# The Role of Relevance in Fair Ranking

Aparna Balagopalan Massachusetts Institute of Technology USA aparnab@mit.edu Abigail Z. Jacobs University of Michigan USA azjacobs@umich.edu Asia J. Biega Max Planck Institute for Security and Privacy Germany asia.biega@mpi-sp.org

# ABSTRACT

Online platforms mediate access to opportunity: relevance-based rankings create and constrain options by allocating exposure to job openings and job candidates in hiring platforms, or sellers in a marketplace. In order to do so responsibly, these socially consequential systems employ various fairness measures and interventions, many of which seek to allocate exposure based on *worthiness*. Because these constructs are typically not directly observable, platforms must instead resort to using proxy scores such as *relevance* and infer them from behavioral signals such as searcher clicks. Yet, it remains an open question whether relevance fulfills its role as such a worthiness score in high-stakes fair rankings.

In this paper, we combine perspectives and tools from the social sciences, information retrieval, and fairness in machine learning to derive a set of desired criteria that relevance scores should satisfy in order to meaningfully guide fairness interventions. We then empirically show that not all of these criteria are met in a case study of relevance inferred from biased user click data. We assess the impact of these violations on the estimated system fairness and analyze whether existing fairness interventions may mitigate the identified issues. Our analyses and results surface the pressing need for new approaches to relevance collection and generation that are suitable for use in fair ranking.

# **CCS CONCEPTS**

• Information systems → Retrieval models and ranking; • Human-centered computing → ranking, fairness.

# **KEYWORDS**

relevance, ranking, measurement theory, fair exposure, validity

#### **ACM Reference Format:**

Aparna Balagopalan, Abigail Z. Jacobs, and Asia J. Biega. 2023. The Role of Relevance in Fair Ranking. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23), July 23–27, 2023, Taipei, Taiwan.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3539618.3591933

# **1** INTRODUCTION

Search systems generate rankings that do the sorting and provision of candidates to hire, of jobs to apply to, of housing options and of commercial products or services. As these rankings mediate access to opportunity [10, 61], various fairness measures and interventions haven been proposed to prevent harmful outcomes, such



This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGIR <sup>1</sup>23, July 23–27, 2023, Taipei, Taiwan © 2023 Copyright held by the owner/author(s). ACM ISBN 978-1-4503-9408-6/23/07. https://doi.org/10.1145/3539618.3591933 as systematic penalization of marginalized groups or new entrants to online marketplaces [8, 15, 28, 29, 48, 69, 77]. Such treatments characterize rankings as a causal intervention on the allocation of exposure of ranked items to searchers [38, 57] and direct mediators of fairness [10, 61], in line with social and legal perspectives on fairness-related harms [69, 77].

Many existing fairness definitions propose to explicitly distribute exposure in rankings as a function of worthiness or deservedness, e.g., how qualified for a job a candidate is [10, 13, 17, 24, 38, 47, 61]. In the fair ranking literature to date, however, the concept of deservedness has remained underdefined and has instead been represented by relevance: Fairness definitions, such as equity of attention [10] or fairness of exposure [61], mandate that item exposure allocations in a ranking be proportional to the item's relevance to the ranking query. In socially consequential search systems, relevance thus now plays a doubly important role in the distribution of the key livelihood-impacting resources - inequitable estimation of relevance then becomes inequitably allocated exposure. Yet, the construct of relevance in information retrieval has originally been developed to guide ordering of items to create results of high utility to searchers. Whether the score can be repurposed as a notion of worthiness in fair rankings remains an open question that we focus on in this paper.

We unite perspectives on measurement theory from the social sciences [34], information retrieval, and fairness in machine learning to critically examine the role of relevance in equitable information access. We begin with a conceptual analysis of relevance, examining it from historical and definitional angles, and probing whether it can reasonably guide fairness interventions as a proxy for worthiness. The analysis allows us to derive desiderata for relevance in fair rankings: a set of conditions (credibility, consistency, stability, comparability, and availability) the scores should satisfy in order to be a reliable component of exposure-based fairness measures.

We then empirically test whether the derived criteria are met in practice in a case study of relevance inferred from biased user click data. We approach these questions as questions of long-term dynamics in a complex system that combines algorithmic and human agents; as such we conduct a simulation study with both synthetic and real world datasets, similar to past efforts [11, 20, 25].

Click-based approaches, measuring relevance via user browsing models (i.e., conditional probabilities of relevance given a model of user clicking behavior [22, 68]), are indeed a common way to gather relevance at scale as it is not readily available otherwise. Models of click behavior have been shown in prior work to sometimes poorly represent searcher behavior, reflect biases, or capture only narrow conceptions of relevance. Our results furthermore show that they can violate some of the desiderata of relevance for fair ranking. In particular, click-based relevance scores violate the criteria of availability and comparability – as the feedback is collected only for the top-ranked items [22], accurate may not be available for most items in the corpus, and the resulting scores may not allow for accurate proportional relevance comparisons between individuals and groups. This paper makes the following contributions: (1) We probe the role of relevance as a proxy for worthiness in fair exposure allocation, and derive a set of five desired criteria for relevance to be a valid and reliable proxy. (2) We consider click-based relevance rankings as a case-study and observe that some desired criteria may not be met in practice. (3) We explore factors such as class imbalance or fairness interventions that may mitigate the impact of desiderata being unmet in practice, and find that impact can vary from dataset to dataset. In summary, our analyses and results demonstrate that relevance needs to be critically examined as a valid and reliable target for fair exposure allocation. The results also surface the pressing need for new approaches to relevance collection and generation (e.g., inference of relevance as calibrated probabilities), while providing a set of criteria for viable solutions.

## **2 FAIR RANKING: THE ROLE OF RELEVANCE**

The probability ranking principle [54] states that an optimal ranking sorts items in order of relevance–where relevance is assumed to be unbiased and consistent. Yet, despite this 'optimal' ranking of items, the outcomes may be unfair to different groups of users [10, 13, 14, 17, 24, 47, 61, 71, 75]. For instance, we may observe an unfair ranking when top search results for a recruiter feature primarily male candidates, despite the presence of qualified female candidates in the selection pool. Many papers have proposed methods to quantify and mitigate unfairness by allocating ranking exposure in proportion to "worthiness" of allocation [10, 13, 17, 24, 47, 61, 75]. Yet, the concept of worthiness remains undefined and often the relevance of an item is used as a proxy [10, 53, 61]. In this section, we examine the role of relevance in fair ranking and probe its validity as a proxy for worthiness.

## 2.1 Goals of Fair Exposure Allocation

All ranking systems and operationalizations of fairness express a normative goal. In fair ranking, different interpretations of fairness thus build from differing normative theories of discrimination underlying each framework [7]. For example, for a goal of equal opportunity, under the view that similar items should attain equal attention, one might allocate exposure in proportion to the notion of an item's attention-deservedness or "worthiness" at ranking time. Alternatively, the objective of demographic parity would be to have all groups realize their potential equally well; in implementation, this could be achieved by equalizing the levels of exposure per group. The ideal exposure rates for items may also be externally determined based on domain knowledge or through legal regulations: This includes, for example, the Rooney Rule [18]. We direct the interested reader to recent work by Zehlike et al. [76] which delineates the normative dimensions underlying several fair ranking techniques.

In many fairness metrics, the *relevance* of a ranked item in response to a query is considered to be a proxy for its worthiness, and the attention an item receives from searchers should ideally be proportional to its relevance [10, 53, 61].

# 2.2 Measuring and Mitigating (Un)Fairness

2.2.1 *Measuring Fairness.* In this paper, we focus on the family of exposure-based or attention-based metrics that allocate exposure in proportion to relevance [10, 13, 17, 24, 47, 61, 75].

The exposure obtained by group  $G_k$  under ranking P is defined as the mean exposure obtained by all items in  $G_k$ :

$$Exposure(G_k|P) = \frac{1}{\|G_k\|} \sum_{d_i \in G_k} Exposure(d_i|P)$$
(1)

Similarly, relevance of group  $G_k$  is defined as the mean inferred relevance of all items in it:

$$Relevance(G_k) = \frac{1}{\|G_k\|} \sum_{d_i \in G_k} Relevance(d_i|P)$$
(2)

Since we consider the case of a binary protected attribute in our experiments, for simplicity all the metrics have been defined using two groups (but could be extended). Three commonly-used metrics are:

*Demographic fairness parity.* Ratio between exposure of groups, agnostic to base rates of relevance [47]), measuring how much the exposure allocated to different groups varies:

$$\frac{Exposure(G_0|P)}{Exposure(G_1|P)}$$
(3)

*Exposure fairness:* Ratio between exposure-relevance ratios of groups measuring how much more exposure each group receives proportional to its relevance [47]:

$$\frac{Exposure(G_0|P)}{Relevance(G_0|P)} / \frac{Exposure(G_1|P)}{Relevance(G_1|P)}$$
(4)

*Individual fairness:* Sum of absolute difference in the exposure and relevance scores for individual items, amortized over a sequence of rankings [1,M], with N items [10].

$$\sum_{i=1}^{N} \left| \sum_{j=1}^{M} Exposure(d_i|j) - \sum_{j=1}^{M} Relevance(d_i|j) \right|$$
(5)

To estimate exposure for different positions, a weight denoting the attention a ranking position obtains from searchers on average is required. We use a log-decaying attention model for top-k positions, similar to the discounts in NDCG [32]. All relevance and exposure scores are [0,1]-normalized following Biega et al. [10]. While we focus on the metrics listed above in our paper, we highlight that additional exposure-based fairness metrics have been proposed in recent literature, and direct the interested reader to work by Raj and Ekstrand [53] for a comparitive analysis.

2.2.2 Fairness Interventions. Several methods have been proposed to train ranking systems that maximize fairness according to the metrics described above. These approaches are performed during pre-processing [4, 63] (e.g., Sonoda [63] consider the order of the training data), as in-processing [47, 72] (with a fairness constraint while training a learning to rank system), or as post-processing [10, 31, 43, 61] (re-ranking items so that system fairness is high). Notably, the expected relevance scores in the above fairness formulations may not be known a priori, and in practice one would work with *estimates* from some predictive model. In particular, in the systems we consider, these scores may be obtained from the ranking model itself.

# 2.3 Probing Relevance as a Proxy for Worthiness

2.3.1 *Defining Worthiness.* Fairness interventions seek to produce rankings that are more equitable, i.e., rankings that are capturing something not totally captured by inferred relevance. Several recent papers explicitly or implicitly refer to various notions of "merit" or "worthiness" as a basis for their intervention:

- "allocation of exposure based on merit (i.e, relevance)" [70],
- "We define the merit of a document as a function of its relevance to the query [...] and we state that each document in the candidate set should get exposure proportional to its merit" [62],
- "relevance can be thought of as a proxy for worthiness in the context of a given search task" [10],
- "to implement fairness constraints based on merit, we need to explicitly estimate relevance to the user as a measure of merit" [47].

In this paper, we refer to *worthiness* as the underlying construct that fair rankings aim to infer and operationalize within their (possibly adjusted) relevance scores. Yet, in practice, this construct eludes an easy definition. Zehlike et al. [76] note that among fairness metrics, some consider a notion of merit when measuring disparities in exposure while mostly understanding "merit as the utility score *Y* at face value. However, ...the understanding of merit depends on worldviews and on one's conception of equal opportunity."

Under one interpretation, the worthiness of an item, such as a job seeker's profile on an online hiring platform, can be understood as the value from allocating attention to the item. Yet, a key question remains: the value for whom? Job seekers might receive value from being allocated attention to searchers likely to hire them. Searchers, on the other hand, might receive value from being exposed to job seekers who are qualified and likely to stay with the employer longterm. Worthiness scores based on the value for different stakeholders might thus diverge.

Another interpretation of worthiness, rather than tying the score to the allocation value, might look at inherent deservedness of the ranked items in a given application domain. The difference between these two interpretations is akin to the difference between outcome and process fairness; see, e.g., the discussion by Barocas et al. [7] on merit vs. desert for allocation.

Having laid out some of the definitional nuances of worthiness, the remainder of this section will focus on our paper's key question: Is relevance an adequate and sufficient proxy for worthiness?

2.3.2 A Definitional Perspective on Relevance. Relevance is a construct central to all search systems. However, it is important to note that a single definition might not apply to all scenarios. While the concept of relevance has been examined in detail before as a construct that guides information seeking behaviors [12, 40], several valid definitions exist, including as a user-dependent and measurable construct of information need, utility or usefulness of the viewed and assessed information object, topicality, and others [12, 56]. Accordingly, different definitions of relevance have been instrumented by information systems over time [56]. For example, in search systems, relevance may refer to the topical match between a query and a web-page [73] or binary "appropriateness" scores [41], both often crowd-sourced and averaged [6].

In an ideal system, item ranking order should correspond to the latent relevance and satisfy a users's information need [56]. Needless to say, a single-dimentional notion of relevance used during system training might not satisfy all specifications and user needs. Relevance is a convoluted relation involving a given user's information need and the properties of ranked items – objective scores of relevance may not exist, and instead are only defined within the context of a particular stakeholder's frame of reference [56]. Despite this complexity, relevance is often infferred using common *click measurement models* that make simplifying definitional assumptions about relevance.

2.3.3 Measuring Relevance: From Definitions to Practice. Constructs such as relevance cannot be operationalized directly and are instead inferred from measurements of observable properties thought to be related to them via a measurement model [34]. Behaviorist search systems operationalize relevance as user engagement with ranked content, implicitly assuming that the more (latently) relevant the content is, the higher the engagement. Thus a common strategy is to estimate relevance of an item to a given query using behavioral data in the form of clicks. Clicks are translated to relevance through a

user browsing model. Such models typically also account for various measurement issues, such as interaction and cognitive biases of the searchers, or the physical limitations of a ranking infrastructure.

Another approach to measurement of relevance is crowdsourcing. Here, "ground-truth" relevance judgments [56] are collected from human annotators for a list of items given a query. Because of how complex the concept of relevance is, it is however hard for external annotators to estimate another searcher's relevance [5]. Moreover, the annotation process is costly and impossible to reliably scale. Indeed, relevance judgments elicited from annotators are used primarily for system evaluation, not training [65].

There are several assumptions that have been traditionally made in relevance measurements (summarized from Saracevic [56]): that relevance is binary (ranked items can be relevant or not relevant), that relevance judgments of different items are independent, that relevance has a non-dynamic nature (i.e., values do not change over time), that relevance judgments exhibit low variance across users in non-personalized systems, and others. In operationalizing relevance, from the theoretical understanding to the model underlying data collection, it is crucial to note that several of these assumptions—for example, the binary nature of relevance—are still prevalent.

2.3.4 Issues with Relevance as a Worthiness Score. We hypothesize that using relevance scores as a proxy for worthiness to fairly allocate exposure has thus far been a choice of convenience. We base this claim on the fact that the construct of relevance has primarily been investigated in the context of accurate ranking construction [12, 40, 49, 56], not as a worthiness score to guide fair exposure allocations.

Additionally, limitations of relevance scores [23] (such as unavailability for all ranking subjects or thebinary nature of elicited judgments), of their measurement process (such as misspecifications in browsing model parameters [39]), and the consequences thereof, have also primarily been studied in the context of ranking utility. In the context of fairness, to address the unavailability of relevance scores, Kırnap et al. [42] proposed a methodology to estimate fairness metrics using incomplete relevance judgments. However, these approximations are useful for system evaluation, not for training algorithms that fairly allocate exposure. Similarly, fair ranking evaluation benchmarks rely on exhaustive annotations that are infeasible at the scale of real systems [9].

In addition to the *incompleteness issues*, using relevance as a proxy for worthiness involves making certain implicit assumptions about their relationship: for example, that the relative ordering of items according to worthiness is consistent with the ordering by relevance. Moreover, high variance or noisiness in relevance scores may also cause fairness measurements to be less robust or more arbitrary [19]. It is essential to elucidate and test the assumptions of how the two concepts relate to each other if relevance is to be a *valid* and *reliable* candidate for approximating worthiness.

Attending to the two-sided nature of search systems helps us further understand the way that values and positionality are embedded in relevance: Relevance defined as utility to the searcher/consumer and relevance defined as utility to the items/providers can be in conflict in fair ranking systems [47, 61]. Moreover, search systems themselves may have additional goals as well, such as maximizing the time users spend on the platform. It is thus crucial to surface whose judgment of worthiness a given relevance proxy represents.

Table 1 summarizes several common relevance measurement approaches and explores their potential issues as proxies of worthiness.

Aparna Balagopalan, Abigail Z. Jacobs, and Asia J. Biega

Worthiness Proxy	Positionality of Value	Potential Issues with Measurement		
Crowd-sourced labels [6, 23, 41]	Depends on labeling instruction and perspective of labelers	May be unavailable for many items, or binary valued		
Topical match between query and item [73]	Depends on attributes used	May be noisy, depends on similarity metric		
ML-based predicted relevance [47]	Searcher, if inferred using interaction data	Potentially noisy, may not generalize to new queries		
Click-based relevance [1, 39]	Searcher	Binary-valued, prone to biases, may depend on size of click logs		

Table 1: Potential limitations with using relevance scores as proxies for worthiness. Here, the "positionality of value" denotes the position from which ideal "worthiness" is judged (i.e., whose value from the ranking is considered; see Section 2.3.1 for details).

## **3 CASE STUDY: RELEVANCE FROM CLICKS**

In this paper, we focus on relevance inferred from clicks as a case study. An overview of the inference process shown in Figure 1.<sup>1</sup>

# 3.1 Ranking with Click-Based Relevance

Unbiased learning-to-rank is a well studied problem in information retrieval to learn accurate rankings from biased training data [1, 33, 39, 64, 67, 68]. Some are online ranking algorithms that involve interventions and collecting real-time user feedback [50, 59, 74]. Many of these systems also rely on a counterfactual model of estimations to debias the ranking loss. One common approach is to use Inverse Propensity Weighting or Score (IPS) [39] to account for various user cognitive biases in standard ranking losses, such as position bias [68], trust bias [1], or selection bias [46]. Other work proposes jointly training a ranking model and an examination propensity model [3, 33].

Here we introduce a prototypical learning-to-rank framework with click-based relevance, including the learning algorithms, inference methods, and evaluation metrics.



Figure 1: Relevance inferred using searcher clicks. A searcher inputs a query, with a latent notion of relevance (R) of results. Then, they examine (E) and click (C) on the ranked results, thereby allocating exposure to the items. However, they may exhibit cognitive biases (such as examining and clicking on top-ranked items with higher probability). In order to infer relevance of items from these click logs, the data has to be debiased. Browsing models perform this operation by parameterizing the examination probability and relevance of ranked items, and then defining a model linking the click probability with these for a query and item. With such a model, the probability of an item being relevant can be estimated or relevance can be inferred. Then, to predict the relevance for an unseen query/item, a machine learning model is trained, with loss terms influenced by parameters from the browsing model (e.g., with propensity-weighting). Finally, the quality and fairness of exposure allocation is measured.

<sup>1</sup>Code: https://github.com/Aparna-B/FairRankingRelevance

3.1.1 From Clicks to Relevance via Browsing Models. Click or browsing models - derived using empirical studies of user information retrieval/search behavior - make it possible to simulate ranking systems, approximate item relevance, and evaluate search systems [16]. The relevant-intent hypothesis states that an item is clicked if and only if it is relevant and examined, which is then used to infer the latent relevance using algorithms such as Inverse Propensity Scores. The main variables being modelled with different assumptions in click or user browsing models is the probability of examination of items in a ranked list:  $o_r$ @rank r is a binary variable indicating whether an item at position-r is examined or not, and the relevance of items being ranked. Most browsing models parameterize the examination probability with known distributions identified using domain knowledge. When the browsing model changes, the probability of examination at a specific rank  $p(o_r = 1)$ , and the model for user behaviour changes. As a high-level summary, a browsing model is a graphical model with two unseen nodes - examination and relevance - and an observable node of clicks. (For a comprehensive overview of click models, see [16].)

In this paper, we use a position-biased browsing model of user behavior. This model assumes that examination probability is constant given a specific rank *r*. (We also note that browsing models are also used in other applications, such as debiasing user logs.)

3.1.2 Unbiased Learning to Rank with Propensity-Weighting. The Propensity-Weighted ranking algorithm infers propensity weights to minimize a specific, pre-defined loss function l(f) and supervise a ranking model f. In most ranking systems, the expected loss can be expressed as a sum of loss terms over relevant items in a list (or a pair of items). The ranking algorithm proposed by Joachims et al. [39] uses Inverse Propensity Score (IPS) weighting to account for the biased nature of clicks (e.g., position bias as described in Section 3.1.1), and weight each items's loss correspondingly. In IPS weighting, each item's loss is inversely weighted with its examination probability if it is clicked. Note that only items that are clicked contribute to this IPS-weighted loss. In the next section, we describe how the ranking model f is trained and optimized.

3.1.3 Learning a Ranking Model. This is the machine learning model f that predicts relevance used to produce a ranking. We consider models optimized with a propensity-weighted listwise loss using the outputs of deep neural network models, following Joachims et al. [39]. We optimize a list-wise softmax-based cross-entropy loss while training the ranking model [2, 3]. This loss is further weighted by propensity and optimized during model training, with best parameters chosen based on validation set performance. In the next section, we list evaluation metrics for assessing the utility of ranking systems.

*3.1.4 Performance Evaluation.* Finally, the quality of the ranking algorithm is evaluated. Typically, the ranking produced by the model is compared against a ground truth ranking (assumed to be an unbiased gold standard). In this paper we focus on the Normalized Discounted

Cumulative Gain (NDCG), a normalized measure of ranking quality favoring occurrence of more relevant documents in the higher ranking positions through a logarithmic discount [35].

# 4 DESIDERATA OF RELEVANCE IN FAIR RANKING

In this section, we describe a set of desiderata that the relevance scores ought to satisfy to be meaningful and useful for fairly allocating exposure. Drawing on domain knowledge about relevance-based ranking systems and the assumptions within, these desiderata offer a pathway towards *valid* and *reliable* measures of relevance, where the properties of proxy scores should match their theoretical ideal across a range of qualitative dimensions [34].

In the definitions below, let *w* denote true relevance (worthiness) scores, and *r* denote inferred relevance scores, i.e., the measurements of *w* that will be used to allocate exposure. *K* denotes the number of items to be ranked.

*Credibility.* For inferred relevance scores to be useful, they must behave as we expect. In a ranking setup, we can ask: do items with high true relevance (worthiness) have higher inferred relevance scores across runs? (And lower, lower?) Absence of credibility would imply that the inferred relevance scores are not capturing how the 'true' relevance scores would be expected to, and so lacks face validity and content validity—and may even lack reliability across similar inputs to the model [34]. Thus, it might be an unsuitable proxy under views of fairness such as "equal opportunity" where exposure should ideally be allocated in proportion to worthiness.

**Definition 4.1** (Credibility). If *i* and *j* are two items to be ranked, credibility necessitates, in expectation:

$$w_i \ge w_j \equiv r_i \ge r_j \tag{6}$$

*Consistency.* A core assumption in unbiased learning-to-rank is that of consistency: in the limit of sufficient data, the estimated expected relevance scores converge. Thus, meeting this statistical property necessitates the criterion: do the inferred relevance scores converge? If this criterion remains unmet, then fair ranking metrics relying on relevance may not converge either.

**Definition 4.2** (Consistency). In the limit of sufficient data, the predicted scores converge within an acceptable error range. Let *N* denote the number of expected training steps (e.g., the size of the training set), and  $\hat{r}_i$  the predicted relevance score for item *i* at the *N*th training step. If  $S_n = \frac{1}{K} \sum_{i=1}^{K} [(r_i^n - \hat{r}_i)^2]$  measures the mean squared deviation between relevance scores obtained by training the relevance prediction model for *n* steps and  $\{\hat{r}_i\}_{i=1}^K$ , and  $\epsilon$  is the acceptable error value (close to 0), then consistency requires:

$$\exists N_0 \text{ such that } S_n \leq \epsilon \,\forall n \geq N_0 \tag{7}$$

Stability. Relevance ought to have test-retest reliability, i.e., that measurements from the same model for the same input do not vary more than an expected limit (e.g., due to unavoidable stochasticity). In the ranking setup, the criterion to test then is: do the inferred relevance scores vary more than an established limit across runs with minor variations in stochastic parameters or initialization? This is related to robustness and how noisy estimates of fairness would be.

**Definition 4.3** (Stability). If *n* denotes the number of experimental runs with variations in unimportant parameters, and  $\epsilon$  is the

acceptable variation across runs, then:

$$\frac{1}{K} \sum_{j=1}^{K} \left[ \frac{1}{n} \sum_{i=1}^{n} \left[ \left( r_{ij} - \mu_j \right)^2 \right] \right] \le \epsilon$$
(8)

where  $\mu_j = \frac{1}{n} \sum_{i=1}^{n} r_{ij}$ , and  $r_{ij}$  denotes relevance for the *j*th item under the *i*th run.

*Comparability.* Several fairness interventions suggest that exposure should be allocated in proportion to inferred relevance [10, 61], with the underlying assumption that this proportionality would hold for true relevance (worthiness) of items/groups. Thus we must verify that the relative ratios of relevance scores match the corresponding ratios of true worthiness, i.e., that our model has structural validity that these relative properties behave as expected [34]. For group fairness, the average relevance is taken as the aggregate across all items belonging to a group. For individual fairness, this criterion is evaluated at the per-item level.

**Definition 4.4** (Comparability). If *i* and *j* are two items or groups of items (e.g., defined on the basis of a sensitive attribute) to be ranked, comparability necessitates:

$$\frac{w_i}{w_j} \approx t \equiv \frac{r_i}{r_j} \approx t \tag{9}$$

Availability. Past research in Information Retrieval (IR) has established that searchers exhibit position-bias in viewing and inspecting ranked results [22]. As a result, this variance in attention across different positions is modeled while allocating exposure for an item or group of ranked items. Then, exposure is allocated in proportion to uncorrected (i.e., without any bias-correction) relevance scores. Thus, an underlying assumption is that relevance scores are *available* for all items, and are *unbiased* estimates of true relevance (or worthiness). The availability criterion tests this assumption that unbiased inferred relevance scores are available for all items—without which it would be difficult to claim validity or reliability of relevance scores.

**Definition 4.5** (Availability). If f is a function that outputs 1 if f(x) is defined and 0 if not, then,

$$f(r_i) = 1 \,\forall i \in \{1, 2...K\} \tag{10}$$

Additionally, we define that for the availability property to be met, the distribution of w across items should be statistically indistinguishable from that of r.

Note that some of these properties are *inter-related*. For example, availability may be a necessary condition for comparability to be tested. Thus, it may be possible to design a fixed order in which these properties could be tested.

#### 5 EXPERIMENTAL SETUP

We describe the setup for testing if desiderata described in Section 4 are met in practice when relevance is inferred using searcher clicks.

# 5.1 Data

Datasets in our experiments consist of a list of items which are to be ranked in response to a query. Each item is associated with some features and a relevance label. We assume that the ranking with highest-utility is one where items are arranged in decreasing order of these ground-truth relevance labels. Items also have labels denoting (protected) group membership. We split each dataset in 70-10-20 proportions as train, validation, and test splits. The details of each dataset are below and summarized in Table 2. SIGIR '23, July 23-27, 2023, Taipei, Taiwan

Dataset	d	Relevance Distribution, % (4=highest)				
		0	1	2	3	4
synth-normal	2	0.13	11.00	61.51	26.65	0.71
synth-pareto	2	86.78	11.39	1.58	0.21	0.04
fairtrec	3	50.00	50.00	-	-	-

Table 2: Datasets statistics. *d* denotes the number of features. Each column *i* indicates the relative frequency of relevance grade *i* in a given dataset (4: most relevant, 0: least relevant). Each dataset contains 50000 samples

5.1.1 Synthetic Data. To experiment under controlled conditions, we synthesize two datasets from a graphical model using the synthetic data generation methodology proposed by Yang et al. [71]. In each case, the correlation Directed Acyclic Graph (DAG) consists of four attributes: continuous attributes for 2 features and a relevance label, and one binary group membership attribute. All three continuous attributes follow a Pareto or Normal distribution (see appendix for details). After sampling, we discretize all relevance scores to 5 grades uniformly based on the value, similar to standard learning-to-rank setups. We consider binary protected groups for all fairness analyses (e.g., men and women). We sample equal number of items from both groups. We sample N=50,000 datapoints from both synthetic distributions.

**Pareto distribution:** We sample relevance from a Pareto distribution with P(2.0,1.0). This dataset is referred to as synth-pareto.

**Normal distribution:** Relevance is sampled from a normal distribution with mean and standard deviation set to:  $\mu = 2, \sigma = 1$  respectively. This dataset is referred to as synth-normal.

5.1.2 FairTREC 2021. The FairTREC 2021 [27] dataset consists of Wikimedia articles ranked in response to text queries. In total, there are 50 train and 50 test queries. To align the setup of this dataset with our study, we choose a single query for analysis. We select a train-set query with atleast 25,000 number of relevance annotations available (train query ID 6). Additionally, only positive relevance annotations are available in the FairTREC dataset. We make the assumption that all the other Wikimedia articles in the corpus are not relevant to the query. Finally, we subsample the dataset to a similar size as the synthetic datasets to obtain a dataset with 50,000 items. Each item in this dataset thus consists of a query, article text, and a binary relevance judgment. To obtain a representation for ranking, we utilize a pre-trained checkpoint of a cross-encoder model trained on the ms-marco  $[21]^2$  task. Finally, we embed the high-dimensional representations into three dimensions using Principal Component Analysis (PCA). This is the final "item representation". Additionally, the geographical location metadata is utilized as the sensitive attribute for fair ranking. We binarize these categories into "majority" location (Europe only) and all others (including no available metadata). This dataset is severely imbalanced in terms of group membership with a 90% majority in group distribution.

## 5.2 Simulation Protocol

We experiment in a single-query setup without personalization. In each round, we generate synthetic clicks to simulate a set of users clicking on items based on the position-biased examination probability and relevance of each item to be ranked, using a position-biased click model (as specified in Sec. 3.1.1; with a positive click probability of 1, negative click probability of 0.1, and maximum relevance grade of 4). Then, the average click through-rate of each item is accumulated to compute its propensity and the ranking model is optimized with inverse-propensity weighted optimization (Sec. 3.1.2). In our experiments, we use pre-computed propensity weights using result randomization (i.e., a pre-trained propensity model) for the positionbiased click model<sup>3</sup> while training the ranking model. The metric optimized in each step is the list-wise ranking loss with a selectionbias cutoff (set to 10). The final ranking system output is a list of items, ranked in decreasing order of expected utility to users. This is repeated for 500 iterations in total during training, and performance is evaluated, with models checkpointed every 50 iterations. Model loss is optimized with stochastic gradient descent with a learning rate of 0.01.<sup>4</sup> Performance on the test set is evaluated with the best model checkpoint (based on validation performance). This process is repeated with 10 random seeds, and averaged for all metrics. Note that there are several sources of randomness in this setup, including the items that are clicked, the initialization of parameters in the ranking model, and sampling of data batches during training.

## 6 EXPERIMENTS AND RESULTS

In this section, we seek to answer the question: To what extent can click-based relevance be used to guide fair exposure allocation?

### 6.1 Studying Properties of Relevance

We empirically test if the desired criteria for relevance are met in practice on three datasets. We assume that the true relevance labels are proportional to true worthiness, and we refer to ground truth relevance judgment labels as true relevance, and inferred relevance labels as predicted relevance. We note that both the tests for operationalizing the formal criteria (see Section 4) and the thresholds for meeting these criteria are themselves choices. For example, credibility can be measured using the mean deviation between the inferred relevance scores and true relevance scores for point-estimates, or by measuring the overall Spearman correlation between them (which may be closer to the rank ordering definition). Further, to say that a criterion is met, one may set acceptance thresholds, e.g., that a given criterion might be met only if a specific correlation is greater than 0.5. Thus we emphasize that meeting these criteria may be a continuum, not a binary outcome. Summary of findings is in Table 3. 6.1.1 Credibility. We study if relevance inferred from clicks under this prototypical setup is credible. Specifically, we test if the top-items across multiple runs of the learning-to-rank system (i.e., different seeds) correspond to top items (i.e., the ordering is retained). We construct boxplots showing the distribution of relevance scores (with full list size being the size of the test set; all softmax scores within a given ranking for ease of interpretation) for each relevance judgment label grade. Ideally, we expect that items with higher relevance grades would occupy lower ranks by having high inferred relevance scores. We assert that the relevance scores are not credible if the population medians corresponding to different relevance grades are not significantly different at a level of 0.05.

**Results.** From Figure 2, we observe that for increasing values of true relevance judgments, the medians of the inferred relevance scores for items are monotonically increasing for all the datasets. With a Kruskal Wallis-H test, we observe that these medians are significantly different (p < 0.001). This lets us conclude that the relevance scores are indeed reliable across runs for all three datasets. However,

<sup>&</sup>lt;sup>2</sup>https://huggingface.co/cross-encoder/ms-marco-TinyBERT-L-2-v2

<sup>&</sup>lt;sup>3</sup>From the ULTRA-pytorch simulation repository.

<sup>&</sup>lt;sup>4</sup>We validated that an Adagrad optimizer produced similar train loss.

The Role of Relevance in Fair Ranking



Figure 2: Distributions of inferred vs true graded relevance scores. Items with high true relevance judgments must have high inferred relevance for the *credibility* criterion to be met. We observe this to be true on average (i.e., ordering of medians is monotonic, though there is considerable overlap between inferred scores for different relevance grades).



Figure 3: Consistency of predictions. Relevanceconsistently converges, as measured using mean squared error (MSE).



Figure 4: Distributions of [0,1]-normalized inferred and true relevance scores: the inference process does not preserve distributions.

post-hoc tests may be required to test if medians of specific groups (e.g. those with relevance grades of 2 vs 3) are significantly different.

*6.1.2 Consistency.* We measure the mean squared error between the predicted relevance scores at iteration *i* and the final iteration (i.e., *i* = 500) for *i* = 10,20,30...490. Note that we use the logit-score predictions, before the *softmax* operation. These scores are computed on the validation rather than test set (since test set is unseen throughout training). Then, we assess if this mean squared error converges. Based on the definition for consistency (see Section 4), we need to choose a  $\epsilon$  value to assess consistency. Here, we set  $\epsilon = 0.10$ 

**Results.** We observe that the mean squared error converges to the final value for all datasets, though with some noise in fairtrec. As a result, the consistency criterion is satisfied all datasets on the validation set. (Consistency also met for datasets on the train set.)

6.1.3 Stability. We measure the deviation in inferred relevance values across multiple runs. We compute the average standard deviation in relevance measurement of a specific item across the ten random seeds, and average this value across all items. To ensure that all values can be compared across datasets, we scale predicted scores for each item in given list by subtracting the mean and dividing by standard deviation, all derived from the predictions for all items in that list.

**Results.** Among the synthetic datasets, we find that the mean variation is is between 0.17-0.18. Since all relevance scores have been normalized with unit-variance, this corresponds to less than one standard deviation of the full relevance distribution. The results for the fairtrec dataset show similar trends (variances scores ~0.41). As a result, we conclude that the relevance scores are reasonably stable for all datasets.

6.1.4 Comparability. We compute two metrics to assess the comparability criterion: (1) the Spearman correlation between predicted and true relevance scores for all items, (2) the ratio between average relevance scores for groups for true and predicted relevance scores. We assess that the criterion is met if the correlation is  $\geq$  0.3 from (1) and the difference in ratios from (2) is not more than 0.05. Note that (1) is closely tied to individual fairness since it considers prediction for all items. On the other hand, (2) is directly consequential to group-level exposure fairness: specifically, exposure is allocated in

Aparna Balagopalan, Abigail Z. Jacobs, and Asia J. Biega

proportion to expected group relevance (either predicted or true) as per several fair ranking formulations [10, 37, 53].

**Results.** Spearman correlation between true and predicted relevance scores is 0.09 (0.09 for subgroups) for synth-normal. For synth-pareto, the Spearman correlation between true and predicted relevance scores is 0.23 (0.20, 0.26 for female and male subgroup respectively). This indicates that the comparability criterion of relevance scores is better for one subgroup here. In contrast, the trends observed on the fairtrec dataset are different: the correlation is 0.27 for all items (subgroup correlations of 0.27 and 0.29).

The ratio between normalized average relevance for subgroups using true relevance scores are close to 1 for both synthetic datasets (by construction). Ratios between average inferred group-relevance scores are also close to 1.00 for synth-normal and synth-pareto respectively. The difference for fairtrec is similar, with the two ratios being 1.04 and 1.02 with true and inferred relevance respectively. Thus, the target for exposure ratios may change slightly on using true vs predicted relevance scores. In this case, as per the criteria set, the comparability criterion is not satisfied for all datasets. The comparability criterion related to group-fairness are met on all datasets. 6.1.5 Availability. We analyze the distribution of inferred relevance scores and compare them to the distribution of true relevance labels. Results. By design, since a machine learning model can make predictions for all items, the continuous relevance score predictions are available for all items. However, when we visualize the histogram of predicted and true relevance (plots in Figure 4), we observe that the nature of distributions is not retained. We computed the two-sample Kolmogorov-Smirnov test [44] for goodness of fit with null hypothesis being that the two samples - true and inferred relevance scores arise from the same distribution. We find that the null hypothesis can be rejected in all cases with  $p \ll 0.001$ . As a result, we conclude that the availability criterion is not satisfied for at least subset of ranked items. This makes as sense as available "worthiness" judgements here are graded, while predictions are continuous.

Dataset	Credibility	Consistency	Stability	Comparability	Availability
fairtrec synth-normal synth-pareto	\$ \$ \$	5 5 5	\ \ \	X X X	× × ×

Table 3: Empirical tests of relevance score desiderata on three datasets. In this table, we summarize whether the five desired criteria are empirically met in the datasets.

# 6.2 Impact of Violating Criteria in Practice

While relevance is the construct measured in ranking systems, the user-facing output is a ranked list of items (which is a function of the inferred relevance). Hence, system evaluations form an important part of optimizing ranking systems. We evaluate ranking systems along two axes: (1) quality or utility of ranking, (2) fairness of ranking. We measure the quality of rankings using NDCG@10 score for all datasets, and obtain scores of 0.90 (near perfect score)<sup>5</sup>, 0.29, and 0.26 for fairtrec, synth-pareto, and synth-normal respectively.

To test if violating the criteria matters in practice, we measure exposure fairness using the metrics defined in Sec. 2, all computed either using true, *graded* or inferred *continuous* relevance scores on the test set. From Table 4b and Table 4a, we observe that fairness assessments may vary depending on whether true or inferred relevance scores are used, and the size of difference is often dataset-dependent. In particular, the difference does not seem particularly stark for the synthetic datasets. However, on computing Wilcoxon signed-rank tests paired by the random seed (which influences the stochastic initialization of the prediction model), the exposure fairness metrics are significantly different for all datasets. Similar trends are observed for individual fairness: the difference is significantly different, where pairing is performed at item-level (i.e., 10000 items). We observe that these results are sensitive to normalization of relevance scores.

Thus, fairness metrics vary depending on whether true or predicted relevance scores are used. Notably, the comparability criterion is also violated (especially for individual fairness) on all three datasets (Section 6.1). Analyzing the degree to which a criterion is violated may yield interesting insights. For example, the difference in ratios of group-level relevance assessed during comparability evaluation is generally small for the synthetic datasets (Sec. 6.1.4), as is the difference in group-level exposure fairness (Table 4b). Thus, the difference between fairness metrics using true and predicted relevance may be closely tied to the comparability criterion of predicted scores.

#### 6.3 Mitigating Measurement Issues

In this section, we highlight examples of conditions and interventions that might mitigate relevance measurement issues.

6.3.1 *Mechanisms: Fairness Interventions.* We study a post-hoc fairness processing setup with two algorithms proposed by Geyik et al. [31]. Both fair re-ranking algorithms (DetCons, DetConstSort) rearrange the top-k items in a ranked list such that the distribution of groups matches a user-specified distribution in a greedy manner (see Geyik et al. [31] for details). With a group-fairness lens, we set the desired proportion of a group among the top-k items to be the average relevance of items in that group divided by the sum of average relevance of items for all groups.

We study the impact of these interventions on estimated grouplevel exposure fairness. Do the interventions *mitigate* the measurement issues in fairness? From Table 4c, the fairness interventions are successful in reducing system unfairness (ratios are closer to 1 than pre-intervention scores in Table 4b). In some of the six cases (two algorithms, three datasets), the re-ranking algorithm modifies the system fairness to be indistinguishable from the pre-intervention value with true relevance scores (e.g., the fairtrec+DetCons experiment, and the synth-pareto+DetConstSort experiment). We observe that results are sensitive to the top "k" chosen for re-ranking. Thus, fairness interventions may be a mechanism to reduce the error in fairness assessment, but impact may vary based on the algorithm.

6.3.2 Data: Imbalanced Groups. During synthetic data generation (Sec. 5.1), we ensured equal group sample sizes. While class imbalance is not explicitly a component of the measurement model of relevance, it is often an artefact of ranking system design, and therefore could affect system fairness evaluation. Here, we consider various levels of data imbalance in the datasets, and study the resultant effects on exposure fairness—i.e., does imbalanced training data mitigate or hide the effects of relevance desiderata violations? We simulate imbalance ranging from 50-50% (i.e., no imbalance) to 90-10% (i.e., high imblance) by subsampling the synthetic datasets to 25,000 points.

For both datasets, we observe that the degree of imbalance impacts the difference between exposure fairness metrics as computed with true and predicted relevance: the difference ranges from about -0.10 to 0.12, and about -0.03 to 0.00 for synth-pareto and synth-normal respectively. A negative value here implies that the exposure fairness assessed using the predicted relevance score has

<sup>&</sup>lt;sup>5</sup>We use pre-trained ranking models for embeddings; modifying this led to similar results.

	Individual Fairness		Exposure Fairness		Exposure fairness + intervention	
Dataset	True	Predicted	True	Predicted	DetConstSort	DetCons
fairtrec	19.703 ± 0.01	$19.776 \pm 0.01^{*}$	$0.942 \pm 1.03$	$0.953 \pm 1.03^{*}$	$1.023 \pm 1.11^*$	$1.020 \pm 0.93$
synth-normal	$19.971 \pm 0.03$	$19.966 \pm 0.03^{*}$	$3.420 \pm 0.31$	$3.407 \pm 0.31^{*}$	$3.131 {\pm} 0.34^*$	$1.420 \pm 0.18$ *
synth-pareto	$19.523\pm0.03$	$19.949 \pm 0.03^{\ast}$	$2.180 \pm 1.01$	$2.210\!\pm\! 1.01^*$	$2.154 {\pm} 0.91$	$1.901 \pm 0.69$
(a)		(b)		(c)		

Table 4: (a) Individual fairness, (b) Group Exposure fairness measured on the test set, averaged across 10 runs, using true vs predicted relevance scores for the same set of rankings, (c) Impact of group-level fairness interventions. We observe that when inferred relevance scores are used, fairness metric values are significantly different (paired Wilcoxon test; p < 0.05; not significant with unpaired tests), though the size of the difference is often small. The deviation reported for individual fairness measurements corresponds to variance across items. (c) displays fairness with predicted relevance after unfairness mitigation algorithms intervene.

a lower value than that using true relevance. When fairness measurements here are greater than 1 (true for most of imbalance ratios here), this implies that the system appears more fair when predicted relevance is used. Thus, at specific imbalance ratios, the system may appear *more fair* than it is. In summary, group sample size imbalance may impact the inferred relevance scores and thus system fairness.

# 7 FUTURE PERSPECTIVES FOR IR

A turn towards examining fairness [14] and equity [10] in ranking has inspired new research agendas in IR. Meanwhile attending to the two-sided nature of ranking platforms has broadened the desiderata of fair ranking systems [26, 51, 52]. Past works propose interventions to fairly allocate opportunities and avoid exacerbating social biases based on some latent, complex notion of worthiness [10, 30, 45, 55, 60, 66]. Ultimately these interventions rely on *relevance* as a convenience proxy to allocate exposure. Our work points to a pressing need for information retrieval to reengage with research on relevance in the context of fair ranking.

Critically assessing limitations of relevance as a proxy for worthiness. Relevance can be justified as a proxy for worthiness; however, this justification must be done by meaningfully establishing the validity and reliability of relevance in a given setting. For the goal of fair ranking, this requires showing that the properties of the relevance scores and resulting rankings are aligned with the intended fairness goals. In our paper, we demonstrate an example of validating several such properties. Concrete use cases (e.g., hiring) and intended normative goals should guide the acceptable limitations of relevance.

The properties proposed in this paper were conceptualized based on extensive discussions between the authors of this paper using interdisciplinary lenses of machine learning, sociotechnical systems, measurement theory, and IR. Based on the application at hand, some of these properties may be more important than others. We also emphasize that additional properties may still need to be identified and tested for in various use-cases. Still, we believe that the proposed framework of elucidation followed by examination of properties offers a step towards assessing if a specific proxy (such as relevance) can be reliably used as a target for fair ranking. In concurrent work with an alternate view, Schumacher et al. [58] derive desired properties for group fairness metrics and highlight similar nuances of metrics that rely on relevance, further validating our findings. In contrast, we derive desiderata for *relevance* to guide fair exposure allocation.

**Defining 'worthiness' for different ranking application domains.** Different theoretical conceptualizations of worthiness will correspond to different normative goals and theories of justice. For instance, one could distinguish between the fairness of process (where the worthiness score might encode the merit of the ranked subjects) and the fairness of consequence (where the worthiness score might correspond to the utility of a subject being top-ranked). The design goals of a system will also reflect their application domains: a marketplace of online sellers might think differently about worthiness than a hiring platform. Precisely connecting different conceptualizations to how they are operationalized in fair ranking systems will aid researchers, developers, and auditors to enhance system equity.

**New methods for obtaining worthiness scores.** Beyond the need for new definitions, this paper points to the importance of considering how worthiness is operationalized in the context of fair ranking as well as revealing the need for new operationalizations. Potential approaches might vary from finding better proxies, developing new methods for direct worthiness crowdsourcing, finding new ways of accounting for annotator biases, calibrating predictions of relevance from browsing models, or proposing new approaches that go beyond existing IR relevance methodologies.

# 8 CONCLUSIONS

Ranking systems mediate opportunity in a variety of high impact settings [36]. In this paper, we leveraged an interdisciplinary perspective to show how fair ranking systems may fail to engage with the intended goals of fair ranking, relying instead on a tenuous assumption of relevance as worthiness. Drawing on domain knowledge about relevance-based fair ranking, we derived a set of criteria that relevance must satisfy in order to be a valid and reliable target to guide fair exposure allocation: credibility, consistency, stability, comparability, and availability. Using click-based relevance as a case study, we tested if these criteria can be empirically met on three datasets. We observe that a subset of these criteria may be unfulfilled in practice. Beyond click-based relevance, similar criteria will have to be tested for other forms of relevance estimation (e.g., crowdsourced relevance judgments). The contributions of our work are in establishing a measurement theory-based framework for thinking about the role of relevance in fair exposure allocation settings. More broadly, our work highlights the need for novel approaches to generate and collect relevance scores in a valid, reliable manner in fair ranking settings. Together, this reveals several open problems in relevance measurement at scale for fair exposure allocation.

## ACKNOWLEDGMENTS

This paper is part of the FINDHR (Fairness and Intersectional Non-Discrimination in Human Recommendation) project that received funding from the European Union's Horizon Europe research and innovation program under grant agreement No 101070212. SIGIR '23, July 23-27, 2023, Taipei, Taiwan

Aparna Balagopalan, Abigail Z. Jacobs, and Asia J. Biega

## REFERENCES

- Aman Agarwal, Xuanhui Wang, Cheng Li, Michael Bendersky, and Marc Najork. 2019. Addressing trust bias for unbiased learning-to-rank. In WWW. 4–14.
- [2] Qingyao Ai, Keping Bi, Jiafeng Guo, and W Bruce Croft. 2018. Learning a deep listwise context model for ranking refinement. In ACM SIGIR. 135–144.
- [3] Qingyao Ai, Keping Bi, Cheng Luo, Jiafeng Guo, and W Bruce Croft. 2018. Unbiased learning to rank with unbiased propensity estimation. In SIGIR. 385–394.
- [4] Abolfazl Asudeh, HV Jagadish, Julia Stoyanovich, and Gautam Das. 2019. Designing fair ranking schemes. In SIGMOD. 1259–1276.
- [5] Peter Bailey, Nick Craswell, Ian Soboroff, Paul Thomas, Arjen P de Vries, and Emine Yilmaz. 2008. Relevance assessment: are judges exchangeable and does it matter. In SIGIR. 667–674.
- [6] David La Barbera, Kevin Roitero, Gianluca Demartini, Stefano Mizzaro, and Damiano Spina. 2020. Crowdsourcing truthfulness: The impact of judgment scale and assessor bias. In European Conference on Information Retrieval. Springer, 207–214.
- [7] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. Fairness and Machine Learning. fairmlbook. org.
   [8] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt,
- [8] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, Ed H Chi, et al. 2019. Fairness in recommendation ranking through pairwise comparisons. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 2212–2220.
- [9] Asia J Biega, Fernando Diaz, Michael D Ekstrand, Sergey Feldman, and Sebastian Kohlmeier. 2021. Overview of the TREC 2020 Fair Ranking Track. arXiv preprint arXiv:2108.05135 (2021).
- [10] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. In SIGIR. 405–414.
- [11] Eszter Bokányi and Anikó Hannák. 2020. Understanding inequalities in ride-hailing services through simulations. *Scientific reports* 10, 1 (2020), 1–11.
- [12] Pia Borlund. 2003. The concept of relevance in IR. Journal of the American Society for information Science and Technology 54, 10 (2003), 913–925.
- [13] Amanda Bower, Hamid Eftekhari, Mikhail Yurochkin, and Yuekai Sun. 2020. Individually Fair Rankings. In ICLR.
- [14] Carlos Castillo. 2019. Fairness and transparency in ranking. In ACM SIGIR Forum, Vol. 52. ACM New York, NY, USA, 64–71.
- [15] Mingming Cheng and Carmel Foley. 2018. The sharing economy and digital discrimination: The case of Airbnb. Int. J. Hosp. Mgmt. 70 (2018), 95–98.
- [16] A Chuklin, I Markov, and M de Rijke. 2015. Click models for web search. Synthesis lectures on information concepts, retrieval, and services 7, 3 (2015), 1–115.
- [17] Leonid Churilov and Andy Flitman. 2006. Towards fair ranking of Olympics achievements: The case of Sydney 2000. *Comput. & Op. Res.* 33, 7 (2006), 2057–82.
  [18] Brian W Collins. 2007. Tackling unconscious bias in hiring practices: The plight
- [18] Brian W Comms. 2007. Tacking unconscious bias in niring practices: The pright of the Rooney rule. NYUL Rev. 82 (2007), 870.
- [19] A Feder Cooper, Solon Barocas, Christopher De Sa, and Siddhartha Sen. 2023. Variance, Self-Consistency, and Arbitrariness in Fair Classification. arXiv:2301.11562 (2023).
- [20] Michael D Cooper. 1973. A simulation model of an information retrieval system. Information Storage and Retrieval 9, 1 (1973), 13–32.
- [21] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Jimmy Lin. 2021. Ms marco: Benchmarking ranking models in the large-data regime. In Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. 1566–1576.
- [22] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In WSDM. 87–94.
- [23] Carlos A Cuadra and Robert V Katter. 1967. Opening the black box of 'relevance'. Journal of Documentation 23, 4 (1967), 291–303.
- [24] Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette.
   2020. Evaluating stochastic rankings with expected exposure. In CIKM. 275–284.
   [25] Michael D Ekstrand, Allison Chaney, Pablo Castells, Robin Burke, David Rohde,
- [25] Michael D Ekstrand, Allison Chaney, Pablo Castells, Robin Burke, David Rohde, and Manel Slokom. 2021. SimuRec: Workshop on Synthetic Data and Simulation Methods for Recommender Systems Research. In *RecSys.* 803–805.
- [26] Michael D Ekstrand, Anubrata Das, Robin Burke, Fernando Diaz, et al. 2022. Fairness in information access systems. Foundations and Trends in Information Retrieval 16, 1-2 (2022), 1–177.
- [27] Michael D. Ekstrand, Graham McDonald, Amifa Raj, and Isaac Johnson. 2022. Overview of the TREC 2021 Fair Ranking Track. In *The Thirtieth Text REtrieval* Conference (TREC 2021) Proceedings.
- [28] Francesco Fabbri, Francesco Bonchi, Ludovico Boratto, and Carlos Castillo. 2020. The effect of homophily on disparate visibility of minorities in people recommender systems. In Proceedings of the International AAAI Conference on Web and Social Media, Vol. 14. 165–175.
- [29] Anna Farmaki and Stella Kladou. 2020. Why do Airbnb hosts discriminate? Examining the sources and manifestations of discrimination in host practice. *Journal of Hospitality and Tourism Management* 42 (2020), 181–189.
- [30] Andres Ferraro, Xavier Serra, and Christine Bauer. 2021. What is fair? Exploring the artists' perspective on the fairness of music streaming platforms. In *IFIP* Conference on Human-Computer Interaction. Springer, 562–584.
- [31] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairnessaware 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. 2221–2231.
   [32] Avijit Ghosh, Ritam Dutt, and Christo Wilson. 2021. When fair ranking meets
- [32] Avijit Ghosh, Ritam Dutt, and Christo Wilson. 2021. When fair ranking meets uncertain inference. In SIGIR. 1033–1043.

- [33] Ziniu Hu, Yang Wang, Qu Peng, and Hang Li. 2019. Unbiased LambdaMART: An unbiased pairwise learning-to-rank algorithm. In WWW. 2830–2836.
- [34] Abigail Z Jacobs and Hanna Wallach. 2021. Measurement and fairness. In Proc. 2021 ACM Conference on Fairness, Accountability, and Transparency. 375–385.
- [35] Kalervo Järvelin and Jaana Kekäläinen. 2017. IR evaluation methods for retrieving highly relevant documents. In ACM SIGIR Forum, Vol. 51. 243–250.
   [36] T Joachims, L Granka, B Pan, H Hembrooke, and G Gay. 2017. Accurately inter-
- [30] J Joannis, D Gaina, J Fan, Thenholooke, and G Gay. 2017. Actuality metpreting clickthrough data as implicit feedback. In ACM SIGIR Forum, Vol. 51. 4–11.
   [37] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, Filip Radlinski,
- and Geri Gay. 2007. Evaluating the accuracy of implicit feedback from clicks and query reformulations in web search. *TOIS* 25, 2 (2007), 7–es.
   [38] Thorsten Joachims, Ben London, Yi Su, Adith Swaminathan, and Lequn Wang.
- [36] Thorsten Joachims, Jeri London, Trot, Auth Swaminatrian, and Lequi Wang.
   2021. Recommendations as treatments. AIMagazine 42, 3 (2021), 19–30.
   [39] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased
- [40] Yuri Kagolovsky and Jochen R Möhr. 2001. A new approach to the concept of
- "relevance" in information retrieval (IR). In MEDINFO 2001. IOS Press, 348–352.
   [41] Jaana Kekäläinen. 2005. Binary and graded relevance in IR evalua-
- [41] Jaana Kekäläinen. 2005. Binary and graded relevance in IR evaluations—comparison of the effects on ranking of IR systems. Information processing & management 41, 5 (2005), 1019–1033.
- [42] Ömer Kırnap, Fernando Diaz, Asia Biega, Michael Ekstrand, Ben Carterette, and Emine Yilmaz. 2021. Estimation of Fair Ranking Metrics with Incomplete Judgments. In Proceedings of the Web Conference 2021. 1065–1075.
- [43] Till Kletti, Jean-Michel Renders, and Patrick Loiseau. 2022. Pareto-Optimal Fairness-Utility Amortizations in Rankings with a DBN Exposure Model. In SIGIR 2022-45th International ACM SIGIR Conference on Research and Development in Information Retrieval.
- [44] Frank J Massey Jr. 1951. The Kolmogorov-Smirnov test for goodness of fit. Journal of the American statistical Association 46, 253 (1951), 68–78.
- [45] 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 Proceedings of the 27th acm international conference on information and knowledge management. 2243–2251.
- [46] Massimo Melucci. 2016. Impact of query sample selection bias on information retrieval system ranking. In DSAA. IEEE, 341–350.
- [47] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling fairness and bias in dynamic learning-to-rank. In SIGIR. 429–438.
- [48] Mohammadmehdi Naghiaei, Hossein A. Rahmani, and Yashar Deldjoo. 2022. CP-Fair: Personalized Consumer and Producer Fairness Re-Ranking for Recommender Systems (SIGIR '22). 770–779.
- [49] Harrie Oosterhuis. 2022. Reaching the End of Unbiasedness: Uncovering Implicit Limitations of Click-Based Learning to Rank. In Proceedings of the 2022 ACM SIGIR International Conference on Theory of Information Retrieval. 264–274.
- [50] Harrie Oosterhuis and Maarten de Rijke. 2018. Differentiable unbiased online learning to rank. In CIKM. 1293–1302.
- [51] Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. 2022. Fair ranking: a critical review, challenges, and future directions. In 2022 ACM Conference on Fairness, Accountability, and Transparency. 1929–1942.
- [52] Milo Phillips-Brown. 2023. Algorithmic neutrality. arXiv preprint arXiv:2303.05103 (2023).
- [53] Amifa Raj and Michael D Ekstrand. 2022. Measuring Fairness in Ranked Results: An Analytical and Empirical Comparison. In Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. 726–736.
- [54] Stephen E Robertson. 1977. The probability ranking principle in IR. J Doc (1977).
   [55] Piotr Sapiezynski, Wesley Zeng, Ronald E Robertson, Alan Mislove, and Christo Wilson. 2019. Quantifying the impact of user attentionon fair group representation in ranked lists. In Companion proceedings of the 2019 world wide web conference. 553–562.
- [56] Tefko Saracevic. 2016. The Notion of Relevance in Information Science: Everybody knows what relevance is. But, what is it really? *Synthesis lectures on information* concepts, retrieval, and services 8, 3 (2016), i–109.
- [57] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. 2016. Recommendations as treatments: Debiasing learning and evaluation. In *ICML*. 1670–1679.
- [58] Tobias Schumacher, Marlene Lutz, Sandipan Sikdar, and Markus Strohmaier. 2022. Properties of Group Fairness Metrics for Rankings. arXiv preprint arXiv:2212.14351 (2022).
- [59] Anne Schuth, Harrie Oosterhuis, Shimon Whiteson, and Maarten de Rijke. 2016. Multileave gradient descent for fast online learning to rank. In WSDM. 457–466.
- [60] Dougal Shakespeare, Lorenzo Porcaro, Emilia Gómez, and Carlos Castillo. 2020. Exploring artist gender bias in music recommendation. arXiv preprint arXiv:2009.01715 (2020).
- [61] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In KDD. 2219–2228.
- [62] Ashudeep Singh and Thorsten Joachims. 2019. Policy learning for fairness in ranking. Advances in neural information processing systems 32 (2019).
- [63] Ryosuke Sonoda. 2021. A Pre-processing Method for Fairness in Ranking. arXiv preprint arXiv:2110.15503 (2021).
- [64] Ali Vardasbi, Maarten de Rijke, and Ilya Markov. 2020. Cascade model-based propensity estimation for counterfactual learning to rank. In SIGIR. 2089–2092.

The Role of Relevance in Fair Ranking

SIGIR '23, July 23-27, 2023, Taipei, Taiwan

- [65] Ellen M Voorhees, Donna K Harman, et al. 2005. TREC: Experiment and evaluation in information retrieval, vol. 1.
- [66] Johannes Wachs, Anikó Hannák, András Vörös, and Bálint Daróczy. 2017. Why do men get more attention? Exploring factors behind success in an online design community. In Eleventh International AAAI Conference on Web and Social Media.
- [67] Xuanhui Wang, Michael Bendersky, Donald Metzler, and Marc Najork. 2016. Learning to rank with selection bias in personal search. In SIGIR. 115-124.
- [68] X Wang, N Golbandi, M Bendersky, D Metzler, and M Najork. 2018. Position bias estimation for unbiased learning to rank in personal search. In *WSDM*. [69] Lu Xian and Abigail Z Jacobs. 2022. Harms as process: opportunity harms in the
- US mortgage market. *preprint* (2022). [70] Himank Yadav, Zhengxiao Du, and Thorsten Joachims. 2020. Fair learning-to-rank
- from implicit feedback. In SIGIR.
- [71] Ke Yang, Joshua R Loftus, and Julia Stoyanovich. 2020. Causal intersectionality for fair ranking. arXiv preprint arXiv:2006.08688 (2020).

- [72] Tao Yang and Qingyao Ai. 2021. Maximizing Marginal Fairness for Dynamic Learning to Rank. In WebConf. 137–145.
- Dawei Yin, Yuening Hu, Jiliang Tang, Tim Daly, Mianwei Zhou, Hua Ouyang, [73] Jianhui Chen, Changsung Kang, Hongbo Deng, Chikashi Nobata, et al. 2016. Ranking relevance in yahoo search. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 323-332.
- [74] Yisong Yue and Thorsten Joachims. 2009. Interactively optimizing information retrieval systems as a dueling bandits problem. In ICML. 1201-1208.
- [75] Meike Zehlike, Ke Yang, and Julia Stoyanovich. 2021. Fairness in Ranking: A Survey. arXiv preprint arXiv:2103.14000 (2021).
- [76] Meike Zehlike, Ke Yang, and Julia Stoyanovich. 2022. Fairness in ranking, part ii: Learning-to-rank and recommender systems. *Comput. Surveys* 55, 6 (2022), 1–41.
   [77] Elana Zeide. 2022. The Silicon Ceiling: How Algorithmic Assessments Construct
- an Invisible Barrier to Opportunity. UMKC Law Rev. (2022).