



Joint Multisided Exposure Fairness for Recommendation

Haolun Wu^{*†}
McGill University
Montréal, Canada
haolun.wu@mail.mcgill.ca

Bhaskar Mitra^{*}
Microsoft
Montréal, Canada
bmitra@microsoft.com

Chen Ma[†]
City University of Hong Kong
Hong Kong SAR
chenma@cityu.edu.hk

Fernando Diaz
Canadian CIFAR AI Chair & Google
Montréal, Canada
diazf@acm.org

Xue Liu
McGill University
Montréal, Canada
xueliu@cs.mcgill.ca

Abstract

Prior research on exposure fairness in the context of recommender systems has focused mostly on disparities in the exposure of individual or groups of items to individual users of the system. The problem of how individual or groups of items may be systemically under or over exposed to groups of users, or even all users, has received relatively less attention. However, such systemic disparities in information exposure can result in observable social harms, such as withholding economic opportunities from historically marginalized groups (*allocative harm*) or amplifying gendered and racialized stereotypes (*representational harm*). Previously, Diaz et al. [17] developed the *expected exposure* metric—that incorporates existing user browsing models that have previously been developed for information retrieval—to study fairness of content exposure to individual users. We extend their proposed framework to formalize a family of exposure fairness metrics that model the problem jointly from the perspective of both the consumers and producers. Specifically, we consider group attributes for both types of stakeholders to identify and mitigate fairness concerns that go beyond individual users and items towards more systemic biases in recommendation. Furthermore, we study and discuss the relationships between the different exposure fairness dimensions proposed in this paper, as well as demonstrate how stochastic ranking policies can be optimized towards said fairness goals.

CCS Concepts

• Information systems → Recommender systems.

Keywords

Fairness-aware Recommendation; Multisided Fairness; Exposure Fairness; Recommender System

^{*}Both authors contributed equally to this research.

[†]Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '22, July 11–15, 2022, Madrid, Spain

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-8732-3/22/07...\$15.00

<https://doi.org/10.1145/3477495.3532007>

ACM Reference Format:

Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. 2022. Joint Multisided Exposure Fairness for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*, July 11–15, 2022, Madrid, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3477495.3532007>

1 Introduction

Online information access systems, like recommender systems and search, mediate what information gets exposure and thereby influence their consumption at scale. There is a growing body of evidence [17, 27, 51] that information retrieval (IR) algorithms that narrowly focus on maximizing ranking utility of retrieved items may disparately expose items of similar relevance from the collection. Such disparities in exposure outcome raise concerns of algorithmic fairness and bias of moral import [20], and may contribute to both *representational* and *allocative* harms [14]. Representational harms may manifest by reinforcing negative stereotypes—e.g., disproportionately suggesting arrest record searches for black-identifying first names in online ad delivery systems [54]—and perpetuating inequities in representation of women and other historically marginalized peoples in different occupational roles [27]. Similarly, allocative harms may occur when disparate exposure in retrieved results lead to unfair allocation of economic opportunities [51]—e.g., in IR systems that aim to match employers with potential candidate employees, and for content-producers on the web who depend on search traffic for ad-based monetization.

The *probability ranking principle* [50], which suggests that retrieved items should be ordered in decreasing probability of relevance, has been a guiding principle in many IR model development. In contrast, by viewing IR systems as mediators of exposure, Diaz et al. [17] propose the *principle of equal expected exposure* that requires IR systems to provide equal exposure in expectation to items of comparable relevance for a given information need. Furthermore, they define the expected exposure metric as a deviation between actual system exposure and ideal target exposure for individual (or groups of) items to individual users. While Diaz et al. [17] consider group attributes only on the producer-side, Burke [7] and Ekstrand et al. [18] argue that it is important to consider group attributes on both the consumer-side and producer-side, and that the joint consideration of group attributes on both sides can reveal new fairness concerns. However, these previous works do not provide a formal formulation of joint multisided exposure (JME) fairness,

which is exactly where our current work fits in. We extend the expected exposure metric proposed by Diaz et al. [17] to incorporate group attributes on both the consumer and producer side and formally define a broader set of fairness concerns that we believe should receive consideration for developing fair recommendation algorithms.

These distinct fairness concerns may be best explained through specific examples. Consider a job portal that recommends potential employment opportunities to job seekers. In this recommendation scenario, the employers represent the producer-side and the job seekers consume the recommendations produced by the system. A responsible recommendation algorithm in this context must grapple with several fairness concerns. Firstly, any individual user of the system must have fair exposure to all relevant jobs that they may qualify for. We refer to this as individual-user to individual-item fairness, or II-F to be more concise. If we group similar job roles by salary into, say, high and low paying job groups, then the system must also ensure that an individual job candidate does not disproportionately get recommended either low or high paying jobs. We classify such fairness concerns as individual-user to group-of-items fairness, or IG-F in short. If we now group the job seekers—for example, based on race or gender—then the system must also guarantee that qualified candidates get fair exposure to all relevant individual job opportunities independent of their demographic attributes. We refer to this as group-of-users to individual-item fairness, or GI-F. Extending the framework further, we consider the scenario where job seekers are grouped by demographic attributes, like race and gender, and employment opportunities by salary. This raises a group-of-users to group-of-items fairness (GG-F) consideration such as ensuring that women or racially marginalized groups are not disproportionately recommended low-paying jobs. Another interesting example of the GG-F concern involves grouping both job seekers and employers by demographic, such as race, where an unfair recommender system may, for example, disproportionately recommend Black employees to Black-owned businesses and similarly match White employees with White employers, and codify *algorithmic segregation* much like what Benjamin [2] calls the “New Jim Code”. Finally, the recommender system should also ensure that any individual jobs or groups of jobs (e.g., grouped by race or gender attributes of business owners) are not systemically under or over exposed to all users, which we refer to as all-users to individual-item fairness (AI-F) and all-users to group-of-items fairness (AG-F), respectively. For example, an unfair recommender system may systemically under-expose jobs at businesses owned by historically marginalized groups which according to our definition would fall under AG-F concerns.

While II-F and IG-F has been previously studied by Diaz et al. [17], to the best of our knowledge, this is the first work to formally define and study a family of exposure fairness considerations jointly from the perspective of both the consumers and producers. In summary, the key contributions of this work are as follows:

- (1) We formalize and compare a family of JME-fairness measures that deal with different types of systemic biases in content exposure from recommender systems.
- (2) We demonstrate how each of the six JME-fairness metrics can be decomposed into their corresponding disparity and relevance components, and study their trade-offs.

- (3) We demonstrate how stochastic ranking policies can be optimized towards specific JME-fairness goals.

2 Related work

2.1 Fairness in recommendation

Fairness in recommendation has received increasing attention in the IR community recently. Although there is no unified notion of fairness, one of the criteria is to model from the perspective of different stakeholders [18], which may include among others: (i) consumers (i.e., users who seek content), (ii) producers (i.e., users who create or publish content), (iii) information subjects (when retrieved items correspond to individuals—e.g., recommending candidates for jobs), and (iv) other side stakeholders who may be impacted even if they do not directly interact with the retrieval system (e.g., historically marginalized groups who may suffer from representational harms by information access platforms).

Even for a single stakeholder, different fairness dimensions are important to consider that if unchecked may result in different harmful outcomes. For example, for consumers it is important to consider fairness in quality of service, such as ensuring that consumers belonging to different demographic groups experience comparable retrieval quality. Ekstrand et al. [19] and Neophytou et al. [37] demonstrate that recommender performance can vary across demographic groups and Mehrotra et al. [34] report similar observations in the context of web search, another important retrieval scenario. On similar lines, Wu et al. [60] explicitly model the consumer-sided fairness in quality of service as the difference of Normalized Discounted Cumulative Gain (NDCG) between group of users with different demographic attributes which they then try to minimize during the model optimization process.

Apart from quality of service, fairness concerns may also arise with respect to what content consumers are exposed to [17, 51]. For example, when retrieved results correspond to economic opportunities, unfair distribution of exposure can result in allocative harm. Take the case of a job recommender system, in addition to ensuring that different demographic groups of users receive results of comparable relevance, it is also critical that there are no systemic disparities in exposure to high and low paying jobs across demographics [7]. Previous works have modeled fairness from the perspective of fair allocation and utilize some notions in social welfare, such as Envy-freeness [42] and Least Misery [29]. Disparate exposure can similarly raise concerns of stereotyping of consumers, say in case of a news recommender system whose results may reflect historical gender-based stereotyping [59]. Several previous works [5, 48, 65] have explored adversarial approaches, based on domain-confusion [11, 56], to learn representations that conceal information about protected attributes, such as race or gender.

Other works focus on the fairness on the producer side, which is represented by the systematic exposure disparities [3, 17, 21, 52] across content providers. This line of research mainly aims to ensure the equity of exposure for individual producers or group of producers with sensitive attributes. For instance, Zehlike et al. [64] model producer-sided fairness in a top- k ranking problem and guarantee a minimum proportion of items from the protected group in every prefix of the top- k ranking. Biega et al. [3] aim to achieve amortized fairness of attention by making exposure proportional to

relevance through integer linear programming. Singh and Joachims [51] later propose a more general framework which can achieve individual fairness and group fairness on the producer side concurrently. In the learning-to-rank [30] setting, Singh and Joachims [52] propose a policy learning approach for optimizing ranking models while satisfying fairness of exposure constraints, while Diaz et al. [17] develop a direct supervised approach using Gumbel sampling. Recently, Oosterhuis [39] proposes computationally efficient optimization of Plackett-Luce ranking models for fairness.

It is important to ensure fairness for multiple stakeholders in online platforms for avoiding the super-star economy [1, 35] and building a healthy marketplace, as well as to ensure that improving fairness for one set of stakeholders do not negatively impact the utility of the other sides. Several previous works [8, 35, 53, 60] have explored multisided fairness definitions within recommendation tasks. However, among these related works, with the exception of Burke et al. [8], the remaining focus on the fairness of quality of service, and not of exposure, on the consumer side. The specific problem of two-sided exposure fairness has received limited attention in the community, despite the fact that it has been previously suggested by Burke [7] and Ekstrand et al. [18]. Our work thus aims to fill that gap by presenting a formal formulation that extends previous works on user browsing model based exposure fairness [17].

As the literature on the fairness of recommendation has recently grown significantly, we point our readers to [18] for a broader overview of fairness concerns and mitigation approaches in information access systems.

2.2 Ranking with stochastic policy

Most of the early works in learning-to-rank for IR [30] focus on deterministic ranking policies that produce static ordering of items given a user (in case of recommendation) or a query (in case of search). Motivated by Pandey et al. [41], who first proposed to introduce randomization in ranking, several works have employed randomization as a means to collect unbiased implicit feedback from user behavior data [26, 40, 45, 46, 57] and for training unbiased ranking models on biased user feedback [25]. Stochastic ranking policies can also be employed to improve the diversity of retrieved results [47] and—as is more relevant to our current work—to ensure fairer exposure of information content [17, 52, 62]. Motivated by these use cases, several recent works [6, 17, 40, 52, 62] have explored optimization of stochastic ranking policies. In section 6, we will demonstrate how stochastic ranking policies for recommendation can be optimized towards JME-fairness.

3 Preliminaries

Before we formally define the different JME-fairness metrics, we go over some of the fundamental concepts in this section that have already been developed in existing literature that we build on in subsequent sections.

3.1 Exposure and user browsing models

Simplified models of how users inspect and interact with retrieved results provide a useful tool for metric development [10, 63] and for estimating relevance from historical logs of user behavior data [9].

These user browsing models provide a mechanism to estimate the probability of exposure of an item d in a retrieved ranked list of items σ . For example, the user browsing model behind the rank-biased precision (RBP) metric [36] assumes that the probability of the exposure event ϵ for d depends only on its rank $\rho_{d,\sigma}$ in σ and decreases exponentially as we go down the ranked list.

$$p_{RBP}(\epsilon|d, \sigma) = \gamma^{(\rho_{d,\sigma}-1)}, \quad (1)$$

where the γ is the patience factor and controls how deep in the ranking the user is likely browse.

In this work, we employ the RBP user browsing model although other models can also be employed in its place.

3.2 Stochastic ranking and expected exposure

Diaz et al. [17] define a stochastic ranking policy π as a probability distribution over all permutations of items in the collection. In the recommendation scenario, let \mathcal{D} be the collection of all items and \mathcal{U} the set of all users of the retrieval system. Now, given a stochastic ranking policy π_u conditioned on user $u \in \mathcal{U}$, we can compute the expected value of the probability that an item $d \in \mathcal{D}$ is exposed to the user as follow:

$$p(\epsilon|d, \pi_u) = \mathbb{E}_{\sigma \sim \pi_u} [p(\epsilon|d, \sigma)]. \quad (2)$$

Assuming the RBP user browsing model, we can further compute $p(d|\sigma)$ based on eq. (1). Furthermore, for notational convenience, let $\mathbf{E} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{D}|}$ be the expected exposure matrix, such that $E_{ij} = p(\epsilon|\mathcal{D}_j, \pi_{\mathcal{U}_i})$. In the rest of this paper, we refer to the expected exposure \mathbf{E} corresponding to a stochastic ranking policy π as determined by a retrieval system as *system exposure*. Similarly, *target exposure* is defined as the expected exposure \mathbf{E}^* corresponding to an ideal stochastic ranking policy π^* , whose behavior may be dictated by some desirable principle, such as the *equal expected exposure principle* [17]. The deviation of \mathbf{E} from \mathbf{E}^* gives us a quantitative measure of the suboptimality of the retrieval system under consideration. Finally, we define *random exposure* as the expected exposure $\tilde{\mathbf{E}}$ corresponding to a stochastic ranking policy $\tilde{\pi}$ which is defined by a uniformly random distribution over all permutations of items in the collection.

4 Joint multisided exposure (JME) fairness

Having introduced the expected exposure framework in the previous section, we now formalize six different JME-fairness metrics. Each of these metrics addresses a distinct aspect of fairness, covering the disparity in how individual items and groups of items are exposed to individual users, groups of users, and all users.

4.1 Metric definitions

Individual-user-to-individual-item fairness (II-F). The II-F metric, previously proposed by Diaz et al. [17], measures the disparity between the system and target exposure at the level of individual users and individual items. Using similar notations as before, we have:

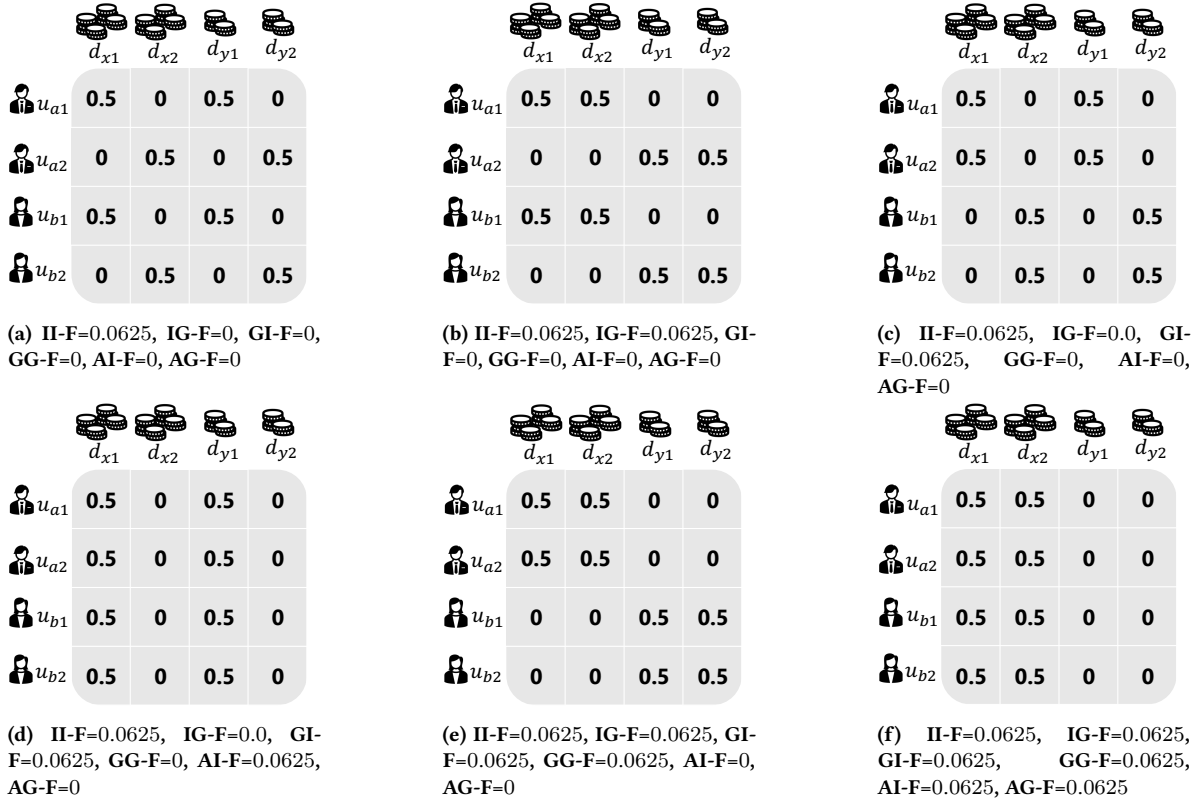


Figure 1: A toy example of a job recommendation system demonstrating the different fairness concerns. We present six systems that have comparable II-F metric values but differ in their fairness as measured along the other five JME-fairness dimensions. The values in the tables represent the probability of exposure of an individual job to an individual candidate.

$$\text{II-F} = \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{d \in \mathcal{D}} \sum_{u \in \mathcal{U}} (p(\epsilon|d, u) - p^*(\epsilon|d, u))^2 \quad (3)$$

$$= \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} (\mathbf{E}_{ij} - \mathbf{E}_{ij}^*)^2. \quad (4)$$

The key motivation of this work lies in the observation that the deviation between the system and target exposure may not be distributed uniformly across different user sub-populations and item groups. As we will demonstrate in section 4.2, the II-F metric cannot distinguish the scenario where the system-to-target exposure deviations systematically and disproportionately impact specific user and/or item groups from the scenario where these deviations are comparable across groups.

Individual-user-to-group-of-items fairness (IG-F). We first introduce group attributes on the item-side and present the IG-F metric which is concerned with whether groups of items are over or under exposed to individual users. We achieve this by making couple of minor modifications to eq. (3): (i) replacing $p(\epsilon|d, u)$ and $p^*(\epsilon|d, u)$ with $p(\epsilon|D, u)$ and $p^*(\epsilon|D, u)$, respectively, where $D \in \mathcal{G}_d$ denotes a group of items and \mathcal{G}_d is the set of all item groups, and (ii) averaging the deviations across groups of items instead of individual items.

$$\text{IG-F} = \frac{1}{|\mathcal{G}_d|} \frac{1}{|\mathcal{U}|} \sum_{D \in \mathcal{G}_d} \sum_{u \in \mathcal{U}} (p(\epsilon|D, u) - p^*(\epsilon|D, u))^2 \quad (5)$$

$$= \frac{1}{|\mathcal{G}_d|} \frac{1}{|\mathcal{U}|} \sum_{D \in \mathcal{G}_d} \sum_{i=1}^{|\mathcal{U}|} \left(\sum_{j=1}^{|\mathcal{D}|} p(D_j|D) (\mathbf{E}_{ij} - \mathbf{E}_{ij}^*) \right)^2. \quad (6)$$

Here, $p(D_j|D)$ can be defined as a uniform probability distribution over all items in a group, or when appropriate a popularity weighted distribution over items can also be employed.

In previous work, Diaz et al. [17, §4.3] have also considered group attributes on the producer-side, and their proposed formulation is in fact identical to eq. (6). However, they provide limited motivation for their specific formulation and as a minor contribution eq. (5) gives us a probabilistic interpretation of this measure.

Group-of-users-to-individual-item fairness (GI-F). Next, we introduce group attributes on the user-side which gives us the GI-F metric that measures the over or under exposure of individual items to groups of users. Similar to the way we define the IG-F metric, the GI-F metric can be defined as follows, where $U \in \mathcal{G}_u$ denote a group of users and \mathcal{G}_u the set of all user groups:

$$\text{GI-F} = \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{G}_u|} \sum_{d \in \mathcal{D}} \sum_{U \in \mathcal{G}_u} (p(\epsilon|d, U) - p^*(\epsilon|d, U))^2 \quad (7)$$

$$= \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{G}_u|} \sum_{j=1}^{|\mathcal{D}|} \sum_{U \in \mathcal{G}_u} \left(\sum_{i=1}^{|U|} p(U_i|U) (E_{ij} - E_{ij}^*) \right)^2. \quad (8)$$

Consequently, $p(U_i|U)$ can be defined as a uniform probability distribution over all users in a group, or could be proportional to their usage of the recommender system.

Group-of-users-to-group-of-items fairness (GG-F). Having introduced group attributes for users and items separately, we now change our focus to exposure disparities that emerge when we look at group attributes for both the users and items jointly. Using similar notations as before, we can write:

$$\begin{aligned} \text{GG-F} &= \frac{1}{|\mathcal{G}_d|} \frac{1}{|\mathcal{G}_u|} \sum_{D \in \mathcal{G}_d} \sum_{U \in \mathcal{G}_u} (p(\epsilon|D, U) - p^*(\epsilon|D, U))^2 \quad (9) \\ &= \frac{1}{|\mathcal{G}_d|} \frac{1}{|\mathcal{G}_u|} \sum_{D \in \mathcal{G}_d} \sum_{U \in \mathcal{G}_u} \left(\sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|U|} p(D_j|D) p(U_i|U) (E_{ij} - E_{ij}^*) \right)^2. \quad (10) \end{aligned}$$

Of all six fairness metrics defined in this section, the GG-F metric is particularly interesting as all the other metrics can be thought of specific instances of GG-F. For example, if we define the group attributes for users in GG-F such that each group contains only one user and every user belongs to only one group then we recover the IG-F metric. A similar trivial definition of groups on the item-side gives us the GI-F metric. Consequently, if this trivial definition of groups is applied to both the users and items, we get the II-F metric. Another trivial, but conceptually interesting, definition of the user group may involve a single group to which all users belong. Under this setting, depending on group definition on the item-side, we can recover the AI-F and AG-F metrics that we describe next.

All-users-to-individual-item fairness (AI-F). A recommender system may systemically under or over expose an item to all users. To quantify this kind of systemic disparities we define the AI-F metric which computes the mean deviation between overall system exposure $p(\epsilon|d)$ and target exposure $p^*(\epsilon|d)$ for items:

$$\text{AI-F} = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} (p(\epsilon|d) - p^*(\epsilon|d))^2 \quad (11)$$

$$= \frac{|\mathcal{D}|}{|\mathcal{U}|} \left(\sum_{i=1}^{|\mathcal{U}|} p(\mathcal{U}_i) (E_{ij} - E_{ij}^*) \right)^2. \quad (12)$$

As earlier, $p(\mathcal{U}_i)$ can either be uniform or weighted by usage.

All-users-to-group-of-items fairness (AG-F). Finally, the AG-F metric is concerned with systemic under or over exposure of groups of items to all users and is defined as follows:

$$\text{AG-F} = \frac{1}{|\mathcal{G}_d|} \sum_{D \in \mathcal{G}_d} (p(\epsilon|D) - p^*(\epsilon|D))^2 \quad (13)$$

$$= \frac{1}{|\mathcal{G}_d|} \sum_{D \in \mathcal{G}_d} \left(\sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} p(D_j|D) p(\mathcal{U}_i) (E_{ij} - E_{ij}^*) \right)^2. \quad (14)$$

We have formally defined six JME-fairness metrics—II-F, IG-F, GI-F, GG-F, AI-F, and AG-F. Readers should note that for all six metrics, a *lower value* is more desirable and corresponds to a more fair recommendation. Next, we discuss the relationship and distinction between these different fairness notions to help the reader develop a more intuitive understanding of our proposed framework of metrics for exposure fairness.

4.2 Relationship between metrics

If we look closely at eqs. (4), (6), (8), (10), (12) and (14), we notice that all six JME-fairness metrics consider the system-to-target exposure deviations $(E_{ij} - E_{ij}^*)$ for individual items to individual users. However, they differ in how they aggregate these differences across users and items leading to interesting relationships and distinctions between these metrics.

Firstly, it is easy to demonstrate that a system that is fair with respect to II-F will also be fair along the other five JME-fairness dimensions. This is because II-F can be zero if and only if $(E_{ij} - E_{ij}^*) = 0, \forall i \in [1, \dots, |\mathcal{U}|], j \in [1, \dots, |\mathcal{D}|]$, and when that is true then IG-F, GI-F, GG-F, AI-F, and AG-F will also be zero. However, the reverse does not hold—*i.e.*, a system that may be considered fair based on one or more metrics from the set IG-F, GI-F, GG-F, AI-F, and AG-F may not necessarily be II-fair. Similarly, it also holds that if a system is either IG-fair or GI-fair then it must also be GG-fair. To demonstrate this, we have to consider that IG-F is zero if and only if $(\sum_{j=1}^{|\mathcal{D}|} p(D_j|D) (E_{ij} - E_{ij}^*)) = 0, \forall i \in [1, \dots, |\mathcal{U}|]$ which implies that $(\sum_{i=1}^{|\mathcal{U}|} p(U_i|U) \sum_{j=1}^{|\mathcal{D}|} p(D_j|D) (E_{ij} - E_{ij}^*)) = 0, \forall U_i \in U, U \in \mathcal{G}_u$. Similarly, we can also prove the relationship between GI-F and GG-F. As before, a system that is GG-fair, however, may not necessarily be IG-fair or GI-fair. Furthermore, we can also show that a system will necessarily be AI-fair if it is GI-fair, and similarly both GG-fairness and AI-fairness independently imply AG-fairness. We do not include the full proof for each of these relationships due space constraints but leave them as an useful exercise for the reader.

Next, we consider the question that if II-fairness implies fairness along the other five dimensions then why it is necessary to develop these other fairness notions at all. The answer lies in the fact that two systems that are comparable according to the II-F metric, may demonstrate very different levels of unfairness along the other dimensions. To motivate this more intuitively, let us consider the toy example of a job recommender system in fig. 1. In this example, there are four candidates ($u_{a1}, u_{a2}, u_{b1}, u_{b2}$) and four jobs ($d_{x1}, d_{x2}, d_{y1}, d_{y2}$), and all four jobs are relevant to each of the four candidates. The candidates belong to two groups a (u_{a1} and u_{a2}) and b (u_{b1} and u_{b2})—*e.g.*, based on gender—and similarly the jobs belong to two groups x (d_{x1} and d_{x2}) and y (d_{y1} and d_{y2}), say based on whether they pay high or low salaries. Furthermore, we assume that the recommender system displays only one result at a time and our simple user model assumes that the user always inspects the displayed result—*i.e.*, the probability of exposure is 1 for the displayed item and 0 for all other items for a given impression. In this setting, an ideal recommender would expose each of the four jobs to each candidate with a probability of 0.25.

The six fictitious recommender systems shown in fig. 1, demonstrate comparable disparity according to the II-F metric but different

Table 1: Decomposing each of the six JME-fairness metrics into their disparity (*-D) and relevance (*-R) components.

Disparity	Relevance
$\Pi\text{-D} = \frac{1}{ \mathcal{D} } \frac{1}{ \mathcal{U} } \sum_{j=1}^{ \mathcal{D} } \sum_{i=1}^{ \mathcal{U} } E_{ij}^{\delta^2}$	$\Pi\text{-R} = \frac{1}{ \mathcal{D} } \frac{1}{ \mathcal{U} } \sum_{j=1}^{ \mathcal{D} } \sum_{i=1}^{ \mathcal{U} } 2E_{ij}^{\delta} E_{ij}^{\Delta}$
$\text{IG-D} = \frac{1}{ \mathcal{G}_d } \frac{1}{ \mathcal{U} } \sum_{D \in \mathcal{G}_d} \sum_{i=1}^{ \mathcal{U} } \left(\sum_{j=1}^{ \mathcal{D} } p(D_j D) E_{ij}^{\delta} \right)^2$	$\text{IG-R} = \frac{1}{ \mathcal{G}_d } \frac{1}{ \mathcal{U} } \sum_{D \in \mathcal{G}_d} \sum_{i=1}^{ \mathcal{U} } \left(\sum_{j=1}^{ \mathcal{D} } 2p(D_j D) E_{ij}^{\delta} E_{ij}^{\Delta} \right)^2$
$\text{GI-D} = \frac{1}{ \mathcal{D} } \frac{1}{ \mathcal{G}_u } \sum_{j=1}^{ \mathcal{D} } \sum_{U \in \mathcal{G}_u} \left(\sum_{i=1}^{ \mathcal{U} } p(U_i U) E_{ij}^{\delta} \right)^2$	$\text{GI-R} = \frac{1}{ \mathcal{D} } \frac{1}{ \mathcal{G}_u } \sum_{j=1}^{ \mathcal{D} } \sum_{U \in \mathcal{G}_u} \left(\sum_{i=1}^{ \mathcal{U} } 2p(U_i U) E_{ij}^{\delta} E_{ij}^{\Delta} \right)^2$
$\text{GG-D} = \frac{1}{ \mathcal{G}_d } \frac{1}{ \mathcal{G}_u } \sum_{D \in \mathcal{G}_d} \sum_{U \in \mathcal{G}_u} \left(\sum_{j=1}^{ \mathcal{D} } \sum_{i=1}^{ \mathcal{U} } p(D_j D) p(U_i U) E_{ij}^{\delta} \right)^2$	$\text{GG-R} = \frac{1}{ \mathcal{G}_d } \frac{1}{ \mathcal{G}_u } \sum_{D \in \mathcal{G}_d} \sum_{U \in \mathcal{G}_u} \left(\sum_{j=1}^{ \mathcal{D} } \sum_{i=1}^{ \mathcal{U} } 2p(D_j D) p(U_i U) E_{ij}^{\delta} E_{ij}^{\Delta} \right)^2$
$\text{AI-D} = \sum_{j=1}^{ \mathcal{D} } \left(\sum_{i=1}^{ \mathcal{U} } p(\mathcal{U}_i) E_{ij}^{\delta} \right)^2$	$\text{AI-R} = \sum_{j=1}^{ \mathcal{D} } \left(\sum_{i=1}^{ \mathcal{U} } 2p(\mathcal{U}_i) E_{ij}^{\delta} E_{ij}^{\Delta} \right)^2$
$\text{AG-D} = \frac{1}{ \mathcal{G}_d } \sum_{D \in \mathcal{G}_d} \left(\sum_{j=1}^{ \mathcal{D} } \sum_{i=1}^{ \mathcal{U} } p(D_j D) p(\mathcal{U}_i) E_{ij}^{\delta} \right)^2$	$\text{AG-R} = \frac{1}{ \mathcal{G}_d } \sum_{D \in \mathcal{G}_d} \left(\sum_{j=1}^{ \mathcal{D} } \sum_{i=1}^{ \mathcal{U} } 2p(D_j D) p(\mathcal{U}_i) E_{ij}^{\delta} E_{ij}^{\Delta} \right)^2$

levels of unfairness along the other dimensions. For example figs. 1a, 1c and 1d correspond to systems that are IG-fair because every candidate is exposed to at least one high-paying and one low-paying job, whereas in figs. 1b, 1e and 1f every candidate is exposed to either only high-paying or only low-paying jobs raising IG-unfairness concerns. Similarly, we can observe that figs. 1a and 1b are GI-fair because each group of candidates are exposed to all four jobs, while consequently the other four systems demonstrate GI-unfairness. Among these six examples, only figs. 1e and 1f demonstrate GG-unfairness. In fig. 1e, men are recommended high-paying jobs while women are recommended low-paying jobs, and in fig. 1f all candidates are exposed only to high-paying jobs leading to no candidates being referred to low-paying jobs. Finally, figs. 1d and 1f demonstrate AI-unfairness because they include individual jobs that are under or over exposed to all candidates, and specifically in the latter instance we also observe AG-unfairness as the low-paying jobs are under-exposed to all candidates as previously noted. It is precisely these differences that motivate us to develop these six JME-fairness metrics which we believe measures different systemic unfairness in recommendation outcomes, and therefore should be considered in complementary combination for more holistic fairness analysis.

4.3 Disparity and relevance

Diaz et al. [17] make an interesting observation that as an artifact of using squared error to compute the deviation between the system and target exposure. The II-F metric can be decomposed into a *disparity* component (II-D), a *relevance* component (II-R), and a remaining component that is a system-independent constant (II-C).

$$\begin{aligned}
\text{II-F} &= \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} (E_{ij} - E_{ij}^*)^2 \\
&= \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} E_{ij}^2 \quad \left. \vphantom{\sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|}} \right\} \text{II-D} \\
&\quad - \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} 2E_{ij} E_{ij}^* \quad \left. \vphantom{\sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|}} \right\} \text{II-R} \\
&\quad + \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} E_{ij}^{*2} \quad \left. \vphantom{\sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|}} \right\} \text{II-C}
\end{aligned}$$

This decomposition is useful as it represents an inherent trade-off between the disparity and relevance as stochasticity is introduced into the model. Increasing randomization in the model has the desirable property of reducing disparity (II-D) but also the undesirable effect of reducing relevance (II-R). In this framing, relevance is maximized when a static ranking model is employed and disparity is minimized when the ranking model is fully stochastic.

A similar decomposition into disparity and relevance components would also be useful for the new metrics proposed in this paper. When considering group attributes, say on the item-side, it should be self-evident that a fully stochastic ranking model would recommend items not necessarily with uniform probability over the groups, but in proportion to the group sizes in the collection. In section 3.2, we defined random exposure \tilde{E} as the expected exposure matrix corresponding to such a fully stochastic model. It makes intuitive sense to define disparity not in terms of the flatness of the system exposure distribution but as a deviation of the system exposure from random exposure, especially when the distribution under consideration is over groups of items rather than individual items. Therefore, we propose a slight modification to the decomposition of the original II-F metric, which is then consistent with how we decompose the other five JME-fairness metrics into their respective components. To do so, we first rewrite eq. (4) as follows:

$$\begin{aligned}
\text{II-F} &= \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} (E_{ij} - E_{ij}^*)^2 \\
&= \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} ((E_{ij} - \tilde{E}_{ij}) - (E_{ij}^* - \tilde{E}_{ij}^*))^2. \quad (15)
\end{aligned}$$

For notational brevity, let $E_{ij}^{\delta} = E_{ij} - \tilde{E}_{ij}$ and $E_{ij}^{\Delta} = E_{ij}^* - \tilde{E}_{ij}^*$. Based on eq. (15), we now redefine II-D and II-R as:

$$\text{II-D} = \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} E_{ij}^{\delta^2} \quad (16)$$

$$\text{II-R} = \frac{1}{|\mathcal{D}|} \frac{1}{|\mathcal{U}|} \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{U}|} 2E_{ij}^{\delta} E_{ij}^{\Delta}. \quad (17)$$

Consequently, the new definition of II-D would produce the same system ordering as the original definition by Diaz et al. [17]. Due

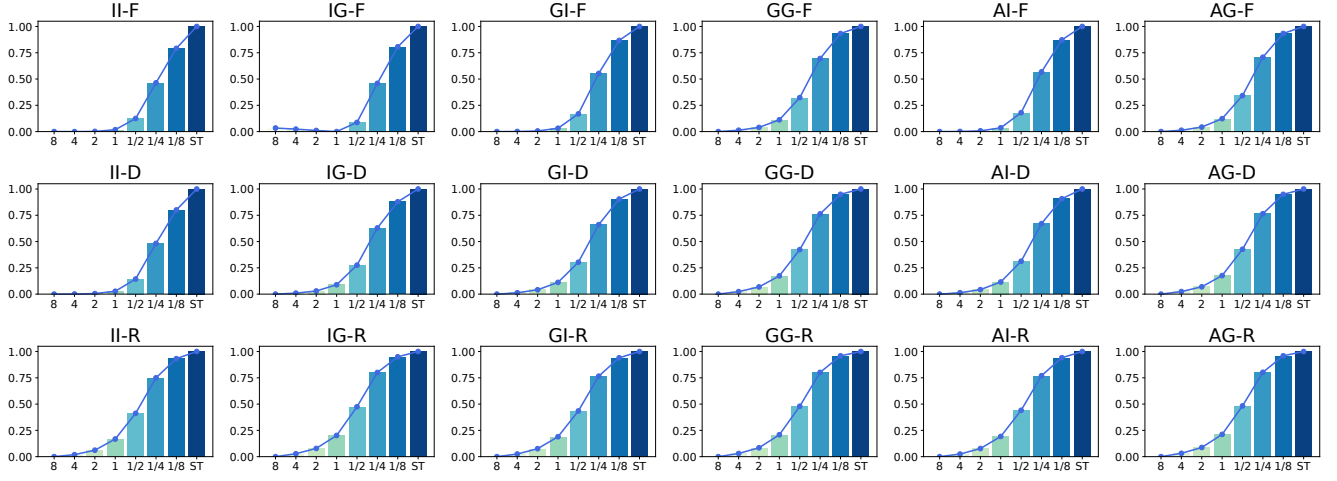


Figure 2: Behavior of JME-fairness metrics for a stochastic ranking policy—generated by randomizing the BPRMF model using Plackett-Luce—on the *MovieLens1M* dataset. The first row shows the impact of different stochasticity on the overall fairness metrics, while the second and third row show the impact on their corresponding disparity and relevance components. The x-axis shows the values of β , where a larger value indicates more randomization.

to space constraints, we do not include the full proof here. Next, employing a similar strategy we can decompose the other five JME-fairness metrics into their disparity and relevance components as shown in table 1. In section 5.2, we will analyse how stochasticity trades-off between the disparity and relevance components of all six JME-fairness metrics.

5 Metric analysis

In this section, we analyze the proposed six JME-fairness metrics with respect to the trade-offs involved between their respective disparity and relevance components (intra-metric analysis) and the relationship between the different JME-fairness metrics themselves (cross-metric analysis). Specifically, we will study the following research questions:

- **RQ1:** What is the impact of stochasticity on each of the six JME-fairness metrics, and their corresponding disparity and relevance components?
- **RQ2:** How do the different JME-fairness metrics, and their respective components, correlate with each other and what does that tell us about the relationships between these different fairness dimensions?

5.1 Method

As aforementioned, we are interested in studying stochastic ranking policies with respect to their ability to distribute exposure across (individual or groups of) items. Towards that goal, we first introduce a method for generating stochastic ranking policies based on a deterministic ranking model. Given a deterministic ranker, and corresponding estimated relevance scores for items with respect to user u , we utilize the Plackett-Luce (PL) model [32, 43] to sample multiple rankings. The PL model depends on Luce’s axiom of choice which states that the odds of choosing an item over another do not depend on the set of items from which the choice is

made [31, 32]. Specifically, the PL model constructs a ranking by iteratively sampling items without replacement from the collection with probability distribution $p(d|u)$ defined as below:

$$p(d|u) = \frac{\exp(Y_{d,u}/\beta)}{\sum_{d' \in \mathcal{D}} \exp(Y_{d',u}/\beta)}, \quad (18)$$

where $Y_{d,u}$ is the relevance score estimated by the deterministic ranker for item d with respect to user u . The parameter β is the softmax temperature. A larger β corresponds to more stochasticity in the ranking. For example, when $\beta = 8$, the probability distribution over all permutations is almost uniform and the stochastic policy approaches a fully random ranking model. As a corollary, when β decreases the stochastic policy converges to the deterministic ranking policy, which is a ranking of items sorted by their estimated relevance score $Y_{d,u}$ in descending order for each user u .

We generate stochastic ranking policies by applying this post-processing step, with different values of β , over a set of trained recommendation models that are publicly available¹ for the *MovieLens1M* dataset. There are 6,040 users and 3,706 items in the *MovieLens1M* dataset. For the group information, we use the gender and age attributes on the user side, and the genre attributes on the item side. We rerank the top 100 items and sample 100 rankings for each user during evaluation. We employ the RBP user browsing model and set the patience factor $\gamma = 0.8$. We select different values for β in the range of 1/8 to 8 for introducing different degree of stochasticity in our ranking, and compare with the deterministic ranking policy.

5.2 Results

RQ1: Impact of Stochasticity. To investigate the impact of stochasticity, we first visualize how different values of β influences the different JME-fairness metrics and their corresponding components in fig. 2. Our analysis is based on a stochastic ranking policy that

¹We obtain the set of runs from: <https://github.com/dvalcarce/evalMetrics>.

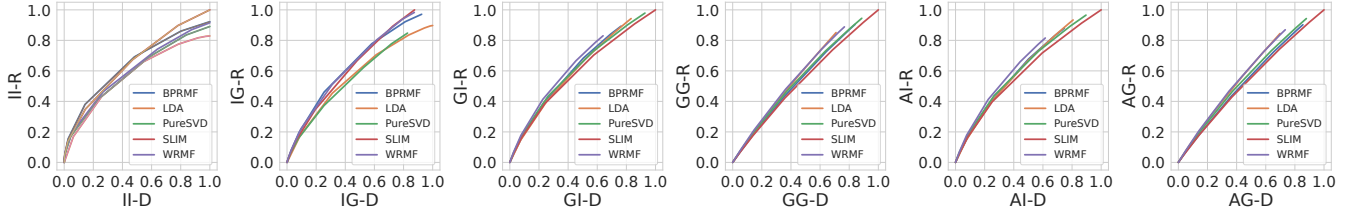


Figure 3: Curves for disparity-relevance trade-off across six different fairness dimensions. We introduce different level of stochasticity on the top of the static rankings from six different recommendation models.

Table 2: The AUC of the disparity-relevance trade-off curves for different models across the six fairness dimensions. The highest value in each column is bolded, while the second highest is underlined.

Model	II	IG	GI	GG	AI	AG
BPRMF	<u>0.6331</u>	0.4774	0.2904	0.2953	0.2712	0.2814
LDA	0.5664	0.4088	0.2837	<u>0.3164</u>	0.2687	<u>0.3115</u>
PureSVD	0.5830	0.4102	<u>0.2921</u>	0.3030	<u>0.2755</u>	0.2942
SLIM	0.6408	0.4654	0.2776	0.2851	0.2605	0.2752
WRMF	0.5996	<u>0.4769</u>	0.3135	0.3186	0.2957	0.3139

uses the BPRMF model as the underlying deterministic ranker, although we have confirmed that we obtain similar results when considering other ranking models. For consistency, we normalize each of the metric values between 0 and 1 using min-max normalization. Our first key observation is that as β increases, the values on all metrics tend to decrease. This is expected given that a larger β corresponds to a more random ranking policy, where the original static relevance estimates have a smaller influence, which consequently results in both lower disparity and relevance. As a corollary, the static ranking policy derived directly from the deterministic BPRMF model has the highest relevance and disparity for all six fairness metrics.

For each of the JME-fairness metrics, these results imply the existence of a disparity-relevance trade-off as the stochastic ranking policy cannot achieve both low disparity and high relevance simultaneously. To visualize this more clearly, we present the disparity-relevance trade-off curves for all six JME-fairness metrics in fig. 3. We study five different models in this experiment: BPRMF [49], LDA [4], PureSVD [24], SLIM [38], and WRMF [55], employing the same randomization strategy on top for all models. As before, we normalize the disparity and relevance metric values in the range of 0 to 1. A point closer to the lower left represents a ranking policy with full stochasticity, while the right-most point for each curve corresponds to the respective static ranking policy—*i.e.*, without any randomization.

For a given metric, different models achieve different levels of disparity-relevance trade-off, and which model achieves the best trade-off may differ across the different JME-fairness metrics. For a given metric, this can be best quantified by computing the area under the disparity-relevance curve for each model, as shown in table 2. For example, WRMF model achieves the best trade-off for GI-F, GG-F, AI-F and AG-F, but is worse for II-F compared to SLIM and BPRMF.

RQ2: Metric Correlation. The intra-metric analysis in the previous subsection demonstrates how all six JME-fairness metrics trades-off between their corresponding disparity and relevance components. Next, we shift our focus to cross-metric analysis to understand how a recommender system optimized for one of the JME-fairness metrics may fare on the other measures. To study this, we evaluate 15 different recommendation models—with 7 different levels of stochasticity in each case—against each of the six JME-fairness metrics. For each metric, this gives us 15×7 values each corresponding to a recommender system instance—*i.e.*, a combination of model and stochasticity level. Now for every pair of JME-fairness metrics, we compute the Kendall rank correlation [28] to quantify the agreement between the two metrics with respect to the ordering of the recommender system instances, as presented in fig. 4. We perform the analysis twice using gender and age as demographic attributes for the user-side, respectively.

We observe that the correlation between the II-F metric and the other five JME-fairness measures is typically low, and the same holds when we compare the corresponding disparity and relevance components across these measures. This provides evidence to our claim that each of these metrics quantifies different notion of unfairness, and that a system instance that may perform well on one fairness dimension may be suboptimal for another. A more nuanced interpretation of these correlation matrices requires us to recall that in section 4.1 we argued that the other five JME-fairness metrics can be considered as a specific instance of the GG-F metric based on the group definitions. For example, on the user dimension, if our group definition involves a very small number of groups then the GG-F metric should intuitively display high correlation with the AG-F metric. On the other hand, if we consider many small groups of users in context of GG-F, then we would expect the correlation between GG-F and IG-F to be higher. In fig. 4, the group definition based on both gender and age correspond to the scenario where we have a small number of groups which may explain why we see strong correlation between GG-F and AG-F—and between age and gender we see that the correlation is higher for gender because it involves relatively fewer groups compared to age. Effectively, GG-F captures a notion of unfairness that lies on spectrum between what IG-F and AG-F measures. The correlation matrix therefore is consistent with the behavior we would expect by comparing eqs. (4), (6), (8), (10), (12) and (14).

6 Optimization for JME-fairness

In the previous section, we demonstrated that a recommender system that performs well on one JME-fairness metric may underperform on another. This highlights a potential opportunity to

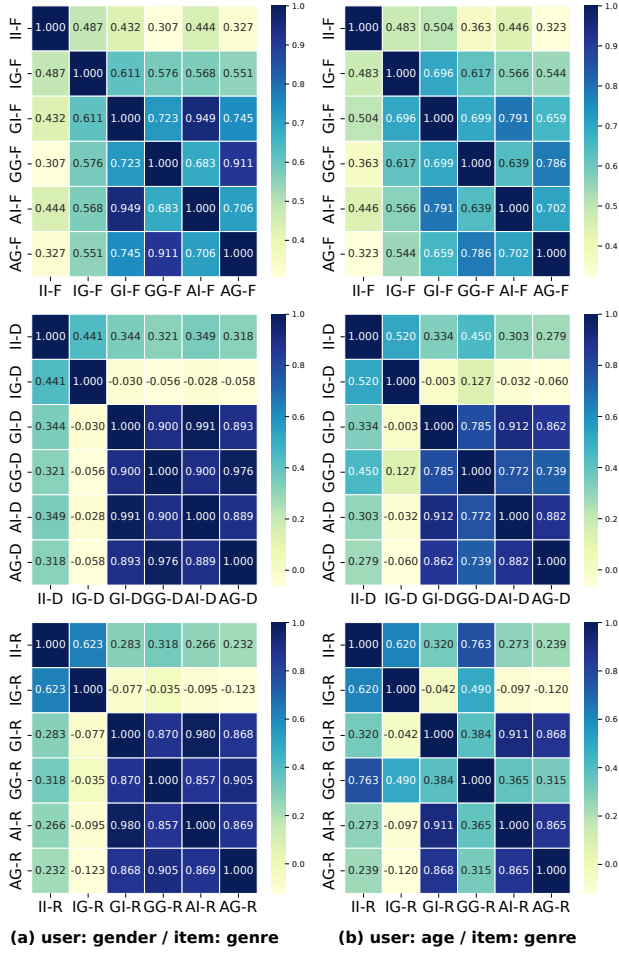


Figure 4: The Kendall rank correlation between different metrics and their two components across six different fairness dimensions.

optimize multiple JME-fairness metrics simultaneously. Ideally, we expect that optimizing for multiple JME-fairness measures allows us to explicitly trade-off between these various fairness concerns.

6.1 Algorithm

To optimize towards a JME-fairness objective, we adopt the approach proposed by Diaz et al. [17]. Let $Y \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{D}|}$ be the estimated relevance between users and items as predicted by a model $f_\theta : \mathcal{U} \times \mathcal{D} \rightarrow \mathbb{R}$, parameterized by Θ . A stochastic ranking policy can now be defined that sampling items without replacement from the following probability distribution as given by the Plackett-Luce model:

$$p(\mathcal{D}_j | \mathcal{U}_i) = \frac{\exp(Y_{ij})}{\sum_k \exp(Y_{ik})}. \quad (19)$$

Sampling from the above distribution is a non-differentiable step that creates a roadblock for gradient-based optimization. However,

following the same strategy as Diaz et al. [17], we can reparameterize the probability distribution by adding independently-drawn noise ζ sampled from the Gumbel distribution to Y_{ij} , and sorting the items by their “noisy” probability distribution $\tilde{p}(\mathcal{D}_j | \mathcal{U}_i)$:

$$\tilde{p}(\mathcal{D}_j | \mathcal{U}_i) = \frac{\exp(Y_{ij} + \zeta_j)}{\sum_k \exp(Y_{ik} + \zeta_k)}. \quad (20)$$

The sorting step itself is also non-differentiable, but we can compute the smooth rank [44, 61] for each item in the ranking as follows:

$$\rho_{\mathcal{D}_j, \pi_{\mathcal{U}_i}} = \sum_{k \in [1..|\mathcal{D}|], k \neq j} \left(1 + \exp \left(\frac{\tilde{p}(\mathcal{D}_j | \mathcal{U}_i) - \tilde{p}(\mathcal{D}_k | \mathcal{U}_i)}{\tau} \right) \right)^{-1}, \quad (21)$$

where the temperature τ is a hyperparameter that controls the smoothness of the approximated ranks. We have now computed the rank position of an item \mathcal{D}_j in the ranking with respect to user \mathcal{U}_i in a differentiable way. Next, we can compute the system exposure using a user browsing model like RBP as $E_{ij} = (1 - \gamma) \cdot \gamma^{\rho_{\mathcal{D}_j, \pi_{\mathcal{U}_i}}}$. To derive the expected exposure, we average the system exposure over 100 different sampled rankings. Finally, having estimated the system exposure E we can now compute its deviation from target exposure E^* using different JME-fairness metric definitions which are themselves differentiable.

Readers should note that merely optimizing towards a JME-fairness metric, such as GG-F, is not sufficient for a recommender system to learn a relevance function between individual users and items. Therefore, it is important to combine the II-F objective, which is appropriate for relevance modeling, with the other desired JME-fairness objective(s). In this work, we combine the GG-F objective linearly with the II-F as follows:

$$\mathcal{L} = \text{II-F} + \alpha \cdot \text{GG-F}, \quad (22)$$

where α is a scaling factor that trades-off between the two objectives. Similarly, other combinations of JME-fairness notions can also be developed.

6.2 Experiment

Datasets. Since our proposed JME-fairness metrics require meaningful attribute information on both the user-side and item-side, we select the *MovieLens100K* and *MovieLens1M* [23] as our datasets. Both of these datasets contain user-to-movie relevance information collected from the *MovieLens* website. These two datasets provide 100 thousand and 1 million user-movie interactions, respectively, with the user metadata (gender, age, occupation group) and movie genres. For each user, we randomly select 20% of the rated items as ground truth for testing. The remaining 70% and 10% data constitutes the training and validation set.²

Implementation Details. We choose Matrix Factorization (MF) as our base model to conduct our experiments. MF is widely used in recommendation, due to its simplicity and efficiency. We optimize the model using the Adam optimizer with the Xavier initialization [22]. The model parameters Θ contain the user embeddings and item embeddings, where the embedding size is fixed to 64.

²The code and datasets are available at <https://github.com/haolun-wu/JMEFairness>.

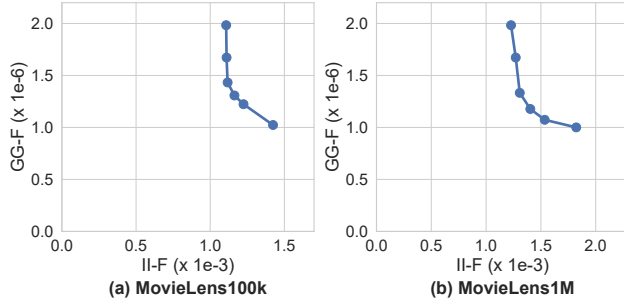


Figure 5: The trade-off between II-F and GG-F when directly optimizing the recommender model towards a combination of both objectives.

In order to obtain the user-item relevance matrix Y , we use the inner-product to compute the similarity between each pair of users and items. The batch size for training is set to be 32. The learning rate and the regularization hyper-parameter are selected from $\{1e^{-2}, 1e^{-3}, 1e^{-4}\}$. The scaling factor α is selected from $\{0, 1, 5, 10, 20, 50\}$. The temperature τ in eq. (21) and the patience parameter γ for the RBP user browsing model are set to be 0.1 and 0.8, respectively.

6.3 Results

Figure 5 shows the trade-off between II-F and GG-F that we achieve by varying α . On both datasets, we initially see a sharp fall in GG-F, which is desirable, with only a small increase in II-F, which depending on the application may be a fair trade-off against the GG-F improvements. For example, if we compare training the model with $\alpha = 0$ (i.e., training on II-F only) and $\alpha = 1$, we see a statistically significant improvement (drop) in GG-F according to the student’s t-test ($p < 0.01$), while the corresponding degradation (increase) in II-F is not statistically significant under the same test, on both datasets. Furthermore, the degradation in NDCG between $\alpha = 0$ and $\alpha = 1$ is also small on both datasets, as shown in table 3, and the difference is not statistically significant under the student’s t-test ($p < 0.01$). These empirical results portray the opportunity to optimize recommender systems along multiple exposure fairness dimensions, while maintaining high recommendation quality.

7 Discussion and Conclusion

In IR, many measures of retrieval quality focus primarily on the position of relevant items in static ranked lists. By explicitly measuring deviations in system and target exposure, expected exposure metrics, in contrast, provide a useful framework to consider other dimensions of ranking evaluations, such as diversity and fairness, in addition to relevance. While, this framework have previously been employed in the context of individual users to individual and groups of items exposure fairness, we argue in this work that joint consideration of group attributes on both user-side and item-side allows us to study other forms of systematic unfairness of social and moral import in recommendation. In this setting, the choice of group attributes on both sides is an important consideration and must be informed by historical and social contexts as well as critical scholarship in the area of socioeconomic justice. For example, in the

Table 3: NDCG@50 for the recommender system corresponding to different values of α .

Dataset	$\alpha = 0$	$\alpha = 1$	$\alpha = 5$	$\alpha = 10$	$\alpha = 20$	$\alpha = 50$
<i>ML100K</i>	0.3703	0.3692	0.3684	0.3680	0.3674	0.3662
<i>ML1M</i>	0.2741	0.2736	0.2725	0.2712	0.2693	0.2684

context of job recommendation, such considerations may include an understanding of historical and ongoing pay and other forms of workplace discrimination based on gender and race [16, 58]. While, the analysis presented in this paper has considered group attributes for users (and correspondingly for items) along single dimensions, it is noteworthy that our formalization does not make any such assumptions. In fact, in the job recommendation scenario, it is meaningful to consider, say, both gender and race attributes of applicants jointly with grouping of jobs by salary and demographic attributes of corresponding business-owners in the GG-F metric formulation. Incorporating multiple group dimensions also raises additional considerations, such as of intersectional fairness [13, 15]. We have not yet analysed the appropriateness and sufficiency of our proposed framework for quantifying unfairness that may result from membership in multiple historically marginalized demographics but we believe this is an exciting area for future studies.

A historical perspective on fairness and justice is also critical to ensuring appropriate operationalization of our proposed framework. For example, readers should note that computing target exposure assumes the availability of *true relevance* labels. However, all real-world recommendation datasets likely suffer from historical biases [33] and therefore computing target exposure based on these *observed relevance* labels is also likely to be unfair to different demographics. In some cases, this gap between true and observed relevance may be mitigated by modifying how target exposure is computed—e.g., for GG-F it may be reasonable to employ E^{\sim} instead of E^* as the ideal target exposure distribution over item groups. The application of our proposed framework of metrics necessitates such careful deliberation to prevent the undesirable case of bias-in-bias-out.

Our work also has implications for the design and deployment of stochastic ranking models. Considering multiple fairness metrics, such as those proposed in this work, requires trading-off different forms of unfairness—e.g., II-F vs. GG-F. This may have implications for model optimization as well as the calibration of stochasticity in the model. While both our current work and previous literature on expected exposure metrics have considered stochastic models that randomize the output rankings by sampling noise from model-independent distributions, future development of stochastic ranking models may also want to consider approaches where the randomization is informed by the model’s own epistemic and aleatoric uncertainty [12]. On a concluding note, while our work has developed the framework of JME-fairness in the context of recommendation task, we believe our formulations can easily be extended to other IR tasks, such as search by replacing users with queries, without loss of generality.

References

- [1] Nina Baranchuk, Glenn MacDonald, and Jun Yang. 2011. The economics of super managers. *The Review of Financial Studies* 24, 10 (2011), 3321–3368.

- [2] Ruha Benjamin. 2019. Race after technology: Abolitionist tools for the new jim code. *Social Forces* (2019).
- [3] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. ACM, New York, NY, USA, 405–414.
- [4] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (2003), 993–1022.
- [5] Avishek Joey Bose and William L. Hamilton. 2019. Compositional Fairness Constraints for Graph Embeddings. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 715–724.
- [6] Sebastian Bruch, Shuguang Han, Michael Bendersky, and Marc Najork. 2020. A stochastic treatment of learning to rank scoring functions. In *Proc. WSDM*. 61–69.
- [7] Robin Burke. 2017. Multisided fairness for recommendation. *arXiv preprint arXiv:1707.00093* (2017).
- [8] Robin Burke, Nasim Sonboli, and Aldo Ordonez-Gauger. 2018. Balanced neighborhoods for multi-sided fairness in recommendation. In *Conference on Fairness, Accountability and Transparency*. 202–214.
- [9] Aleksandr Chuklin, Ilya Markov, and Maarten de Rijke. 2015. Click models for web search. *Synthesis lectures on information concepts, retrieval, and services* 7, 3 (2015), 1–115.
- [10] Aleksandr Chuklin, Pavel Serdyukov, and Maarten de Rijke. 2013. Click model-based information retrieval metrics. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. 493–502.
- [11] Daniel Cohen, Bhaskar Mitra, Katja Hofmann, and W Bruce Croft. 2018. Cross domain regularization for neural ranking models using adversarial learning. In *Proc. SIGIR*. 1025–1028.
- [12] Daniel Cohen, Bhaskar Mitra, Oleg Lesota, Navid Rekabsaz, and Carsten Eickhoff. 2021. Not All Relevance Scores are Equal: Efficient Uncertainty and Calibration Modeling for Deep Retrieval Models. In *Proc. SIGIR*. ACM.
- [13] Patricia Hill Collins and Sirma Bilge. 2020. *Intersectionality*. John Wiley & Sons.
- [14] Kate Crawford. 2017. The trouble with bias. In *Conference on Neural Information Processing Systems, invited speaker*.
- [15] Kimberlé W Crenshaw. 2017. *On intersectionality: Essential writings*. The New Press.
- [16] Greg Daugherty. 2021. The history of the gender wage gap in America. *Investopedia* (2021). <https://www.investopedia.com/history-gender-wage-gap-america-5074898>
- [17] Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. 2020. Evaluating stochastic rankings with expected exposure. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 275–284.
- [18] Michael D Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. 2021. Fairness and Discrimination in Information Access Systems. *arXiv preprint arXiv:2105.05779* (2021).
- [19] Michael D Ekstrand, Mucun Tian, Ion Madrazo Azpiazu, Jennifer D Ekstrand, Oghenemaro Anuyah, David McNeill, and Maria Soledad Pera. 2018. All the cool kids, how do they fit in?: Popularity and demographic biases in recommender evaluation and effectiveness. In *Conference on fairness, accountability and transparency*. PMLR, 172–186.
- [20] Batya Friedman and Helen Nissenbaum. 1996. Bias in computer systems. *ACM Transactions on Information Systems (TOIS)* 14, 3 (1996), 330–347.
- [21] Yingqiang Ge, Juntao Tan, Yan Zhu, Yinglong Xia, Jiebo Luo, Shuchang Liu, Zuohui Fu, Shijie Geng, Zelong Li, and Yongfeng Zhang. 2022. Explainable Fairness in Recommendation. In *Proc. of SIGIR*. ACM.
- [22] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS (JMLR Proceedings, Vol. 9)*. JMLR.org, 249–256.
- [23] F. Maxwell Harper and Joseph A. Konstan. 2016. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4 (2016), 19:1–19:19.
- [24] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *ICDM*. IEEE Computer Society, 263–272.
- [25] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased learning-to-rank with biased feedback. In *Proc. WSDM*. 781–789.
- [26] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2018. Unbiased Learning-to-Rank with Biased Feedback. In *IJCAI*. ijcai.org, 5284–5288.
- [27] Matthew Kay, Cynthia Matuszek, and Sean A Munson. 2015. 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*. 3819–3828.
- [28] Maurice George Kendall. 1948. Rank correlation methods. (1948).
- [29] Xiao Lin, Min Zhang, Yongfeng Zhang, Zhaoquan Gu, Yiqun Liu, and Shaoping Ma. 2017. Fairness-Aware Group Recommendation with Pareto-Efficiency. In *RecSys*. ACM, 107–115.
- [30] Tie-Yan Liu. 2011. Learning to rank for information retrieval. (2011).
- [31] R. Duncan Luce. [n.d.]. The choice axiom after twenty years. *Journal of Mathematical Psychology* ([n. d.]).
- [32] R. Duncan Luce. 1959. *Individual Choice Behavior: A Theoretical analysis*. Wiley, New York, NY, USA.
- [33] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [34] Rishabh Mehrotra, Ashton Anderson, Fernando Diaz, Amit Sharma, Hanna Wallach, and Emine Yilmaz. 2017. Auditing search engines for differential satisfaction across demographics. In *Proc. WWW*. 626–633.
- [35] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a Fair Marketplace: Counterfactual Evaluation of the trade-off between Relevance, Fairness & Satisfaction in Recommendation Systems. In *CIKM*. ACM, 2243–2251.
- [36] Alistair Moffat and Justin Zobel. 2008. Rank-Biased Precision for Measurement of Retrieval Effectiveness. *ACM Trans. Inf. Syst.* 27, 1, Article 2 (Dec. 2008), 27 pages. <https://doi.org/10.1145/1416950.1416952>
- [37] Nicola Neophytou, Bhaskar Mitra, and Catherine Stinson. 2022. Revisiting Popularity and Demographic Biases in Recommender Evaluation and Effectiveness. In *Proc. ECIR*.
- [38] Xia Ning and George Karypis. 2011. SLIM: Sparse Linear Methods for Top-N Recommender Systems. In *ICDM*. IEEE Computer Society, 497–506.
- [39] Harrie Oosterhuis. 2021. Computationally Efficient Optimization of Plackett-Luce Ranking Models for Relevance and Fairness. In *Proc. SIGIR*.
- [40] Harrie Oosterhuis and Maarten de Rijke. 2020. Policy-Aware Unbiased Learning to Rank for Top-k Rankings. In *SIGIR*. ACM, 489–498.
- [41] Sandeep Pandey, Sourashis Roy, Christopher Olston, Junghoo Cho, and Soumen Chakrabarti. 2005. Shuffling a Stacked Deck: The Case for Partially Randomized Ranking of Search Engine Results. In *Vldb*. ACM, 781–792.
- [42] Gourab K. Patro, Arpita Biswas, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. 2020. FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms. In *WWW*. ACM / IW3C2, 1194–1204.
- [43] Robin L Plackett. 1975. The analysis of permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 24, 2 (1975), 193–202.
- [44] Tao Qin, Tie-Yan Liu, and Hang Li. 2010. A general approximation framework for direct optimization of information retrieval measures. *Information retrieval* 13, 4 (2010), 375–397.
- [45] Filip Radlinski and Thorsten Joachims. 2006. Minimally Invasive Randomization for Collecting Unbiased Preferences from Clickthrough Logs. *CoRR abs/cs/0605037* (2006).
- [46] Filip Radlinski and Thorsten Joachims. 2007. Active exploration for learning rankings from clickthrough data. In *KDD*. ACM, 570–579.
- [47] Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. 2008. Learning diverse rankings with multi-armed bandits. In *Proc. ICDL*. 784–791.
- [48] Navid Rekabsaz, Simone Kopeinik, and Markus Schedl. 2021. Societal Biases in Retrieved Contents: Measurement Framework and Adversarial Mitigation for BERT Rankers. *arXiv preprint arXiv:2104.13640* (2021).
- [49] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI*. AUAI Press, 452–461.
- [50] Stephen E Robertson. 1977. The probability ranking principle in IR. *Journal of documentation* (1977).
- [51] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *KDD*. ACM, 2219–2228.
- [52] Ashudeep Singh and Thorsten Joachims. 2019. Policy Learning for Fairness in Ranking. In *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 5427–5437.
- [53] Tom Sühr, Asia J Biega, Meike Zehlike, Krishna P Gummadi, and Abhijnan Chakraborty. 2019. Two-sided fairness for repeated matchings in two-sided markets: A case study of a ride-hailing platform. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 3082–3092.
- [54] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Commun. ACM* 56, 5 (2013), 44–54.
- [55] Gábor Takács, István Pilászy, Botyán Németh, and Domonkos Tikk. 2009. Scalable Collaborative Filtering Approaches for Large Recommender Systems. *J. Mach. Learn. Res.* 10 (2009), 623–656.
- [56] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. 2014. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474* (2014).
- [57] Xuanhui Wang, Michael Bendersky, Donald Metzler, and Marc Najork. 2016. Learning to Rank with Selection Bias in Personal Search. In *SIGIR*. ACM, 115–124.
- [58] Ward Williams. 2021. Wage gaps by race. *Investopedia* (2021). <https://www.investopedia.com/wage-gaps-by-race-5073258>
- [59] Chuhan Wu, Fangzhao Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. 2021. Fairrec: fairness-aware news recommendation with decomposed adversarial learning. AAAI.

- [60] Haolun Wu, Chen Ma, Bhaskar Mitra, Fernando Diaz, and Xue Liu. 2021. Multi-FR: A Multi-Objective Optimization Method for Achieving Two-sided Fairness in E-commerce Recommendation. *arXiv:2105.02951 [cs.IR]*
- [61] Mingrui Wu, Yi Chang, Zhaohui Zheng, and Hongyuan Zha. 2009. Smoothing DCG for learning to rank: a novel approach using smoothed hinge functions. In *CIKM*. ACM, 1923–1926.
- [62] Himank Yadav, Zhengxiao Du, and Thorsten Joachims. 2019. Policy-Gradient Training of Fair and Unbiased Ranking Functions. *arXiv preprint arXiv:1911.08054* (2019).
- [63] Emine Yilmaz, Milad Shokouhi, Nick Craswell, and Stephen Robertson. 2010. Expected browsing utility for web search evaluation. In *Proceedings of the 19th ACM international conference on Information and knowledge management*. 1561–1564.
- [64] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. FA*IR: A Fair Top-k Ranking Algorithm. In *CIKM*. ACM, 1569–1578.
- [65] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. 2018. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. 335–340.