A Unified Meta-learning Framework for Fair Ranking with Curriculum Learning

Yuan Wang, Zhiqiang Tao, Yi Fang

Abstract—In recent information retrieval systems, it is observed that the datasets used to train machine learning models can be biased, leading to systematic discrimination against certain demographic groups, which means the ranking utility of specific groups is often lower than others in a biased dataset. Training models on these datasets will further decrease the exposure of the minority groups. To address this problem, we propose a Meta Curriculum-based Fair Ranking framework (MCFR) which could alleviate the data bias issue through the weighted loss using gradient-based learning to learn. Specifically, we optimize a meta learner from a sampled dataset (*meta-dataset*), and meanwhile train a ranking model on the whole (*biased*) dataset. The meta-dataset is sampled with a curriculum learning scheduler to guide the meta learner's training to gradually mitigate the skewness towards biased attributes. The meta learner serves as a weighting function to make the ranking loss focus more on the minority group. We formulate the proposed MCFR as a bilevel optimization problem and solve it using gradients through gradients. Extensive experiments on real-world datasets demonstrate that our approach can be used as a generic framework to work with various ranking losses and fairness metrics.

Index Terms—Fairness-aware Search, Meta-learning, Learning-to-rank, Curriculum Learning.

1 Introduction

Fairness in search engines is an important topic, which focuses on training an unbiased ranking model towards protected attributes. Typically, when a user query is given, the ranking model predicts relevant scores among candidate items and returns items with the highest scores to users. The data-driven ranking model is usually trained with large datasets, and thus the ranker will learn user/item patterns from the training dataset and make predictions based on them. However, in many cases, the systematic biases such as exposure bias [1] in the dataset will cause unfairness to the ranking model. The historical discrimination against the socially underrepresented group [2] will make its way into the model as the pattern will be observed during the training process. Such an unfairness problem could be summarized as the disparate exposure [1], leading to a negative impact on many real-world ranking problems.

Disparate exposure is prevalent in information retrieval. For instance, expert search and job recommendation systems historically underrepresented minority groups like females and African Americans. Consequently, traditional learning to rank (LTR) models, such as ListNet [3], often rank these groups lower due to data biases. Fig. 1 shows ranking scores from different models on four datasets, highlighting this unfairness. Disparate exposure implies uneven group visibility in algorithm outcomes, especially linked to attributes like gender or race, distinct from biases like selection or conformity, which challenge algorithmic fairness and efficiency.

To reduce disparate exposure in a ranking context, many research works have been proposed recently by designing fairness-aware algorithms, which can be divided into two categories: 1) the score-based models and 2) the supervised-learning models. The score-based models [9], [10], [11], [12], [13], [14] compute the ranking scores on the fly for a given candidates list and return the sorted candidates as the model outcome. The supervised-learning models generally solve ranking as a prediction problem and focus on different mitigation strategies, such as the post- [15], in- [6], [16], [17], [18], [19], [20], [21], and pre-processing [22], [23], [24] in model training. Although the in-processing models have achieved promising performance on both fairness and ranking metrics, learning on biased datasets is still under-explored and challenging, due to the unbalanced distributions of protected attributes in the public training datasets.

One possible way to alleviate system discrimination inherited from data bias is dynamically re-weighting the minority groups to contribute more penalties in computing a ranking loss. To this end, meta-learning [25] emerges as an effective way to enable a learning-to-weight approach by leveraging a small, unbiased dataset – meta dataset. For the fairness-aware ranking problem, we propose to mitigate the exposure issue in the biased dataset by learning a weighting model (meta-learner) to re-weight the loss of the ranking model on the biased dataset. The meta-learner will be optimized on the meta dataset (unbiased), and the weighted loss on the training dataset (biased) will be used to optimize the ranking model. However, due to the distribution shift between the biased and unbiased datasets, it is non-trivial to directly train the meta-learner and base learner on these two datasets where a large training loss may impair ranking utility and burden convergence speed.

We propose to adopt curriculum learning to gradually increase the difficulty of training meta-learners to address the above challenge. Specifically, we define the difficulty as the exposure of the protected groups in a dataset. We first

- *Yi Fang is the corresponding author.*
- *Yuan Wang and Yi Fang are with the Department of Computer Science and Engineering, Santa Clara University, Santa Clara, CA.*
  E-mail: {ywang4, yfang}@scu.edu
- *Zhiqiang Tao is with the School of Information, GCCIS, Rochester Institute of Technology, Rochester, NY.*
  E-mail: zhiqiang.tao@rit.edu

Manuscript received 4 Jan. 2023; revised 20 Jan. 2024; accepted 2 Mar. 2024.

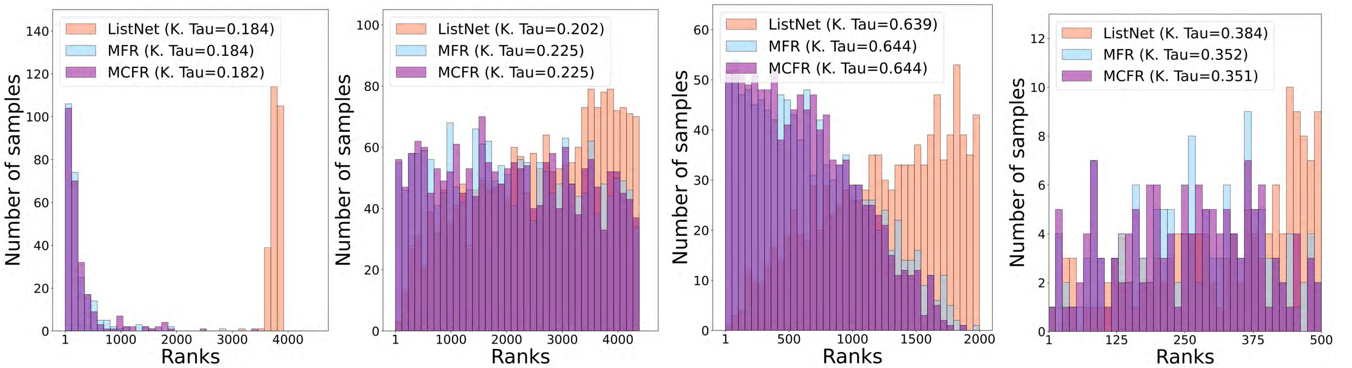

(a) African American (b) Female (c) African American (d) Female

Fig. 1: Illustration of the predicted rankings distribution of two protected attributes on four datasets – *(a) Law Student (gender)* [4], *(b) Law Student (race)* [4], *(c) COMPAS* [5], *and (d) Engineering Student* [6]. We report Kendall's Tau [7] as the ranking performance. MCFR and MFR [8] improve the protected attributes' ranking while realizing competitive ranking performance compared with ListNet [3], demonstrating that our approach could increase the exposure of the minority.

randomly sample a meta dataset that has the same exposure as the training dataset. Then, we continually increase the protected groups' exposure in the meta dataset by sampling more candidates from this group at each ongoing epoch until a uniform distribution (equal exposure) is achieved over sensitive attributes. Intuitively, this incremental concept learning [26] is a good fit to solve the distribution shift problem, because meta-learners are trained with samples from the biased dataset at the early epochs, which means there is less distribution shift between the meta-dataset and training dataset. The experimental results demonstrate the effectiveness of curriculum learning and the improved data efficiency during training.

In this study, we propose a unified meta-learning framework with curriculum learning to formulate the fairness-aware ranking task as a bilevel optimization problem where the upper level focuses on learning-to-weight to mitigate the biased exposure of protected attributes, and the lower level solves learning-to-rank with a dynamic loss governed by a meta learner. Specifically, we alleviate the data bias issue for the protected groups through an automatically weighted loss. The contributions of this work are as follows.

- We propose a novel Meta Curriculum-based Fair Ranking framework, namely (MCFR), which addresses the data bias by automatically re-weighting the ranking losses. The proposed MCFR is formulated as a bilevel optimization problem and solved using gradients through gradients.
- The proposed fair ranking algorithm marries in-processing methods with pre-processing techniques by seamlessly incorporating curriculum learning into the construction process of meta datasets.
- We develop MCFR as a general framework applicable to various ranking loss functions and fairness metrics. A systematic empirical study has been provided to show the versatility of the proposed framework over different ranking and fairness criteria.
- Experiments on public datasets show our method matches existing ranking performance and enhances fairness metrics. Additionally, evaluations confirm MCFR improves fairness with less training data and achieves comparable convergence times.

This paper is a substantial extension of our previous work [8]. The updated version integrates curriculum learning into the MFR model, offering the first fair ranking framework to utilize both pre-processing and in-processing methods. This new approach enhances the model's adaptability and robustness by allowing for a broader range of loss functions and dynamically adjusting meta-datasets during training. Additionally, our framework demonstrates data efficiency in comparative experiments. We've also conducted more comprehensive tests, incorporating additional baseline models and performing an ablation study on various fairness terms and ranking losses. Lastly, we've updated the manuscript to include more recent related works, providing a fuller understanding of fairness in ranking.

The rest of this paper is organized as follows. Section 2 introduces the related works. Section 3 elaborates the proposed Meta Curriculum-based Fair Ranking (MCFR) framework in detail. Section 4 shows the experimental setting, implementation details, evaluation results, and performance analysis. Section 5 finally concludes the study.

2 Related Work

The proposed MCFR is related to the general fairness issue in ranking, the meta-learning on fairness problems, and the curriculum learning. We discuss in detail the recent works in the following subsections from these research areas.

2.1 Fairness on Ranking

Zehlike et al. [27] categorized fair ranking models into score-based and supervised learning models. Score-based models modify score outcomes or distributions for enhanced fairness. Notable contributions include works by Yang et al. [9], [10], Celis et al. [11], Stoyanovich et al. [12], Kleinberg et al. [14], and Asudeh et al. [13].

Supervised fairness models in ranking span pre-processing, in-processing, and post-processing approaches. Pre-processing models, exemplified by Lahoti et al. [15], work on deriving fair training data. In-processing models, such as Zehlike et al.'s DELTR [6], address fairness during training, focusing on exposure bias. Similarly, Beutel et al. [16] introduced a pairwise ranking loss function with fairness regularizer, while Ma et al. [17] tackled fairness in query generation. Haak et al. [18] aimed at search query bias identification, and Chu et al. [19] highlighted biases in neural architecture search evaluations. Importantly, Chen et al. [21] proposed a meta-learning-based debiasing framework for recommendations. Post-processing models, conversely, refine model outputs post-training for fairness. Among these, Zehlike et al.'s works [22], [23] like FA*IR ensure representation of protected groups and offer continuous fairness interpolation. Additionally, Biega et al. [24] developed an algorithm optimizing the equity of user attention through relevance loss function.

Existing fairness ranking models utilize traditional machine learning, whereas our method employs meta-learning. We structure MCFR as a bilevel optimization solved via gradient-based techniques. Our model trains a meta-learner on a uniformly sampled meta-dataset, enabling the ranking model to learn unbiasedly from skewed data. Unlike previous works that fit into processing categories, we introduce a combined pre-processing and in-processing framework.

2.2 Meta-Learning on Fairness

Meta-learning is a field of study that aims to improve the learning ability of models by adapting to new tasks or environments, and it could be divided into two main categories: model-based [28], [29] and learning algorithm-based [30]. In addition to tasks such as few-shot learning [31], continual learning [32], and hyperparameter optimization [33], fairness is an important field.

Zhao et al. [34] presented the Follow the Fair Meta Leader (FFML) that learns an online fair classification model's primal, delivering both accuracy and fairness. In a subsequent work, Zhao et al. [35] emphasized the Primal-Dual Fair Meta-learning, targeting the optimal initialization of the base model's weights to rapidly adjust to new fairness tasks. They further advanced their research in [36], creating a few-shot discrimination prevention model for unbiased multi-class classification, rooted in the MAML framework. Concurrently, Slack et al. [37] introduced Fair-MAML, designed to derive fair models from minimal data for emerging tasks. This model, like Zhao's, is built upon the MAML framework but incorporates fairness regularization and a specific fairness hyperparameter. On recommender systems, Chen et al. [21] applied meta-learning principles on the AutoDebias framework. This framework is tailored to confront various biases, from selection to position bias. Thus, meta-learning's application in fairness ranking is emerging, and our MCFR proficiently targets this area.

2.3 Curriculum Learning

Bengio et al. [26] proposed the first curriculum learning approach, which orders the training examples based on the difficulty, and Hacohen et al. [38] applied it on deep learning model training. Generally, the curriculum learning could be classified into two types: data-level curriculum learning and model-level curriculum learning. The data-level curriculum learning aims to learn the data from easy to difficult in terms of a certain difficulty measurement, and the model-level curriculum learning would learn increase the model complexity as the training steps increase. Previously, the curriculum learning is applied to many interesting tasks such as relation extraction [39], transfer learning [40], domain adaptation [41], class-incremental learning [42], stochastic optimization [43], daily schedule recommendation [44], and etc. While curriculum learning has been successfully applied to many tasks, the meta-learning is one of the under-explored areas. Recently, the work on the relation extraction [39] combine the meta-learning and curriculum learning together to quickly adapt model parameters to new tasks. On few-shot classification, [45] proposed a simple and novel curriculum schedule that decreases the number of support size through the training. Also, the hardness aware meta-learning concept was applied in next point-of-interest recommendation task [46]. Among the works mentioned above, the curriculum has not be applied to fairness ranking problem as the measure of difficulty is not general. To the best of our knowledge, we propose the first curriculum learning framework for fairness ranking problem.

3 Meta Curriculum-based Fair Ranking

In this section, we will explain the proposed Meta Curriculum-based Fair Ranking framework in detail. In the MCFR framework, we will train an unbiased ranking model by using a meta-leaner to re-weight the ranking losses. We formulate it as a bilevel optimization problem and solve it using gradients through gradients. We also show that the framework could be trained with various ranking loss functions and fairness terms. Finally, we describe the design of the curriculum sampling strategy for meta dataset.

To address bias in datasets, traditional methods have utilized pre-processing, in-processing, or post-processing techniques [27], [47]. Our model combines pre-processing and in-processing, introducing the Meta Curriculum-based Fair Ranking framework. We derive a smaller dataset for meta-learner training, which assigns weights to emphasize the protected group during training. Curriculum learning adjusts this dataset's distribution ratio over epochs, facilitating smoother meta-learner training. This integrates ranking loss with fairness regularization, using the meta-learner to guide model training, as depicted in Fig. 2.

3.1 Problem Setting

We denote the set of queries in the training dataset as $\mathcal{Q}^{train}$ with the size $|\mathcal{Q}^{train}| = m$ and the set of items $\mathcal{D}^{train}$ with $|\mathcal{D}^{train}| = n$. Each query q in the $\mathcal{Q}^{train}$ has a list of item candidates $d^{(q)}$ from $\mathcal{D}^{train}$. Each pair of query and item is represented as a feature vector $x_i^{(q)}$ and is associated with the relevance score $y_i^{(q)}$. In the dataset, the candidates D have a binary attribute that specifies whether the candidate d belongs to the protected group or the non-protected group. For example, the binary attribute could represent gender or race, and systematic bias exists during the dataset collection.

3.2 A Unified MCFR Framework

To address the fairness problem, we train a meta learner on the meta-dataset which could help train a fair ranking model with the biased training dataset. We have the

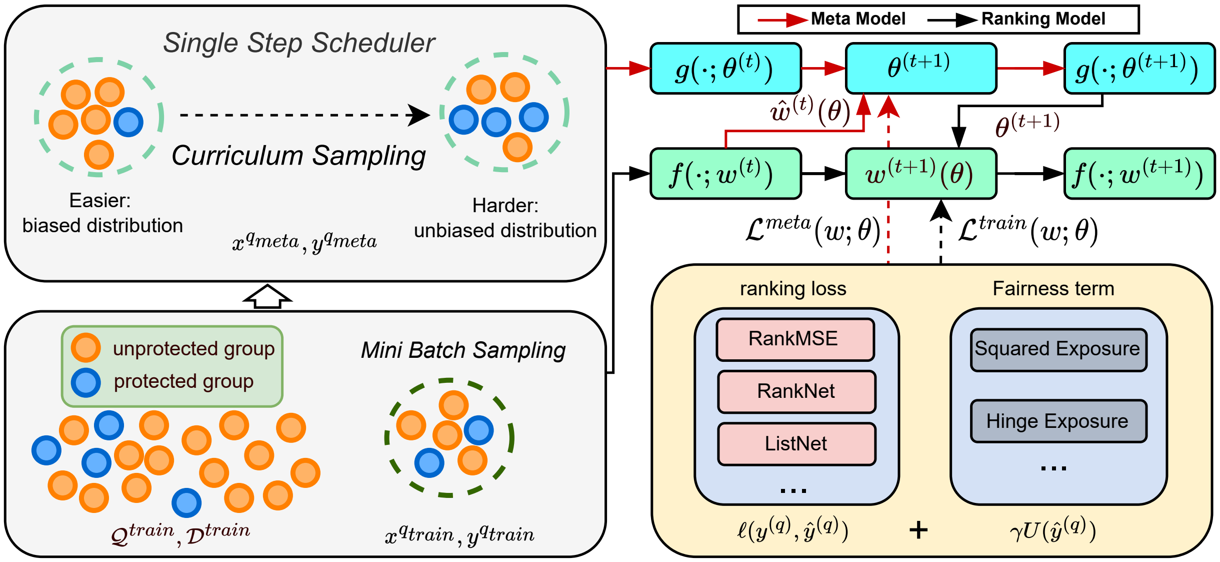

Fig. 2: MCFR learning algorithm flowchart (steps 4 and 6 in Algorithm 1). Note that $f(\cdot; w)$ is the ranking model, $g(\cdot; \theta)$ is the meta learner, b is the batch size for the training dataset, c is the batch size for the meta-dataset, and α and β are the learning rates. At each iteration, we firstly update θ in the meta learner using Eq. (8) with the meta-dataset sampled from the curriculum sampling with update of sampling difficulty at each epoch, and then we update w in the ranking model using Eq. (9) with the training dataset.

ranking model $f(x^{(q)}; w)$ and w is the learnable parameters of f, and we denote the output of the model as $\hat{y}^{(q)} = f(x^{(q)}; w)$. Generally, the model parameter w is optimized by $\min_w \frac{1}{m} \sum_{i=1}^{m} \mathcal{L}(y_i^{(q)}, \hat{y}_i^{(q)})$ which could minimize any given ranking loss function $\mathcal{L}$ such as pairwise loss and listwise loss. However, these loss functions treat $\mathcal{L}$ of each sample equally so that the ranking model will be unfair as there is a heavy data bias issue towards minority groups in the training dataset. To mitigate this problem, we introduce a meta learner $g(\cdot; \theta)$ with the learnable parameters θ to adaptively tune loss weights for each sample to achieve a fair exposure over diversity, and we could rewrite the training loss as the following:

$$\mathcal{L}^{train}(w; \theta) = \frac{1}{m} \sum_{i=1}^{m} \phi_i \mathcal{L}_i(w) = \frac{1}{m} \sum_{i=1}^{m} \phi_i \mathcal{L}(y_i^{(q)}, \hat{y}_i^{(q)}), \tag{1}$$

where $\hat{y}_i^{(q)} = f(x_i^{(q)}; w)$ denotes the model output, and $\phi_i \in [0, 1]$ denotes the i-th sample's loss weight given by the aforementioned meta learner $g(\cdot; \theta)$. Notably, $\mathcal{L}^{train}(w; \theta)$ governed by the meta learner's output weights depends on a fixed θ and is used for updating the ranking model's parameter w. In short, we write $\mathcal{L}_i(w)$ as the original loss value of the i-th training data sample output from the ranking loss $\mathcal{L}$. For the meta learner g, we use a multi-layer Perceptron network as proposed in [48], which takes loss values as input and output weighted loss as

$$\phi_i = g\big(\mathcal{L}_i(w); \theta\big) = g\Big(\mathcal{L}_i\big(y^{(q)}, f(x^{(q)}; w)\big); \theta\Big), \tag{2}$$

where i is the sample from the training dataset or the meta-dataset. We use `sigmoid` as the last-layer's activation function. Then we define a meta training loss function as

$$\mathcal{L}^{meta}\big(w(\theta)\big) = \frac{1}{s} \sum_{i=1}^{s} \mathcal{L}_i\big(w(\theta)\big), \tag{3}$$

Algorithm 1 Parameter update algorithm of MCFR

Input: A batch of training data $x^{q_{train}}, y^{q_{train}}$, a batch of meta-dataset $x^{q_{train}}, y^{q_{train}}$, ranking model's parameter $w^{(t)}$, and the meta learner's parameter $\theta^{(t)}$.
Output: Ranking model's parameter update $w^{(t+1)}$
1: Update $\hat{w}^{(t)}(\theta)$ by Eq. (5) with $\{x^{q_{train}}, y^{q_{train}}\}$.
2: Update $\theta^{(t+1)}$ by Eq. (8) with $\{x^{q_{meta}}, y^{q_{meta}}\}$.
3: Update $w^{(t+1)}$ by Eq. (9) with $\{x^{q_{train}}, y^{q_{train}}\}$.

where $s = |\mathcal{Q}^{meta}|$. The goal of the meta learner $g(\cdot; \theta)$ is to leverage the meta-dataset to learn how to re-weight the loss values to train the model $f(\cdot; w)$ on the biased dataset, indicating the relationship that the meta-learner plays a pivotal role in directing the tuning of the ranking model's parameters, inherently making w a function of θ. Since w is a function of θ, we naturally formulate the proposed MCFR as a bilevel optimization problem and give the objective function as

$$\min_{\theta} \mathcal{L}^{meta}\big(w^*(\theta)\big) \text{ s.t. } w^*(\theta) = \arg\min_{w} \mathcal{L}^{train}(w; \theta). \tag{4}$$

As illustrated in Fig. 2, our proposed MCFR model takes advantage of the sampled meta-dataset to learn an unbiased ranking model. The meta-dataset guide the meta learner to reweight the training loss, which helps the ranking model to focus on the candidates from the protected group.

3.3 Parameter Update

Since we formulate the framework as a bilevel optimization problem, it could be challenging as calculating the optimal parameters requires two nested loops of optimization. Following the well-known MAML works [48], [51], [52], we adopt an online strategy with a single optimization loop to update the ranking model and meta-learner parameters to guarantee the training efficiency.

	Type	Formula
Fairness	Hinge Exposure [6]	$U(\hat{y}^{(q)}) = \max(0, \text{Exposure}(G_0\|P) - \text{Exposure}(G_1\|P))^2$
	Squared Exposure	$U(\hat{y}^{(q)}) = (\text{Exposure}(G_0\|P) - \text{Exposure}(G_1\|P))^2$
Ranking	RankMSE [49]	$\ell(y^{(q)}, \hat{y}^{(q)}) = \frac{1}{n}\sum_{i=1}^{n}(y_i^{(q)}, \hat{y}_i^{(q)})^2$
	RankNet [50]	$\ell(y^{(q)}, \hat{y}^{(q)}) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=i}^{n}\log(1+\exp^{-(y_i^{(q)}-\hat{y}_j^{(q)})})$
	ListNet [3]	$\ell(y^{(q)}, \hat{y}^{(q)}) = -\sum_{i=1}^{n} P_{y^{(q)}}(i)\log P_{\hat{y}^{(q)}}(i)$

TABLE 1: Summary of ranking and fairness terms used in the loss function. The loss function used in the framework is $\mathcal{L}(y^{(q)}, \hat{y}^{(q)}) = \ell(y^{(q)}, \hat{y}^{(q)}) + \gamma U(\hat{y}^{(q)})$, and we can insert the above exposure terms and ranking loss terms as needed. Note that n denotes the number of candidates per query.

We update the parameters of the ranking network using the gradient decent on a batch of a training data with the loss function in Eq. (1), and we define the update of $w^{(t)}$ as:

$$\hat{w}^{(t)}(\theta) = w^{(t)} - \alpha\frac{1}{b}\sum_{i=1}^{b} g\big(\mathcal{L}_i^{train}(w^{(t)}); \theta^{(t)}\big)\nabla_w \mathcal{L}_i^{train}(w^{(t)}), \tag{5}$$

where t is each step of the update, and $w^{(t)}$ is the ranking model parameters at the step t. To obtain optimal parameters w^* and θ^*, we minimize the training loss by

$$w^*(\theta) = \arg\min_{w} \mathcal{L}^{train}(w;\theta) = \frac{1}{m}\sum_{i=1}^{m}\phi_i \mathcal{L}_i^{train}(w), \tag{6}$$

and the loss for the meta learner by

$$\theta^* = \arg\min_{\theta} \mathcal{L}^{meta}\big(w^*(\theta)\big) = \frac{1}{s}\sum_{i=1}^{s}\mathcal{L}_i^{meta}\big(w^*(\theta)\big). \tag{7}$$

Then given $\hat{w}^{(t)}(\theta)$ from Eq. (5), we update θ with the loss of the ranking model on the meta-dataset as the following:

$$\theta^{(t+1)} = \theta^{(t)} - \beta\frac{1}{c}\sum_{i=1}^{c}\nabla_\theta \mathcal{L}_i^{meta}\big(\hat{w}^{(t)}(\theta)\big), \tag{8}$$

where β is the learning rate, and c is the batch size of the meta-dataset. Then we update w as the following:

$$w^{(t+1)}(\theta) = w^{(t)} - \alpha\frac{1}{b}\sum_{i=1}^{b}\phi_i\nabla_w \mathcal{L}_i^{train}(w^{(t)}), \tag{9}$$

where α is the learning rate and b is the batch size of the training dataset. We adopt an alternating optimization strategy [48], [51], [52] to implement Eq. (8) and Eq. (9) instead of using nested optimization loops. The one step update algorithm is summarised in Alg. 1.

3.4 Ranking and Fairness Loss

The proposed MCFR serves as a unified framework that aims to improve both the ranking and fairness metrics, given any ranking and fairness objectives. To achieve this goal, we propose to include two terms in the loss functions similar to some in-processing fairness methods such as DELTR [6], and we develop our loss functions with the ranking term and fairness term given by:

$$\mathcal{L}(y^{(q)}, \hat{y}^{(q)}) = \ell(y^{(q)}, \hat{y}^{(q)}) + \gamma U(\hat{y}^{(q)}), \tag{10}$$

where $U(\hat{y}^{(q)})$ is the fairness term, $\ell(y^{(q)}, \hat{y}^{(q)})$ is the ranking loss term, and $\gamma > 0$ is a balancing parameter.

3.4.1 *Ranking Terms*

For the ranking loss, we use the following loss functions in the experiments: RankMSE [49], RankNet [50], and ListNet [3]. **RankMSE** is a pointwise loss which is based on least mean squared regression. **RankNet** proposed the first pairwise cross entropy loss which consider the preference relationships between documents. However, it is not possible to correctly predict the document order in all cases. **ListNet** aims to directly compute the ranking loss with each query and their candidates list instead of computing pairwise loss one pair by one pair.

It is worth noting that other ranking losses are also applicable in MCFR as we provide a general framework to improve the ranking metrics.

3.4.2 *Fairness Terms*

In this work, we focus on disparate exposure for the fairness term. For candidates D, there are two different groups: the non-protected group G_0 and the protected group G_1. The candidates from G_1 belong to a discriminated group such as female and African American and have significant disadvantages in the datasets. Then following the definition of Singh, et al [1], the exposure of a candidate d in a ranked list generated by a probabilistic ranking P is given by:

$$\text{Exposure}(x_i^{(q)}|P) = \sum_{a=1}^{n} P_{i,a}\cdot v_a, \tag{11}$$

where v_a is the position bias of position a. We then follow the implementation of Zelike, el al [6] to only consider the position bias of position 1 with v_1. Then the average exposure of candidates in each group G could be written as:

$$\text{Exposure}(G|P) = \frac{1}{|G|}\sum_{x_i^{(q)}\in G}\text{Exposure}(x_i^{(q)}|P). \tag{12}$$

With the exposure term defined above, we can introduce the fairness measure by minimizing the difference between the $\text{Exposure}(G_0|P)$ and $\text{Exposure}(G_1|P)$. In the experiments, we use two exposure measurements. Hinge Exposure calculates hinge squared loss from the exposure difference between two groups, while Square Exposure computes the squared exposure difference.

The ranking loss terms and exposure terms could be used in an arbitrary combination, and our framework could improve both the fairness and ranking metrics given different combinations. The ranking terms and fairness terms are summarised in Table 1.

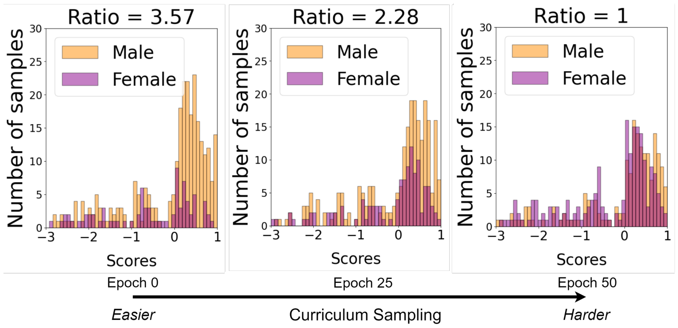

Fig. 3: Curriculum sampling strategy illustrated on the Engineering Student (Gender) dataset. We use the same ratio between the unprotected group and protected group in the meta-dataset as the training dataset at the beginning training epoch. We gradually decrease the ratio as the training epoch increase until the ratio becomes 1 which shows a balanced meta-dataset.

3.5 Curriculum Sampling

The training data shows systematic bias, with fewer candidates from protected groups than unprotected ones. To address this issue, we trained a meta learner using an unbiased meta-dataset since real unbiased data is rare. While AutoDebias [21] previously tackled a similar issue for recommendation systems, it does not fit our ranking-focused needs. Another approach, used in MFR [8], equally samples candidates from each group. However, this method creates a meta-dataset that may fall short of accurately capturing the real biased data. For tasks like ranking, where the order and relevance of items are crucial, this mismatch in the data distribution can significantly hinder the model's ability to provide fair and effective rankings in practical applications, biased situations. To this end, we adopt curriculum learning [26], a method that starts with easier, less biased samples and gradually introduces more complex ones. This mimics natural learning, helping the model adapt better and become more robust. It's designed to ease the model into understanding and correcting biases, ensuring it performs well and fairly in real-world applications, even with the underlying biases in the data it was trained on.

In detail, we want to downsample the meta-dataset with the similar distribution as the training dataset at the early training epochs, and we gradually change the ratio of the number of candidates from the protected and unprotected groups to 1.0. Since we could not collect a real unbiased dataset, we define 1.0 to be the unbiased ratio of the number of candidates from the two different groups ($d^{(q)}_{\text{unprotected}}$ vs $d^{(q)}_{\text{protected}}$), which means there is an equal number of candidates from each group. Here the downsampling ratio is defined as $r = |d^{(q)}_{\text{unprotected}}|/|d^{(q)}_{\text{protected}}|$. The underlying assumption behind this curriculum sampling strategy is that it is easier to train the model when the meta-dataset and training dataset have similar distribution and that it is difficult to optimize the parameters in the ranking model when the meta learner sees a very different meta-dataset compared to the training dataset. As shown in Fig. 3, we illustrate the change in the distribution of two groups in the meta-dataset at different training epochs.

To train the meta learner, we use the curriculum sampled data $\{x^{q_{meta}}, y^{q_{meta}}\}$. The meta-dataset represents the meta-knowledge of the true distribution of the protected group and the other group, where $|\mathcal{Q}^{meta}| = s \ll m$ and $|\mathcal{D}^{meta}| = o \ll n$. In the meta-dataset, we denote the feature vector of each item as $x^{(q_{meta})}$ and the relevance score as $y^{(q_{meta})}$ given a query q_{meta} from $\mathcal{Q}^{meta}$. Similar to $\mathcal{L}_i^{train}(w)$, we denote $\mathcal{L}_i^{meta}(w(\theta))$ as the loss value for each meta-dataset sample. Thus we define CurriculumSampling($\mathcal{Q}^{train}, \mathcal{D}^{train}$, b, t) as the following:

$$r^{(t)} = r - t \times (r - 1.0)/T, \tag{13}$$

where $r^{(t)}$ is the ratio of sampled candidates for each group for each query. Note that this is a single step scheduler as the ratio $r^{(t)}$ is updated at each epoch. After executing CurriculumSampling at each epoch, the sampling meta-dataset $\{x^{q_{meta}}, y^{q_{meta}}\}$ should have the property that $|d^{(q)}_{\text{unprotected}}|/|d^{(q)}_{\text{protected}}| = r^{(t)}$. Intuitively, the CurriculumSampling decreases the ratio epoch by epoch from the biased ratio to 1.0.

As described in Section 3.2, the meta-dataset is an important part of the model training as it is the key data to guide the meta learner. Since the meta learner aims to reweight the loss for the ranking model, how well the meta learner is trained determine the performance of the ranking model. With the curriculum sampling, we decrease the training difficulty of the meta learner compared to MFR [8] which only uses one sampled unbiased dataset. The meta learner could progressively be trained with a more unbiased meta-dataset as the epoch increases, which could improve the meta learner's performance and lead to a better overall performance for the ranking model. The whole training process is summarized in Algorithm 2.

Our framework provides flexibility to solve different ranking problems as ListNet [3] may not work for all ranking problems. In other cases, the fairness terms could also be switched by using different fairness metrics or a different formula to compute the disparate exposure. As the exposure issue is not the only fairness problem, the MCFR is capable of being optimized with other fairness terms such as position bias and conformity bias.

Algorithm 2 The MCFR Learning Algorithm

Input: Training dataset $\mathcal{Q}^{train}$, $\mathcal{D}^{train}$, batch size b, c, max iterations T.
Output: Ranking model's parameter $w^{(T)}$
1: Initialize ranking model's parameter $w^{(0)}$ and the meta learner's parameter $\theta^{(0)}$.
2: **for** $t = 0$ **to** $T - 1$ **do**
3: $\{x^{q_{meta}}, y^{q_{meta}}\} \leftarrow$ CurriculumSampling($\mathcal{Q}^{train}, \mathcal{D}^{train}$, b, t).
4: $\{x^{q_{train}}, y^{q_{train}}\} \leftarrow$ SampleMiniBatch($\mathcal{Q}^{train}, \mathcal{D}^{train}, c$).
5: Update $w^{(t+1)}$ by Alg. 1
6: **end for**

4 Experiments

In the experiments, we train and evaluate the model on four real-world public datasets. We study both the ranking and fairness metrics of our approach compared to other baseline models. We also conduct an ablation study for the

	W3C Experts (gender)	Engineering Students (high school type)	Engineering Students (gender)	Law Students (gender)	Law Students (race)	COMPAS (race)
#items/query	200	480.6	480.6	21791	19567	6889
#protected/query	21.5	167.6	97.6	9537	1282	3528

TABLE 2: Summary of dataset statistics. We report the average counts of total and unprotected items per query for the W3C Experts and Engineering Students datasets. We provide the exact item counts for the Law Students and COMPAS datasets, each of which contains only one query.

	W3C Experts (gender)		Engineering Students (high school type)		Engineering Students (gender)	
	Precision@10	Fairness	Kendall's Tau	Fairness	Kendall's Tau	Fairness
ListNet [3]	0.178	0.759	0.390	1.070	**0.384**	0.858
LambdaMART [53]	0.095	0.738	0.355	1.002	0.326	0.907
DELTR [6]	**0.180**	0.827	**0.391**	1.075	0.370	0.976
FA*IR pre [22]	**0.180**	0.770	0.374	1.020	0.360	0.942
FA*IR post [22]	**0.180**	0.827	**0.391**	1.075	0.370	0.976
AutoDebias [21]	0.033	0.829	0.372	0.955	0.372	0.955
FairGBM [20]	0.087	**0.941**	0.338	0.909	0.336	0.892
MFR	0.126	0.830	**0.391**	1.086	0.352	1.052
MCFR	0.118	0.843	0.390	**1.088**	0.350	**1.055**

	Law Students (gender)		Law Students (race)		COMPAS (race)	
	Kendall's Tau	Fairness	Kendall's Tau	Fairness	Kendall's Tau	Fairness
ListNet [3]	0.202	0.931	**0.184**	0.853	0.639	0.836
LambdaMART [53]	0.199	0.979	0.156	0.847	0.542	0.956
DELTR [6]	0.188	0.993	0.130	1.014	0.576	0.970
FA*IR pre [22]	0.203	0.931	0.161	0.895	0.557	1.039
FA*IR post [22]	0.182	0.965	0.140	0.944	0.557	1.040
AutoDebias [21]	0.222	0.894	0.135	1.009	0.644	1.136
FairGBM [20]	0.141	0.998	0.210	1.116	0.550	0.917
MFR	**0.225**	1.015	**0.184**	1.654	**0.644**	1.138
MCFR	**0.225**	**1.023**	0.182	**1.671**	**0.644**	**1.144**

TABLE 3: Experimental results with hinge exposure [6]. To measure fairness, we compute the exposure ratio between the protected and the non-protected group, so the values greater than 1.0 indicate greater visibility for the protected group and vice versa. For the ranking metric, higher Kendall's Tau / Precision@10(P@10) scores indicate better performance. The bold text indicates the model with the best performance, and the results show that the MCFR model is better on the fairness metrics with comparable performance on the ranking metrics against other state-of-the-art models.

effectiveness of our framework by changing the ranking loss term and the disparate exposure term. We repeat the experiment on the same datasets with different settings of loss functions, and we evaluate the proposed framework by comparing it with the baseline models. In the analysis, the following questions are answered:

- What is the proposed MCFR's performance compared to the baseline models?
- Could MCFR improve both the ranking and fairness metrics in different loss functions?
- What are the effects of the curriculum sampling?

4.1 Experimental setting

We train and evaluate the model on four real-world public datasets: (i) Engineering Student; (ii) Law Student, (iii) W3C Experts; (iv) COMPAS (Correctional Offender Management Profiling for Alternative Sanctions). The statistics of each dataset are summarized in Table 2.

W3C experts Dataset This dataset originates from TREC 2005 Enterprise Track [54]. It involves searching for experts based on a topic, using features from their emails. We designate gender as the protected attribute, with technical topics as queries. In this context, females are the protected group, and males are non-protected. Each query has 200 items, averaging 21.5 from the protected group. Given that the original dataset ranks retrieved experts equally, we adopt the DELTER experiments' setting [6], categorizing expert candidates as: male experts, female experts, male non-experts, and female non-experts. For candidate features, we utilize the Elasticsearch Learning to Rank Plug-in[1] for all query-candidate pair text features.

Law Student Dataset This dataset [4] was collected to determine if the LSAT (Law School Admission Test in the US) is biased against ethnic minorities. The dataset contains information from first-year law students, and the protected attributes are gender and race. The query is academic year, and the task is to retrieve students with good LSAT scores. Since our problem setting is focused on one protected attribute at a time, we have two datasets: Law Students (gender) and Law Students (race). In the Law Students (gender) dataset, females are the protected group among 21,791 candidates, with 9,537 being female. In the Law Students (race) dataset, African Americans are the protected

1. https://elasticsearch-learning-to-rank.readthedocs.io/en/latest/

	Exposure Type	W3C Experts (gender)		Engineering Students (high school type)		Engineering Students (gender)	
		Precision@10	**Fairness**	**Kendall's Tau**	**Fairness**	**Kendall's Tau**	**Fairness**
RankMSE	n/a	0.121	0.770	0.187	0.800	0.376	0.836
MFR	Hinge	0.115	0.781	0.384	1.049	0.357	1.010
MCFR	Hinge	0.115	0.782	0.384	1.052	0.353	1.020
MFR	Squared	0.115	0.780	0.384	1.045	0.360	0.982
MCFR	Squared	0.115	0.782	0.384	1.045	0.360	0.990
	Exposure Type	**Law Students (gender)**		**Law Students (race)**		**COMPAS (race)**	
		Kendall's Tau	**Fairness**	**Kendall's Tau**	**Fairness**	**Kendall's Tau**	**Fairness**
RankMSE	n/a	0.213	0.874	0.190	0.847	0.493	0.768
MFR	Hinge	0.225	0.910	0.191	0.847	0.634	0.911
MCFR	Hinge	0.226	0.920	0.190	0.851	0.634	0.911
MFR	Squared	0.223	1.010	0.139	0.992	0.633	0.911
MCFR	Squared	0.225	1.023	0.138	0.996	0.630	0.928

(a) RankMSE [49]

	Exposure Type	W3C Experts (gender)		Engineering Students (high school type)		Engineering Students (gender)	
		Precision@10	**Fairness**	**Kendall's Tau**	**Fairness**	**Kendall's Tau**	**Fairness**
RankNet	n/a	0.121	0.770	0.131	0.806	0.190	0.800
MFR	Hinge	0.121	0.774	0.126	0.925	0.188	0.810
MCFR	Hinge	0.123	0.775	0.131	0.867	0.186	0.820
MFR	Squared	0.121	0.774	0.126	0.925	0.188	0.810
MCFR	Squared	0.121	0.774	0.131	0.867	0.186	0.812
	Exposure Type	**Law Students (gender)**		**Law Students (race)**		**COMPAS (race)**	
		Kendall's Tau	**Fairness**	**Kendall's Tau**	**Fairness**	**Kendall's Tau**	**Fairness**
RankNet	n/a	0.093	0.942	0.105	0.866	0.128	0.768
MFR	Hinge	0.131	1.033	0.140	1.284	0.373	0.839
MCFR	Hinge	0.132	1.036	0.152	1.370	0.375	0.840
MFR	Squared	0.173	1.033	0.105	0.866	0.352	0.832
MCFR	Squared	0.220	1.050	0.105	0.866	0.352	0.832

(b) RankNet [50]

	Exposure Type	W3C Experts (gender)		Engineering Students (high school type)		Engineering Students (gender)	
		Precision@10	**Fairness**	**Kendall's Tau**	**Fairness**	**Kendall's Tau**	**Fairness**
ListNet	n/a	0.178	0.759	0.390	1.070	0.384	0.858
MFR	Hinge	0.126	0.830	0.391	1.086	0.352	1.052
MCFR	Hinge	0.118	0.843	0.390	1.088	0.350	1.055
MFR	Squared	0.118	0.803	0.330	1.005	0.358	1.006
MCFR	Squared	0.118	0.803	0.341	1.005	0.342	1.018
	Exposure Type	**Law Students (gender)**		**Law Students (race)**		**COMPAS (race)**	
		Kendall's Tau	**Fairness**	**Kendall's Tau**	**Fairness**	**Kendall's Tau**	**Fairness**
ListNet	n/a	0.202	0.931	0.184	0.853	0.639	0.836
MFR	Hinge	0.225	1.015	0.184	1.654	0.644	1.138
MCFR	Hinge	0.225	1.023	0.182	1.671	0.644	1.144
MFR	Squared	0.223	1.010	0.113	1.166	0.340	0.828
MCFR	Squared	0.225	1.014	0.079	1.115	0.632	1.068

(c) ListNet [3]

TABLE 4: Ablation study results. We conduct experiments on all 6 combinations of ranking loss terms and fairness terms on MFR and MCFR. For ListNet, RankMSE and RankNet models, they serves as the baseline model considering no fairness term during the training. The results show that MCFR could generally improve the fairness metrics with comparable ranking performance given different ranking loss and fairness terms.

group out of 19,567 candidates, with 1,282 from this group.

Engineering Students This dataset [6] contains information on first-year students at a Chilean university. The qualification features include admission test results in mathematics, language, and science, the students' high school grades, and the number of credits taken at the university. The task is to predict academic performance, and the protected attributes are high school type and gender. Similarly, we have two datasets: Engineering Students (high school type) and Engineering Students (gender). For Engineering Students datasets, one focuses on high school type, with public high school students as the protected group, averaging 167.6 out of 480.6 items per query. The other considers gender, with females as the protected group, averaging 97.6 out of 480.6 items per query.

COMPAS COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) is a commercial algorithm for scoring a criminal defendant's likelihood of recidivism. In the COMPAS dataset [5], it has been observed that the algorithm is biased towards African American candidates. In this dataset, the task is to predict the recidivism score, and the protected attribute is race. There are 6,889 candidates in total, and 3,528 are African Americans.

4.1.1 Baselines

We integrated several baseline models in our implementation. ListNet [3] introduces a listwise loss function. LambdaMART [53] combines MART and LambdaRank, transforming ranking tasks with gradient boosting decision trees. DELTR [6] offers an LTR strategy with listwise fairness metrics. FA*IR [22] applies pre and post-processing techniques for enhanced fairness. AutoDebias [21] presents a debiasing method for recommendation systems. FairGBM [20] delivers a fairness-centric classification model for GBDT, while MFR [8] employs meta-learning for fair LTR. Notably, only ListNet and LambdaMART focus solely on ranking metrics, with DELTR and MFR emphasizing fairness-aware ranking.

4.1.2 Implementation Details

To split the datasets, we have 50 queries for training and 10 queries for testing in the W3C dataset, 4 queries for training and 1 query for testing in the Engineering Students dataset, and 80% for training and 20% for testing in the Law Students dataset and the COMPAS dataset. We use Precision@10 (P@10) [55] for the W3C dataset and Kendall's Tau [7] for other datasets to evaluate the ranking performance. Kendall's Tau assesses the correlation between two ranking sets, calculating the difference between the number of concordant and discordant pairs divided by the total number of pairs. It ranges from -1 to 1, indicating perfect agreement, no correlation, or perfect disagreement in the rankings, respectively. In details, the Kendall's Tau is calculated as the following:

$$\text{Kendall's Tau} = \frac{p-q}{\sqrt{(p+q+t)\times(p+q+u)}}, \tag{14}$$

where p is the number of concordant pairs, q the number of discordant pairs, t the number of ties in the ground truth rankings, and u the number of ties in the predicted rankings. To measure fairness, we compute the exposure ratio between the protected and the non-protected group [6]. Thus, in the fairness metric, values greater than 1.0 indicate greater visibility for the protected group and vice versa.

In the training, we set the update frequency of the weighting model parameter θ to be per 2 steps, the optimizer to be SGD [56], the momentum to be 0.98, the learning rate to be 0.022, the hidden layer dimension to be 30, and the number of hidden layers to be 3. For the ranking model, we set the learning rate to be 0.005, the optimizer to be SGD, the momentum to be 0.95, and the weight decay to be 0.005. We set different values for γ and training epoch for different dataset: W3C dataset uses γ = 500 and 100 epochs, Engineering Students (high school) uses γ = 5,000 and 280 epochs, Engineering Students (Gender) uses γ = 400 and 150 epochs, Law Students(gender) uses γ = 1,200 and 550 epochs, Law Students (race) uses γ = 50,000 and 110 epochs, and COMPAS (race) uses γ = 2,500 and 45 epochs.

In the ablation study to evaluate the effectiveness of our framework, we use the same hyperparameters as described above for other ranking losses such as RankMSE and RankNet. In the experiment, we collect results with all combinations of ranking losses and fairness terms.

4.2 Fair Ranking Performance

In Table 3, we detail the performance of both baseline and fair ranking models trained with hinge exposure. The proposed MCFR outperforms other baseline models in fairness metrics across all datasets. When compared to ListNet and LambdaMART, models like DELTR, MFR, FA*IR, AutoDebias, FairGBM, and MCFR show enhanced results due to the inclusion of fairness measures during training. Notably, MCFR's use of curriculum sampling for the meta-dataset allows it to surpass MFR in fairness metrics, as the meta-learner adeptly adjusts the loss distribution. During MCFR training, curriculum sampling creates the meta-dataset for the Meta Model. The W3C dataset's limited items from the protected group hinder significant distribution shifts in meta-dataset sampling, affecting its ranking performance. This constraint primarily contributes to the decreasing ranking performance observed in the model trained on W3C data. Except on the W3C dataset, MCFR has competitive results on the ranking metrics compared to the other baseline models, indicating that training MCFR does not focus solely on the fairness metrics. For ListNet, the results are also expected, as they only optimize for ranking metrics and have better performance in ranking metrics on Engineering Students (gender) and Law Students (race). Since AutoDebias and FairGBM are tailored for recommendation and classification tasks respectively, their limited performance on ranking problems is as expected. In Fig. 1, we also plot the histogram of ranks on the protected attributes from the different models. From the plot, we can see that the distribution of predicted ranks shifts from right to left, indicating that the MCFR model generally ranks items from the protected group higher compared to ListNet and MFR. In the plot, 1 on the x-axis indicates the top rank, and more candidates falling in bins on the left means the candidates receive higher ranks. In ranking algorithms, MCFR enhances visibility for underrepresented protected groups. However, fairness doesn't mean maximizing exposure for them at the expense of the non-protected group's visibility.

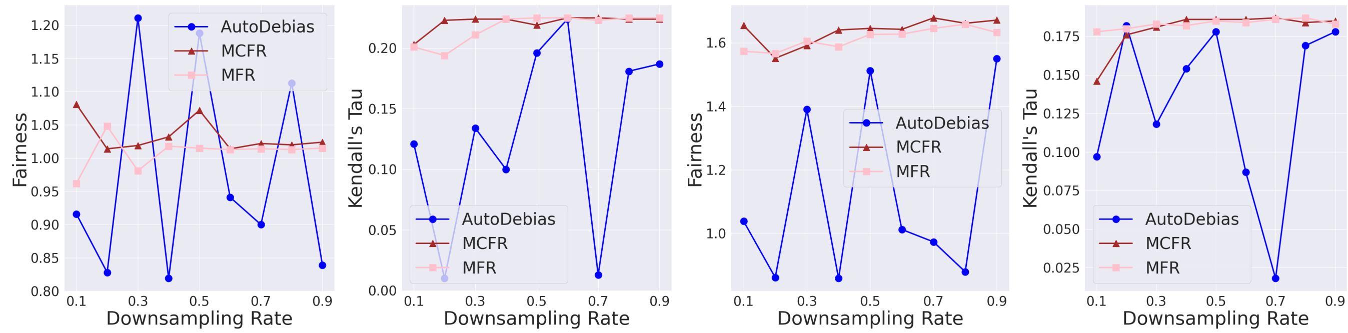

(a) Fairness Metric (gender) (b) Ranking Metric (gender) (c) Fairness Metric (race) (d) Ranking Metric (race)

Fig. 4: Evaluation results on the down-sampling experiments. We conduct the experiment on Law Students (gender) and Law students (race) datasets, and we down-sample the training data from the rate of 0.1 to 0.9. The results show that MCFR has better data efficiency as it could achieve better fairness metrics with similar ranking performance than MFR and AutoDebias at different down-sampling rate.

	W3C Experts (gender)	Engineering Students (high school type)	Engineering Students (gender)	Law Students (gender)	Law Students (race)	COMPAS (race)
DELTR	43.69	14.09	40.92	14.35	17.70	19.67
MFR	21.16	15.24	17.24	51.29	49.72	76.88
MCFR	171.37	92.57	91.64	294.42	293.92	352.96

TABLE 5: Experimental results on total convergence time in seconds. It shows the total convergence time for different algorithms (DELTR, MFR, and MCFR) across various datasets or scenarios. Based on the table, the MCFR framework generally has comparable convergence time than the other two algorithms.

4.3 Ablation Studies

In Table 4, we present the ablation study results for MCFR, which offers flexibility in choosing loss functions and fairness terms. As a generalized framework, MCFR consistently enhances both ranking and fairness metrics across various loss functions and exposure formulas. We employed RankMSE, RankNet, and ListNet as representatives for pointwise, pairwise, and listwise losses, which serve as baseline models in Table 4(a), 4(b), and 4(c).

4.3.1 Ranking Terms Analysis

First, we analyze the performance of MCFR using different ranking terms in loss functions. When using ListNet, MCFR has worse ranking performance on the W3C Experts (gender) and Engineering Students (gender) datasets than the ListNet model. On other datasets, MCFR and the ListNet model have similar ranking performance. Note that on the Law Students (gender) dataset, MCFR also improves the ranking metrics. When using RankMSE, a similar pattern is observed. On RankNet, MCFR achieves similar ranking performance on the W3C Experts (gender) dataset and improves the ranking metrics on the Law Students (gender) and Law Students (race) datasets, in addition to the fairness metrics. The consistent improvement in ranking metrics shows that the proposed MCFR is a generalized framework that can adapt to many ranking loss functions.

4.3.2 Fairness Terms Analysis

Second, we evaluate different fairness terms in loss functions. When using ListNet as the ranking loss term, MCFR greatly improves the fairness metrics on the W3C Experts (gender) and Engineering Students (gender) datasets. On other datasets, MCFR outperforms the ListNet model on the fairness metrics with similar ranking performance. When using RankMSE, MCFR also improves the fairness metrics on the Law Students (gender) and Law Students (race) datasets. We see that MCFR can improve the fairness metrics with various ranking loss terms.

4.3.3 Curriculum Sampling Analysis

Moreover, we compare the performance of MCFR and MFR to show the effectiveness of curriculum learning using different losses. Note that in MFR, we use the same settings in loss functions as in MCFR to have a fair comparison. When using the Hinge exposure, MCFR usually has better fairness performance with minor trade-offs in ranking metrics, except on the W3C Experts (gender) dataset using ListNet. While using the Squared exposure, except on the Law Students (race) dataset, MCFR improves both ranking and fairness metrics compared to MFR. These results demonstrate the effectiveness of curriculum learning.

4.3.4 Data Efficiency

To assess curriculum learning's effect on data efficiency, we compare with MCFR, MFR, and AutoDebias using down-sampled training data, varying from 10% to 90% of the original data. Figure 4 illustrates how MCFR outperforms MFR and AutoDebias across most sampling rates in fairness for gender-related data, achieving fair metrics close to 1.0 while maintaining high ranking performance. MCFR demonstrates superior fairness with reduced training data. For race-related data, MCFR achieves better ranking performance and higher fairness metrics, indicating our curriculum strategy effectively enhances fairness of the protected groups even with less data.

4.3.5 Training and Inference Efficiency

To enhance ranking fairness with MCFR, we sought a balance between fairness and efficiency. As shown in Table 5, MCFR has a training complexity comparable to methods like DELTR, and the curriculum sampling extends the training time linearly with sampling rounds. Notably, during the inference, MCFR, MFR, and DELTR will show consistent efficiency since these algorithms share the same base ranking model with the same number of parameters and layers and there is only one forward pass for predictions. Table 5 shows MCFR's extended convergence time due to curriculum sampling and added epochs. MCFR's fairness benefits are clear, yet we value efficiency in time-sensitive applications. Overall, these results demonstrate that the curriculum learning in MCFR enhances fairness without compromising ranking performance, also making training more efficient.

5 Conclusion and Future Work

In this study, we introduced the Meta Curriculum-based Fair Ranking (MCFR) framework to address data bias in search problems. By employing a meta-learner trained on a curriculum-learning-sampled meta-dataset, our approach re-weights the training loss from the target ranker on biased data. This re-weighted loss aids in developing an unbiased ranking model, enhancing exposure for minority groups. Comparative experiments on real-world datasets confirm MCFR's superiority over fair ranking models lacking meta-learning and curriculum learning components. Future work will explore handling multiple protected attributes and expanding MCFR's applicability to diverse ranking tasks and datasets, probing its adaptability and potential constraints.

Acknowledgments

This work is supported in part by the Ciocca center research award and also by the Responsible AI program at Santa Clara University.

References

[1] A. Singh and T. Joachims, "Fairness of exposure in rankings," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY, USA: ACM, 2018, p. 2219–2228.

[2] B. Friedman and H. Nissenbaum, "Bias in computer systems," *ACM Trans. Inf. Syst.*, vol. 14, no. 3, p. 330–347, jul 1996.

[3] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li, "Learning to rank: from pairwise approach to listwise approach," in *Proceedings of the 24th international conference on Machine learning*, 2007, pp. 129–136.

[4] L. F. Wightman, "Lsac national longitudinal bar passage study," *LSAC Research Report Series*, 1998.

[5] M. Barenstein, "ProPublica's COMPAS Data Revisited," *arXiv e-prints*, p. arXiv:1906.04711, Jun 2019.

[6] M. Zehlike and C. Castillo, *Reducing Disparate Exposure in Ranking: A Learning To Rank Approach*. New York, NY, USA: ACM, 2020, p. 2849–2855.

[7] W. R. Knight, "A computer method for calculating kendall's tau with ungrouped data," *Journal of the American Statistical Association*, vol. 61, no. 314, pp. 436–439, 1966.

[8] Y. Wang, Z. Tao, and Y. Fang, "A meta-learning approach to fair ranking," in *The 45th International ACM SIGIR Conference on Research & Development in Information Retrieval*. New York, NY, USA: ACM, 2022, p. 2539–2544.

[9] K. Yang and J. Stoyanovich, "Measuring fairness in ranked outputs," in *Proceedings of the 29th International Conference on Scientific and Statistical Database Management*. New York, NY, USA: ACM, 2017.

[10] K. Yang, V. Gkatzelis, and J. Stoyanovich, "Balanced ranking with diversity constraints," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 7 2019, pp. 6035–6042.

[11] L. E. Celis, D. Straszak, and N. K. Vishnoi, "Ranking with fairness constraints," in *45th International Colloquium on Automata, Languages, and Programming, 2018, July 9-13, 2018, Prague, Czech Republic*, I. Chatzigiannakis, C. Kaklamanis, D. Marx, and D. Sannella, Eds., vol. 107. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, pp. 28:1–28:15.

[12] J. Stoyanovich, K. Yang, and H. Jagadish, "Online set selection with fairness and diversity constraints," in *Proceedings of the EDBT Conference*, 2018.

[13] A. Asudeh, H. V. Jagadish, J. Stoyanovich, and G. Das, "Designing fair ranking schemes," in *Proceedings of the 2019 International Conference on Management of Data*. New York, NY, USA: ACM, 2019, p. 1259–1276.

[14] J. Kleinberg and M. Raghavan, "Selection Problems in the Presence of Implicit Bias," in *9th Innovations in Theoretical Computer Science Conference*, ser. Leibniz International Proceedings in Informatics, A. R. Karlin, Ed., vol. 94. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, pp. 33:1–33:17.

[15] P. Lahoti, G. Weikum, and K. P. Gummadi, "ifair: Learning individually fair data representations for algorithmic decision making," *2019 IEEE 35th International Conference on Data Engineering*, pp. 1334–1345, 2019.

[16] A. Beutel, J. Chen, T. Doshi, H. Qian, L. Wei, Y. Wu, L. Heldt, Z. Zhao, L. Hong, E. H. Chi, and C. Goodrow, "Fairness in recommendation ranking through pairwise comparisons," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. New York, NY, USA: ACM, 2019, p. 2212–2220.

[17] H. Ma, S. Guan, C. Toomey, and Y. Wu, "Diversified subgraph query generation with group fairness," in *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. New York, NY, USA: ACM, 2022, p. 686–694.

[18] F. Haak and P. Schaer, "Auditing search query suggestion bias through recursive algorithm interrogation," in *14th ACM Web Science Conference 2022*. New York, NY, USA: ACM, 2022, p. 219–227.

[19] X. Chu, B. Zhang, and R. Xu, "Fairnas: Rethinking evaluation fairness of weight sharing neural architecture search," in *International Conference on Computer Vision*, 2021.

[20] A. Cruz, C. G. Belém, J. Bravo, P. Saleiro, and P. Bizarro, "FairGBM: Gradient boosting with fairness constraints," in *The Eleventh International Conference on Learning Representations*, 2023.

[21] J. Chen, H. Dong, Y. Qiu, X. He, X. Xin, L. Chen, G. Lin, and K. Yang, "Autodebias: Learning to debias for recommendation," in *The 44th International ACM SIGIR Conference on Research & Development in Information Retrieval*. New York, NY, USA: ACM, 2021, p. 21–30.

[22] M. Zehlike, F. Bonchi, C. Castillo, S. Hajian, M. Megahed, and R. Baeza-Yates, "Fa*ir: A fair top-k ranking algorithm," in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. New York, NY, USA: ACM, 2017, p. 1569–1578.

[23] M. Zehlike, P. Hacker, and E. Wiedemann, "Matching code and law: Achieving algorithmic fairness with optimal transport," *Data Min. Knowl. Discov.*, vol. 34, no. 1, p. 163–200, jan 2020.

[24] A. J. Biega, K. P. Gummadi, and G. Weikum, "Equity of attention: Amortizing individual fairness in rankings," in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. New York, NY, USA: ACM, 2018, p. 405–414.

[25] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70. PMLR, 06–11 Aug 2017, pp. 1126–1135.

[26] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th Annual International Conference on Machine Learning*. New York, NY, USA: ACM, 2009, p. 41–48.

[27] M. Zehlike, K. Yang, and J. Stoyanovich, "Fairness in ranking, part i: Score-based ranking," *ACM Comput. Surv.*, vol. 55, no. 6, dec 2022.

[28] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*. JMLR.org, 2017, p. 1126–1135.

[29] M. Andrychowicz, M. Denil, S. G. Colmenarejo, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, and N. de Freitas, "Learning to learn by gradient descent by gradient descent," in *Proceedings of the 30th International Conference on Neural Information Processing*

Systems. Red Hook, NY, USA: Curran Associates Inc., 2016, p. 3988–3996.

[30] A. Antoniou, H. Edwards, and A. Storkey, "How to train your MAML," in *International Conference on Learning Representations*, 2019.

[31] M. A. Jamal and G.-J. Qi, "Task agnostic meta-learning for few-shot learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 11 719–11 727.

[32] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," *Neural Networks*, vol. 113, pp. 54–71, 2019.

[33] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil, "Bilevel programming for hyperparameter optimization and meta-learning," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1568–1577.

[34] C. Zhao, F. Chen, and B. Thuraisingham, "Fairness-aware online meta-learning," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. New York, NY, USA: ACM, 2021, p. 2294–2304.

[35] C. Zhao, F. Chen, Z. Wang, and L. Khan, "A primal-dual subgradient approach for fair meta learning," in *2020 IEEE International Conference on Data Mining*. IEEE, 2020, pp. 821–830.

[36] C. Zhao and F. Chen, "Unfairness discovery and prevention for few-shot regression," in *2020 IEEE International Conference on Knowledge Graph*, 2020, pp. 137–144.

[37] D. Slack, S. A. Friedler, and E. Givental, "Fairness warnings and fair-maml: Learning fairly with minimal data," in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: ACM, 2020, p. 200–209.

[38] G. Hacohen and D. Weinshall, "On the power of curriculum learning in training deep networks," in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, vol. 97. PMLR, 2019, pp. 2535–2544.

[39] T. Wu, X. Li, Y.-F. Li, G. Haffari, G. Qi, Y. Zhu, and G. Xu, "Curriculum-meta learning for order-robust continual relation extraction," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, 2021, pp. 10 363–10 369.

[40] D. Weinshall, G. Cohen, and D. Amir, "Curriculum learning by transfer learning: Theory and experiments with deep networks," in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, vol. 80. PMLR, 2018, pp. 5235–5243.

[41] Y. Shu, Z. Cao, M. Long, and J. Wang, "Transferable curriculum for weakly-supervised domain adaptation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 4951–4958.

[42] Y. Yang, Z. Sun, H. Zhu, Y. Fu, Y. Zhou, H. Xiong, and J. Yang, "Learning adaptive embedding considering incremental class," *IEEE Transactions on Knowledge and Data Engineering*, no. 01, pp. 1–1, sep 5555.

[43] B. Han, I. W. Tsang, X. Xiao, L. Chen, S. Fung, and C. P. Yu, "Privacy-preserving stochastic gradual learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 33, no. 08, pp. 3129–3140, aug 2021.

[44] W. Huang, J. Liu, T. Li, T. Huang, S. Ji, and J. Wan, "Feddsr: Daily schedule recommendation in a federated deep reinforcement learning framework," *IEEE Transactions on Knowledge and Data Engineering*, no. 01, pp. 1–1, nov 5555.

[45] E. stergiadis, P. Agrawal, and O. Squire, "Curriculum meta-learning for few-shot classification," in *Fifth Workshop on Meta-Learning at the Conference on Neural Information Processing Systems*, 2021.

[46] Y. Chen, X. Wang, M. Fan, J. Huang, S. Yang, and W. Zhu, "Curriculum meta-learning for next poi recommendation," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 2692–2702.

[47] M. D. Ekstrand, A. Das, R. Burke, F. Diaz *et al.*, "Fairness in information access systems," *Foundations and Trends® in Information Retrieval*, vol. 16, no. 1-2, pp. 1–177, 2022.

[48] J. Shu, Q. Xie, L. Yi, Q. Zhao, S. Zhou, Z. Xu, and D. Meng, *Meta-Weight-Net: Learning an Explicit Mapping for Sample Weighting*. Red Hook, NY, USA: Curran Associates Inc., 2019.

[49] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.

[50] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender, "Learning to rank using gradient descent," in *Proceedings of the 22nd International Conference on Machine Learning*. New York, NY, USA: ACM, 2005, p. 89–96.

[51] Z. Tao, Y. Li, B. Ding, C. Zhang, J. Zhou, and Y. Fu, "Learning to mutate with hypergradient guided population," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 17 641–17 651.

[52] X. Yang, J. Wang, X. Zhao, S. Li, and Z. Tao, "Calibrate automated graph neural network via hyperparameter uncertainty," in *CIKM*, 2022.

[53] C. J. Burges, "From ranknet to lambdarank to lambdamart: An overview," *Learning*, vol. 11, no. 23-581, p. 81, 2010.

[54] N. Craswell, A. P. De Vries, and I. Soboroff, "Overview of the trec 2005 enterprise track." in *Trec*, vol. 5, 2005, pp. 1–7.

[55] C. Goutte and E. Gaussier, "A probabilistic interpretation of precision, recall and f-score, with implication for evaluation," in *Advances in Information Retrieval*, D. E. Losada and J. M. Fernández-Luna, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 345–359.

[56] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *International conference on machine learning*. PMLR, 2013, pp. 1139–1147.

Yuan Wang received the B.S. degree in Department of Mathematics and Computer Science and the M.S. degree in Department of Computer Science and Engineering, Santa Clara University, Santa Clara, California, US, in 2018 and 2020, respectively. Currently, he is a PhD student in Department of Computer Science and Engineering, Santa Clara University. His research interests include fairness ranking and deep learning.

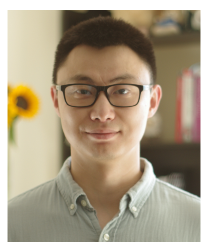

Zhiqiang Tao (M'22) received the B.Eng. degree in software engineering from the School of Computer Software, and the M.Eng. degree in computer science from the School of Computer Science and Technology, Tianjin University, Tianjin, China, in 2012 and 2015, respectively. He received the Ph.D. degree from the Department of Electrical and Computer Engineering, Northeastern University, Boston MA, 2020. He has been a Tenure-Track Assistant Professor at the School of Information, Rochester Institute of Technology, since 2022. Before that, he was an Assistant Professor at the Department of Computer Science and Engineering, Santa Clara University, from 2020 to 2022. He has served as the Associate Editor of Neurocomputing (2022-) and IEEE TCSVT (2023-).

Yi Fang is an Associate Professor in the Department of Computer Science and Engineering at Santa Clara University. He received his Ph.D. degree in Computer Science from Purdue University, West Lafayette, in 2012. His research interests lie in information retrieval, data mining, and machine learning. He served as the Chair of the Steering Committee for ACM International Conference on the Theory of Information Retrieval (ICTIR) from 2022 to 2023, the General co-Chair for ICTIR 2019, and Program co-Chair for ICTIR 2021. He has served as Senior Program Committee members/Area Chairs in SIGIR, WWW, ACL, EMNLP, NAACL, CIKM, WSDM, SIGIR-AP, and AAAI. He received the Outstanding Service Award at ACM CIKM 2013 and IEEE ICME 2013.