
Fairly Recommending with Social Attributes: A Flexible and Controllable Optimization Approach

Jinqiu Jin^{1,*} Haoxuan Li^{2,*} Fuli Feng^{1,†} Sihao Ding¹ Peng Wu³ Xiangnan He¹

¹University of Science and Technology of China

²Peking University ³Beijing Technology and Business University

jjq20021015@mail.ustc.edu.cn, hxli@stu.pku.edu.cn, fulifeng93@gmail.com,
dsihao@mail.ustc.edu.cn, pengwu@btbu.edu.cn, xiangnanhe@gmail.com

Abstract

Item-side group fairness (IGF) requires a recommendation model to treat different item groups similarly, and has a crucial impact on information diffusion, consumption activity, and market equilibrium. Previous IGF notions only focus on the *direct* utility of the item exposures, *i.e.*, the exposure numbers across different item groups. Nevertheless, the item exposures also facilitate utility gained from the neighboring users via social influence, called *social* utility, such as information sharing on the social media. To fill this gap, this paper introduces two social attribute-aware IGF metrics, which require similar user social attributes on the exposed items across the different item groups. In light of the trade-off between the direct utility and social utility, we formulate a new multi-objective optimization problem for training recommender models with flexible trade-off while ensuring controllable accuracy. To solve this problem, we develop a gradient-based optimization algorithm and theoretically show that the proposed algorithm can find Pareto optimal solutions with varying trade-off and guaranteed accuracy. Extensive experiments on two real-world datasets validate the effectiveness of our approach. Our codes are available at https://github.com/mitao-cat/nips23_social_igf.

1 Introduction

Developing fair recommendation algorithms is crucial to perform reliable information search and decision making, which prevents users’ as well as platforms’ interests from being sacrificed [1]. One particular aspect of recommendation fairness is item-side group fairness (IGF), which requires a recommendation model to treat different item groups similarly, and the items are grouped based on attributes such as category or brand [2, 3, 4]. Existing IGF frameworks primarily focus on the *direct* utility of item exposures, requiring a similar number of exposures for different item groups, and can be broadly categorized into two forms: Statistical Parity (SP) [5, 6] and Equal Opportunity (EO) [7].

Nevertheless, the existing IGF notions overlook the *social* utility of item exposures, as users with different social attributes may produce varying utilities due to the social influence among users. For example, a user’s friends may view a recommended micro-video due to sharing activities or public record in the user’s timeline. This can lead to unfairness problem in which some item groups benefit more than others though the existing IGF notions are satisfied. Figure 1a illustrates a toy example where item group A receives the same number of exposures as group B, thus satisfying the requirement of the previous SP, while items in group A are recommended to users with greater numbers of friends. This motivates us to consider both direct and social utility for IGF notions.

*Equal contribution.

†Corresponding author.

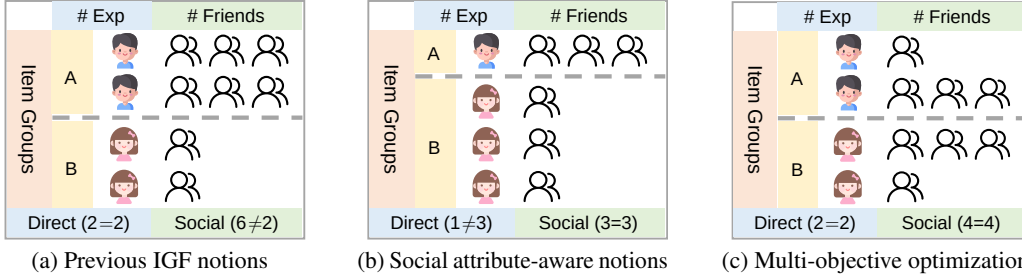


Figure 1: A toy example with item groups A and B to illustrate direct and social utilities of the IGF notions, where "# Exp" is the number of exposed users across the item group, and "# Friends" is the number of friends of the exposed users. The multi-objective optimization considers both utilities.

In this paper, we first introduce two social attribute-aware IGF metrics, namely Neighborhood SP (NSP) and Neighborhood EO (NEO), which require that users exposed to different item groups have similar total social utility. As shown in Figure 1b, item groups A and B have the same total social utility, which is measured by the number of friends in this toy example. Compared to the previous SP and EO notions, NSP and NEO require similarity in the neighbors of the exposed users, rather than simply requiring the similar number of exposed users across the item groups. However, optimizing only the social attribute-aware IGF metrics may result in varying direct utilities across different item groups, *e.g.*, item groups A and B in Figure 1b have the different number of exposed users. Therefore, it is reasonable to consider both direct and social utility when developing fair recommender models.

To this end, we next formulate a multi-objective optimization problem for training the recommender models with consideration of both direct and social utility, and further require flexibility and controllability of the trade-offs. Specifically, motivated by the previous study [8], we incorporate pre-defined preference regions with varying trade-offs between the direct and social utility into the multi-objective optimization problem, enabling the flexibility to choose the desired trade-offs in practice. Meanwhile, we also incorporate a customizable accuracy constraint for the controllability of the optimization, which leads to a guaranteed recommendation accuracy while satisfying the IGF notions.

We further propose a **S**ocial attribute-aware **F**lexible fair recommendation training algorithm with controllable Accuracy (SoFA), for solving the above multi-objective optimization problem. Given a pre-defined preference region, the proposed gradient-based algorithm updates the model parameters within this region to achieve the desired trade-off. On the other hand, when the accuracy loss exceeds a given threshold, the model parameters update direction is changed to achieve the minimum accuracy guarantee. We also theoretically show that SoFA can find Pareto optimal solutions with varying trade-off, meaning the direct and social utility of the trained model is not dominated by any other solutions. Extensive experiments are conducted on two real-world datasets with user social attributes, and the experimental results validate the effectiveness of the proposed training algorithm.

The main contributions of this paper are summarized as follows:

- We propose two social attribute-aware IGF metrics, named NSP and NEO, to study the item exposure utility gained from the user social network.
- We formalize a multi-objective optimization problem to achieve flexible trade-off between the direct utility and social utility with controllable accuracy sacrifice of the recommendation.
- We further propose a gradient-based optimization algorithm to solve the above problem, called SoFA, and theoretically show that SoFA can find Pareto optimal solutions with varying trade-off.
- We conduct extensive experiments on two real-world datasets, revealing that methods ignoring social influence would lead to unfair exposure, and validating the effectiveness of our proposal.

2 Preliminaries

In this section, we present the background about personalized recommendation, item-side group fairness including SP and EO notions, and multi-objective optimization.

Personalized Recommendation. Personalized recommendation aims to select a small subset of items to each user u based on historical data, with users denoted as $\mathcal{U} = \{u\}$ and items denoted as

$\mathcal{I} = \{i\}$. In our study, we focus on a common recommendation scenario known as collaborative filtering (CF) with implicit feedback, where the historical data contains the items that each user has interacted with [9, 10, 11, 12]. Bayesian Personalized Ranking (BPR) [9] is one of the representative CF algorithm modeling the interaction probabilities, by minimizing the pairwise ranking loss

$$\min_{\theta} \mathcal{L}_{\text{BPR}}(\theta) = - \sum_{(u,i,j) \in \mathcal{D}} \ln \sigma(\hat{y}_{u,i} - \hat{y}_{u,j}) + \frac{\lambda_{\theta}}{2} \|\theta\|_{\text{F}}^2, \quad (1)$$

where \mathcal{D} denotes the historical data for training, including user u , interacted item i , and no interacted item j . $\sigma(\cdot)$ is the Sigmoid function, θ represents the model parameters, and λ_{θ} is the regularization hyper-parameter to prevent overfitting. The predicted user preference scores $\hat{y}_{u,i}$ and $\hat{y}_{u,j}$ for positive item i and negative item j can be computed by using the widely-adopted matrix factorization [13].

Item-Side Group Fairness (IGF). The IGF requires fairness of recommendation holds across a set of item groups $\mathcal{G} = \{g_1, g_2, \dots, g_A\}$, where each item $i \in \mathcal{I}$ belongs to one or more groups. The group of an item is correspond to attributes such as genre, brand, or other item characteristics, and denote $G_{g_a}(i)$ as the indicator function of whether item i belongs to group g_a , i.e., $G_{g_a}(i) = 1$ if $i \in g_a$, otherwise $G_{g_a}(i) = 0$. IGF notions aim to ensure fairness across different groups without over-recommending or under-recommending any specific group [14]. The existing IGF notions mainly focus on the exposure levels across different groups, and can be classified into positive rate-based metrics and confusion matrix-based metrics [1]. One particular positive rate-based metric is *Statistical Parity* (SP) [5, 6], which aims to ensure that each group has an equal likelihood of being recommended and requires the following quantities to be similar across different item groups

$$\mathbb{P}(R | g = g_a) = \frac{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} G_{g_a}(i) \cdot \hat{Y}(u, i)}{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} G_{g_a}(i)}, \quad \text{for all } a = 1, \dots, A, \quad (2)$$

where $\hat{Y}(u, i) = \sigma(\hat{y}_{u,i})$ estimates the interaction probability for user u to item i , and $\sum_{i \in \mathcal{I}} G_{g_a}(i)$ is the total number of items belonging to group g_a . Instead, confusion matrix-based metrics further consider the ground truth labels. For instance, *Equal Opportunity* (EO) [7] aims to ensure the same true positive rate across different item groups

$$\mathbb{P}(R | g = g_a, y = 1) = \frac{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} G_{g_a}(i) \cdot Y(u, i) \cdot \hat{Y}(u, i)}{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} G_{g_a}(i) \cdot Y(u, i)}, \quad \text{for all } a = 1, \dots, A, \quad (3)$$

where $Y(u, i)$ is the ground truth label for user u and item i , and $\sum_{i \in \mathcal{I}} G_{g_a}(i) Y(u, i)$ computes the total number of items that user u has interacted with in group g_a . To keep the same value scale for different metrics, we evaluate the IGF of a recommendation model by computing the relative standard deviation of the group probabilities

$$\begin{aligned} \text{SP} &= \text{rsd}(\mathbb{P}(R | g = g_1), \dots, \mathbb{P}(R | g = g_A)), \\ \text{EO} &= \text{rsd}(\mathbb{P}(R | g = g_1, y = 1), \dots, \mathbb{P}(R | g = g_A, y = 1)), \end{aligned} \quad (4)$$

where $\text{rsd}(\cdot) = \text{std}(\cdot) / \text{mean}(\cdot)$ calculates the relative standard deviation. The lower the SP and EO values, the fairer the recommender system satisfies the IGF fairness notions.

Multi-Objective Optimization. In many real-world tasks, there might be multiple optimization objectives rather than single one, while these objectives may collaborate or conflict with each other [15, 16]. Multi-objective optimization [17] aims to find a set of solutions that effectively balance the trade-off among these objectives, and it has been widely used in areas such as reinforcement learning [18] and E-commerce [19]. Typically, a multi-objective optimization problem is

$$\min_{\theta} \mathcal{L}(\theta) = (\mathcal{L}_1(\theta), \mathcal{L}_2(\theta), \dots, \mathcal{L}_M(\theta))^{\text{T}}, \quad (5)$$

where $(\mathcal{L}_1(\theta), \dots, \mathcal{L}_M(\theta))$ states the multiple objectives. To solve the multi-objective optimization problem, several population-based and evolutionary algorithm-based methods [20, 21, 22] have been proposed. Nevertheless, they are not efficient in handling large-scale datasets in recommendation. Alternatively, multi-objective gradient descent approaches [23, 24] address this problem by leveraging Karush-Kuhn-Tucker (KKT) conditions [25] to find a gradient direction that reduces all objective values simultaneously. For example, gradient-based steepest descent methods [23] are proposed to seek a descent direction \mathbf{d}_t by solving the following optimization problem

$$(\mathbf{d}_t, \alpha_t) = \arg \min_{\mathbf{d}, \alpha \in \mathbb{R}} \alpha + \frac{1}{2} \|\mathbf{d}\|^2, \quad \text{s.t. } \nabla \mathcal{L}_i(\theta)^{\text{T}} \mathbf{d} \leq \alpha, \quad i = 1, \dots, M, \quad (6)$$

then updates the model parameters using gradient descent $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \eta \mathbf{d}_t$ with step size η . Theoretically, the solutions of the descent direction guarantee that all objective values will decrease, thus achieving the Pareto optimality [26]. To obtain these descent direction solutions, one can employ the widely-adopted multiple gradient descent algorithm (MGDA) [24].

3 Social Attributes-Aware IGF

Existing IGF notions such as SP and EO define the utility of recommending an item to a user only depends on the exposure and interaction numbers. For example, if item i_1 and item i_2 are recommended to user u_1 and user u_2 respectively, then both i_1 and i_2 have the same utility when $\hat{Y}(u_1, i_1) = \hat{Y}(u_2, i_2)$. However, they overlook the social attributes of users. Considering that the user social network plays a crucial role in click or conversion behaviors in real-world recommendation scenarios [27], recommending to users with distinct social attributes may lead to diverse item utilities.

To bridge this gap, we propose social attribute-aware IGF metrics as the extension of previously widely used SP and EO notions. Denote the neighbors (such as friends or followers) of user u as \mathcal{N}_u , for each user $v \in \mathcal{N}_u$, let $R_v(i) > 0$ be the utility of item i gained from user v via the social influence of user u , when the item i is recommended to the user u . Therefore, for positive rate-based metrics and confusion matrix-based metrics, the following quantities are required to be similar

$$\begin{aligned} \text{NR}(g = g_a) &= \frac{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} G_{g_a}(i) \cdot \hat{Y}(u, i) \sum_{v \in \mathcal{N}_u} R_v(i)}{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} G_{g_a}(i)}, \quad \text{for all } a = 1, \dots, A, \\ \text{NR}(g = g_a, y = 1) &= \frac{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} G_{g_a}(i) \cdot Y(u, i) \cdot \hat{Y}(u, i) \sum_{v \in \mathcal{N}_u} R_v(i)}{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} G_{g_a}(i) \cdot Y(u, i)}, \quad \text{for all } a = 1, \dots, A, \end{aligned} \quad (7)$$

where the sum of utilities from the social network $\sum_{v \in \mathcal{N}_u} R_v$ are multiplied by the interaction probability $\hat{Y}(u, i)$ to measure the total social utility. Similar to Eq. (4), we then compute the *Neighborhood SP* and *Neighborhood EO* by the relative standard deviations

$$\begin{aligned} \text{NSP} &= \text{rsd}(\text{NR}(g = g_1), \dots, \text{NR}(g = g_A)), \\ \text{NEO} &= \text{rsd}(\text{NR}(g = g_1, y = 1), \dots, \text{NR}(g = g_A, y = 1)). \end{aligned} \quad (8)$$

Further Discussion. For the existing and the proposed IGF notions, the former including SP and EO focus on the direct utility obtained through recommendation exposures to users, whereas the latter including SP and EO emphasize the social utility gained from the user social network. Since these metrics emphasize different aspects, optimizing one IGF notion alone does not guarantee the optimality of the other IGF notion. Figures 1a and 1b illustrate this with a toy example: if we solely optimize SP, it does not guarantee the achievement of the other IGF notion, *i.e.*, NSP, and vice versa. Since both IGF metrics are important, it is desirable to optimize both SP (EO) and NSP (NEO) simultaneously from the prospective of multi-objective optimization, as illustrated in Figure 1c.

4 Methodology

4.1 Multi-Objective Optimization Problem Formulation

The optimization objective for training a fair recommendation model under IGF notions is to minimize both SP (EO) and NSP (NEO), so as to obtain similar direct utility and similar social utility simultaneously. Building upon prior studies [8, 28], we formulate the fair recommendation training task as a multi-objective optimization problem, where each optimization objective corresponds to a pre-defined IGF metric. Without loss of generality, let the number of IGF metrics used as optimization objectives be M , then we aim to obtain Pareto optimal solutions among these IGF metrics.

Pareto Dominance. A solution $\boldsymbol{\theta}_1$ dominates another solution $\boldsymbol{\theta}_2$ if $\mathcal{L}_i(\boldsymbol{\theta}_1) \leq \mathcal{L}_i(\boldsymbol{\theta}_2)$, $\forall i = 1, \dots, M$ and $\mathcal{L}(\boldsymbol{\theta}_1) = (\mathcal{L}_1(\boldsymbol{\theta}_1), \dots, \mathcal{L}_M(\boldsymbol{\theta}_1))^T \neq \mathcal{L}(\boldsymbol{\theta}_2) = (\mathcal{L}_1(\boldsymbol{\theta}_2), \dots, \mathcal{L}_M(\boldsymbol{\theta}_2))^T$.

Pareto Optimality. A solution $\boldsymbol{\theta}$ is a Pareto optimal solution if $\boldsymbol{\theta}$ is not dominated by any other solutions. The set of the Pareto optimal solutions is called as Pareto optimal set.

It is now attractive to formulate the trade-off between different IGF metrics as a multi-objective optimization problem, then obtain a Pareto-optimal solution between these metrics. However, due to the well-known accuracy sacrifice for improving fairness [14], directly solving the above problem without considering recommendation accuracy may lead to a significant reduction in recommendation quality. To ensure controllable recommendation accuracy during the optimization process, we introduce an accuracy constraint that penalizes instances with the accuracy loss (*e.g.*, BPR loss) exceeding a pre-defined threshold. Formally, the multi-objective optimization problem is

$$\min_{\theta} \mathcal{L}(\theta) = (\mathcal{L}_1(\theta), \mathcal{L}_2(\theta), \dots, \mathcal{L}_M(\theta))^T, \text{ s.t. } \mathcal{L}_{\text{BPR}}(\theta) \leq \xi, \quad (9)$$

where ξ is the pre-defined accuracy threshold, and ξ should be set to exceed the loss of a recommendation model trained using only BPR loss to ensure the existence of feasible solutions.

Furthermore, in real-world recommendation scenarios, the two types of IGF metrics that consider direct and social effects, respectively, may require different proportions of trade-offs. For example, a recommender system designed for social media applications may place more emphasis on the social utility than the direct utility. Therefore, it is meaningful to find flexible Pareto-optimal solutions to the multi-objective optimization problem with varying trade-offs. Following the previous work [8], we decompose the optimization problem into N subproblems using a set of pre-defined unit preference vectors $\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N \in R_+^M$, and define the N preference regions $\Omega_1, \Omega_2, \dots, \Omega_N$ as follows

$$\Omega_n = \{\mathcal{L}(\theta) \in R_+^M \mid \mathbf{s}_j^T \mathcal{L}(\theta) \leq \mathbf{s}_n^T \mathcal{L}(\theta), \forall j = 1, \dots, N\}, \quad (10)$$

where $\mathcal{L}(\theta) \in \Omega_n$ if and only if $\mathcal{L}(\theta)$ forms the smallest acute angle with \mathbf{s}_n , resulting in the largest inner product $\mathbf{s}_n^T \mathcal{L}(\theta)$ among all $\mathbf{s}_1^T \mathcal{L}(\theta), \mathbf{s}_2^T \mathcal{L}(\theta), \dots, \mathbf{s}_N^T \mathcal{L}(\theta)$. To ensure that the solution $\mathcal{L}(\theta)$ obtained from the optimization phase falls within the preference region Ω_n , we further impose a constraint that penalizes the distance of the obtained solution from this preference region

$$\begin{aligned} \min_{\theta} \mathcal{L}(\theta) &= (\mathcal{L}_1(\theta), \mathcal{L}_2(\theta), \dots, \mathcal{L}_M(\theta))^T \\ \text{s.t. } \mathcal{G}_j(\theta) &= (\mathbf{s}_j - \mathbf{s}_n)^T \mathcal{L}(\theta) \leq 0, \forall j = 1, \dots, N, \\ \mathcal{L}_{\text{BPR}}(\theta) &\leq \xi. \end{aligned} \quad (11)$$

By letting the preference vectors be uniformly distributed in R_+^M , we can obtain flexible Pareto optimal solutions with varying trade-offs. Each subproblem corresponds to a unique Pareto-optimal solution in a specific preference region, which can reflect different attention to the direct and social utilities. Finally, we can select the desired Pareto optimal solution for diverse recommendation scenarios to satisfy both the flexible trade-offs and the controllable prediction accuracy.

4.2 Gradient-Based Flexible and Controllable Fair Recommendation Training

Finding the Initial Solution. For more efficient finding a solution to the multi-objective problem that falls within the specified preference region, we first find a feasible initial solution that satisfies all the constraints of that preference region. Specifically, for a randomly generated initial solution θ_r , we define the set of indices that violate the constraints as $I_\epsilon(\theta_r) = \{j = 1, \dots, N \mid \mathcal{G}_j(\theta_r) \geq -\epsilon\}$, where ϵ is a small value to deal with solutions near the boundary. Then we can compute a descending direction \mathbf{d}_r that reduces all the values of $\{\mathcal{G}_j(\theta_r) \mid j \in I_\epsilon(\theta_r)\}$ by solving the optimization problem

$$(\mathbf{d}_r, \alpha_r) = \arg \min_{\mathbf{d}, \alpha \in R} \alpha + \frac{1}{2} \|\mathbf{d}\|^2, \text{ s.t. } \nabla \mathcal{G}_j(\theta_r)^T \mathbf{d} \leq \alpha, j \in I_\epsilon(\theta_r). \quad (12)$$

The gradient-based update rule can be expressed as $\theta_{r_{t+1}} = \theta_{r_t} + \eta_r \mathbf{d}_{r_t}$, where η_r denotes the step size, and will stop when a feasible solution is found or the maximum iteration number is reached.

Solving the Subproblem. Given the feasible initial solution θ_0 , we next propose a gradient-based learning approach to solve the multi-objective optimization problem in Eq. (11). Specifically, we compute the update direction \mathbf{d}_t from θ_t to θ_{t+1} by solving the following optimization problem

$$\begin{aligned} (\mathbf{d}_t, \alpha_t) &= \arg \min_{\mathbf{d}, \alpha \in R} \alpha + \frac{1}{2} \|\mathbf{d}\|^2 \\ \text{s.t. } \nabla \mathcal{L}_i(\theta_t)^T \mathbf{d} &\leq \alpha, i = 1, \dots, M, \\ \nabla \mathcal{G}_j(\theta_t)^T \mathbf{d} &\leq \alpha, j \in I_\epsilon(\theta_t), \\ \nabla \mathcal{L}_{\text{BPR}}(\theta_t)^T \mathbf{d} &\leq \alpha, \text{ if } \mathcal{L}_{\text{BPR}}(\theta_t) \geq \xi. \end{aligned} \quad (13)$$

The Lagrange function of the above optimization problem is

$$\begin{aligned} \mathcal{L}(\mathbf{d}, \alpha, \alpha_i, \beta_j, \lambda) &= \alpha + \frac{1}{2} \|\mathbf{d}\|^2 + \sum_{i=1}^M \alpha_i (\nabla_{\boldsymbol{\theta}_t} \mathcal{L}_i(\boldsymbol{\theta}_t)^T \mathbf{d} - \alpha) \\ &+ \sum_{j \in I_\epsilon(\boldsymbol{\theta}_t)} \beta_j (\nabla_{\boldsymbol{\theta}_t} \mathcal{G}_j(\boldsymbol{\theta}_t)^T \mathbf{d} - \alpha) + \lambda \cdot \mathbb{I}(\mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi) (\nabla_{\boldsymbol{\theta}_t} \mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t)^T \mathbf{d} - \alpha), \end{aligned} \quad (14)$$

where $\alpha_i \geq 0, \beta_j \geq 0, \lambda \geq 0$ are the Lagrange multipliers. By requiring that the derivatives of $\mathcal{L}(\mathbf{d}, \alpha, \alpha_i, \beta_j, \lambda)$ with respect to both \mathbf{d} and α be zero, we have the following equations

$$\begin{aligned} \mathbf{d} &= - \sum_{i=1}^M \alpha_i \nabla_{\boldsymbol{\theta}_t} \mathcal{L}_i(\boldsymbol{\theta}_t) - \sum_{j \in I_\epsilon(\boldsymbol{\theta}_t)} \beta_j \nabla_{\boldsymbol{\theta}_t} \mathcal{G}_j(\boldsymbol{\theta}_t) - \lambda \cdot \mathbb{I}(\mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi) \nabla_{\boldsymbol{\theta}_t} \mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t), \\ \sum_{i=1}^M \alpha_i + \sum_{j \in I_\epsilon(\boldsymbol{\theta}_t)} \beta_j + \lambda \cdot \mathbb{I}(\mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi) &= 1. \end{aligned} \quad (15)$$

Next, we compute the minimal value of Eq. (14) by substituting the corresponding terms in Eq. (15), and obtain the dual problem of Eq. (13) as $\max_{\alpha_i, \beta_j, \lambda} \min_{\mathbf{d}, \alpha} \mathcal{L}(\mathbf{d}, \alpha, \alpha_i, \beta_j, \lambda)$, equals to

$$\begin{aligned} \min_{\alpha_i, \beta_j, \lambda} \frac{1}{2} \left\| \sum_{i=1}^M \alpha_i \nabla_{\boldsymbol{\theta}_t} \mathcal{L}_i(\boldsymbol{\theta}_t) + \sum_{j \in I_\epsilon(\boldsymbol{\theta}_t)} \beta_j \nabla_{\boldsymbol{\theta}_t} \mathcal{G}_j(\boldsymbol{\theta}_t) + \lambda \cdot \mathbb{I}(\mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi) \nabla_{\boldsymbol{\theta}_t} \mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \right\|^2, \\ \text{s.t. } \sum_{i=1}^M \alpha_i + \sum_{j \in I_\epsilon(\boldsymbol{\theta}_t)} \beta_j + \lambda \cdot \mathbb{I}(\mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi) &= 1, \end{aligned} \quad (16)$$

which can be efficiently solved by the gradient-based optimization methods such as MGDA [24].

Denote the solutions of Eq. (13) and Eq. (16) as (\mathbf{d}^*, α^*) and $(\alpha_i^*, \beta_j^*, \lambda^*)$, respectively. To obtain the solution of the dual problem, according to the KKT conditions, we have

$$\begin{aligned} \alpha_i^* (\nabla_{\boldsymbol{\theta}_t} \mathcal{L}_i(\boldsymbol{\theta}_t)^T \mathbf{d}^* - \alpha^*) &= 0, \quad i = 1, \dots, M, \\ \beta_j^* (\nabla_{\boldsymbol{\theta}_t} \mathcal{G}_j(\boldsymbol{\theta}_t)^T \mathbf{d}^* - \alpha^*) &= 0, \quad \forall j \in I_\epsilon(\boldsymbol{\theta}_t), \\ \lambda \cdot \mathbb{I}(\mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi) (\nabla_{\boldsymbol{\theta}_t} \mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t)^T \mathbf{d}^* - \alpha^*) &= 0, \quad \text{if } \mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi. \end{aligned} \quad (17)$$

By adding all equations in Eq. (17) together, we have $\alpha^* = -\|\mathbf{d}^*\|^2$. In this way,

- If $\boldsymbol{\theta}_t$ is Pareto optimal, then no other solution in its neighborhood can achieve better objective values, and we obtain the solution $\mathbf{d}^* = \mathbf{0}$, indicating that no direction can simultaneously improve all objective values.
- If $\boldsymbol{\theta}_t$ is not Pareto optimal, we have the following conclusions from Eq. (13)

$$\begin{aligned} \nabla \mathcal{L}_i(\boldsymbol{\theta}_t)^T \mathbf{d}^* \leq \alpha^* \leq -\|\mathbf{d}^*\|^2 < 0, \quad i = 1, \dots, M, \\ \nabla \mathcal{G}_j(\boldsymbol{\theta}_t)^T \mathbf{d}^* \leq \alpha^* \leq -\|\mathbf{d}^*\|^2 < 0, \quad j \in I_\epsilon(\boldsymbol{\theta}_t), \\ \nabla \mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t)^T \mathbf{d}^* \leq \alpha^* \leq -\|\mathbf{d}^*\|^2 < 0, \quad \text{if } \mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi, \end{aligned} \quad (18)$$

therefore \mathbf{d}^* will be a descent direction that at least decreases all the IGF losses simultaneously, as well as the recommendation accuracy loss when $\mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t) \geq \xi$.

Finally, the model parameters are updated as $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \eta \mathbf{d}_t$, where η denotes the step size. Remarkably, throughout the optimization procedure, the values of all IGF objectives and the distance between the obtained solution and the preference region decrease, indicating the Pareto optimality within the preference region. Moreover, to tackle the problem of high-dimensional parameter $\boldsymbol{\theta}$ in the multi-objective optimization, we propose to optimize its dual problem rather than directly solving the primal problem. As a result, the numbers of parameters (α_i, β_j) in the dual problem is significantly reduced to be same as the numbers of IGF objectives and constraints, respectively, which greatly increases the scalability of the proposed algorithm to the large-scale recommendation datasets. We summarize the whole gradient-based fair recommendation training process in Algorithm 1.

Algorithm 1: Social-Aware Flexible Fair Recommendation with Controllable Accuracy (SoFA)

Input: Boundary threshold ϵ , accuracy threshold ξ , step size η_r and η .

- 1 Set the fairness preference vectors $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N\}$;
- 2 **for** $n = 1$ to N **do**
- 3 Randomly initialize the model parameters $\boldsymbol{\theta}_r^{(n)}$;
- 4 Find the nearest feasible solution $\boldsymbol{\theta}_0^{(n)}$ from $\boldsymbol{\theta}_r^{(n)}$ for region Ω_n ;
- 5 **for** $t = 0$ to T **do**
- 6 **if** $\mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t^{(n)}) \leq \xi$ **then**
- 7 Obtain $\alpha_i, i = 1, \dots, M$, and $\beta_j, j \in I_\epsilon(\boldsymbol{\theta}_t^{(n)})$ by solving the dual optimization;
- 8 Compute the update direction
 $\mathbf{d}_t^{(n)} = -\sum_{i=1}^M \alpha_i \nabla_{\boldsymbol{\theta}_t^{(n)}} \mathcal{L}_i(\boldsymbol{\theta}_t^{(n)}) - \sum_{j \in I_\epsilon(\boldsymbol{\theta}_t^{(n)})} \beta_j \nabla_{\boldsymbol{\theta}_t^{(n)}} \mathcal{G}_j(\boldsymbol{\theta}_t^{(n)})$;
- 9 **else**
- 10 Obtain $\alpha_i, i = 1, \dots, M, \beta_j, j \in I_\epsilon(\boldsymbol{\theta}_t^{(n)})$, and λ by solving the dual optimization;
- 11 Compute the update direction $\mathbf{d}_t^{(n)} =$
 $-\sum_{i=1}^M \alpha_i \nabla_{\boldsymbol{\theta}_t^{(n)}} \mathcal{L}_i(\boldsymbol{\theta}_t^{(n)}) - \sum_{j \in I_\epsilon(\boldsymbol{\theta}_t^{(n)})} \beta_j \nabla_{\boldsymbol{\theta}_t^{(n)}} \mathcal{G}_j(\boldsymbol{\theta}_t^{(n)}) - \lambda \nabla_{\boldsymbol{\theta}_t^{(n)}} \mathcal{L}_{\text{BPR}}(\boldsymbol{\theta}_t^{(n)})$;
- 12 **end**
- 13 Update the model parameters $\boldsymbol{\theta}_{t+1}^{(n)} = \boldsymbol{\theta}_t^{(n)} + \eta \mathbf{d}_t^{(n)}$;
- 14 **end**
- 15 **end**

5 Experiments

In this section we aim to answer the following research questions:

- **RQ1:** How does our proposed method perform compared with the existing approaches?
- **RQ2:** Do the preference regions in our method lead to more flexible Pareto solutions? Does the accuracy constraint result in more accurate recommendations while achieving IGF?
- **RQ3:** How does varying the accuracy thresholds affect the performance of our method?

5.1 Experiment Setups

Datasets and Pre-processing. The experiments are performed on two benchmark recommendation datasets that include user social attributes and item features:

- **KuaiRec** is a short-video recommendation dataset on the video-sharing platform Kuaishou³, which contains interactions between users and videos. In our experiments, we select user clicks spanning 10 consecutive days as the outcome of interest [29] and apply 10-core filtering to eliminate users and items with minimal interactions, which results in a pre-processed dataset consisting of 7,010 users, 1,405 items, and 863,819 interactions. For each user, the number of followers of that user is treated as the user social attribute, and we define the social utility derived from the user social attribute as the total fans count of that user. As to videos, we use tags to categorize them into six item groups.

- **Epinions** is a dataset derived from a trust network, comprising users' rating scores for products. Following previous work [10], we only keep the interactions with ratings greater than 3 and also adopt the 10-core filtering setting. Subsequently, we obtain a dataset consisting of 11,710 users, 13,682 items, and 215,262 interactions. For each user, the dataset includes their trust relations with other users, while for each product, it specifies its category. We let user's social utility equal to the number of trustors, and classify the products into six groups based on their categories.

Throughout our experiment, we randomly split the interactions of each dataset into training set (60%), validation set (20%), and testing set (20%). We use the validation set for tuning hyper-parameters and the testing data for evaluation the prediction and fairness performance of the trained models.

³<https://www.kuaishou.com/>.

Table 1: Performance comparison using SP and NSP as IGF notions, where SoFA is implemented with five preference regions. The best and second best results are bolded and underlined, respectively.

	KuaiRec				Epinions			
	N@5 \uparrow	SP \downarrow	NSP \downarrow	FISP \downarrow_{deg}	N@5 \uparrow	SP \downarrow	NSP \downarrow	FISP \downarrow_{deg}
BPRMF	0.2426	0.0966	0.1119	0.1037 49.2°	0.0443	0.0252	0.0286	0.0268 48.6°
+ SP Reg	0.2389	<u>0.0062</u>	0.0168	<u>0.0091</u> 69.7°	<u>0.0450</u>	0.0140	0.0196	0.0163 54.5°
+ NSP Reg	0.2279	0.0366	0.0142	0.0205 21.2°	0.0378	0.0224	0.0188	0.0205 40.0°
+ SP&NSP Reg	0.2369	0.0090	0.0245	0.0132 69.8°	0.0448	0.0154	0.0205	0.0176 53.2°
+ SP Post	0.2412	0.0388	0.0545	0.0454 54.5°	0.0445	0.0141	0.0196	0.0164 54.2°
+ NSP Post	0.2348	0.0844	0.0311	0.0455 20.3°	0.0398	0.0212	<u>0.0185</u>	0.0197 41.2°
+ SP&NSP Post	0.2405	0.0817	0.0562	0.0666 34.5°	0.0443	0.0152	0.0207	0.0175 53.7°
MOOMTL	0.2229	0.0069	0.0238	0.0107 73.8°	0.0446	0.0138	0.0193	0.0161 54.4°
SoFA _{region 0}	0.2349	0.0296	0.0096	0.0145 18.0°	0.0364	0.0909	0.0294	0.0445 17.9°
SoFA _{region 1}	0.2376	0.0179	<u>0.0105</u>	0.0133 30.4°	0.0441	0.0326	0.0225	0.0266 34.6°
SoFA _{region 2}	0.2329	0.0103	0.0146	0.0121 54.8°	0.0451	0.0153	0.0210	0.0177 53.9°
SoFA _{region 3}	<u>0.2413</u>	0.0074	0.0194	0.0107 69.1°	0.0427	<u>0.0118</u>	0.0177	0.0142 56.3°
SoFA _{region 4}	0.2402	0.0046	0.0227	0.0077 78.5°	0.0185	0.0095	0.0314	<u>0.0146</u> 73.1°

Table 2: Performance comparison using EO and NEO as IGF notions, where SoFA is implemented with five preference regions. The best and second best results are bolded and underlined, respectively.

	KuaiRec				Epinions			
	N@5 \uparrow	EO \downarrow	NEO \downarrow	F1EO \downarrow_{deg}	N@5 \uparrow	EO \downarrow	NEO \downarrow	F1EO \downarrow_{deg}
BPRMF	0.2426	0.0189	0.1031	0.0319 79.6°	0.0443	0.1182	0.2059	0.1502 60.1°
+ EO Reg	0.2340	0.0044	0.1082	0.0085 87.7°	0.0425	0.1126	0.1956	0.1429 60.1°
+ NEO Reg	0.2329	0.0147	0.0992	0.0256 81.6°	0.0430	0.1289	0.1557	0.1410 50.4°
+ EO&NEO Reg	0.2346	0.0054	0.1035	0.0102 87.0°	0.0432	0.1136	0.1465	0.1280 52.2°
+ EO Post	0.2371	0.0061	0.1082	0.0116 86.8°	0.0437	0.1100	0.1850	0.1379 59.3°
+ NEO Post	0.2280	0.0143	0.0672	0.0236 78.0°	0.0427	0.1274	0.1565	0.1405 50.8°
+ EO&NEO Post	<u>0.2413</u>	0.0133	0.1039	0.0235 82.7°	0.0436	0.1131	0.1630	0.1335 55.3°
MOOMTL	0.2332	<u>0.0031</u>	0.1003	<u>0.0061</u> 88.2°	0.0444	0.1093	<u>0.1449</u>	0.1246 53.0°
SoFA _{region 0}	0.1635	0.1454	<u>0.0557</u>	0.0806 21.0°	0.0224	0.3273	0.1494	0.2051 24.5°
SoFA _{region 1}	0.2077	0.1135	0.0558	0.0749 26.2°	0.0353	0.1752	0.1394	0.1552 38.5°
SoFA _{region 2}	0.2254	0.0610	0.0528	0.0566 40.9°	0.0448	<u>0.1094</u>	0.1473	<u>0.1256</u> 53.4°
SoFA _{region 3}	0.2241	0.0284	0.0713	0.0406 68.3°	0.0443	0.1160	0.1604	0.1347 54.1°
SoFA _{region 4}	0.2352	0.0023	0.0685	0.0045 88.1°	<u>0.0447</u>	0.1157	0.1574	0.1334 53.7°

Evaluation Protocols. We evaluate the recommendation accuracy with NDCG@5 [30]. For evaluating IGF, we report both the previous fairness metrics and social attribute-aware metrics, where lower value indicates fairer performance. We study two settings with (SP, NSP) and (EO, NEO) as the IGF objectives, respectively. Under each setting, the overall fairness is assessed using the harmonic-mean of the two IGF metrics denoted as FISP. In addition, to quantify the trade-offs between the two IGF metrics, we calculate the angles formed by these metrics via $deg = \arctan(NSP/SP)$, where a lower angle degree signifies a greater emphasis on the social utility.

Hyper-Parameter Settings. For all compared methods, we use a pre-trained BPRMF [9] as the backbone with fixed optimal batch size and regularization coefficient, then fine-tune the model by using the baselines methods for achieving IGF notions. For the proposed SoFA, we set $N = 5$ as the number of preference vectors that equally divide the first quadrant, and tune the learning rate for finding the initial solution within $\{0.3lr, lr, 3lr\}$, where lr denotes the learning rate for the pre-trained BPRMF. This optimization step stops once a feasible solution is found. The learning rate for all methods is searched within $\{0.03lr, 0.1lr, 0.3lr, lr, 3lr, 10lr\}$, and the coefficients of IGF terms for regularization-based and post-processing baselines are searched within $\{0.1, 0.2, 0.5, 1, 2, 5\}$. The early stopping strategy is employed when the FISP or F1EO value does not decrease over 5 epochs.

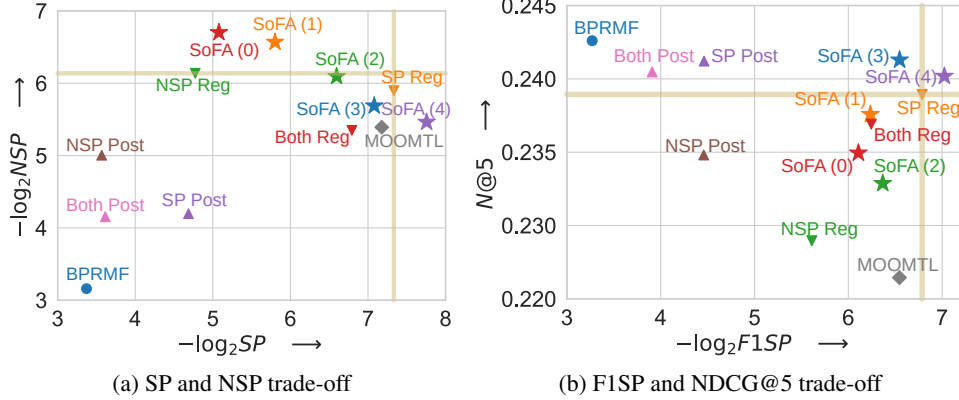


Figure 2: The trade-offs on KuaiRec between (a) IGF metrics, *i.e.*, SP and NSP, and (b) overall fairness and accuracy, *i.e.*, F1SP and NDCG, where yellow lines refer to the results when only optimizing one IGF metric, and scatters located on the upper right indicate better overall performance.

Table 3: Ablation studies of SoFA on the accuracy constraint, IGF objectives, and preference region.

	KuaiRec				Epinions					
	N@5 \uparrow	SP \downarrow	NSP \downarrow	F1SP \downarrow deg	N@5 \uparrow	SP \downarrow	NSP \downarrow	F1SP \downarrow deg		
SoFA _{region 0}	0.2349	0.0296	0.0096	0.0145	18.0°	0.0364	0.0909	0.0294	0.0445	17.9°
w/o con	0.2181	0.0415	0.0124	0.0191	16.6°	0.0339	0.0814	0.0260	0.0395	17.7°
SoFA _{region 1}	0.2376	0.0179	0.0105	0.0133	30.4°	0.0441	0.0326	0.0225	0.0266	34.6°
w/o con	0.2254	0.0211	0.0121	0.0154	29.8°	0.0429	0.0371	0.0267	0.0311	35.8°
SoFA _{region 2}	0.2329	0.0103	0.0146	0.0121	54.8°	0.0451	0.0153	0.0210	0.0177	53.9°
w/o con	0.2117	0.0085	0.0131	0.0103	57.0°	0.0445	0.0163	0.0224	0.0189	53.9°
SoFA _{region 3}	<u>0.2413</u>	0.0074	0.0194	0.0107	69.1°	0.0427	0.0118	<u>0.0177</u>	0.0142	56.3°
w/o con	0.2212	0.0077	0.0216	0.0113	70.4°	0.0417	0.0131	0.0222	0.0165	59.4°
SoFA _{region 4}	0.2402	0.0046	0.0227	<u>0.0077</u>	78.5°	0.0185	<u>0.0095</u>	0.0314	0.0146	73.1°
w/o con	0.2193	0.0069	0.0256	0.0109	74.9°	0.0169	0.0093	0.0313	<u>0.0144</u>	73.4°
w/o SP	0.2220	0.0206	0.0064	0.0098	17.4°	0.0365	0.0216	0.0147	0.0175	34.3°
w/o NSP	0.2415	0.0050	0.0173	<u>0.0077</u>	74.0°	<u>0.0447</u>	0.0148	0.0200	0.0170	53.5°
w/o region	0.2412	<u>0.0046</u>	0.0193	0.0074	76.7°	<u>0.0447</u>	0.0141	0.0197	0.0164	54.4°

5.2 Performance Comparison (RQ1)

We compare our proposed SoFA with the vanilla BPRMF [9] and several baseline methods to achieve recommendation fairness. These baselines include: (1) Regularization-based method [31], which incorporates IGF metrics as regularization terms into the BPR loss. (2) Post-processing method [32], which reconstructs the recommendation results while imposing fairness objectives as constraints. (3) Multi-objective optimization-based method [28], which optimizes the formulated multi-objective problem without the preference regions and recommendation accuracy constraint. We adopt three versions of regularization-based methods and post-processing methods, which considers only direct utility (SP or EO), only social utility (NSP or NEO), and both utilities, respectively.

Tables 1 and 2 show the overall performance of compared methods under the (SP, NSP) and (EO, NEO) settings, respectively, from which we have the following observations: (1) In all cases, SoFA outperforms other baselines in terms of accuracy and fairness, which validates its effectiveness in fairly recommending with accuracy guarantee. (2) SoFA can effectively find solutions in different preference regions, showing the ability to accommodate customized trade-off between the direct utility and social utility as the IGF metrics. Figure 2a visualizes the trade-off between SP and NSP, showing that SoFA is able to find pareto optimal solutions between the IGF metrics. Figure 2b suggests that F1SP and NDCG do not always exhibit a conflicting trade-off, meaning that enhancing fairness does not necessarily sacrifice recommendation accuracy in certain scenarios.




Figure 3: Effects of varying accuracy thresholds in the optimization on KuaiRec with 100 repeats.

5.3 In-depth Analysis

Ablation Study (RQ2). We conduct the following ablations to evaluate the effectiveness of specific designs in SoFA: (1) *w/o con*, which trains the model without the recommendation accuracy constraint. (2) *w/o region*, which trains the model without specifying the preference region. (3) *w/o SP*, which trains the model without the SP fairness loss, thus does not require the preference region specification. (4) *w/o NSP*, similar to *w/o SP* but trains the model without the NSP fairness loss. Table 3 shows the performance of SoFA and its ablated versions, we find that training the model without the accuracy constraint encounters significant drop in recommendation accuracy compared to SoFA, indicating the effectiveness of the proposed accuracy constraint. In addition, removing the specification of the preference region may not hurt the performance in terms of accuracy and fairness, because the preference region do not directly contribute to the recommendation accuracy and overall fairness. However, it fails to obtain a recommendation model with flexibility to trade-off these IGF metrics.

Effects of Accuracy Requirement (RQ3). To further investigate the effect of recommendation accuracy constraint, we evaluate SoFA with relaxed constraints by increasing the threshold ξ in Eq. (9). In particular, we compare the default choice of the accuracy threshold with 1.1 times and 1.2 times of that value, as well as the optimization without the recommendation accuracy constraint. We conduct 100 runs for each optimization process on KuaiRec with varying preference regions. Figure 3 shows the accuracy and fairness performance, and we have observations as below: (1) NDCG has a clear decrease trend as the threshold increases, which further reveals that constraining the recommendation loss can control the accuracy of the obtained solution. (2) FISP improves as the threshold decreases in most regions, which indicates that the accuracy constraint may also benefit the fairness. Therefore, we suggest to use the original BPRMF loss as the accuracy threshold in practice.

6 Conclusion

This study addresses a new problem in IGF motivated by the effect of user social influence on item utility. First, we propose two IGF notions, namely NSP and NEO, as extensions to the existing IGF notions that only considers the direct utility of item exposures. Next, we formulate a multi-objective optimization problem with pre-defined preference regions to ensure the flexibility of the trade-off between these IGF metrics. Then, we incorporate a recommendation accuracy constraint to control the accuracy sacrifice for satisfying the fairness requirements. We further propose an algorithm to solve the optimization problem and theoretically demonstrate its Pareto optimality. Extensive experiments on two real-world datasets validate the effectiveness of our method. One interesting future research direction is to generalize the user social influence studies in this paper to ranking-based IGF settings.

Acknowledgement

This work is supported by the National Key Research and Development Program of China (No. 2022YFB3104701) and the National Natural Science Foundation of China (No. 12301370).

References

- [1] Simon Caton and Christian Haas. Fairness in Machine Learning: A Survey. *ACM Computing Surveys*, 2020.
- [2] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. Towards a Fair Marketplace: Counterfactual Evaluation of the Trade-Off between Relevance, Fairness & Satisfaction in Recommendation Systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 2243–2251, 2018.
- [3] Ziwei Zhu, Xia Hu, and James Caverlee. Fairness-Aware Tensor-Based Recommendation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1153–1162, 2018.
- [4] Tao Qi, Fangzhao Wu, Chuhan Wu, Peijie Sun, Le Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. ProFairRec: Provider Fairness-Aware News Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1164–1173, 2022.
- [5] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 214–226, 2012.
- [6] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P Gummadi. Fairness Constraints: A Flexible Approach for Fair Classification. *The Journal of Machine Learning Research*, 20(1):2737–2778, 2019.
- [7] Moritz Hardt, Eric Price, and Nati Srebro. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, 2016.
- [8] Xi Lin, Hui-Ling Zhen, Zhenhua Li, Qing-Fu Zhang, and Sam Kwong. Pareto Multi-Task Learning. In *Advances in Neural Information Processing Systems*, 2019.
- [9] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian Personalized Ranking from Implicit Feedback. *arXiv preprint arXiv:1205.2618*, 2012.
- [10] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web*, pages 173–182, 2017.
- [11] Yehuda Koren, Steffen Rendle, and Robert Bell. Advances in Collaborative Filtering. *Recommender Systems Handbook*, pages 91–142, 2021.
- [12] Le Wu, Xiangnan He, Xiang Wang, Kun Zhang, and Meng Wang. A Survey on Accuracy-Oriented Neural Recommendation: From Collaborative Filtering to Information-Rich Recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 35(5):4425–4445, 2022.
- [13] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix Factorization Techniques for Recommender Systems. *Computer*, 42(8):30–37, 2009.
- [14] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, Juntao Tan, Shuchang Liu, and Yongfeng Zhang. Fairness in Recommendation: Foundations, Methods, and Applications. *ACM Transactions on Intelligent Systems and Technology*, 14(5):1–48, 2023.
- [15] Sebastian Ruder. An Overview of Multi-Task Learning in Deep Neural Networks. *arXiv preprint arXiv:1706.05098*, 2017.
- [16] Yu Zhang and Qiang Yang. A Survey on Multi-Task Learning. *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609, 2021.
- [17] Kalyanmoy Deb and Kalyanmoy Deb. Multi-Objective Optimization. In *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques*, pages 403–449. Springer, 2013.

- [18] Xiao Lin, Hongjie Chen, Changhua Pei, Fei Sun, Xuanji Xiao, Hanxiao Sun, Yongfeng Zhang, Wenwu Ou, and Peng Jiang. A Pareto-Efficient Algorithm for Multiple Objective Optimization in E-commerce Recommendation. In *Proceedings of the 13th ACM Conference on Recommender Systems*, pages 20–28, 2019.
- [19] Kristof Van Moffaert and Ann Nowé. Multi-Objective Reinforcement Learning Using Sets of Pareto Dominating Policies. *The Journal of Machine Learning Research*, 15(1):3483–3512, 2014.
- [20] Eckart Zitzler. *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*, volume 63. Shaker Ithaca, 1999.
- [21] Abdullah Konak, David W Coit, and Alice E Smith. Multi-Objective Optimization Using Genetic Algorithms: A Tutorial. *Reliability Engineering & System Safety*, 91(9):992–1007, 2006.
- [22] Ioannis Giagkiozis, Robin C Purshouse, and Peter J Fleming. An Overview of Population-Based Algorithms for Multi-Objective Optimization. *International Journal of Systems Science*, 46(9):1572–1599, 2015.
- [23] Jörg Fliege and Benar Fux Svaiter. Steepest Descent Methods for Multicriteria Optimization. *Mathematical Methods of Operations Research*, 51:479–494, 2000.
- [24] Jean-Antoine Désidéri. Multiple-Gradient Descent Algorithm (MGDA) for Multiobjective Optimization. *Comptes Rendus Mathématique*, 350(5-6):313–318, 2012.
- [25] William Karush. Minima of Functions of Several Variables with Inequalities as Side Constraints. *M. Sc. Dissertation. Dept. of Mathematics, Univ. of Chicago*, 1939.
- [26] Vilfredo Pareto et al. *Manual of Political Economy*. 1971.
- [27] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph Neural Networks for Social Recommendation. In *The World Wide Web Conference*, pages 417–426, 2019.
- [28] Ozan Sener and Vladlen Koltun. Multi-Task Learning as Multi-Objective Optimization. In *Advances in Neural Information Processing Systems*, 2018.
- [29] Chongming Gao, Shijun Li, Wenqiang Lei, Jiawei Chen, Biao Li, Peng Jiang, Xiangnan He, Jiabin Mao, and Tat-Seng Chua. KuaiRec: A Fully-observed Dataset and Insights for Evaluating Recommender Systems. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 540–550, 2022.
- [30] Kalervo Järvelin and Jaana Kekäläinen. Cumulated Gain-Based Evaluation of IR Techniques. *ACM Transactions on Information Systems*, 20(4):422–446, 2002.
- [31] Dimitris Sacharidis, Carine Pierrette Mukamakuza, and Hannes Werthner. Fairness and Diversity in Social-Based Recommender Systems. In *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, pages 83–88, 2020.
- [32] Preetam Nandy, Cyrus Diccio, Divya Venugopalan, Heloise Logan, Kinjal Basu, and Nouredine El Karoui. Achieving Fairness via Post-Processing in Web-Scale Recommender Systems. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 715–725, 2022.