



FARA: Future-aware Ranking Algorithm for Fairness Optimization

Tao Yang
University of Utah
Salt Lake City, Utah
taoyang@cs.utah.edu

Zhenduo Wang
University of Utah
Salt Lake City, Utah, USA
zhenduow@cs.utah.edu

Zhichao Xu
University of Utah
Salt Lake City, Utah
zhichao.xu@utah.edu

Qingyao Ai*
DCST, Tsinghua University
Quan Cheng Laboratory, Zhongguancun Laboratory
Beijing, China
aiqy@tsinghua.edu.cn

Abstract

Ranking systems are the key components of modern Information Retrieval (IR) applications, such as search engines and recommender systems. Besides the ranking relevance to users, the exposure fairness to item providers has also been considered an important factor in ranking optimization. Many fair ranking algorithms have been proposed to jointly optimize both ranking relevance and fairness. However, we find that most existing fair ranking methods adopt greedy algorithms that only optimize rankings for the next immediate session or request. As shown in this paper, such a myopic paradigm could limit the upper bound of ranking optimization and lead to suboptimal performance in the long term.

To this end, we propose **FARA**, a novel Future-Aware Ranking Algorithm for ranking relevance and fairness optimization. Instead of greedily optimizing rankings for the next immediate session, FARA plans ahead by jointly optimizing multiple ranklists together and saving them for future sessions. Specifically, FARA first uses the Taylor expansion to investigate how future ranklists will influence the overall fairness of the system. Then, based on the analysis of the Taylor expansion, FARA adopts a two-phase optimization algorithm where we first solve an optimal future exposure planning problem and then construct the optimal ranklists according to the optimal future exposure planning. Theoretically, we show that FARA is optimal for ranking relevance and fairness joint optimization. Empirically, our extensive experiments on three semi-synthesized datasets show that FARA is efficient, effective, and can deliver significantly better ranking performance compared to state-of-the-art fair ranking methods. We make our implementation public at https://github.com/Taosheng-ty/QP_fairness/.

CCS Concepts

• Information systems → Learning to rank.

*Corresponding authors



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

CIKM '23, October 21–25, 2023, Birmingham, United Kingdom
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0124-5/23/10.
<https://doi.org/10.1145/3583780.3614877>

Keywords

Fair Ranking, Position Bias, Exposure, Exposure Fairness

ACM Reference Format:

Tao Yang, Zhichao Xu, Zhenduo Wang, and Qingyao Ai. 2023. FARA: Future-aware Ranking Algorithm for Fairness Optimization. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3583780.3614877>

1 INTRODUCTION

Ranking systems are one of the important cornerstones of information retrieval (IR). Existing ranking systems are usually constructed to optimize ranking relevance with the Probability Ranking Principle (PRP) [34] where items of greater likely relevance should be ranked higher. The PRP is a user-centered ranking strategy that helps save users energy and time since users could satisfy their needs with the top-ranked items [18]. However, recent research has shown that, besides users, item providers also draw utility from ranking systems, and the PRP could result in severe unfairness for item providers [4, 37]. Particularly, the PRP always assigns a few top items with high-rank positions. Those top items usually get the majority of exposure while other items rarely get exposure, although other items might still be relevant [19, 29, 40, 49]. The unbalanced exposure leads to unfair opportunities and unfair economic gains for item providers. Such unfairness will eventually force unfairly treated providers to leave the system, and fewer options will be left for users [51]. Therefore, IR researchers have argued that ranking relevance and fairness are both important for modern ranking systems [37, 38]. Many fair ranking algorithms have been proposed to optimize both of them jointly [29, 55].

However, existing fair algorithms are mostly greedy algorithms and could only deliver suboptimal ranking performance in the long run. In particular, existing fair ranking algorithms [4, 24, 26, 37, 47] usually behave greedily to sequentially produce the locally optimal ranklist for the next immediate session without being aware of the influence of future sessions¹ (see more discussion in §2). The unawareness could lead to unmitigated ranking conflict between relevance and fairness optimization. For example, imagine a case that there are in total 3 items in consideration, item *A*, item *B*, and

¹In this paper, we define a session as a query issued by a user.

item C , where item A is the most relevant one and item C is the least relevant one. Ranklist $[A, B, C]$ is the ranklist to maximize ranking relevance. We now consider a scenario where item C is severely unfairly treated in history. To optimize exposure fairness, we need to allocate item C more exposure by boosting item C to a higher position. However, If we try to greedily boost item C within the next immediate session, it is highly likely that item C will be boosted to the first rank to get the maximum exposure and the result ranklist is $[C, A, B]$. However, Ranklist $[C, A, B]$ is of poor ranking relevance due to the ranking conflict that the least relevant item (item C) is put on the most important rank (the first rank).

Intuitively, the ranking conflict can be smoothed if we plan ahead and jointly optimize multiple future sessions' ranklists together instead of greedily optimizing the next immediate session. For example, the multiple ranklists after joint optimization can be $[[A, C, B], [A, C, B], \dots]$, where item C is smoothly boosted in multiple ranklists and the most relevant item, i.e., item A , is still ranked the highest. Based on the above idea, we propose **FARA**, a novel **Future-Aware Ranking Algorithm** for relevance and fairness optimization. Briefly, FARA precomputes and jointly optimizes multiple ranklists together and saves them for future use. Particularly, to be able to plan for the future, FARA first uses the Taylor expansion to investigate how future ranklists will influence fairness. Then, based on the influence, FARA uses a two-phase optimization to jointly optimizes multiple ranklists together for future use. In phase 1, we solve an exposure planning problem and get the optimal future planning for item exposure. In phase 2, we construct the optimal ranklists according to the optimal future planning for item exposure. We prove FARA's optimum in terms of ranking relevance and fairness joint optimization in § 5. Extensive experiments on three semi-synthesized datasets also demonstrate FARA's effectiveness and efficiency compared to existing fair ranking algorithms (§ 6).

2 RELATED WORK

Ranking Fairness: Due to the importance of rankings for providers (sellers, job seekers, content creators, etc.), [9, 17, 37, 53, 54], there has been growing interest in ranking fairness for providers [5, 12, 15, 20, 25, 32, 41, 41, 47]. However, the definitions of ranking fairness vary a lot in the existing literature, and there exists no universal definition. At a high level, existing fairness definitions can be grouped into probability-based fairness and exposure-based fairness [29, 55]. Probability-based fairness [3, 6, 13] usually requires a minimum number or proportion of protected (e.g., race, gender) items to be distributed evenly across a ranklist. However, only considering the number or proportion of items in a ranklist neglects the fact that different ranks usually have different importance. To address this, exposure-based fairness [4, 8, 11, 37, 48] assigns values to each ranking position based on the expected user attention or click probability. Exposure-based fairness argues that total exposure is a limited resource for a ranking system and advocates for fair distribution of exposure among items to ensure fairness for item providers [29]. In this paper, we limit our discussion of fair ranking algorithms within the scope of exposure-based fairness.

Fair Ranking Algorithms: Recently, a few ranking algorithms [10, 15, 23, 29, 35, 41, 55] have been proposed to achieve exposure-based fairness. In this work, we classify them as **open-loop** algorithms or **feedback-loop** algorithms depending on whether historically

Table 1: A summary of notations in this paper.

$d, q, D(q)$	For a query q , $D(q)$ is the set of candidates items. $d \in D(q)$ is an item.
e, r, c	All are binary random variables indicating whether an item d is examined ($e = 1$), perceived as relevant ($r = 1$) and clicked ($c = 1$) by a user respectively.
R, P_i, E, π	$R = P(r = 1 d)$, is the probability of an item d perceived as relevant. $P_i = P(e = 1 rank(d \pi) = i)$ is the examination probability of item d when it is put in i^{th} rank in a ranklist π . E is item's accumulated examination probability (see Eq.7).
k_s, k_c	Users will stop examining items lower than rank k_s due to selection bias (see Eq. 3). k_c is the cutoff prefix to evaluate Cum-NDCG and $k_c \leq k_s$.

generated ranklists are used to correct ranking scores. For open-loop fair algorithms [15, 26, 36–38, 42, 46], each item usually has a static and fixed ranking score once the ranking model is optimized. Then ranklists are stochastically sampled for each session according to the static ranking scores. Various techniques have been used to optimize the static ranking model, such as linear programming [15, 37], policy gradient [38], differentiable PL model optimization [26]. However, given the fact that ranking scores are static, open-loop algorithms are usually not robust. To improve ranking robustness, feedback-loop algorithms [4, 24, 48, 52] dynamically take historical ranklists as input to correct items' scores. For example, Morik et al. [24] proposes to use a proportional controller to boost ranking scores of historically unfairly treated items.

3 BACKGROUND AND PRIOR KNOWLEDGE

In this section, we provide readers with background knowledge of the paper. A summary of notations we use throughout the paper can be found in Table 1.

Ranking Services Workflow: We take web search as an example to detail the ranking service workflow. At time step t , a ranking session starts when user t issues a query q . For query q , there exist candidate items provided by item providers. With the query and candidate items, the ranking system first estimates each item's relevance and then constructs a ranklist of candidate items by optimizing certain ranking objectives. Then the ranking system presents the ranklist to users and collects users' feedback (e.g., clicks) which can be utilized to update the relevance estimator.

Partial and Biased Feedback: Relevance estimation is usually updated using users' feedback. However, such feedback is usually a noisy and biased indicator of relevance since users only provide meaningful feedback for items they have examined. If we consider user clicks as the main signal for user feedback, then we have

$$c = \begin{cases} r, & \text{if } e = 1 \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where e, r, c are binary random variables indicating whether an item is examined, perceived as relevant, and clicked, respectively. With User Examination Hypothesis [33], we can model users' click probability as

$$P(c = 1) = P(e = 1)P(r = 1) \quad (2)$$

For the rest of the paper, we use $R = P(r = 1)$ to simplify the notation. Although there exist several types of biases in the examination probability $P(e = 1)$, we focus on two most important ones: positional bias and selection bias.

Positional Bias [7]: Examination probability is decided by the rank (also called position) and drops along ranks. Particularly, the examining probability is denoted as $P_{rank(d|\pi)}$, where $rank(d|\pi)$ is item d 's rank in ranklist π .

Selection Bias [27, 28]: This bias exists when not all of the items are selected to be shown to users, or some lists are so long that no users will examine every item in them. Assuming the items ranked lower than rank k_s will not be examined [27], we model this as:

$$P(e = 1|d, \pi) = \begin{cases} P_{rank(d|\pi)}, & \text{if } rank(d|\pi) \leq k_s \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

Ranking Utility Measurement: Here we introduce the evaluation of ranking performance from both the user side and the provider side, which in later sections will guide the ranking optimization.

The User-side Utility: User-side utility measures a ranking system's ability to put relevant items on higher ranks. A popular user-side utility measurement is DCG [16]. Specifically, given a query q and a ranklist π , $DCG@k_c$ is defined as,

$$DCG@k_c(\pi) = \sum_{i=1}^{k_c} R(\pi[i], q) \lambda_i = \sum_{i=1}^{k_c} R(\pi[i], q) P_i \quad (4)$$

where $\pi[i]$ indicates the i^{th} ranked item in ranklist π , $R(\pi[i], q)$ indicates $\pi[i]$'s relevance to query q , cutoff k_c indicates the prefix we want to evaluate, λ_i indicates the weight we put on i^{th} rank. λ_i is usually monotonically decreases as i increases, e.g., λ_i is usually set to $\frac{1}{\log_2(i+1)}$ [16]. In this paper, following previous works in [37], we set λ_i as the examination probability P_i at the i^{th} rank when computing DCG . Furthermore, based on DCG , we can measure multiple ranklists with the cumulative NDCG,

$$eff.@k_c = cum-NDCG@k_c(q, t) = \sum_{\tau=1}^t \gamma^{t-\tau} \frac{DCG@k_c(\pi_\tau)}{DCG@k_c(\pi^*)} \quad (5)$$

where $0 \leq \gamma \leq 1$ is a constant discount factor and t is the current time step. π^* is the ideal ranklist constructed by ranking items according to true relevance, and we use $DCG@k_c(\pi^*)$, referred to as $IDCG$, to normalize $DCG@k_c(\pi_\tau)$. By ignoring γ , we can get the average NDCG as,

$$\begin{aligned} aver-NDCG@k_c &= \frac{\sum_{\tau=1}^t DCG@k_c(\pi_\tau)}{t DCG@k_c(\pi^*)} \\ &= \frac{\sum_{d \in D(q)} R(d) (\sum_{\tau=1}^t \sum_{j=1}^{k_c} P_j \mathbb{1}_{\pi_i[j]=d})}{t DCG@k_c(\pi^*)} \\ &= \frac{\sum_{d \in D(q)} R(d) E^t@k_c(d)}{t DCG@k_c(\pi^*)} \end{aligned} \quad (6)$$

where $E^t@k_c(d)$ is the cumulative exposure at top k_c ranks,

$$E^t@k_c(d) = \sum_{\tau=1}^t \sum_{j=1}^{k_c} P_j \mathbb{1}_{\pi_i[j]=d} \quad (7)$$

where $\pi_i[j]$ indicates the j^{th} item in ranklist π_i , $\mathbb{1}$ is an indicator function which means we only accumulate item d 's exposure. To

simplify notations, we use $E^t(d)$ to denote $E^t@k_s(d)$ and $eff.@k_s$, where k_s is the ranklist length introduced in Eq. 3.

The Provider-side Utility (Fairness): As items' rankings can have significant effects on their providers' profit, it is important to create a fair ranking environment. To evaluate whether exposure is fairly allocated to users, we use the negative exposure disparity between item pairs as fairness measurement [26],

$$\begin{aligned} unfair.(q, t) &= \frac{1}{n(n-1)} \sum_{d_x \in D(q)} \sum_{d_y \in D(q)} \left(E^t(d_x) R(d_y) - E^t(d_y) R(d_x) \right)^2 \\ fair.(q, t) &= -unfair.(q, t) \end{aligned} \quad \begin{matrix} (8) \\ (9) \end{matrix}$$

where $D(q)$ is the set of candidate items that will construct the ranklist of query q and $n = |D(q)|$. The intuition of the above fairness measurement is that the optimal fairness can be achieved when items exposure is proportional to their relevance, i.e., $\forall d_x, d_y \in D(q)$, $\frac{E^t(d_x)}{R(d_x)} = \frac{E^t(d_y)}{R(d_y)}$. In other words, exposure fairness means we should let items of similar relevance get similar exposure. In this paper, we choose the exposure fairness evaluation proposed by Oosterhuis [26] instead of the original evaluation proposed in [37]. The reason for the choice is that the fairness evaluation in [37] needs to divide the exposure of an item by its relevance, i.e., $\frac{E^t(d)}{R(d)}$, which has zero denominator problem when item d is irrelevant, and $R(d)$ is near zero. This paper uses average unfairness across different queries to evaluate a ranking algorithm. We also refer to the average unfairness as the unfairness tolerance.

4 PROPOSED METHOD

Most existing fair algorithms are greedy algorithms, i.e., they sequentially construct the locally optimal ranklist for the next immediate session. Therefore they usually fail to optimize the construction procedure if we expect multiple sessions will come for the same query in the future. To mitigate this gap and reach a global optimal for a query, we propose to (i) **plan ahead and precompute multiple ranklists for future use** (§ 4.1) and (ii) **jointly optimize those ranklists together to maximize both fairness and ranking relevance** (§ 4.2 & § 4.3). We hypothesize that jointly optimizing multiple ranklists can construct better ranklists compared to sequentially greedily optimizing one single ranklist at each time step. Such hypothesis is verified by both the theoretical analysis in § 5 and the empirical results in § 6.2.

4.1 Future-aware Ranking Objective

We first propose a ranking fairness objective to plan and optimize the future ΔT ranklists for a query q . Specifically, when we are at time step $t + 1$, the objective is to pre-compute the optimal ranklists $\mathcal{B}^* = [\pi_{t+1}, \dots, \pi_{t+\Delta T}]$ that can maximize the marginal fairness $\Delta fair.(q, t, t + \Delta T)$,

$$\mathcal{B}^* = \arg \max_{\mathcal{B}=[\pi_{t+1}, \dots, \pi_{t+\Delta T}]} \Delta \widehat{fair}.(q, t, t + \Delta T) \quad (10)$$

$\Delta \widehat{fair}.(q, t, t + \Delta T) = \widehat{fair}.(q, t + \Delta T) - \widehat{fair}.(q, t)$, and \widehat{fair} is the estimated fairness. \widehat{fair} is calculated with Eq. 9 by substituting true relevance R with the estimated relevance, denoted as \hat{R} , since true relevance is mostly not available during optimization. Here maximizing the marginal fairness $\Delta \widehat{fair}.(q, t, t + \Delta T)$ is equivalent

to maximizing final fairness $\widehat{fair}(q, t + \Delta T)$ at time step $t + \Delta T$. The reason for the equivalence is that ranklists $\mathcal{B} = [\pi_{t+1}, \dots, \pi_{t+\Delta T}]$ can only influence the marginal fairness from timestep $t + 1$ to timestep $t + \Delta T$ rather than fairness before timestep t . Here, fairness evaluation in Eq. 9 is a direct objective in our method.

To the best of our knowledge, there is no trivial algorithm to get the optimal ranklists \mathcal{B}^* due to the discontinuity of ranking problem [21]. One example of discontinuity is that increasing an item's ranking score may not change the output ranklist, and fairness stays the same unless the increased score can surpass another item's score, and fairness will experience a sudden change. To alleviate the discontinuity of $\mathcal{B}^* \leftarrow \arg\max_{\mathcal{B}} \widehat{fair}$, we propose a novel two-phase solution path by introducing a continuous variable, ΔE ,

$$\mathcal{B}^* \xleftarrow[\text{phase2}]{\text{Vertical Allocation}} \Delta E^*(d) \forall d \in D(q) \xleftarrow[\text{phase1}]{\arg\max_{\Delta E}} \widehat{fair}. \quad (11)$$

where $\Delta E(d)$, also referred to as the planning exposure, is the marginal (or incremental) exposure we plan to assign to item d within the next ΔT timesteps. $\Delta E^*(d)$ is the optimal marginal exposure. We found that introducing ΔE helps to effectively maximize \widehat{fair} in phase 1 of our solution. In phase 2, we construct the optimal ranklists \mathcal{B}^* by allocating the optimal exposure ΔE^* to each item with a vertical allocation method (more details are in §4.3).

To get ΔE^* , we carry out a Taylor series expansion to investigate how future exposure will influence the fairness objective,

$$\begin{aligned} \widehat{fair}(q, t, t + \Delta T) &= \sum_{d \in D(q)} \frac{\partial \widehat{fair}}{\partial E(d)} \Delta E(d) \\ &+ \frac{1}{2} \sum_{d_x \in D(q)} \sum_{d_y \in D(q)} \frac{\partial^2 \widehat{fair}}{\partial E(d_x) \partial E(d_y)} \Delta E(d_x) \Delta E(d_y) \\ &= \sum_{d \in D(q)} \widehat{G}(d) \Delta E(d) \\ &- \frac{1}{2} \sum_{d_x \in D(q)} \sum_{d_y \in D(q)} \widehat{H}(d_x, d_y) \Delta E(d_x) \Delta E(d_y) \\ &= \vec{G} \cdot \Delta \vec{E} - \frac{1}{2} \Delta \vec{E}^T \cdot \widehat{H} \cdot \Delta \vec{E} \end{aligned} \quad (12)$$

where $\Delta E(d) = E^{t+\Delta T}(d) - E^t(d)$, the exposure increments. \widehat{G} and \widehat{H} are the first and the second order derivative, i.e., the gradient vector and the Hessian matrix, respectively,

$$\begin{aligned} \widehat{G}(d) &= \frac{4}{n(n-1)} \left(\widehat{R}(d) \sum_l E^t(l) \widehat{R}(l) - E^t(d) \sum_h \widehat{R}^2(h) \right) \\ \widehat{H}(d_x, d_y) &= \frac{4}{n(n-1)} \left(\sum_{d \in D(q)} \widehat{R}^2(d) \mathbb{1}_{x=y} - \widehat{R}(d_x) \widehat{R}(d_y) \right) \end{aligned}$$

By observing Eq. 8, we could derive two facts about the above second-order expansion in Eq. 12. (i) The above second-order expansion is not an approximation, but equality since (un)fairness in Eq. 8 is defined as a polynomial of E with a degree of two, and its derivative of order higher than two is zero. (ii) Since (un)fairness in Eq. 8 is defined as a sum of squares, the second order derivative \widehat{H} is semi-definite. Being equality, Eq. 12 allows us to correctly estimate future fairness given marginal exposure ΔE even when we consider a long-term future (large ΔT and ΔE). Based on the

correct future fairness estimation, it is possible to find the optimal marginal exposure planning, denoted as ΔE^* , that can maximize future fairness. Since \widehat{H} is semi-definite, Quadratic Programming-based (QP) optimization is valid to find ΔE^* . We give the specific QP problem formulation to find ΔE^* in § 4.2 and leave constructing optimal ranklists \mathcal{B}^* from ΔE^* in § 4.3.

4.2 Phase 1: Future Exposure Planning

When giving the QP problem formulation, we noticed that existing ranking fairness optimization usually considers two settings: (i) **the post-processing setting** [4, 37, 38] where relevance is assumed to be known or well estimated in advance; and (ii) **the online setting** [24, 48] where fairness is optimized while relevance is still being learned. To consider both settings, we first illustrate the QP problem formulation in the post-processing setting in § 4.2.1 and then extend it to work in the online setting in § 4.2.2.

4.2.1 The post-processing setting. To get the optimal exposure planning ΔE^* , we propose the following QP formulation with $\Delta E(d) \forall d \in D(q)$ as decision variables,

$$\max_{\Delta E} \widehat{fair}(q, t, t + \Delta T) \quad (13a)$$

$$\text{s.t.} \quad \sum_{d \in D(q)} \Delta E(d) = \sum_{i=1}^{\Delta T} \sum_{j=1}^{k_s} P_j \quad (13b)$$

$$\sum_{d \in D(q)} \Delta E(d) \widehat{R}(d) \geq (1 - \alpha) \sum_{i=1}^{\Delta T} \sum_{j=1}^{k_s} P_j \widehat{R}(d_{order_j}) \quad (13c)$$

$$\Delta E(d) \geq 0, \forall d \in D(q) \quad (13d)$$

$$\Delta E(d) \leq \sum_{i=1}^{\Delta T} P_1, \forall d \in D(q) \quad (13e)$$

where k_s is the length of ranklists, P_j is the examining probability at rank j . Eq.13b indicates that the sum of items' marginal exposure should equal the sum of the ΔT ranklists' exposure. In Eq.13c, we introduce the *NDCG* constraint, where $\widehat{R}(d_{order_j})$ means the j^{th} largest estimated relevance. According to Eq.(4&6), $\sum_{d \in D(q)} \Delta E(d) \widehat{R}(d)$ indicates the *DCG* and $\sum_{i=1}^{\Delta T} \sum_{j=1}^{k_s} P_j \widehat{R}(d_{order_j})$ indicates *IDCG*. Then, it is straightforward that $(1 - \alpha)$ indicates the minimum *NDCG* requirement we want to guarantee. In the post-processing setting, relevance is assumed to be given or already well-estimated prior to ranking optimization. Eq.13d indicates that an item's marginal exposure $\Delta E(d)$ should be no less than 0. In Eq.13e, $\Delta E(d)$ should be more than the accumulation of the first rank's exposure in the ΔT ranklists because an item should be unique in a ranklist and there are ΔT ranklists under consideration. Here we assume that the first rank's exposure is the largest and exposure drops from the top to the lower ranks.

4.2.2 The online setting. In the online setting, ranklists are optimized while relevance is still being learned. How to actively explore items and get more accurate relevance for ranking optimization is critical. Yang et al. [50] show that a more accurate relevance estimation for an item can be achieved by exposing an item more because more exposure leads to more interaction with users. Based on this, we do explorations by setting a minimum exposure requirement

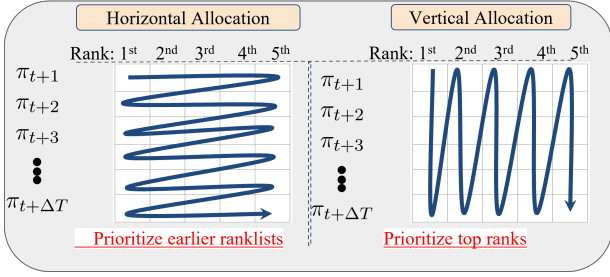


Figure 1: The ranklist construction order of the horizontal allocation and vertical allocation.

for items and propose the following QP formulation,

$$\max_{\Delta E} \quad \Delta \widehat{fair}.(q, t, t + \Delta T) - \beta \sum_{d \in D(q)} s(d) \quad (14a)$$

$$s.t. \quad Eq. (13b, 13c, 13d, 13e) \quad (14b)$$

$$s(d) \geq 0, \forall d \in D(q) \quad (14c)$$

$$s(d) + \Delta E(d) + E^t(d) \geq E_{min}, \forall d \in D(q) \quad (14d)$$

where β indicates the importance of exploration, $s(d) \forall d \in D(q)$ are slack variables to encourage exploration. In other words, both $s(d)$ and $\Delta E(d) \forall d \in D(q)$ are decision variables in the setting. In Eq. 14d, E_{min} is the minimum exposure requirement, $E^t(d)$ is item d 's exposure accumulated till time step t , and $\Delta E(d)$ is the marginal exposure we plan to allocate to item d within the next ΔT steps. $s(d)$ can be interpreted as the additional exposure still needed to satisfy the minimum exposure requirement after the next ΔT steps. When $E^t(d) \geq E_{min}$ for item d , i.e., minimum exposure requirement is already satisfied for item d , it is straightforward that $s(d)$ will be 0 and will not contribute to the ranking objective in Eq. 14a. When $E^t(d) < E_{min}$, i.e. item d does not meet the minimum exposure requirement, the objective in Eq. 14a will try to minimize $s(d)$. In other words, $\Delta E(d)$ will be boosted in order to satisfy the constraint in Eq. 14d. With more exposure, item d will be explored more. In this paper, we refer to the introduction of s as **Exploration**. And we treat β in Eq. 14a and E_{min} as hyper-parameters to control the degree of exploration.

As quadratic programming has been well studied, there are many available existing solvers. In this paper, we use quadratic programming library qpsolvers² within python to solve Equation 13 and Equation 14 to get the optimal exposure planning ΔE^* .

4.3 Phase 2: Ranklists Construction

Following the solution path in Eq. 11, the next step is to construct the optimal ranklists \mathcal{B}^* according to ΔE^* as ΔE^* has been solved in Phase 1. Here, we should allocate each item exactly its optimal exposure ΔE^* within \mathcal{B}^* . However, we find that the allocation solution of \mathcal{B} is not unique, and they share the same aver-NDCG@ k_s (see Theorem 5.3). Therefore, we additionally aim to find the optimal \mathcal{B}^* that can optimize **all top ranks' effectiveness**, i.e., aver-NDCG@ k_c , $\forall k_c \leq k_s$. Optimizing top ranks' effectiveness is important since users usually pay more attention to top ranks.

Inspired by [51], we propose a vertical exposure allocation method in Algorithm 1 to construct the optimal \mathcal{B}^* based on ΔE^* . The difference between a vertical allocation and a horizontal allocation

Algorithm 1: Vertical Allocation

```

1 Input: The optimal exposure planning  $\Delta E$ , the number of
   planning sessions to consider  $\Delta T$ , the ranked list length  $k_s$ ,
   relevance estimation  $\hat{R}$ ;
2 Initialize: ranking lists  $\mathcal{B}^*(i) \leftarrow [\emptyset]$  for  $i \in \text{range}(\Delta T)$ , set
    $\tilde{E}(d) \leftarrow 0 \quad \forall d \in D$ ;
3 for  $rnk \in [1, 2, \dots, k_s]$  do
4   for  $sess \in [1, 2, \dots, \Delta T]$  do
5      $Set1 \leftarrow \{d | \Delta E(d) - \tilde{E}(d) \geq P_{rnk}\}$ ;
6      $Set2 \leftarrow \{d | d \notin \mathcal{B}^*(sess)\}$ ;
7     if  $Set1 \cap Set2 = \emptyset$  then
8        $Candidates \leftarrow Set2$ 
9     else
10       $Candidates \leftarrow Set1 \cap Set2$ 
11       $d^* \leftarrow \underset{d \in Candidates}{\operatorname{argmax}} \hat{R}(d)$ ;
12       $\mathcal{B}^*(sess).append(d^*)$ ;
13       $\tilde{E}(d^*) \leftarrow \tilde{E}(d^*) + P_{rnk}$ ;
14 Output:  $\mathcal{B}^*$ ;

```

is the ranklist construction order. As shown in Fig. 1, a horizontal allocation prioritizes earlier ranklists and first fills out all ranks of the i^{th} ranklist π_i before filling out π_{i+1} . However, a vertical allocation prioritizes top ranks and fills out the i^{th} ranked items of all ranklists before filling out any $(i+1)^{th}$ ranked item. Since top ranks are usually more important, our proposed Algorithm 1 adopts a vertical allocation to fill out \mathcal{B}^* . In our proposed Algorithm 1, \mathcal{B}^* is the generated ΔT ranklists and $\mathcal{B}^*(sess, rnk)$ denote the rnk^{th} rank of the $sess^{th}$ ranklist. To fill out $\mathcal{B}^*(sess, rnk)$, the proposed vertical allocation first generates a feasible candidate set, i.e., $Candidates$, which contains items that are not selected for this session before ($d \notin \mathcal{B}(sess)$) but still have planned exposure left ($\Delta E(d) - \tilde{E}(d) \geq P_{rnk}$). Here, examination probability P_{rnk} serves as the margin, $\tilde{E}(d)$ stores the actual exposure item d receives. Within $Candidates$, our algorithm selects the most relevant item from the candidate set to fill out $\mathcal{B}^*(sess, rnk)$. Algorithm 1 is theoretically justified to accurately allocate exposure ΔE^* (see Theorem 5.1) and can construct optimal \mathcal{B}^* for aver-NDCG@ k_c $\forall k_c \leq k_s$ (see Theorem 5.2).

Although inspired by the vertical method in [51], the proposed vertical allocation is different from it. Yang et al. [51] focus on a certain share of exposure to be guaranteed and have a complicated 3-step procedure, i.e., allocation, appending, and resorting, which cannot be used to allocate ΔE^* for our problem. The proposed allocation algorithm in this paper uses up all ΔE^* , and the allocation procedure is less complicated and more straightforward than those introduced by Yang et al. [51].

4.4 FARA: Future-Aware Ranking Algorithm

Combining Phase 1 and Phase 2, we propose a future-aware ranking algorithm for fairness optimization, **FARA**, detailed in Algorithm 2. FARA serves users in an online manner where we pre-compute \mathcal{B}^* , the next ΔT ranklists, and randomly pop out one ranklist from \mathcal{B}^* when needed. When \mathcal{B}^* is used up and empty, We will re-compute

²<https://pypi.org/project/qpsolvers/>

Algorithm 2: FARA: Future-aware Ranking Algorithm

```

1 Input: The number of planning sessions to consider  $\Delta T$ ,
  fairness-effectiveness tradeoff parameter  $\alpha$ . And in the
  online setting, we need to additionally give the exploration
  parameters  $\beta$  and  $E_{min}$  (see Eq. 14);
2 Initialize time step  $t \leftarrow 0$ , initialize an empty dictionary
   $\mathbb{B} = \{\}$  to store ranked lists, items' exposure  $E \leftarrow 0$  and
  cumulative click  $cumC \leftarrow 0$ ;
3 while True do
4    $t \leftarrow t + 1$ ;
5   A user issues a query  $q_t$ ;
6   if  $q_t \notin \mathbb{B}$  then
7      $\mathbb{B}[q_t] = []$ , i.e., add an empty list
8   if  $\mathbb{B}[q_t]$  is empty then
9     Get  $\Delta E^*$  by solving Eq. 13 (post-processing) or
      Eq. 14 (online);
10    Get  $\mathcal{B}^*$  with Algorithm 1;
11    Randomly shuffle  $\mathcal{B}$ ;
12     $\mathbb{B}[q_t] \leftarrow \mathcal{B}^*$ ;
13  Pop out a ranking list from  $\mathbb{B}[q_t]$  and present it;
14  Update cumulative click  $cumC$  and items' exposure  $E$ ;
15  Update the relevance estimation  $\hat{R}$  via Eq. 15;

```

\mathcal{B}^* for future ΔT timesteps. $\mathbb{B}[q]$ is used to store \mathcal{B}^* for query q . Besides, FARA does not depend on any specific relevance estimation model, therefore, can be seamlessly integrated into most existing ranking applications. In this paper, we follow works by [50] to use the following unbiased estimator of relevance,

$$\hat{R}(d) = \frac{cumC^t(d)}{E^t(d)} \quad (15)$$

where $cumC^t(d) = \sum_{i=1}^t \sum_{j=1}^{k_s} C_{i,j} \mathbb{1}_{\pi_i[j]=d}$ is the cumulative clicks³. Moreover, it is worth noting that the relevance estimator can be replaced with other relevance estimators as well.

5 THEORETICAL ANALYSIS

THEOREM 5.1. *Algorithm 1 can theoretically guarantee $\Delta E(d) - \tilde{E}(d) \leq P_{k_s}$ for at least $|D| - k_s$ items.*

PROOF. Here we discuss the exposure allocation error bounds in Phase 2 of FARA, i.e., $|\tilde{E}(d) - \Delta E(d)|$, in Algorithm 1. We noticed that there are two possible scenarios of lines 6-12 in Algorithm 1:

Scenario 1: There exists a $(\text{rank}^*, \text{sess}^*)$ pair where $\text{Set1} \cap \text{Set2} = \emptyset$. If this happens, (k_s, sess^*) will also have $\text{Set1} \cap \text{Set2} = \emptyset$ since the size of Set1 and Set2 monotonically decrease for lower rank of the same session. As Set2 is the set of unselected items for a session, we know that there are at least $|D| - k_s$ items in Set2 , i.e., $|\text{Set2}| \geq |D| - k_s$. If $\text{Set1} \cap \text{Set2} = \emptyset$, those $|D| - k_s$ items are not in Set1 . In other words, there are at least $|D| - k_s$ items that satisfy $\Delta E(d) - \tilde{E}(d) < P_{k_s}$.

Scenario 2: $\text{Set1} \cap \text{Set2} \neq \emptyset \forall (\text{rank}, \text{sess})$ pair. Due to the margin P_{rank} in line 5 of Algorithm 1, $\Delta E(d) \geq \tilde{E}(d) \forall d \in D$ should always hold if $\text{Set1} \cap \text{Set2} = \emptyset$ never happens. By considering $\Delta E(d) \geq$

³Given the page limit, we skip the proof of unbiasedness for the above estimator and refer interested readers to [50]

Table 2: Datasets statistics. For each dataset, the table below shows the number of queries, the average number of docs for each query, and the relevance annotation y 's range.

Datasets	#Queries	#Aver. Docs per Query	y 's range
MQ2008	800	20	0 – 2
MSLR-10k	10k	122	0 – 4
Istella-S	33k	103	0 – 4

$\tilde{E}(d) \forall d \in D$ and the identity $\sum_{d \in D(q)} \Delta E(d) \equiv \sum_{d \in D(q)} \tilde{E}(d)$, we would know that $\Delta E(d) = \tilde{E}(d) \forall d \in D$, which means exposures are perfectly allocated according to $\Delta E(d)$.

Combing the two scenarios, the vertical allocation in Algorithm 1 can theoretically guarantee $\Delta E(d) - \tilde{E}(d) \leq P_{k_s}$ for at least $|D| - k_s$ items. Since $\Delta E(d) \gg P_{k_s}$ and $|D| \gg k_s$ when using FARA, we can claim that Algorithm 1 correctly allocates exposure. \square

THEOREM 5.2. *FARA can reach the optimal NDCG with the given exposure planning.*

PROOF. Here we provide theoretical proof that vertical allocation, i.e., phase 2, can optimize effectiveness (*aver-NDCG*) when exposure planning ΔE is given. Specifically, maximizing *aver-NDCG*@ k_c in Eq. 6 is equivalent to

$$\max \sum_{d \in D(q)} R(d) E @ k_c(d) \quad (16a)$$

$$\text{s.t.} \quad \sum_{d \in D(q)} E @ k_c(d) = \text{Const.} \quad (16b)$$

$$0 \leq E @ k_c(d) \leq \Delta E(d) \quad \forall d \in D \quad (16c)$$

where normalization is ignored in Eq. 16a, the sum of top ranks exposure should be a constant in Eq. 16b, and the top ranks exposure should be less than the total exposure planning in Eq. 16c. According to Rearrangement Inequality [14], it is straightforward to know that *aver-NDCG*@ k_c , i.e., Eq. 16a, can be optimized by letting item of greater relevance R get more exposure at top ranks, i.e., greater $E @ k_c$. In other words, we should prioritize letting items of greater relevance R fulfill their exposure planning ΔE at top k_c ranks since $E @ k_c$ is bounded in $[0, \Delta E]$. By assuming that *aver-NDCG* at higher ranks is more important [34], we should maximize *aver-NDCG*@ k_c before maximizing *aver-NDCG*@ $(k_c + 1)$, $\forall 1 \leq k_c < k_s$. As we maximize *aver-NDCG* from top to lower ranks, it is straightforward that the optimal way is to follow a greedy selection strategy to let an item of greater relevance R fulfill its exposure planning ΔE at its highest possible ranks. In Algorithm 1, the proposed vertical allocation exactly follows the above greedy selection strategy to let item of greater relevance R (line 11 in Algorithm 1) fulfill its exposure planning ΔE at the highest possible ranks (setting rank loops as the outer loop in 3 and line 4 of Algorithm 1). So it can reach optimal effectiveness at the top ranks. \square

THEOREM 5.3. *Effectiveness and fairness are fixed when ΔE is fixed.*

PROOF. Given the same exposure planning ΔE , effectiveness (*aver-NDCG*@ k_s) and fairness are fixed since we can substitute exposure planning $\Delta E(d)$ for $E^t @ k_c(d)$ in Eq. 6 and substituting ΔE for E in Eq. 8, respectively. In other words, for any ranklist \mathcal{B}^* , as long as exposure planning ΔE can be accurately allocated in \mathcal{B}^* , the effectiveness and fairness are fixed. \square

6 EXPERIMENTS

6.1 Experimental setup

Datasets: In this work, we use three public Learning-to-Rank (LTR) datasets: MQ2008 [30], MSLR10k⁴ and Istella-S [22]. Datasets' statistics are shown in Table 2. MQ2008 has a three-level relevance judgment (from 0 to 2). MSLR10k and Istella-S have a five-level relevance judgment (from 0 to 4). Queries in each dataset are already divided into training, validation, and test partitions according to a 60%-20%-20% scheme. In this work, we mainly focus on comparison within the LTR tasks. However, the proposed method can be adapted to recommendation tasks, which we leave for future studies.

Baselines: In this paper, we compare the following methods:

- **TopK:** Sort items according to $\hat{R}(d)$
- **RandomK:** Randomly rank items.
- **FairCo** [24]: Fair ranking algorithm based on a proportional controller. $\alpha \in [0.0, 1000.0]$
- **MCFair** [52]: Fair ranking algorithm directly uses gradient as the ranking score. $\alpha \in [0.0, 1000.0]$
- **ILP** [4]: Fair ranking algorithm based on Integer Linear Programming (ILP). $\alpha \in [0.0, 1.0]$
- **LP** [37]: Fair ranking algorithm based on Linear Programming (LP). $\alpha \in [0.0, 1000.0]$
- **MMF** [48]: Similar to FairCo but focus on top ranks fairness. $\alpha \in [0.0, 1.0]$
- **PLFair** [26]: A fair ranking algorithm based on Plackett-Luce optimization. $\alpha \in [0.0, 1.0]$
- **FARA-Horiz.** (ours): A variant of FARA. Compared to FARA, we switch line 3 and line 4 in Algorithm 1 to first iterate the sessions and then iterate the ranks. We refer to the iterations as the **horizontal allocation** paradigm. $\alpha \in [0.0, 1.0]$
- **FARA** (ours): The proposed fair ranking algorithm. $\alpha \in [0.0, 1.0]$.

Among the above ranking algorithm, TopK and RandomK are unfair algorithms, while the others are fair algorithms. While all the fair ranking algorithms aim to maximize effectiveness and fairness, FARA and FARA-Horiz. differ from others by taking a joint optimization across multiple ranklists rather than a traditional greedy optimization approach. For fair ranking algorithms, there exists a tradeoff parameter α , similar to α in Eq. 13, to balance effectiveness and fairness. For fair algorithms, the greater α is, the more we care about fairness while potentially sacrificing more effectiveness. For example, when increasing α in Eq. 13c, FARA can maximize fairness with less effectiveness constraint. For different fair algorithms, α lies in different ranges. For FairCo, MCFair, LP, α are originally within $[0.0, +\infty]$, and we adopt $\alpha \in [0.0, 1000.0]$ which is enough according to our experiments. For ILP, MMF, PLFair, FARA-Horiz. and FARA, $\alpha \in [0.0, 1.0]$. Although the vertical allocation in Algorithm 1 was inspired by [51], [51] cannot be used as a baseline because [51] works with offline ranking services where all user queries are known in advance. However, in this paper, we consider the online services depicted in Algorithm 2.

Ranking Service Simulation: Following the workflow in Algorithm 2, at each time step, a simulated user will issue a query q , which is randomly sampled from the training, validation, or test partition. Corresponding to the query q , a ranking algorithm will

construct a ranklist π of candidate items and present it to the simulated user. To collect users' feedback for the ranked list π , we need to simulate relevance and examination (see Eq.2). Same as [2], the relevance probabilities of each document-query pair (d, q) are simulated with their relevance judgement y as

$$P(r = 1|d, q) = \epsilon + (1 - \epsilon) \frac{2^y - 1}{2^{y_{\max}} - 1}$$

where y_{\max} is the maximum value of relevance judgement y , i.e., 2 or 4 depending on the datasets. Besides relevance, following [24, 28], we simulate users' examination probability as,

$$P(e = 1|d, \pi) = \begin{cases} \frac{1}{\log_2(\text{rank}(d|\pi)+1)}, & \text{if } \text{rank}(d|\pi) \leq k_s \\ 0, & \text{otherwise} \end{cases}$$

For simplicity, we only simulate users' examination behavior on top ranks, and we set k_s to 5 throughout the experiments (refer to Eq. 3 for more details of k_s). With $P(r = 1|d, q, \pi)$ and $P(e = 1|d, \pi)$, we sample clicks with Equation 2. The advantage of the simulation is that it allows us to do online experiments on a large scale while still being easy to reproduce by researchers without access to live ranking systems [28]. For simplicity, same as existing works [24, 28, 48, 50], we assume that users' examination $P(e = 1|d, \pi)$ is known in experiment since many existing works [1, 2, 31, 45] have been proposed to estimate it. Due to different data sizes, we simulate 200k steps for MQ2008 and 4M steps for MSLR10k and Istella-S.

Experiment Settings: We noticed that LP and ILP methods are proposed in the **post-processing setting**, where relevance is already known or well estimated in advance. However, in most real-world settings, ranking optimization and relevance learning are carried out at the same time, which we refer to as the **online setting**. To give a comprehensive comparison, we evaluate ranking methods in both settings. In the post-processing setting, all the ranking methods in Section 6.1 are based on true relevance R , and FARA will set β as 0. In the online setting, all the ranking methods in Section 6.1 are based on the relevance estimation \hat{R} in Eq. 15 to perform ranking optimization. FARA set β to 1 and $E_{\min} = 10$ unless otherwise explicitly specified, as they work well across all our experiments.

Evaluation: We use the cum-NDCG (cNDCG) in Eq. 5 with $\gamma = 0.995$ (same γ adopted in [43, 44]) to evaluate the effectiveness at different cutoffs, $1 \leq k_c \leq 5$. Aside from effectiveness, unfairness defined in Eq. 8 is used for unfairness measurement. We run each experiment five times and report the average evaluation performance on the test partition. We use the Fisher randomization test [39] with $p < 0.05$ to do significant tests. Due to the time cost (see Table 3), we do not run ILP and LP on the larger datasets, MSLR10k and Istella-S, and the performances are not available (NA).

6.2 Results and Analysis

In this section, we first compare the ranking relevance performance given different degrees of fairness requirements. Then we dive deep into our method to offer more insights into FARA's supremacy.

6.2.1 Can FARA reach a better balance between fairness and effectiveness? In Figure 2, we compare ranking methods' effectiveness-fairness balance given different fairness requirements. To generate the balance curves in Figure 2, we incrementally sample α from the minimum value to the maximum value within α 's ranges indicated in Section 6.1. For each method, twenty α are sampled with the

⁴<https://www.microsoft.com/en-us/research/project/mslr/>

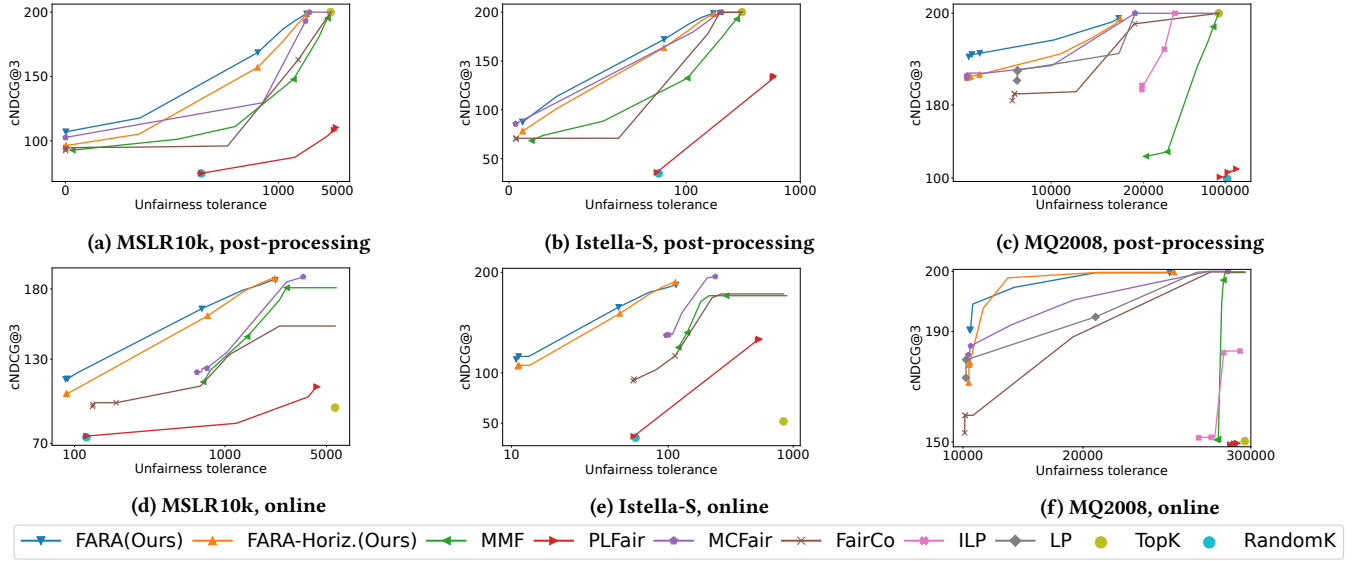


Figure 2: c-NDCG vs. unfairness tolerance (Eq. 8) in the post-processing setting and the online setting. Given the same unfairness, the higher curves or points lie, the better their performances are. Our methods FARA and FARA-Horiz. lie higher than all fair baselines in all figures. ILP and LP are unavailable for MSLR10k and Istella-S due to time costs (refer to Table 4).

Table 3: Comparison of cNDCG@{1,3,5} and unfairness tolerance in the post-processing setting. Significant improvements or degradations with respect to FairCo are indicated with +/- . Within fair algorithms, the best performance with statistical significance is **bolded and underlined**. Here, α is set to the maximum value (see Sec. 6.1 for α 's range) for each fair algorithm respectively, which means that all algorithms are trying their best to optimize ranking fairness (Eq. 9) and the numbers in the table represents their unfairness lower bound. Results are rounded to one decimal place.

Methods	MSLR-10k				Istella-S				MQ2008			
	cNDCG@1	cNDCG@3	cNDCG@5	unfair.	cN@1	cN@3	cN@5	unfair.	cN@1	cN@3	cN@5	unfair.
TopK	200.0 ⁺	200.0 ⁺	200.0 ⁺	4165.0 ⁻	200.0 ⁺	200.0 ⁺	200.0 ⁺	310.1 ⁻	200.0 ⁺	200.0 ⁺	200.0 ⁺	86001.1 ⁻
Randomk	68.0 ⁻	74.7 ⁻	79.7 ⁻	119.0 ⁻	30.2 ⁻	35.7 ⁻	41.1 ⁻	56.7 ⁻	74.0 ⁻	95.7 ⁻	114.6 ⁻	104632.2 ⁻
PLFair	68.2 ⁻	74.8 ⁻	79.9 ⁻	119.6 ⁻	31.8 ⁻	36.2 ⁻	41.5 ⁻	54.6 ⁻	79.2 ⁻	99.3 ⁻	117.2 ⁻	101245.1 ⁻
MMF	84.4	92.8	99.9	8.0 ⁻	62.2	68.8	80.1	6.6 ⁻	132.8 ⁻	162.3 ⁻	172.5 ⁻	20688.7 ⁻
ILP	NA	NA	NA	NA	NA	NA	NA	NA	185.6 ⁺	183.8	186.7	19916.3 ⁻
LP	NA	NA	NA	NA	NA	NA	NA	NA	188.4 ⁺	187.4 ⁺	187.9	9425.7
MCFair	114.8 ⁺	102.7 ⁺	101.0	0.0	113.7 ⁺	85.25 ⁺	81.3	0.4	193.5 ⁺	186.0 ⁺	186.6	9113.7
FairCo	85.5	93.7	100.8	0.0	63.3	69.9	80.4	0.5	179.0	182.0	187.4	9382.0
FARA-Horiz.(Ours)	90.7 ⁺	96.1 ⁺	100.4	0.0	78.5 ⁺	79.8 ⁺	82.6	0.9	187.3 ⁺	186.1 ⁺	187.0	9125.9
FARA(Ours)	<u>129.0⁺</u>	<u>107.0⁺</u>	99.7	0.0	<u>135.5⁺</u>	<u>89.1⁺</u>	82.6	0.9	<u>196.3⁺</u>	<u>190.9⁺</u>	186.8	9129.9

step size as $(\alpha_{max} - \alpha_{min})/20$. After sampling, we perform ranking simulation experiments for each α to get a (cNDCG, unfairness) pair. Then we connect different α 's (cNDCG, unfairness) pair to form a curve for each method respectively in Figure 2. All the curves start from the top right to the bottom left as α increases, which means there exists a tradeoff between fairness and effectiveness (cNDCG). The reason behind this tradeoff is that requiring more fairness will bring more constraints on optimizing effectiveness. Since TopK and RandomK do not have trade-off parameters, both of them only have one single pair of (cNDCG, unfairness), and their performances are shown as single points in Figure 2.

In Figure 2, our methods FARA and FARA-Horiz. outperform all other fair methods since our methods reach the best cNDCG given the same unfairness tolerance. And FARA's supremacy is consistent in both post-processing and online settings. All fair ranking algorithms are effective fair ranking algorithms since they all

show the tradeoff, i.e., higher cNDCG when increasing the unfairness tolerance. For unfair algorithms, TopK performs differently in post-processing and online settings. In the post-processing setting, TopK reaches the highest cNDCG since relevance is known, and ranking relevance is the only consideration. However, in the online setting, TopK can not reach the highest cNDCG. We think the drop in cNDCG is that TopK naively trusts the relevance estimation without any exploration when optimizing ranking relevance. However, fair algorithms are shown to be robust to the online setting since they mostly can reach better cNDCG than Topk when increasing unfairness tolerance. We think the reason for the robustness is that fair algorithms usually rerank items for different sessions to optimize fairness, and such reranking brings explorations.

6.2.2 What is the fairness upper bound that FARA can reach? In Table 3, lower unfairness means higher fairness capacity and fairness upper bound, i.e., the maximum possible fairness one algorithm can

Table 4: The average time (seconds per 1k ranklists) cost with standard deviations in parentheses. Since ILP and LP are time-consuming on large datasets, the time costs on MSLR-10k and Istella-S are estimated by only running 1k steps instead of the total simulation steps indicated in Sec. 6.1.

Algorithms	Datasets		
	MSLR10k	Istella	MQ2008
TopK	0.65 _(0.14)	0.50 _(0.00)	0.55 _(0.10)
Randomk	0.63 _(0.12)	0.57 _(0.04)	0.59 _(0.14)
PLFair	2.24 _(0.04)	3.11 _(0.07)	1.77 _(0.04)
MMF	8.01 _(0.39)	6.57 _(0.23)	1.82 _(0.28)
ILP	1208.90 _(85.80)	1102.30 _(75.20)	19.70 _(1.29)
LP	≥10 days	≥10 days	2.09 _(0.48)
McFair	0.724 _(0.016)	0.660 _(0.025)	0.567 _(0.035)
FairCo	0.73 _(0.04)	0.71 _(0.03)	0.70 _(0.12)
FARA-Horiz.(Ours)	1.00 _(0.17)	0.86 _(0.07)	0.97 _(0.22)
FARA(Ours)	0.91 _(0.07)	0.91 _(0.00)	0.97 _(0.30)

reach. Fair effective ranking algorithms, including FairCo, LP, FARA-Horiz. and FARA, have similar fairness capacity and outperform unfair ranking algorithms in terms of unfairness. The success of FARA-Horiz. and FARA validates the proposed quadratic programming formulation can optimize fairness. Similar cNDCG@5 and unfairness for those effective algorithms are expected according to Theorem 5.3. For other fair ranking algorithms, ILP and MMF, and PLFair show inferior fairness capacity. As for the possible reason, ILP uses the integer linear programming method, which may not be effective in optimizing fairness. MMF actually follows a slightly different definition of fairness which require fairness at any cutoff should be fair, which is more strict than the definition we use in this paper. As for PLFair, PLFair tries to learn the ranking score that optimizes fairness based on the feature representation (the exact setting in original paper [26]). However, the feature representation is initially designed for relevance which makes PLFair suboptimal for fairness optimization. In Table 3, ILP and LP are NA for MSLR10k and Istella-S due to time costs (refer to Table 4). Due to the page limit, we show the ranking performance of the online setting in Fig. (2), instead of in Table 3.

6.2.3 How is FARA’s effectiveness at different cutoffs? In Table 3, we show cNDCG at different cutoffs. Although FairCo, LP, FARA-Horiz. and FARA have similar fairness capacities, FARA significantly outperforms those fair algorithms for cNDCG@1 and cNDCG@3 on all three datasets. Compared to FARA-Horiz., shown in Table 3, FARA still significantly outperforms FARA-Horiz. at top ranks, which shows the necessity of vertical allocation.

6.2.4 How is FARA’s time efficiency? Besides fairness and effectiveness optimization, we also empirically compare the time efficiency. In Table 4, ILP and LP are really time-consuming, especially on large datasets, MSLR10k and Istella-S. Compared to ILP and LP, FARA is more than 1000× time efficient on MSLR10k and Istella-S, although all three of them are programming-based methods. There are two reasons behind FARA’s time efficiency. The first one is that FARA has a much fewer number of decision variables since FARA only has $O(n)$ decision variables, while ILP and LP have $O(n^2)$ decision variables. The second one is that FARA does not need to solve quadratic programming for every time step. By solving quadratic programming once, we can get ΔT ranklists used for




Figure 3: The numbers of planning session ΔT ’s influence on FARA in the post-processing setting on MQ2008. α is set as 1.




Figure 4: Ablation study of exploration in the online setting. The higher curves lie, the better their performances are.

future ΔT sessions. FARA can reach comparable time efficiency with non-programming-based algorithms like TopK, RandomK, and FairCo. Compared with those non-programming-based algorithms, the slightly additional time cost of FARA is acceptable given FARA’s superior ranking performance in Tab. 3 and in Fig. 2.

6.2.5 How does ΔT influence FARA? In Figure 3, we show the results of cNDCG and unfairness by varying the value of ΔT . With greater ΔT , we see a clear boost of cNDCG for FARA, while such a boost does not happen for FARA-Horiz. We know that the proposed vertical allocation is the key reason to have better-ranking relevance when we get the optimal exposure planning ΔE^* , and we theoretically analyze the reason in § 5.2. Besides, as we increase ΔT , unfairness does not vary much, and its value stays close to the minimum unfairness we can achieve in Table 3. We think the reason for the steady value of unfairness is that FARA already reaches the upper limit of fairness when ΔT is small, and it is hard to improve when we increase ΔT .

6.2.6 How does exploration influence FARA in the online setting? To study how the exploration part (the slack variables s in Eq. 14a) influences FARA, we did an ablation study for FARA with or without exploration. Due to the page limit, we only show the ablation results on the larger dataset, i.e., MSLR10k and Istella-S, in Figure 4. The advantage of exploration is two-folded based on Figure 4. Firstly, FARA lies higher than FARA-w/o-Exp. in the figure, which suggests exploration leads to a better effectiveness-fairness balance. Secondly, FARA has a smaller lower bound of unfairness tolerance, which implies exploration enables FARA to have a higher fairness capacity and can meet a more strict fairness requirement.

ACKNOWLEDGEMENTS

This work was supported in part by NSF CCF-2115677 and in part by the School of Computing, University of Utah. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

References

- [1] Aman Agarwal, Ivan Zaitsev, Xuanhui Wang, Cheng Li, Marc Najork, and Thorsten Joachims. 2019. Estimating position bias without intrusive interventions. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. 474–482.
- [2] Qingyao Ai, Keping Bi, Cheng Luo, Jiafeng Guo, and W Bruce Croft. 2018. Unbiased learning to rank with unbiased propensity estimation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 385–394.
- [3] Abolfazl Asudeh, HV Jagadish, Julia Stoyanovich, and Gautam Das. 2019. Designing fair ranking schemes. In *Proceedings of the 2019 international conference on management of data*. 1259–1276.
- [4] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*. 405–414.
- [5] Amin Bigdeli, Negar Arabzadeh, Shirin SeyedSalehi, Morteza Zihayat, and Ebrahim Bagheri. 2022. Gender Fairness in Information Retrieval Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3436–3439.
- [6] L Elisa Celis, Damian Straszak, and Nisheeth K Vishnoi. 2017. Ranking with fairness constraints. *arXiv preprint arXiv:1704.06840* (2017).
- [7] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 international conference on web search and data mining*. 87–94.
- [8] Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. 2020. Evaluating stochastic rankings with expected exposure. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 275–284.
- [9] Michael D Ekstrand, Graham McDonald, Amifa Raj, and Isaac Johnson. 2023. Overview of the TREC 2022 Fair Ranking Track. *arXiv preprint arXiv:2302.05558* (2023).
- [10] Ruoyuan Gao, Yingqiang Ge, and Chirag Shah. 2022. FAIR: Fairness-aware information retrieval evaluation. *Journal of the Association for Information Science and Technology* 73, 10 (2022), 1461–1473.
- [11] Ruoyuan Gao and Chirag Shah. 2021. Addressing bias and fairness in search systems. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*. 2643–2646.
- [12] Yingqiang Ge, Juntao Tan, Yan Zhu, Yinglong Xia, Jiebo Luo, Shuchang Liu, Zuohui Fu, Shijie Geng, Zelong Li, and Yongfeng Zhang. 2022. Explainable Fairness in Recommendation. *arXiv preprint arXiv:2204.11159* (2022).
- [13] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *Proceedings of the 25th acm sigkdd international conference on knowledge discovery & data mining*. 2221–2231.
- [14] Godfrey Harold Hardy, John Edensor Littlewood, George Pólya, György Pólya, et al. 1952. *Inequalities*. Cambridge university press.
- [15] Maria Heuss, Fatemeh Sarvi, and Maarten de Rijke. 2022. Fairness of Exposure in Light of Incomplete Exposure Estimation. *arXiv preprint arXiv:2205.12901* (2022).
- [16] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.
- [17] Thorsten Joachims. 2021. Fairness and Control of Exposure in Two-sided Markets. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval*. 1–1.
- [18] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2017. Accurately interpreting clickthrough data as implicit feedback. In *ACM SIGIR Forum*, Vol. 51. Acm New York, NY, USA, 4–11.
- [19] James Kotary, Ferdinando Fioretto, Pascal Van Hentenryck, and Ziwei Zhu. 2022. End-to-End Learning for Fair Ranking Systems. In *Proceedings of the ACM Web Conference 2022*. 3520–3530.
- [20] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. 2021. Towards personalized fairness based on causal notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1054–1063.
- [21] Tie-Yan Liu et al. 2009. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval* 3, 3 (2009), 225–331.
- [22] Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Fabrizio Silvestri, and Salvatore Trani. 2016. Post-learning optimization of tree ensembles for efficient ranking. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 949–952.
- [23] Masoud Mansoury. 2022. Understanding and mitigating multi-sided exposure bias in recommender systems. *ACM SIGWEB Newsletter Autumn* (2022), 1–4.
- [24] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling Fairness and Bias in Dynamic Learning-to-Rank. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (Virtual Event, China) (SIGIR '20)*. Association for Computing Machinery, New York, NY, USA, 429–438. <https://doi.org/10.1145/3397271.3401100>
- [25] Mohammadmehdi Naghiaei, Hossein A Rahmani, and Yashar Deldjoo. 2022. Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. *arXiv preprint arXiv:2204.08085* (2022).
- [26] Harrie Oosterhuis. 2021. Computationally Efficient Optimization of Plackett-Luce Ranking Models for Relevance and Fairness. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1023–1032.
- [27] Harrie Oosterhuis and Maarten de Rijke. 2020. Policy-aware unbiased learning to rank for top-k rankings. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 489–498.
- [28] Harrie Oosterhuis and Maarten de Rijke. 2021. Unifying online and counterfactual learning to rank: A novel counterfactual estimator that effectively utilizes online interventions. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 463–471.
- [29] Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. 2022. Fair ranking: a critical review, challenges, and future directions. *arXiv preprint arXiv:2201.12662* (2022).
- [30] Tao Qin and Tie-Yan Liu. 2013. Introducing LETOR 4.0 datasets. *arXiv preprint arXiv:1306.2597* (2013).
- [31] Filip Radlinski and Thorsten Joachims. 2006. Minimally invasive randomization for collecting unbiased preferences from clickthrough logs. In *Proceedings of the national conference on artificial intelligence*, Vol. 21. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 1406.
- [32] 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.
- [33] Matthew Richardson, Ewa Dominowska, and Robert Ragno. 2007. Predicting clicks: estimating the click-through rate for new ads. In *Proceedings of the 16th international conference on World Wide Web*. 521–530.
- [34] Stephen E Robertson. 1977. The probability ranking principle in IR. *Journal of documentation* (1977).
- [35] Yuta Saito and Thorsten Joachims. 2022. Fair Ranking as Fair Division: Impact-Based Individual Fairness in Ranking. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1514–1524.
- [36] Ashudeep Singh. 2021. *Fairness of Exposure for Ranking Systems*. Ph.D. Dissertation. Cornell University.
- [37] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2219–2228.
- [38] Ashudeep Singh and Thorsten Joachims. 2019. Policy learning for fairness in ranking. In *Advances in Neural Information Processing Systems*. 5426–5436.
- [39] Mark D Smucker, James Allan, and Ben Carterette. 2007. A comparison of statistical significance tests for information retrieval evaluation. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*. 623–632.
- [40] Anh Tran, Tao Yang, and Qingyao Ai. 2021. ULTRA: an unbiased learning to rank algorithm toolbox. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 4613–4622.
- [41] Nicolas Usunier, Virginie Do, and Elvis Dohmatob. 2022. Fast online ranking with fairness of exposure. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 2157–2167.
- [42] Ali Vardasbi, Fatemeh Sarvi, and Maarten de Rijke. 2022. Probabilistic Permutation Graph Search: Black-Box Optimization for Fairness in Ranking. *arXiv preprint arXiv:2204.13765* (2022).
- [43] Huazheng Wang, Sonwoo Kim, Eric McCord-Snoek, Qingyun Wu, and Hongning Wang. 2019. Variance reduction in gradient exploration for online learning to rank. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 835–844.
- [44] Huazheng Wang, Ramsey Langley, Sonwoo Kim, Eric McCord-Snoek, and Hongning Wang. 2018. Efficient exploration of gradient space for online learning to rank. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 145–154.
- [45] Xuanhui Wang, Nadav Golbandi, Michael Bendersky, Donald Metzler, and Marc Najork. 2018. Position bias estimation for unbiased learning to rank in personal search. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. 610–618.
- [46] Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. 2022. Joint multisided exposure fairness for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 703–714.
- [47] Yao Wu, Jian Cao, Guandong Xu, and Yudong Tan. 2021. Tfrom: A two-sided fairness-aware recommendation model for both customers and providers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1013–1022.
- [48] Tao Yang and Qingyao Ai. 2021. Maximizing Marginal Fairness for Dynamic Learning to Rank. In *Proceedings of the Web Conference 2021*. 137–145.

- [49] Tao Yang, Cuize Han, Chen Luo, Parth Gupta, Jeff M Phillips, and Qingyao Ai. 2023. Mitigating Exploitation Bias in Learning to Rank with an Uncertainty-aware Empirical Bayes Approach. *arXiv preprint arXiv:2305.16606* (2023).
- [50] Tao Yang, Chen Luo, Hanqing Lu, Parth Gupta, Bin Yin, and Qingyao Ai. 2022. Can clicks be both labels and features? Unbiased Behavior Feature Collection and Uncertainty-aware Learning to Rank. In *The 45th International ACM SIGIR Conference on Research & Development in Information Retrieval*.
- [51] Tao Yang, Zhichao Xu, and Qingyao Ai. 2022. Effective Exposure Amortizing for Fair Top-k Recommendation. *arXiv preprint arXiv:2204.03046* (2022).
- [52] Tao Yang, Zhichao Xu, Zhenduo Wang, Anh Tran, and Qingyao Ai. 2023. Marginal-Certainty-aware Fair Ranking Algorithm. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 24–32.
- [53] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. Fa* ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1569–1578.
- [54] Meike Zehlike and Carlos Castillo. 2020. Reducing disparate exposure in ranking: A learning to rank approach. In *Proceedings of The Web Conference 2020*. 2849–2855.
- [55] Meike Zehlike, Ke Yang, and Julia Stoyanovich. 2021. Fairness in ranking: A survey. *arXiv preprint arXiv:2103.14000* (2021).