cause any disparate impact on members of protected groups. For example, consumers of all groups should be served the same distribution of job ads.
2. **Provider Fairness:** A recommendation satisfies provider fairness [5] if all the providers have an equal chance of exposure to the consumers. For example, Spotify recommends both famous and less famous artists publishing a specific genre of music to users who prefer that genre.
3. **Side Stakeholder Fairness:** A recommendation satisfies side stakeholder fairness [5] if it takes into consideration the preferences of side stakeholders. For example, fairly distributing consumer orders and commute distance among drivers in Uber Eats.



## 8 Conclusion

In this short survey, we collected and presented various metrics proposed in the emerging literature on fairness in recommendations. Succinct and distinct categorization of fairness metrics would help people understand the landscape and triage missing gaps, consequently fuelling future research. We are already experiencing an adoption of this research in the industry. For instance, Geyik et al.[13], in a first large-scale deployment, enforced fair ranking in LinkedIn search. We envision such deployments to other major search engines in the future.

## References

1. *Protected Group*. [https://en.wikipedia.org/wiki/Protected\\_group](https://en.wikipedia.org/wiki/Protected_group), accessed: 2020-01-20
2. *What is Search Neutrality?* <https://hackernoon.com/what-is-search-neutrality-d05cc30c6b3e>, accessed: 2020-01-20
3. *Women less likely to be shown ads for high-paid jobs on Google, study shows*. <https://www.theguardian.com/technology/2015/jul/08/women-less-likely-ads-high-paid-jobs-google-study>, accessed: 2020-01-20
4. Abdollahpouri, H., Burke, R., Mobasher, B.: Recommender systems as multistakeholder environments. In: Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization. UMAP '17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3079628.3079657>
5. Abdollahpouri, H., Burke, R.D.: Multi-stakeholder recommendation and its connection to multi-sided fairness. ArXiv **abs/1907.13158** (2019)
6. Amigó, E., Spina, D., Carrillo-de Albornoz, J.: An axiomatic analysis of diversity evaluation metrics: Introducing the rank-biased utility metric. In: The 41st International ACM SIGIR Conference on Research& Development in Information Retrieval. SIGIR '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3209978.3210024>
7. Beutel, A., Chen, J., Doshi, T., Qian, H., Wei, L., Wu, Y., Heldt, L., Zhao, Z., Hong, L., Chi, E.H., Goodrow, C.: Fairness in recommendation ranking through pairwise comparisons. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery& Data Mining. KDD '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3292500.3330745>
8. Burke, R.: Multisided Fairness for Recommendation. arXiv:1707.00093 [cs] (Jul 2017), arXiv: 1707.00093
9. Chakraborty, A., Patro, G.K., Ganguly, N., Gummadi, K.P., Loiseau, P.: Equality of voice: Towards fair representation in crowdsourced top-k recommendations. In: Proceedings of the Conference on Fairness, Accountability, and Transparency. FAT\* '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3287560.3287570>
10. Datta, A., Tschantz, M.C., Datta, A.: Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. ArXiv **abs/1408.6491** (2014)
11. Ferraro, A., Bogdanov, D., Serra, X., Yoon, J.J.: Artist and style exposure bias in collaborative filtering based music recommendations. ArXiv **abs/1911.04827** (2019)

12. Gao, R., Shah, C.: Toward creating a fairer ranking in search engine results. *Information Processing & Management* **57** (2020). <https://doi.org/10.1016/j.ipm.2019.102138>
13. Geyik, S.C., Ambler, S., Kenthapadi, K.: Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '19*, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3292500.3330691>
14. Grimmelmann, J.: *Some skepticism about search neutrality. The Next Digital Decade: Essays on the Future of the Internet* (2011)
15. Heidari, H., Krause, A.: Preventing disparate treatment in sequential decision making. In: *IJCAI* (2018)
16. Ilvento, C., Jagadeesan, M., Chawla, S.: Multi-category fairness in sponsored search auctions. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3351095.3372848>
17. Kuhlman, C., VanValkenburg, M., Rundensteiner, E.: Fare: Diagnostics for fair ranking using pairwise error metrics. In: *The World Wide Web Conference. WWW '19*, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3308558.3313443>
18. Mehrotra, R., McInerney, J., Bouchard, H., Lalmas, M., Diaz, F.: Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management. CIKM '18*, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3269206.3272027>
19. Richardson, M., Dominowska, E., Ragno, R.: Predicting clicks: Estimating the click-through rate for new ads. In: *Proceedings of the 16th International Conference on World Wide Web. WWW '07*, Association for Computing Machinery, New York, NY, USA (2007). <https://doi.org/10.1145/1242572.1242643>
20. Sakai, T., Craswell, N., Song, R., Robertson, S.E., Dou, Z., Lin, C.Y.: Simple evaluation metrics for diversified search results. In: *EVIA@NTCIR* (2010)
21. Serbos, D., Qi, S., Mamoulis, N., Pitoura, E., Tsaparas, P.: Fairness in package-to-group recommendations. In: *Proceedings of the 26th International Conference on World Wide Web. WWW '17*, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE (2017). <https://doi.org/10.1145/3038912.3052612>
22. Singh, A., Joachims, T.: Fairness of exposure in rankings. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18*, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3219819.3220088>
23. Verma, S., Rubin, J.: Fairness definitions explained. In: *Proceedings of the International Workshop on Software Fairness. FairWare '18*, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3194770.3194776>
24. Wan, M., Ni, J., Misra, R., McAuley, J.: Addressing marketing bias in product recommendations. In: *Proceedings of the 13th International Conference on Web Search and Data Mining. WSDM '20*, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3336191.3371855>
25. Wu, Y., Zhang, L., Wu, X.: On discrimination discovery and removal in ranked data using causal graph. In: *Proceedings of the 24th ACM*

- SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3219819.3220087>
26. 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>
  27. 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. CIKM '17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3132847.3132938>