

Policy-Gradient Training of Fair and Unbiased Ranking Functions

Himank Yadav*
himankyadav1@gmail.com
Cornell University
Ithaca, NY, USA

Zhengxiao Du*
zx-du20@mails.tsinghua.edu.cn
Tsinghua University
Beijing, China

Thorsten Joachims
tj@cornell.edu
Cornell University
Ithaca, NY, USA

ABSTRACT

While implicit feedback (e.g., clicks, dwell times, etc.) is an abundant and attractive source of data for learning to rank, it can produce unfair ranking policies for both exogenous and endogenous reasons. Exogenous reasons typically manifest themselves as biases in the training data, which then get reflected in the learned ranking policy and often lead to rich-get-richer dynamics. Moreover, even after the correction of such biases, reasons endogenous to the design of the learning algorithm can still lead to ranking policies that do not allocate exposure among items in a fair way. To address both exogenous and endogenous sources of unfairness, we present the first learning-to-rank approach that addresses both presentation bias and merit-based fairness of exposure simultaneously. Specifically, we define a class of amortized fairness-of-exposure constraints that can be chosen based on the needs of an application, and we show how these fairness criteria can be enforced despite the selection biases in implicit feedback data. The key result is an efficient and flexible policy-gradient algorithm, called FULTR, which is the first to enable the use of counterfactual estimators for both utility estimation and fairness constraints. Beyond the theoretical justification of the framework, we show empirically that the proposed algorithm can learn accurate and fair ranking policies from biased and noisy feedback.¹

CCS CONCEPTS

• Information systems → Learning to rank; • Computing methodologies → Learning from implicit feedback.

KEYWORDS

fairness, bias, learning-to-rank, counterfactual learning

ACM Reference Format:

Himank Yadav, Zhengxiao Du, and Thorsten Joachims. 2021. Policy-Gradient Training of Fair and Unbiased Ranking Functions. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3404835.3462953>

^{*}Equal contribution

¹The code and datasets are available at <https://github.com/him229/fultr>.

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 the author(s) 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 '21, July 11–15, 2021, Virtual Event, Canada.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8037-9/21/07...\$15.00
<https://doi.org/10.1145/3404835.3462953>

1 INTRODUCTION

Implicit feedback from user behavior (e.g., clicks, dwell times, purchases, scroll patterns) [27, 38] is an attractive source of training data for learning-to-rank (LTR), since it gives all users a voice in what the system learns. But does the participatory nature of implicit feedback automatically ensure that the learned ranking policies are fair and lead to desirable effects on the overall health of the ranking system? We argue that both endogenous and exogenous factors can lead to undesirable outcomes as follows.

Endogenous factors are design choices built into the allocation policies of the ranking system. Specifically, deciding which ranking to present to a user implies an explicit choice of how much exposure each item receives – where higher-ranked items receive more exposure and thus more opportunity [16, 18]. Conventional ranking algorithms [40] make this exposure allocation in a haphazard way, which can lead to unfairness and undesirable dynamics [7, 8, 13, 20, 43, 44, 47, 48]. To address this problem, many have noted that there should be an explicit link between the amount of exposure an item receives, and the amount of merit (e.g. relevance) the item has. Specifically, merit-based exposure allocation [7, 8, 43, 44] can ensure that items of similar merit (e.g. relevance) also have similar outcomes (e.g. exposure). The particular choice of allocation policy can provide various fairness guarantees and provide a mechanism for steering the system dynamics (e.g. avoiding superstar economics [30]).

While merit-based exposure allocation addresses an endogenous factor of unfairness, it is important to recognize that an LTR algorithm can still be unfair due to *exogenous* biases in the training data. In particular, position bias [26] makes the use of implicit feedback for LTR challenging. Especially for positive-only feedback like clicks or purchases, highly-ranked items receive more positive feedback due to the increased attention they receive, which further skews the ranking system and affects future rankings. The result can be a dynamic amplification of position bias, leading to phenomena like rich-get-richer [1, 26, 28, 42]. This means the even in the presence of merit-based fairness constraints like [43, 44], naive and biased estimates of merit from implicit feedback signals can still lead to inequities. We, therefore, need to design fair LTR algorithms that ensure unbiased learning despite the biases in the data.

In this paper, we simultaneously address endogenous and exogenous causes of unfairness and present the first LTR algorithm – called FULTR (Fair Unbiased Learning-to-Rank) – for unbiased learning from logged implicit feedback that can enforce merit-based exposure-allocation constraints. We provide a theoretical analysis to show that a naive combination of conventional merit-based fairness constraints and counterfactual relevance estimators cannot guarantee unbiased and fair ranking policies. To overcome this

problem, we define a new type of allocation constraint for groups of items and derive counterfactual estimators [28, 45] that provably correct the position bias. To make the proposed framework operational and practical, we show how to search the space of fairness-constrained ranking policies via a novel policy-gradient approach that is unbiased. Note that FULTR is trained end-to-end, and thus can consider fairness during learning, eliminating failure cases of two-step approaches that rely on post-processing an unfair ranking to make it fair [8, 32, 43].

Beyond the theoretical justification, we also provide an extensive empirical evaluation on real-world datasets. We find that FULTR can effectively optimize utility and fairness over a range of settings even when trained with biased and noisy feedback.

1.1 A Motivating Example

Consider an online movie streaming service which presents a ranked list of movies for any incoming query. Further, consider the case where the movies were produced by only two studios. Studio G_1 produced 30 movies that all have relevance 0.7 (i.e. there is a 70% probability that the user will want to watch each movie). Studio G_2 also produced 30 movies, but they all have slightly lower relevance of 0.69. Ranking purely based on relevance to the user would place movies from G_1 at spots 1 – 30 and movies from G_2 at spots 31 – 60, leading to much higher exposure and revenue for movies from G_1 than movies from G_2 . This can be considered unfair since movies of almost equal merit (i.e. relevance) receive disproportionately worse outcomes (i.e. streaming revenue).

A first thought for addressing this shortcoming may be to enforce various forms of demographic parity between groups of items in the ranking [13, 20, 47, 48]. However, allocating exposure without consideration of merit has other undesirable effects. Assume that there is a third, incompetent, studio G_3 that produced another 30 movies of low relevance 0.1. Clearly, these movies do not deserve as much exposure as those from G_1 and G_2 . We therefore follow an approach that explicitly links exposure to the merit (i.e. relevance) of a group. Specifically, merit-based exposure allocation [7, 8, 43, 44] can ensure that groups of similar merit (e.g. relevance) have similar outcomes (e.g. exposure).

Enforcing merit-based fairness of exposure is a factor endogenous to the ranking system, meaning that it is fully under the control of the system designer. However, it relies on accurate estimates of merit, which can be affected by exogenous factors. Consider, for example, a situation where movies from G_1 have consistently been ranked higher than those from G_2 in the past. If we now naively count the number of clicks as an estimate of merit, this estimate will be biased by how humans browse ranked lists [26]. In particular, the movies from G_2 are likely to get substantially fewer clicks than those from G_1 not because they are substantially worse, but because they were less likely to be discovered by the users. In this way, even an endogenously fair ranking system can become unfair as its past actions create biased estimates that amplify past inequities and that can more generally lead to undesirable system dynamics [28, 30].

Our approach is the first that can simultaneously control both sources of unfairness when learning a ranking policy, as we will detail in the following.

2 RELATED WORK

There have been numerous approaches to defining fairness in different areas of machine learning, including online learning [21, 35], classification [2, 33], regression [6, 39]. We focus on fairness in the relatively under-explored domain of LTR, which has only recently caught attention despite its substantial implications in a broad range of real-world applications. To structure the discussion, we follow the distinction of endogenous and exogenous sources of unfairness introduced above.

Concerning endogenous fairness, several methods have followed the concept of demographic parity which enforces a proportional allocation of exposure between groups without considering merit. This is achieved by either reducing the difference in occurrences of different groups on a subset of the rankings [47] or by placing a limit on the number of items from each group in the top-k positions [13, 20, 48]. In contrast, merit-based fairness of exposure [8, 43] allocates exposure to groups based on their merit rather than purely based on group size. However, these methods assume knowledge of the true relevance labels for all items and do not involve learning, while in practice, true relevance labels are not available. and one needs to use other learning methods to impute relevance labels at prediction time [32]. This leads to a two-step process, where the first regression step is unaware of fairness considerations and fairness is only introduced during post-processing. Instead, our method is trained end-to-end in one step, such that it can, for example, discount features that lead to unfair relevance estimates [44].

More recently, Zehlike and Castillo [49] proposed an LTR method that incorporates an exposure-based fairness regularizer into the ranking loss, but it only considers the exposure at the top-1 position. The limitation was removed by Singh and Joachims [44], who propose an end-to-end LTR method that optimizes both utility and fairness constraints for the full ranking. However, both methods assume that expert-labeled relevances for all items is available during training. This is a weaker assumption than the one made in [8, 43], but it still does not apply to real-world settings where implicit feedback is often used for training.

Our work focuses on exposure-based fairness, but another approach to defining fairness in rankings is through pairwise comparisons [7, 34], without considering exposure as the key criterion. Such comparisons between items count all swaps in the ranking equally and do not reflect that swapping items at the top has a stronger effect on exposure than doing so at lower positions.

Switching to exogenous sources of unfairness, most work on traditional LTR algorithms [10–12, 25] assumes that unbiased relevance judgments by experts are available. However, the field of information retrieval has long been conscious of the biases inherent to implicit feedback and its effect on the ability to rank well [16, 18, 22, 25, 26]. In particular, various studies have demonstrated the presence of position bias, where the quantity and quality of the feedback depend on the rank at which an item is presented.

One approach to modeling and removing position bias is generative click modeling [9, 14, 15]. Click models typically treat relevance as latent variables, and perform inference by maximizing log-likelihood of clicks. These inferred relevances can then be used as a substitute for expert labels in LTR. Unfortunately, most click models suffer from the limitation of requiring large amounts of

repeat impressions for individual query-item pairs, which makes them inapplicable to tail queries.

Recent and more direct approaches to dealing with position bias are counterfactual learning methods [28, 45]. These methods use techniques from causal inference like inverse propensity score (IPS) weighting [41] and do not require repeated queries or latent-variable inference. Instead, they directly optimize over a debiased utility objective that incorporates click data in a principled fashion. Additional algorithms [5, 23] that jointly estimate the propensities and optimize the performance have been proposed. Unlike our proposed FULTR framework, existing counterfactual learning methods do not control for endogenous unfairness. Moreover, we find that naively combining existing counterfactual learning techniques with endogenous fairness constraints fail to control for unfairness.

The only existing method that aims to address both endogenous and exogenous sources of unfairness in rankings is [32]. However, it is the two-step method that learns relevance estimation with unbiased regression objectives first. The regression learning is unaware of fairness constraints and may yield unfairness in relevance estimation. We show empirically that our end-to-end FULTR can achieve better fairness-utility tradeoff than the two-step method. Additionally, it requires interactive experimental control in the form of an online algorithm, whereas FULTR can reuse logged implicit feedback data from past interactions for learning. To the best of our knowledge, no other existing method addresses fairness in rankings while directly learning from logged implicit feedback.

3 LEARNING FAIR RANKING POLICIES FROM IMPLICIT FEEDBACK

In the following, we derive the first learning algorithm for enforcing merit-based fairness of exposure [43] with logged implicit feedback. A key component is a new amortized fairness criterion, and its corresponding disparity measure, for which we show that unbiased counterfactual estimators exist. We analyze the connection between this disparity measure and those in [8, 43]. Furthermore, we quantify the effect of click noise on the disparity measure and propose a noise-corrected estimator of disparity.

The derivation of the algorithm is based on the standard definition of Learning to rank (LTR), which is the problem of learning a ranking policy π for queries q sampled from a distribution $P(q)$. Each query q has a set of candidate items d^q that need to be ranked. Items can represent a wide variety of things depending on the application — web-pages in a search engine, jobs on a job board, or movies in a streaming service. Each item is associated with a feature vector $x_{q,d}$ that describes the match between item d and query q .

In the so-called full-information setting, it is assumed that for a training query q the relevances $\text{rel}_{q,d}$ of all items $d \in d^q$ are known. Instead, implicit feedback only reveals the relevances of part of the candidates. We denote this partial feedback as c_q and call this the partial-information setting [28].

Unlike most existing works, we consider stochastic ranking policies, where $\pi(\sigma|q)$ is a distribution over the rankings σ of the candidates. As will become clear later, stochastic ranking policies have the advantage of providing more fine-grained control of exposure and enabling gradient-based optimization. Note that deterministic

ranking policies are a special case of stochastic ranking policies, where all probability mass lies on a single ranking.

Like in conventional LTR algorithms [10–12, 25], our objective is to learn a policy π from policy space Π with high expected utility

$$U(\pi) = \mathbb{E}_{q \sim P(q)} [U(\pi|q)] = \mathbb{E}_{q \sim P(q)} \mathbb{E}_{\sigma \sim \pi(\sigma|q)} [\Delta(\sigma, \text{rel}_q)],$$

where Δ can be any ranking metric (e.g. DCG [24]).

Utility optimization for the users, however, does not by itself ensure fairness to the items [43]. We, therefore, include an additional constraint that addresses endogenous unfairness by enforcing a merit-based allocation of exposure. For simplicity, consider the case of two groups G_i and G_j . A group is simply a collection of items by any criterion, such as gender, price, brand, etc. To ensure that exposure is allocated fairly between G_i and G_j , we measure unfairness via a disparity measure $D_{ij}(\pi)$ that will be defined in Section 3.2. A perfectly fair ranking policy has zero disparity. More generally, we want to restrict $[D_{ij}(\pi)]^2$ to be at most some threshold δ to get our disparity to be as close to zero as possible. We define our objective by combining the fairness constraint with the conventional utility optimization.

$$\pi^* = \underset{\pi \in \Pi}{\text{argmax}} U(\pi) \text{ s.t. } [D_{ij}(\pi)]^2 \leq \delta. \quad (1)$$

Directly optimizing this objective is not possible, since neither the query distribution $P(q)$ nor the relevances rel_q are known. We therefore take the following Empirical Risk Minimization (ERM) approach, where both $U(\pi)$ and $D_{ij}(\pi)$ are replaced by estimates $\widehat{U}(\pi|Q)$ and $\widehat{D}_{ij}(\pi|Q)$ based on a partial-info training set Q .

$$\widehat{\pi} = \underset{\pi \in \Pi}{\text{argmax}} \widehat{U}(\pi|Q) \text{ s.t. } [\widehat{D}_{ij}(\pi|Q)]^2 \leq \delta. \quad (2)$$

This leads to three key challenges which we address as follows. First, in Section 3.1, we show how to design $\widehat{U}(\pi|Q)$ to get unbiased utility estimates even if we do not have access to true relevance labels rel_q but only have access to biased implicit feedback c_q . Second, in Section 3.2, we show how to define estimators of the fairness disparity $\widehat{D}_{ij}(\pi|Q)$ that ensure merit-based exposure allocation while also being unbiased even with only the implicit feedback c_q . Finally, in Section 4, we show how to efficiently solve the resulting training problem from Equation (2).

3.1 Unbiased Utility Estimator

We consider the class of additive ranking measures that can be expressed as

$$\Delta(\sigma, \text{rel}_q) = \sum_{d \in d^q} [f(\sigma(d)) \cdot \text{rel}_{q,d}],$$

where $\sigma(d)$ denotes the rank of item d in ranking σ and $f(\cdot)$ can be any weighting function. For the DCG metric, $f(\sigma(d)) = 1/\log_2(1 + \sigma(d))$, while for the Average Rank metric, $f(\sigma(d)) = -\sigma(d)$. For simplicity, we assume binary relevances $\text{rel}_{q,d} \in \{0, 1\}$.

In the partial-information setting, the relevances rel_q are not known and we can only observe c_q . What can we infer from c_q about rel_q ? We defer the discussion of noise to Section 3.4 and consider the standard model that users click if and only if the result is relevant (i.e. $\text{rel}_{q,d} = 1$) and examined by the user. We introduce the binary random variable $o_{q,d}$ indicating whether the item d is

examined by the user. Based on this, we can model the "propensity" $p(o_{q,d} = 1|\sigma_q)$, of observing $\text{rel}_{q,d}$ given that ranking σ_q was presented when the implicit feedback was logged. The propensity can be modelled in various ways [19, 28]. Most common is the position-based examination model, where the propensity depends on only the position $\sigma_q(d)$ of d in the ranking. This means $p(o_{q,d} = 1|\sigma_q) = v_{\sigma_q(d)}$, where v_k denotes the examination probability at position k . These position biases v_k can be estimated with swap experiments [28] or intervention harvesting [4]. With knowledge of the propensity, we can use Inverse Propensity Score (IPS) weighting to arrive at the following estimator for Δ [28]

$$\widehat{\Delta}(\sigma, c_q) = \sum_{d: c_{q,d}=1} \frac{f(\sigma(d))}{p(o_{q,d} = 1|\sigma_q)}. \quad (3)$$

The estimator is unbiased if all propensities are bounded away from zero [28]. Furthermore, we can define an estimator of the utility as

$$\widehat{U}(\pi|Q) = \frac{1}{|Q|} \sum_{q \in Q} \widehat{U}(\pi|q) = \frac{1}{|Q|} \sum_{q \in Q} \mathbb{E}_{\sigma \sim \pi(\sigma|q)} [\widehat{\Delta}(\sigma, c_q)]. \quad (4)$$

It is easy to verify that this estimator inherits unbiasedness from $\widehat{\Delta}$. It is also consistent under the same condition on the propensities, thus implying that $\widehat{U}(\pi|Q)$ will converge to $U(\pi)$ as $|Q|$ increases.

3.2 Unbiased Fairness Constraints

To remedy that the utility-maximizing ranking policy can be unfair even if the utility $U(\pi)$ is perfectly known for all $\pi \in \Pi$ [43], we enforce additional criteria of how exposure is allocated based on merit (i.e. relevance). In our training problem from Equation (2), this is implemented through the disparity measure $\widehat{D}_{ij}(\pi|Q)$ in the constraint, which we now define formally.

As a first step, we define the merit of an item d_i as its relevance, i.e. $M_q(d) = \text{rel}_{q,d} \in \{0, 1\}$. The exposure of an item d in ranking σ is defined as the probability of a user examining the item in the ranking. This is identical to the examination probability $p(o_{q,d} = 1|\sigma)$ defined in Section 3.1, but this model is now applied to all rankings instead of just the logged ranking σ_q . The exposure of d under a stochastic ranking policy π for a query q , denoted as $\text{Exp}_q(d|\pi)$, is the expected exposure over all the possible rankings. Formally,

$$\text{Exp}_q(d|\pi) = \mathbb{E}_{\sigma \sim \pi(\sigma|q)} [p(o_{q,d} = 1|\sigma)]$$

Furthermore, the exposure of group G_i is the aggregate of the exposure of the group members, $\text{Exp}_q(G_i|\pi) = \sum_{d \in G_i^q} \text{Exp}_q(d|\pi)$, where $G_i^q = G_i \cap d^q$. Similarly, we define the merit of group G_i for query q as $M_q(G_i) = \sum_{d \in G_i^q} M_q(d)$.

With these definitions in hand, we define our fairness disparity of policy π as $D_{ij}(\pi) = \mathbb{E}_{q \sim P(q)} [D_{ij}(\pi|q)]$, where $D_{ij}(\pi|q)$ measures the disparate exposure of G_i and G_j based on their merit for query q as

$$D_{ij}(\pi|q) = M_q(G_j) \text{Exp}_q(G_i|\pi) - M_q(G_i) \text{Exp}_q(G_j|\pi). \quad (5)$$

As discussed in more detail in the next section, this disparity measure formalizes the principle that proportionality between merit and exposure should hold for a fair ranking [8, 43]. While $\text{Exp}_q(G_i|\pi)$ is an expectation that can be computed, we do have to estimate

$D_{ij}(\pi)$ with respect to the unknown query distribution and the unknown merits. We therefore take the estimator of $D_{ij}(\pi)$ as

$$\widehat{D}_{ij}(\pi|Q) = \frac{1}{|Q|} \sum_{q \in Q} \widehat{D}_{ij}(\pi|q). \quad (6)$$

Furthermore, in the partial-information setting, we can get an unbiased estimate of $M_q(G_i)$ using the IPS estimator $\widehat{M}_q(G_i) = \sum_{d \in G_i^q} \frac{c_{q,d}}{p(o_{q,d}=1|\sigma_q)}$. This means that $M_q(G_i)$ in Equation (5) is replaced with $\widehat{M}_q(G_i)$ to arrive at the following empirical disparity measure

$$\widehat{D}_{ij}(\pi|q) = \widehat{M}_q(G_j) \text{Exp}_q(G_i|\pi) - \widehat{M}_q(G_i) \text{Exp}_q(G_j|\pi). \quad (7)$$

Note that $\widehat{D}_{ij}(\pi|q)$ is unbiased since $\widehat{M}_q(G_i)$ is unbiased and the exposure terms are non-random constants.

$$\begin{aligned} \mathbb{E}_{o_q}[\widehat{D}_{ij}(\pi|q)] &= \mathbb{E}_{o_q}[\widehat{M}_q(G_j)] \text{Exp}_q(G_i|\pi) - \mathbb{E}_{o_q}[\widehat{M}_q(G_i)] \text{Exp}_q(G_j|\pi) \\ &= M_q(G_j) \text{Exp}_q(G_i|\pi) - M_q(G_i) \text{Exp}_q(G_j|\pi) \end{aligned}$$

Since $\widehat{D}_{ij}(\pi|q)$ is unbiased for each query q , the aggregate $\widehat{D}_{ij}(\pi|Q)$ is also unbiased for $D_{ij}(\pi)$, i.e. $\mathbb{E}_q \mathbb{E}_{o_q}[\widehat{D}_{ij}(\pi|Q)] = \mathbb{E}_q [D_{ij}(\pi|q)] = D_{ij}(\pi)$. Furthermore, through the use of Hoeffding bounds, it is possible to show that $\widehat{D}_{ij}(\pi|Q)$ converges to the true disparity $D_{ij}(\pi)$ as $|Q|$ increases.

3.3 Comparison with other Disparity Measures

In this section, we further highlight the differences between our disparity measure from Equation (5) and various other disparity measures.

The Disparate Treatment constraint from [44]

$$\forall q : \frac{\text{Exp}_q(G_i|\pi)}{M_q(G_i)} = \frac{\text{Exp}_q(G_j|\pi)}{M_q(G_j)},$$

expresses that for a given query, the exposure of each group should be proportional to its merit. However, evaluating this constraint requires knowledge of the full relevances which is infeasible to decipher from real-world implicit feedback. Directly replacing $\text{rel}_{q,d}$ with $c_{q,d}$ when evaluating $M_q(G_j)$ leads to biased estimation due to selection bias [28]. To combat this, naively plugging in the IPS estimator $\widehat{M}_q(G_i)$ into the corresponding disparity measure

$$\widehat{D}'_{ij}(\pi) = \frac{\text{Exp}_q(G_i|\pi)}{\widehat{M}_q(G_i)} - \frac{\text{Exp}_q(G_j|\pi)}{\widehat{M}_q(G_j)} \quad (8)$$

cannot eliminate bias effectively, since generally $\mathbb{E}[1/\widehat{M}_q(G_i)] \neq 1/\mathbb{E}[\widehat{M}_q(G_i)]$. For an unbiased estimate of $1/M_q(G_i)$ from c_q , we need to know the joint distribution of $o_{q,d}$ for all the items in G_i^q , which is difficult to model. Instead, the proposed disparity measure in Equation (5) has an unbiased estimator (Equation (7)) to overcome the selection bias, which we find to provide substantial empirical benefits in Section 5.3.

Another proposed disparity measure using implicit feedback replaces query-specific exposure and merit with an amortized notion [8] over the query distribution. The disparity measure

$$D'_{ij}(\pi) = \frac{\mathbb{E}_q[\text{Exp}_q(G_i|\pi)]}{\mathbb{E}_q[M_q(G_i)]} - \frac{\mathbb{E}_q[\text{Exp}_q(G_j|\pi)]}{\mathbb{E}_q[M_q(G_j)]} \quad (9)$$

expresses that for all the queries, the amortized exposure of each group should be proportional to its amortized merit. We will show how $D'_{ij}(\pi)$ is related to our disparity measure $D_{ij}(\pi)$ from Equation (5) under two conditions. First, assume that the total exposure of both groups is a constant $\text{Exp}_q(G_i) + \text{Exp}_q(G_j) = C_E$. In practice, the top items receive most of the users' attention, so the total exposure for each query is relatively stable even if the size of the candidate set varies. Second, assume that the covariance of the exposure of group G_j and the total number of relevant items in both groups, $\text{cov}[\text{Exp}(G_j|\pi), M_q(G_i) + M_q(G_j)]$, is zero. One sufficient condition for this assumption to be approximately satisfied is when the numbers of relevant items do not vary much between queries. Under these two assumptions, the disparity measures $D_{ij}(\pi)$ and $D'_{ij}(\pi)$ are equivalent up to a constant coefficient. This can be shown by transforming the disparity measure $D'_{ij}(\pi)$ as follows

$$\frac{D_{ij}(\pi) + \text{cov}[\text{Exp}_q(G_j|\pi), M_q(G_i)] - \text{cov}[\text{Exp}_q(G_i|\pi), M_q(G_j)]}{\mathbb{E}_q[M_q(G_i)] \mathbb{E}_q[M_q(G_j)]}.$$

The difference between the covariance terms is zero under the two assumptions. Therefore, we have $D_{ij}(\pi) = \mathbb{E}_q[M_q(G_i)] \mathbb{E}_q[M_q(G_j)] D'_{ij}(\pi)$.

However, [8] requires the knowledge of true relevance labels and does not involve learning. Extending this approach to learn relevance labels can suffer from the failure case of a two-step approach that relies on post-processing an unfair relevance estimation to make it fair, since fairness constraints are not taking into account during the learning process. We can get a consistent estimator of $D'_{ij}(\pi)$ by using the IPS estimator of merits and replacing the expectations of exposure and merit with the average over the dataset

$$\widehat{D}'_{ij}(\pi|Q) = \frac{\sum_{q \in Q} \text{Exp}_q(G_i|\pi)}{\sum_{q \in Q} \widehat{M}_q(G_i)} - \frac{\sum_{q \in Q} \text{Exp}_q(G_j|\pi)}{\sum_{q \in Q} \widehat{M}_q(G_j)}.$$

As N increases, $\widehat{D}'_{ij}(\pi|Q)$ will converge to $D'_{ij}(\pi)$. However, the estimator is biased, which means $\mathbb{E}[\widehat{D}'_{ij}(\pi|Q)] \neq D'_{ij}(\pi)$. This bias might increase the error of the estimator. We therefore conclude that our new disparity $D_{ij}(\pi)$ with its unbiased estimator $\widehat{D}_{ij}(\pi)$ in Equations (6) and (7) is preferable.

3.4 Incorporating Click Noise

Up to now, we assumed that positive feedback $c_{q,d} = 1$ reveals the relevance label $\text{rel}_{q,d}$ in a noise-free way – precisely that $c_{q,d} = 1$ if and only if $\text{rel}_{q,d} = 1$ and $o_{q,d} = 1$. In practice, the user might make mistakes when examining the relevances of items, and we now define the following noise model. With $1 \geq \epsilon_+ > \epsilon_- \geq 0$, we have

$$\begin{aligned} P(c_{q,d} = 1 | \text{rel}_{q,d} = 1, o_{q,d} = 1) &= \epsilon_+ \\ P(c_{q,d} = 1 | \text{rel}_{q,d} = 0, o_{q,d} = 1) &= \epsilon_- \end{aligned}$$

If $\epsilon_+ < 1$, users might ignore relevant items even after examination. Otherwise if $\epsilon_- > 0$, users might give false positive feedback to irrelevant items after examination. Fortunately, the IPS estimator of utility is order-preserving for two policies π_1 and π_2 under click

noise [28], namely

$$\mathbb{E}_q[\widehat{U}(\pi_1|Q)] > \mathbb{E}_q[\widehat{U}(\pi_2|Q)] \Leftrightarrow U(\pi_1) > U(\pi_2|q).$$

Therefore, given enough data, the click noise does not affect the ability to find the ranking policy with optimal utility.

In contrast, the disparity estimator in Equations (6) and (7) is not order-preserving when $\epsilon_- > 0$:

$$\begin{aligned} \mathbb{E}_q[\widehat{D}_{ij}(\pi_1|Q) - \widehat{D}_{ij}(\pi_2|Q)] &= (\epsilon_+ - \epsilon_-) \mathbb{E}_q[D_{ij}(\pi_1|Q) - D_{ij}(\pi_2|Q)] \\ &\quad + \epsilon_- \mathbb{E}_q[|G_i^q| \delta \text{Exp}_q(G_i) - |G_j^q| \delta \text{Exp}_q(G_j)], \end{aligned}$$

where $\delta \text{Exp}_q(G_i)$ denotes $\text{Exp}_q(G_i|\pi_1) - \text{Exp}_q(G_i|\pi_2)$. However, if we know the level of negative noise ϵ_- , we can correct the IPS estimator for group disparity as

$$\begin{aligned} \widehat{D}_{ij}(\pi|q, \epsilon_-) &= \widehat{M}_q(G_j) \text{Exp}_q(G_i|\pi) - \widehat{M}_q(G_i) \text{Exp}_q(G_j|\pi) \\ &\quad - \epsilon_- (|G_i^q| \text{Exp}_q(G_i|\pi) - |G_j^q| \text{Exp}_q(G_j|\pi)). \end{aligned} \quad (10)$$

We can prove that the corrected estimator is unbiased up to a constant factor that can be absorbed into the hyperparameter δ of Equation (2)

$$\mathbb{E}_{c_q}[\widehat{D}_{ij}(\pi|q, \epsilon_-)] = (\epsilon_+ - \epsilon_-) D_{ij}(\pi|q).$$

To estimate the probability ϵ_- for false-positive noise, we can use a simple intervention similar to [28, 45]. While identifying relevant items is difficult, identifying irrelevant items is typically easy (e.g., documents that contain none of the query terms). Therefore, we can insert an item d that is known to be irrelevant at position k of the ranking. The expected clickthrough rate for the item is $p(c_{q,d} = 1 | \text{rel}_{q,d} = 0) = v_k \cdot \epsilon_-$. This means that given the position bias and the clickthrough rate estimated from the intervention data, we can compute the noise level ϵ_- . Note that this intervention only needs to be performed on a small set of users and only affects the quality of one position.

4 POLICY-GRADIENT ALGORITHM FOR FAIR LTR

In the previous section, we defined a general framework for learning ranking policies from biased and noisy feedback under amortized fairness constraints. However, we still need an efficient algorithm to search the specific policy space for the solution of the training problem in Equation (2). To this effect, we first define a stochastic ranking policy space based on the Plackett-Luce ranking model as in [44] and then present a policy-gradient algorithm that optimizes the training objective.

4.1 Plackett-Luce Ranking Model

We define a stochastic ranking policy space Π based on the Plackett-Luce model [29, 36]. Specifically, each ranking policy $\pi \in \Pi$ is defined by a scoring function h_θ , which can be any differentiable model with parameters θ . h_θ takes the feature vectors x_q of all the items for the current query q as input and outputs a vector of scores $h_\theta(x_q) = (h_\theta(x_{q,d_1}), h_\theta(x_{q,d_2}), \dots, h_\theta(x_{q,d_{n_q}}))$. Based

on this score vector, the probability $\pi_\theta(r|q)$ of a ranking $r = \langle d_1, d_2, \dots, d_{n_q} \rangle$ is defined as the product of softmax distributions

$$\pi_\theta(r|q) = \prod_{i=1}^{n_q} \frac{\exp(h_\theta(x_{q,d_i}))}{\sum_{j=i}^{n_q} \exp(h_\theta(x_{q,d_j}))}. \quad (11)$$

To sample rankings from $\pi_\theta(r|q)$, we can sample items from the distribution $\text{Softmax}(h_\theta(x_q))$ without replacement and rank the items in the order in which they are drawn.

4.2 Policy Gradient Training Algorithm

We use a Lagrange multiplier to solve the constrained optimization problem from Equation (2): $\hat{\pi} = \operatorname{argmax}_\pi \min_{\lambda \geq 0} \widehat{U}(\pi|Q) - \lambda([\widehat{D}_{ij}(\pi|Q)]^2 - \delta)$. Instead of solving the minimization problem w.r.t. λ , we search a specific range of $\lambda \in \{\lambda_1, \dots, \lambda_k\}$. For each λ , we need to solve

$$\hat{\pi}_\lambda = \operatorname{argmax}_\pi \widehat{U}(\pi|Q) - \lambda[\widehat{D}_{ij}(\pi|Q)]^2. \quad (12)$$

Afterwards, we can compute the corresponding $\delta_\lambda = [\widehat{D}_{ij}(\hat{\pi}_\lambda|Q)]^2$ for which the constraint holds and then pick the optimal $\hat{\pi}_\lambda$ that satisfies $\delta_\lambda \leq \delta$ and provides maximal utility $\widehat{U}(\hat{\pi}_\lambda|Q)$.

It remains to find an efficient algorithm for solving the unconstrained optimization problem in Equation (12). We use stochastic gradient descent (SGD) to iteratively update the parameters of the ranking policy. However, it is intractable to compute the expectations $\widehat{U}(\pi|q)$ and $\widehat{D}_{ij}(\pi|q)$ over the exponential space of rankings. Instead, we use the log-derivative trick of the REINFORCE algorithm [46] to compute the gradient of \widehat{U}

$$\begin{aligned} \nabla_\theta \widehat{U}(\pi_\theta|Q) &= \frac{1}{|Q|} \nabla_\theta \sum_q \mathbb{E}_{\sigma \sim \pi(\sigma|q)} [\widehat{\Delta}(\sigma, c_q)] \\ &= \frac{1}{|Q|} \sum_q \mathbb{E}_{\sigma \sim \pi(\sigma|q)} [\nabla_\theta \log \pi_\theta(\sigma|q) \widehat{\Delta}(\sigma, c_q)]. \end{aligned} \quad (13)$$

The gradient of the squared fairness disparity, $\nabla_\theta [\widehat{D}_{ij}(\pi_\theta|Q)]^2$, can also be reformulated this way,

$$\frac{2}{|Q|} \widehat{D}_{ij}(\pi_\theta|Q) \sum_{q \in Q} \mathbb{E}_{\sigma \sim \pi_\theta(\sigma|q)} [\nabla_\theta \log \pi_\theta(\sigma|q) \widehat{\text{diff}}_{ij}(\sigma|q)], \quad (14)$$

where $\widehat{\text{diff}}_{ij}(\sigma|q)$ denotes $\widehat{M}_q(G_j) \text{Exp}_q(G_i|\sigma) - \widehat{M}_q(G_i) \text{Exp}_q(G_j|\sigma)$. The expectations over rankings in Equations (13) and (14) are approximated via $S = 32$ Monte-Carlo samples in our experiments.

To apply SGD, we need to compute $\widehat{D}_{ij}(\pi_{\theta_t}|Q)$ on the whole dataset at each step t , which is quite expensive. As a simple but practically effective solution, we use the running average $\frac{1}{n} \sum_{\tau=0}^{n-1} \widehat{D}_{ij}(\pi_{\theta_{t-\tau}}|q_{t-\tau})$ as an approximation of $\widehat{D}_{ij}(\pi_{\theta_t}|Q)$ in Equation (14). This is biased since the disparities are computed on previous parameters, but it can reduce the variance of the unbiased estimator $\widehat{D}_{ij}(\pi_{\theta_t}|q'_t)$ with q'_t from the query set.

5 EMPIRICAL EVALUATION

We conduct various experiments on two real-world datasets from which we generate synthetic click data following the experiment setup in [3, 28]. This allows us to control the experimental conditions and test multiple data distributions to evaluate robustness.

5.1 Setup

5.1.1 Datasets. We use the Microsoft Learning to Rank Dataset (Fold1), a large-scale LTR dataset [37], and the German Credit Dataset, a credit scoring dataset often used to study algorithmic fairness [17]. The German Credit Dataset contains information about 1000 loan applicants, where each applicant is described by a set of attributes and labeled as creditworthy or non-creditworthy. We randomly split the applicants into train, validation, and test sets with ratio 1:1:1. To define groups, we use the binary feature A43 (indicating whether the purpose of the loan applicant is radio/television) as the group attribute. The ratio of applicants in two groups is around 8:2.

The Microsoft LTR Dataset contains a large number of queries from Bing with manually-judged relevance labels for retrieved web-pages. We adopt the train, validation, and test split provided with the dataset. Since the dataset does not come with any designated groups, we select the QualityScore (feature id 133) as the group attribute and divide items into two groups with the 40th percentile on the attribute as the threshold.

5.1.2 Click Simulation. We generate click data for training and validation from the full-information datasets following [3, 28]. We first train a Ranking SVM with 1 percent of the full-information training data as the logging policy. This logging policy is then used to generate the rankings for which click data is logged. The click data is generated by simulating the position-based examination model. The position bias decays with the presented rank k of the item as $v_k = (1/k)^\eta$. When not stated otherwise, we use $\eta = 1$, $\epsilon_+ = 1$, and $\epsilon_- = 0$. Since the click signal is always binary and cannot reflect the relevance ratings, we binarize relevances in the same way as [28], by assigning $\text{rel}_q(d) = 1$ to all the items that were judged as 3 or 4 and $\text{rel}_q(d) = 0$ to judgments 0, 1, and 2.

We preprocess both datasets into similar formats for LTR. Following [44], we build a query for the German Credit Dataset by sampling 20 individuals from the corresponding set with ratio 9:1 for non-creditworthy individuals to creditworthy individuals respectively. We sample 500 queries from train, validation, and test sets respectively. For MSLR, after the binarization step, the dataset becomes extremely sparse, with only about 2.5% relevant items per query. To better compare different methods and amplify differences, for each query we sample 20 candidate items with 3 relevant items to build a new query.

The performance of the methods is reported on the full-information test set for which all relevance labels are known. Ranking utility and fairness are measured with Average DCG [24] and squared disparity $[D_{ij}(\pi)]^2$ respectively.

5.1.3 Models and Hyperparameters. We train FULTR for two types of models: a linear model and a feed-forward network (one hidden layer with ReLU activation). We use the SGD optimizer for the linear model and the Adam optimizer for the neural network. The learning rate is 0.001. Following [44], we subtract the average reward of the Monte-Carlo samples from the reward to act as a control variate for variance reduction [46]. To encourage exploration and avoid premature convergence to suboptimal policies [31], we add the entropy of the probability distribution $\text{Softmax}(h_\theta(x_q))$ to the objective, with a regularization coefficient γ . We initialize

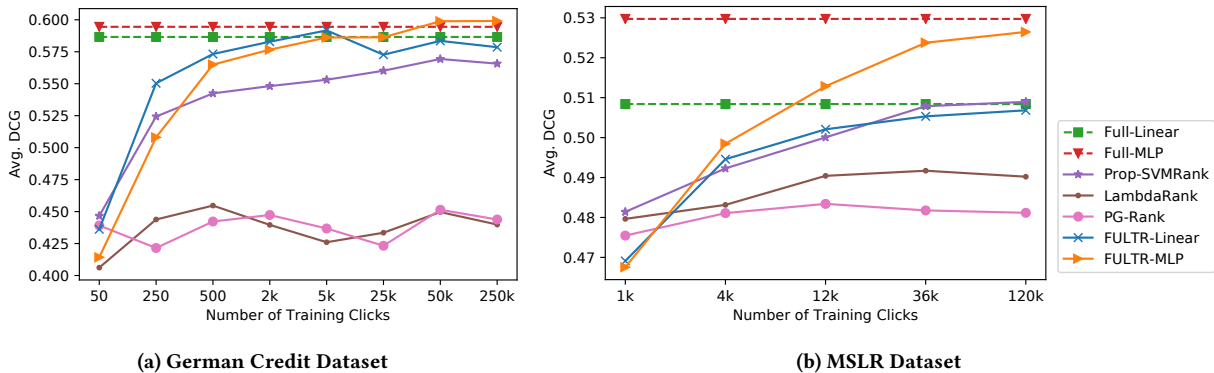


Figure 1: Ranking utility performance in terms of training clicks in the partial-info setting ($\eta = 1, \epsilon_- = 0.1$).

the coefficient as $\gamma = 1.0$ and reduce it by a factor of 3 each time the validation metric has stopped improving. Finally, we add an L2 regularization term and cross-validate for the best regularization coefficient.

5.2 Can FULTR learn accurate ranking policies from biased feedback?

We begin our empirical evaluation by establishing that FULTR without fairness constraints is competitive with conventional LTR methods when learning from partial feedback. Here, we ignore fairness and entirely focus on utility, just like in conventional LTR. For realism, we inject noise $\epsilon_- = 0.1$. We compare FULTR with the following LTR methods based on linear models or neural networks:

- LambdaRank [10]: This is a nonlinear ranking model based on neural networks to optimize DCG.
- Propensity SVM-Rank [28]: This is an unbiased version of SVM-Rank [25] and utilizes the IPS estimator to correct position bias in implicit feedback.
- PG-Rank [44]: This is a full-information policy-gradient method that is training the same model as FULTR-Linear.

Note that the full-information methods LambdaRank and PG-Rank ignore the position bias and treat click signals naively as full relevances. We also add a linear model and a neural network which get to see all true relevance labels as two skylines, called Full-Linear and Full-MLP. They represent the maximum performance we could hope to achieve with FULTR, which has access to only strictly less informative click feedback.

In Figure 1, we show the test-set performance of the conventional LTR methods compared to FULTR. Since the optimal policy is always deterministic, following [44], we use the highest probability ranking for each query when evaluating FULTR. With increasing amounts of click data on the x-axis, FULTR with the neural network and the linear model approaches the skyline performance of Full-MLP and Full-Linear respectively on both datasets. The performance of FULTR-Linear is similar to Propensity SVM-Rank. Baseline methods that naively ignore position bias, namely LambdaRank and PG-Rank, cannot make effective use of the increasing amount of click data, and their learning curves are quite flat. Their performance is well below FULTR given sufficient click data.

Overall, we conclude that FULTR is an LTR algorithm that achieves ranking performance that is at least competitive with existing baselines when learning linear or neural models from partial-information feedback. This implies that FULTR is a promising contender for use in practical applications, and it thus makes sense to further investigate how far it can enforce fairness considerations.

5.3 Can FULTR learn fair ranking policies from biased feedback?

Next, we investigate FULTR’s ability to enforce merit-based fairness of exposure. We compare with the following methods as baselines:

- Group-blind: A variant of FULTR without any fairness constraints, but with the group attribute masked (i.e., fairness through unawareness).
- Fair-PG-Rank [44]: PG-Rank with a fairness constraint.
- Fair-PG-Rank-IPS: Fair-PG-Rank [44], naively using the IPS estimator as in Equation (8).
- Equity of Attention [8]: Post-processing method that optimizes amortized group disparity. We train a regression model $r(x_{q,d})$ to predict item relevances using the position-bias-corrected objective for relevance estimation in [32].

The resulting utility/disparity curves on the test set are shown in Figure 2. First, note that group-blind training does not lead to merit-based fairness of exposure. Second, FULTR is able to effectively trade-off utility and fairness on both datasets. As λ increases, the disparity goes down with an associated drop in DCG as expected. Increasing λ further drives disparity close to zero. The neural network version, FULTR-MLP, achieves better utility-fairness trade-off than the linear version on MSLR dataset. The German Credit dataset was not large enough to train a reliable MLP model.

Third, Fair-PG-Rank cannot effectively reduce amortized disparity. On both datasets, as we vary the coefficient λ , utility drops while disparity increases. We conjecture that this is due to optimizing fairness per query instead of amortized fairness, and the propagation of bias in clicks to the merit-based disparity measure. The naive combination of Fair-PG-Rank and IPS weighting cannot reduce the disparity, since the disparity estimator is still biased.

Fourth, the post-processing method cannot reduce amortized disparity on the MSLR dataset but can approach zero disparity on

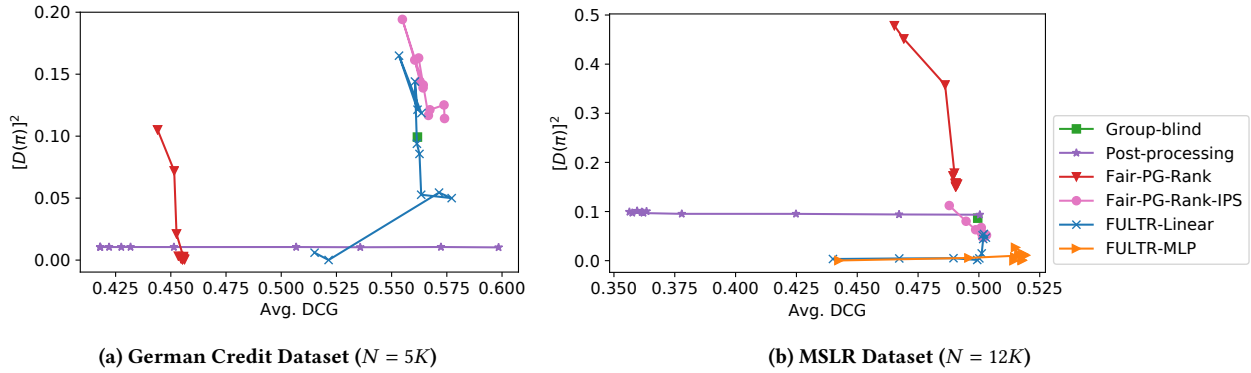


Figure 2: DCG-Disparity curves for FULTR against baselines in the partial-info setting ($\eta = 1, \epsilon_- = 0$).

German Credit dataset. To further analyze the reason, we ran the post-processing method with ground-truth relevances. This infeasible approach achieves disparities close to zero on MSLR, which demonstrates that the problem lies in the regression model. To investigate further, we compute the accumulated regression errors for both groups. We observe that on MSLR, the regression model underestimates the relevance of G_i but estimates the relevance of G_j precisely (error -14.05% vs -0.09%), which means that the regression model is already unfair despite its bias-corrected training objective. On the German Credit dataset, the regression model underestimates the relevance of two groups more uniformly (error -15.86% vs -19.47%). We conclude that the failure of the post-processing on MSLR is due to its two-step nature, where the regression model is trained oblivious to fairness.

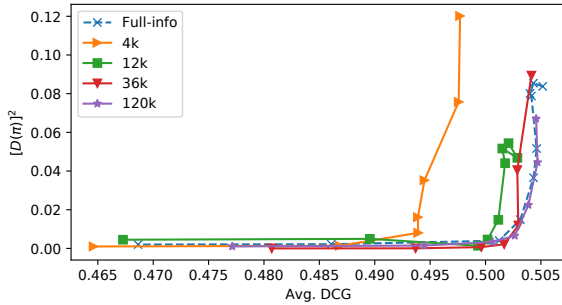


Figure 3: DCG-Disparity curves for FULTR in terms of training clicks on the partial MSLR dataset ($\eta = 1, \epsilon_- = 0$).

5.4 Can FULTR converge to the performance of training on the true relevance labels?

We now explore how ranking utility and fairness of FULTR converge as the amount of click data increases. The utility/disparity curves with various amounts of training clicks are shown in Figure 3. We also show the policy learned by FULTR in the full-info setting, which serves as the skyline that has full knowledge of relevance labels in the training set. With increasing amounts of click data, FULTR approaches the skyline performance of the policy learned

on the full-info data. The policy learned with 120k clicks is almost identical to that of the full-information policy. This demonstrates that the unbiased estimator of FULTR enables it to converge to the full-information policy given enough click data.

5.5 How important is unbiasedness for utility and fairness?

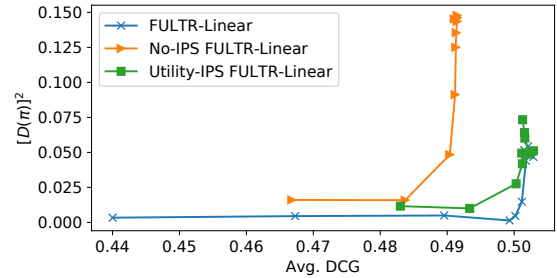


Figure 4: DCG-Disparity curves for FULTR and two variants with unweighted utility and disparity estimators on the partial MSLR dataset ($N = 12K, \eta = 1, \epsilon_- = 0$).

We conduct an ablation study to understand the effect of IPS-weighting on utility and fairness. In particular, we compare the following variants of FULTR in Figure 4:

- No-IPS FULTR-Linear: Both utility and disparity estimators are not IPS-weighted, assuming that implicit feedback reveals full-information relevance labels.
- Utility-IPS FULTR-Linear: The utility estimator is IPS-weighted, but the disparity estimator is unweighted.

With biased estimates of both utility and disparity, the *No-IPS* variant of FULTR achieves suboptimal utility and cannot reduce disparity to zero. The *Utility-IPS* variant can achieve utility similar to that of FULTR, but its fairness disparity remains higher even for large values of λ . This implies eliminating position bias through IPS-weighting is essential for both utility and fairness.

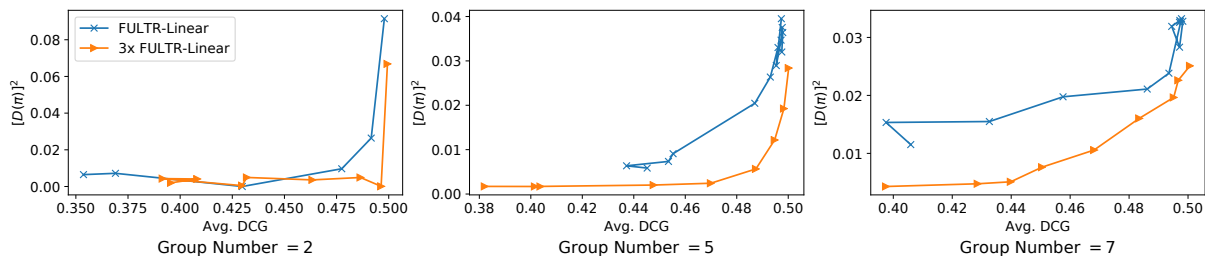


Figure 5: DCG-Disparity curves for different group numbers on the partial MSLR dataset ($N = 4K$ and $N = 12K$, $\eta = 1$, $\epsilon_- = 0$).

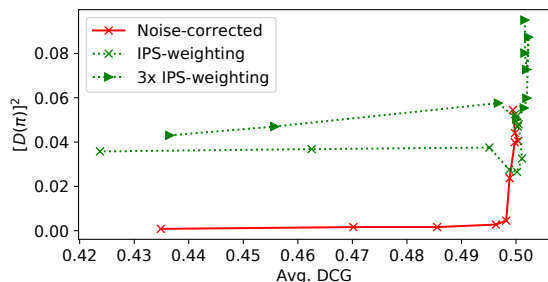


Figure 6: DCG-Disparity curves of noise-corrected estimators against pure IPS-weighting on the partial MSLR dataset with click noise ($N = 36K$ and $N = 120K$, $\eta = 1$, $\epsilon_- = 0.1$).

5.6 Can the disparity estimator adjust to click noise?

Next, we investigate the effectiveness of the noise-corrected disparity estimator in Equation (10). We use partial MSLR with click noise $\epsilon_- = 0.1$. Noise-corrected FULTR estimates ϵ_- via the intervention with 1% of the training data and applies the noise-corrected estimator in Equation (10). Figure 6 shows that the noise-corrected IPS estimator of $D_{ij}(\pi)$ can reduce disparity effectively and achieve zero disparity, while the pure IPS estimator cannot. To verify that this is not due to a lack of data, we increase the amount of training data by a factor of three. However, more data alone does not remedy the problem. This is consistent with the theoretical argument, since more training data can only reduce the variance, while click noise leads to a bias in the disparity estimates.

5.7 Can FULTR ensure fairness between more than two groups?

So far, all experiments are conducted with only two groups. To deal with multiple groups, we can modify the constraint in Equation (2) to

$$\sum_{G_i \neq G_j} \left[\widehat{D}_{ij}(\pi|Q) \right]^2 \leq \delta.$$

To validate the effectiveness of FULTR when dealing with multiple groups, we compare the utility-fairness trade-off for a varying number of groups. To create more than two groups for MSLR, we partition the group attribute into the desired number of intervals with the same numbers of items for each group. Figure 5 shows

the results for 2, 5, and 7 groups. We can observe that with a fixed number of training clicks, the disparity of FULTR increases as the number of groups increases. With 7 groups, a substantial disparity exists even for the largest value of λ . This is due to a lack of data for each group, as the number of groups is increasing while the amount of training data remains fixed. If we increase the number of training clicks, FULTR can again achieve disparity close to zero even with 7 groups. This demonstrates that FULTR can deal with multiple groups, but that an increase in the number of groups also requires an increase in training data.

6 CONCLUSION

We presented a framework FULTR for learning accurate and fair ranking policies from biased feedback. Specifically, we introduced fairness-of-exposure constraints that can allocate amortized exposure to groups of items based on their amortized relevance. Furthermore, we proposed counterfactual estimators of the corresponding utility and disparity measures that remove the effects of position bias and click noise. For efficient training, we derived a policy gradient algorithm that directly optimizes both utility and fairness. Finally, we presented extensive empirical evidence that FULTR can effectively learn ranking policies under fairness constraints despite biased and noisy feedback. Given that FULTR estimates relevance based on position bias across multiple queries, one existing limitation of the framework is its inability to deal with individual-level fairness due to a lack of sufficient data. In the future, it would be interesting to explore ways to scale FULTR to incorporate individual fairness.

ACKNOWLEDGMENTS

This research was supported in part by NSF Awards IIS-1901168 and IIS-2008139. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

REFERENCES

- [1] Lada A Adamic and Bernardo A Huberman. 2000. Power-law distribution of the world wide web. *science* 287, 5461 (2000), 2115–2115.
- [2] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. 2018. A reductions approach to fair classification. *arXiv preprint arXiv:1803.02453* (2018).
- [3] Aman Agarwal, Xuanhui Wang, Cheng Li, Michael Bendersky, and Marc Najork. 2019. Addressing Trust Bias for Unbiased Learning-to-Rank. In *The Web Conference*. 4–14.
- [4] Aman Agarwal, Ivan Zaitsev, Xuanhui Wang, Cheng Li, Marc Najork, and Thorsten Joachims. 2019. Estimating Position Bias without Intrusive Interventions. In *WSDM*. 474–482.

- [5] 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.
- [6] Richard Berk, Hoda Heidari, Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2017. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409* (2017).
- [7] 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. *arXiv preprint arXiv:1903.00780* (2019).
- [8] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *SIGIR*. 405–414.
- [9] Alexey Borisov, Ilya Markov, Maarten de Rijke, and Pavel Serdyukov. 2016. A neural click model for web search. In *WWW*. 531–541.
- [10] Christopher J. C. Burges, Robert Ragno, and Quoc Viet Le. 2006. Learning to Rank with Nonsmooth Cost Functions. In *NIPS*. 193–200.
- [11] Christopher J. C. Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Gregory N. Hullender. 2005. Learning to rank using gradient descent. In *ICML*. 89–96.
- [12] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. 2007. Learning to rank: from pairwise approach to listwise approach. In *ICML*. 129–136.
- [13] L Elisa Celis, Damian Straszak, and Nisheeth K Vishnoi. 2017. Ranking with fairness constraints. *arXiv preprint arXiv:1704.06840* (2017).
- [14] Olivier Chapelle and Ya Zhang. 2009. A dynamic bayesian network click model for web search ranking. In *WWW*. ACM, 1–10.
- [15] 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.
- [16] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *WSDM*. 87–94.
- [17] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- [18] Georges E Dupret and Benjamin Piwowarski. 2008. A user browsing model to predict search engine click data from past observations. In *SIGIR*. 331–338.
- [19] Zhichong Fang, Aman Agarwal, and Thorsten Joachims. 2019. Intervention Harvesting for Context-Dependent Examination-Bias Estimation. In *SIGIR*. 825–834.
- [20] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In *SIGKDD*. 2221–2231.
- [21] Stephen Gillen, Christopher Jung, Michael Kearns, and Aaron Roth. 2018. Online learning with an unknown fairness metric. In *NIPS*. 2600–2609.
- [22] Fan Guo, Chao Liu, and Yi Min Wang. 2009. Efficient multiple-click models in web search. In *WSDM*. 124–131.
- [23] Ziniu Hu, Yang Wang, Qu Peng, and Hang Li. 2018. A Novel Algorithm for Unbiased Learning to Rank. *arXiv preprint arXiv:1809.05818* (2018).
- [24] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems* 20, 4 (2002), 422–446.
- [25] Thorsten Joachims. 2002. Optimizing search engines using clickthrough data. In *SIGKDD*. 133–142.
- [26] 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. *ACM Transactions on Information Systems* 25, 2 (2007), 7.
- [27] Thorsten Joachims, Laura A Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately interpreting clickthrough data as implicit feedback. In *SIGIR*. Vol. 5. 154–161.
- [28] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased Learning-to-Rank with Biased Feedback. In *WSDM*. 781–789.
- [29] R. Duncan Luce. 1959. *Individual Choice Behavior: A Theoretical analysis*. Wiley.
- [30] 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 and Satisfaction in Recommendation Systems.
- [31] Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *ICML*. 1928–1937.
- [32] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling Fairness and Bias in Dynamic Learning-to-Rank. In *SIGIR*. 429–438.
- [33] Harikrishna Narasimhan. 2018. Learning with complex loss functions and constraints. In *AISTATS*. 1646–1654.
- [34] Harikrishna Narasimhan, Andrew Cotter, Maya Gupta, and Serena Wang. 2019. Pairwise Fairness for Ranking and Regression. *arXiv preprint arXiv:1906.05330* (2019).
- [35] Vishakha Patil, Ganesh Ghalme, Vineet Nair, and Y Narahari. 2019. Achieving Fairness in the Stochastic Multi-armed Bandit Problem. *arXiv preprint arXiv:1907.10516* (2019).
- [36] R. L. Plackett. 1975. The analysis of permutations. *Applied Statistics* 24 (1975), 193–202.
- [37] Tao Qin and Tie-Yan Liu. 2013. Introducing LETOR 4.0 Datasets. *arXiv preprint arXiv:1306.2597* (2013).
- [38] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *UAI*. 452–461.
- [39] Ashkan Rezaei, Rizal Fathony, Omid Memarrast, and Brian Ziebart. 2019. Fair Logistic Regression: An Adversarial Perspective. *arXiv preprint arXiv:1903.03910* (2019).
- [40] S. E. Robertson. 1997. *The Probability Ranking Principle in IR*. Morgan Kaufmann Publishers Inc., 281–286.
- [41] PAUL ROSENBAUM and Donald Rubin. 1983. The Central Role of the Propensity Score in Observational Studies For Causal Effects. *Biometrika* 70 (1983), 41–55.
- [42] Matthew J Salganik, Peter Sheridan Dodds, and Duncan J Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *science* 311, 5762 (2006), 854–856.
- [43] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *SIGKDD*. 2219–2228.
- [44] Ashudeep Singh and Thorsten Joachims. 2019. Policy Learning for Fairness in Ranking. In *NeurIPS 2019*. 5427–5437.
- [45] Xuanhui Wang, Michael Bendersky, Donald Metzler, and Marc Najork. 2016. Learning to Rank with Selection Bias in Personal Search. In *SIGIR*. 115–124.
- [46] Ronald J. Williams. 1992. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine Learning* 8 (1992), 229–256.
- [47] Ke Yang and Julia Stoyanovich. 2017. Measuring fairness in ranked outputs. In *International Conference on Scientific and Statistical Database Management (2017)*. 22.
- [48] 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.
- [49] Meike Zehlike and Carlos Castillo. 2018. Reducing disparate exposure in ranking: A learning to rank approach. *arXiv preprint arXiv:1805.08716* (2018).