When Variety Seeking Meets Multi-Sided Recommendation Fairness: A Consistent and Personalized Multi-Objective Optimization Framework

Jiayi Guo*
School of Information Management and Engineering
MoE Key Laboratory of Interdisciplinary Research of Computation and Economics
Shanghai University of Finance and Economics
Shanghai, China
guo.jiayi@sufe.edu.cn

Jiangning He*†
School of Information Management and Engineering
MoE Key Laboratory of Interdisciplinary Research of Computation and Economics
Shanghai University of Finance and Economics
Shanghai, China
he.jiangning@mail.shufe.edu.cn

Chenyan Wang*
Polytechnic Institute
Zhejiang University
Hangzhou, Zhejiang, China
wangcy24@zju.edu.cn

Xinran Wu*
School of Business Administration
Shandong University of Finance and Economics
Jinan, Shandong, China
wuxinran611@gmail.com

Abstract

Recommendation research has evolved from solely improving accuracy to addressing ethical and fairness concerns. While prior works focus on optimizing fairness from either the user or product perspective, recent research emphasizes the importance of multi-sided fairness. This issue is inherently challenging due to the competing goals of different stakeholders. To tackle this challenge, we propose a Consistent and Personalized Fairness Recommendation framework with Multi-Objective Integer Programming (CPFR-MOIP). Our framework introduces two key innovations. First, we develop a novel similarity-based individual fairness metric for the user side and formulate a consistent product-side fairness metric, ensuring that the generated recommendation list aligns with the user preference distribution and the expected product exposure distribution. Second, we incorporate users' variety-seeking levels as a moderating factor to adjust fairness trade-offs and introduce personalized weights to balance user-side and product-side fairness. To effectively solve this optimization problem, we devise an alternating algorithm with theoretical guarantee and demonstrate the Pareto optimality of the obtained solutions. Extensive experiments on two real-world datasets demonstrate that our CPFR-MOIP achieves superior multi-sided fairness while maintaining competitive recommendation accuracy. Furthermore, ablation analysis highlights the advantages of incorporating user variety-seeking levels for personalizing fairness trade-offs. Our work paves the way for more ethical and personalized recommendation systems. The implementation code is available at: https://github.com/P0ise-Wang/CPFR-MOIP.

CCS Concepts

• **Information systems → Recommender systems**.

Keywords

Recommendation systems, Multi-sided fairness, Multi-objective optimization, Re-ranking algorithms

ACM Reference Format:
Jiayi Guo, Jiangning He, Chenyan Wang, and Xinran Wu. 2025. When Variety Seeking Meets Multi-Sided Recommendation Fairness: A Consistent and Personalized Multi-Objective Optimization Framework. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25), November 10–14, 2025, Seoul, Republic of Korea.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3746252.3761024

1 Introduction

Recommendation systems have become an integral part of digital platforms. The majority of research focuses on developing more advanced models to capture user preferences and enhance recommendation accuracy [23]. However, due to inherent data and algorithmic bias [5], optimizing solely for accuracy gives rise to fairness concerns that affect multiple stakeholders. On the product side, this leads to the Matthew effect, where popular items are overexposed and long-tail items are underexposed [17]. On the user side, algorithms tend to favor active users while overlooking the preferences of minority groups [18]. Such multi-sided unfairness not only undermines user trust but also reduces product diversity, significantly threatening the long-term sustainability of platforms [15, 27]. Therefore, it is critical to simultaneously consider and balance multi-sided fairness in recommendation systems.

*All authors have made equal contributions as co-first authors to this work.
†Corresponding author

CIKM '25, Seoul, Republic of Korea.

ACM ISBN 979-8-4007-2040-6/2025/11
https://doi.org/10.1145/3746252.3761024

A widely used approach to improving fairness is post-processing, which re-ranks the recommendation lists generated by initial recommendation models to meet fairness requirements [27]. This *re-ranking* approach offers the advantage of being flexible and model-agnostic [4], allowing fairness to be incorporated with any existing recommendation algorithm. In this line, we focus on solving multi-sided fairness recommendation problem through re-ranking.

Achieving multi-sided fairness can be formulated as a multi-objective optimization task, where the system should simultaneously consider fairness objectives on both the user and product side [18]. This necessitates resolving two fundamental issues:

(1) *How to measure multi-sided fairness?* The fairness criteria are generally categorized into individual fairness and group fairness [4]. On the product side, researchers primarily focus on *group fairness*, which ensures equitable outcomes across predefined product groups (e.g., popular v.s. long-tail) [8, 24]. On the user side, individual fairness focuses on optimizing personalized experiences at a more fine-grained level. *Individual fairness* requires that similar users receive similar treatment [7], ensuring that users with comparable preferences receive recommendations of comparable quality and relevance. Existing research typically implements this by performing pairwise comparisons between users' recommendation lists [25, 30], which is computationally expensive [27]. Hence, how to design an efficient metric for individual fairness with theoretical guarantee is underexplored.

(2) *How to effectively balance the trade-offs between multi-sided fairness?* Prior works in fairness-aware recommendations focus on either user- or product-side fairness [30, 32], while multi-sided fairness recommendations are underexplored. This issue is inherently challenging due to competing interests and goals of different stakeholders [6, 10]. For instance, user-side fairness requires recommendation lists to align closely with user preferences, while product-side fairness focuses on distributing exposure across different product categories. Existing approaches address this trade-off by heuristically considering different fairness rules in re-ranking algorithms [19, 31] or introducing global weights in a multi-objective optimization framework [18]. These strategies are usually system-specific and ignore individuals' diverse sensitivities or preferences regarding different fairness objectives. However, our empirical analysis based on a real-world dataset (Tmall), as shown in Figure 1, reveals a crucial insight: the multi-sided fairness trade-offs vary across users and are moderated by individuals' variety-seeking levels. For users with narrow interests (i.e., low variety-seeking levels), we observe a lower alignment between the user preference distribution and the expected product exposure distribution. This means that maintaining the user-side fairness inevitably hampers product-side fairness, resulting in significant trade-offs. In contrast, for users with high variety-seeking levels, the two fairness objectives tend to align more closely and can be satisfied simultaneously. Diverse recommendations meet users' variety-seeking preferences while also distributing exposure opportunities across different product categories. Therefore, how to adaptively balance multi-sided fairness based on users' variety-seeking levels is worthy of investigation.

To address the research gaps, our contributions are fourfold:

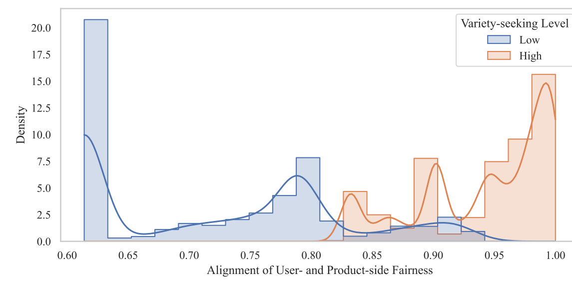

Notations.
User preference distribution: $\boldsymbol{p}$ (Eq.(1)), Expected product exposure distribution: $\boldsymbol{e}$ (Eq.(5)),
Fairness alignment: cosine$(\boldsymbol{p}, \boldsymbol{e})$, Variety-seeking levels: τ (Eq.(7)).

Figure 1: Alignment Between User- and Product-side Fairness Objectives for High and Low Variety-Seeking Users.

(1) We propose a Consistent and Personalized Fairness Recommendation framework with Multi-Objective Integer Programming (CPFR-MOIP), which introduces personalized weights depending on users' variety-seeking levels to adaptively balance the trade-offs between user- and product-side fairness;
(2) This framework features (i) a novel metric for user-side individual fairness that guarantees similar recommendations for similar users with theoretical guarantee, and (ii) a product-side group fairness metric in a consistent form, ensuring that recommendations align with the expected product exposure distribution;
(3) To effectively solve the problem, we propose an alternating algorithm based on KKT conditions and second-order cone programming. Theoretical guarantee of the algorithm and the Pareto optimality of its solutions are also provided;
(4) Extensive experiments on two public datasets demonstrate that our method achieves superior multi-sided fairness while maintaining competitive recommendation accuracy.

2 Related Work

2.1 Literature on Fairness Criteria

To quantify fairness in recommendation systems, a commonly adopted taxonomy involves two main categories: group fairness and individual fairness.

Group fairness requires that recommendation outcomes be equitable across different groups [27], where groups are typically divided by attributes such as popularity and category. Most existing works optimize product-side fairness at the group level [9, 16, 28, 31, 34]. For example, Gomez et al. [9] measure unfairness by evaluating exposure disparities among items from different regions; Wu et al. [31] propose that exposures of each group should be proportional to the quality of products in that group. In line with existing research [8, 28, 31], our study focuses on product-side group fairness by aligning the exposure distribution for recommended product groups with the expected product exposure distribution.

Individual fairness promotes fair treatment on a more fine-grained level (e.g., for each user or item) [27]. On the user side, Wu et al. [31] define fairness as equal recommendation accuracy across users, while Patro et al. [19] propose an envy-free definition that no user should prefer another user's recommendation list over their own. These definitions treat users as homogeneous entities.

Another widely adopted definition of individual fairness is that "similar users should be treated similarly [2, 7]." Based on this definition, Song et al. [25] propose to align graph-based embeddings based on user behavior similarities, while Wu et al. [30] gauge the dissimilarity between the recommendation lists of similar users. These approaches require pairwise comparisons between users' recommendation lists, resulting in high computational complexity. A research gap remains in devising an efficient approach to measuring this similarity-based individual fairness with theoretical guarantees.

2.2 Fairness-aware Recommendation Literature

Existing recommendation research typically enhances fairness in two stages: in-process and post-process [27]. In-process approaches integrate fairness metrics as regularization terms into classic recommendation loss (e.g., BPR [29], MSE [26]), thereby improving fairness during model training. Post-process methods, on the other hand, adjust original recommendation lists based on predefined fairness criteria and select the top-k items to form a re-ranking list. The latter, referred to as *re-ranking* approaches [16, 18, 24], offer a more direct way to optimize fairness [4] and can be flexibly integrated with any underlying recommendation algorithm.

Achieving multi-sided fairness through re-ranking approaches is generally formulated as a multi-objective optimization problem. Some studies solve this problem greedily based on fairness rules [19, 31]. For example, TFROM proposed by Wu et al. [31] sequentially selects the highest-ranked item from the user's original list based on the fairness criteria. An alternative stream formulates this problem as an integer programming problem [6, 18]. For instance, CPFair [18] generates re-ranking lists by minimizing the weighted sum of the fairness deviations from both the consumer and producer perspectives. However, one common limitation of these works is that they adopt global weights or static fairness rules when balancing multiple objectives across all users. As a result, they overlook the fact that users may have diverse fairness sensitivities or preferences on different fairness sides.

Several works on single-sided fairness optimization have introduced adaptive weights to balance recommendation accuracy and fairness. For example, to optimize product-side group fairness, Sonboli et al. [24] recognize that users with diverse preferences are more tolerant of achieving product-side fairness and propose OFAiR using user-specific tolerance weights to balance the trade-offs. Liu et al. [16] focus on a micro-lending scenario to promote loans from uncovered borrower groups and develop PFAR to optimize accuracy and borrower-side fairness using personalized weights. However, these methods concentrate on personalizing the trade-off between accuracy and single-sided fairness and are incapable of addressing the complexities of multi-sided fairness optimization, where different stakeholders have competing interests and objectives [6].

3 Problem Setup

3.1 Preliminary

Let I be the set of all products in the recommendation system. Each product $i \in \mathcal{I}$ is characterized by H attributes. Formally, for each attribute $h = 1, \dots, H$, let $\mathbf{A}^{(h)} \in \{0,1\}^{|\mathcal{I}| \times n^{(h)}}$ denote the attribute value matrix for all products, where $|\mathcal{I}|$ is the number of products and $n^{(h)}$ is the number of possible values for attribute h. The element $a_{ik}^{(h)} \in \mathbf{A}^{(h)}$ indicates whether product i possesses the k-th value on the attribute h. For example, in a movie recommendation system, the attribute "genre" with values "Sci-Fi", "Romance", and "Comedy" can be represented by the matrix $\mathbf{A}^{(\text{genre})}$ as:

$$\mathbf{A}^{(\text{genre})} = \left[\begin{array}{c|ccc} & \text{Sci-Fi} & \text{Romance} & \text{Comedy} \\ \hline \text{Titanic} & 0 & 1 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ \text{Love Actually} & 0 & 1 & 1 \end{array}\right]$$

For each user $u \in \mathcal{U}$, we use $s =< i_1, \dots, i_j, \dots, i_{|s|} >$ to denote their temporally ordered interaction sequence. Here, $i_j \in \mathcal{I}$, $j = 1, \dots, |s|$ and $|s|$ is the length of the sequence. Accordingly, we represent the *user preference distribution* for each attribute h using $\boldsymbol{p}^{(h)} = \left[p_1^{(h)}, \dots, p_k^{(h)}, \dots, p_{n^{(h)}}^{(h)}\right]^T$, where $p_k^{(h)}$ represents the user's preference level for the k-th value and can be estimated by the ratio of interactions with the k-th value in the sequence, i.e.,

$$p_k^{(h)} = \sum_{i \in s} a_{ik}^{(h)} / |s|. \tag{1}$$

Similarly, a recommendation list L can also be characterized at the attribute level. We denote the *recommendation list representation* for attribute h using $\boldsymbol{r}^{(h)} = \left[r_1^{(h)}, \dots, r_k^{(h)}, \dots, r_{n^{(h)}}^{(h)}\right]^T$, where $r_k^{(h)}$ denotes the proportion of the recommended products in list L possessing the k-th value and can be expressed as

$$r_k^{(h)} = \sum_{i \in L} a_{ik}^{(h)} / |L|. \tag{2}$$

3.2 User-side Individual Fairness

Individual fairness, as proposed by [7] and operationalized through Lipschitz property, requires that similar individuals should be treated similarly. In the context of recommender systems, this can be extended by requiring users with similar preferences to obtain similar recommendations [30]. To achieve this, we propose measuring the alignment between the user preference distribution and the recommendation list representation in a shared attribute value space, as formalized in Definition 1.

Definition 1 (Attribute-level User-side Individual Fairness). *Given a user's interaction sequence s and a recommendation list L, the attribute-level user-side individual fairness for attribute h is defined as the cosine similarity between the user preference distribution $\boldsymbol{p}^{(h)}$ and the recommendation list representation $\boldsymbol{r}^{(h)}$:*

$$UF^{(h)} = \frac{\boldsymbol{p}^{(h)} \cdot \boldsymbol{r}^{(h)}}{\|\boldsymbol{p}^{(h)}\| \cdot \|\boldsymbol{r}^{(h)}\|}, \ \forall\ h = 1, \dots, H. \tag{3}$$

This new metric adheres to the individual fairness criterion [7, 20] with the Lipschitz condition (Proposition 1), bridging our metric to foundational fairness frameworks.

Proposition 1. *Denote cosine similarity function to be* $\cos(\boldsymbol{x}_1, \boldsymbol{x}_2) = \frac{\boldsymbol{x}_1^T \boldsymbol{x}_2}{\|\boldsymbol{x}_1\| \|\boldsymbol{x}_2\|}$. *For any given $\boldsymbol{r} \in \mathbb{R}^n$, Lipschitz condition holds:*

$$|\cos(\boldsymbol{p}_1, \boldsymbol{r}) - \cos(\boldsymbol{p}_2, \boldsymbol{r})| \le L \cdot \|\boldsymbol{p}_1 - \boldsymbol{p}_2\|$$

for all $\boldsymbol{p}_1, \boldsymbol{p}_2 \in \mathbb{R}^n$ with $L = \sqrt{n}$.

Moreover, this new metric follows the general principle of fairness. Specifically, high $UF^{(h)}$ values indicate great attribute-level alignment between a user's preference and the recommendation list. Thus, similar users with high $UF^{(h)}$ values receive similar recommendations, according to Proposition 2.

Proposition 2. *For distance metric* $d(\boldsymbol{x}_1, \boldsymbol{x}_2) = \sqrt{1 - \cos(\boldsymbol{x}_1, \boldsymbol{x}_2)}$, *the following inequality holds*

$$d(\boldsymbol{r}_1, \boldsymbol{r}_2) \le d(\boldsymbol{r}_1, \boldsymbol{p}_1) + d(\boldsymbol{p}_1, \boldsymbol{p}_2) + d(\boldsymbol{p}_2, \boldsymbol{r}_2).$$

Proof. The proofs of these two propositions are provided in Appendix A.

3.3 Product-side Group Fairness

Product-side group fairness in recommendation systems ensures equitable exposure across product groups [31]. From a platform perspective, the goal is to ensure that the overall exposure distribution over recommended products aligns with the expected distribution.

We first define the *expected product exposure distribution* based on two notable fairness principles: demographic parity [8] and equal opportunity [11]. *Demographic parity* reflects the prevalence of each attribute value across all products [8]. Formally, for a given attribute h, the *expected product exposure distribution* under demographic parity is denoted as $\boldsymbol{e}^{(h)}_{DP} = \left[e^{(h)}_{DP,1}, \ldots, e^{(h)}_{DP,k}, \ldots, e^{(h)}_{DP,n^{(h)}}\right]^T$, where $e^{(h)}_{DP,k}$ is the proportion of products that take the k-th value on a platform, i.e.,

$$e^{(h)}_{DP,k} = \sum_{i \in \mathcal{I}} a^{(h)}_{ik} / |\mathcal{I}|. \tag{4}$$

Equal opportunity emphasizes that products deemed as "qualified" deserve greater recommendation visibility [11]. We define whether a product is qualified based on user interactions [8]. Then, the *expected product exposure distribution* under equal opportunity for attribute h is denoted as $\boldsymbol{e}^{(h)}_{EO} = \left[e^{(h)}_{EO,1}, \ldots, e^{(h)}_{EO,k}, \ldots, e^{(h)}_{EO,n^{(h)}}\right]^T$, where $e^{(h)}_{EO,k}$ represents the proportion of products possessing the k-th value across all interactions, and s_u is the interaction sequence of user $u \in \mathcal{U}$. Each $e^{(h)}_{EO,k}$ is calculated as

$$e^{(h)}_{EO,k} = \sum_{u \in \mathcal{U}} \sum_{i \in s_u} a^{(h)}_{ik} / \sum_{u \in \mathcal{U}} |s_u|. \tag{5}$$

Based on the expected product exposure distribution under DP or EO principles, we optimize the product-side fairness metric for each recommendation list, as formalized in Definition 2.

Definition 2 (Attribute-level product-side group fairness). *Given a recommendation list L, the attribute-level product-side group fairness for attribute h is defined as the cosine similarity between the expected product exposure distribution* $\boldsymbol{e}^{(h)}$ *(i.e.,* $\boldsymbol{e}^{(h)}_{DP}$ *or* $\boldsymbol{e}^{(h)}_{EO}$*) and the recommendation list representation* $\boldsymbol{r}^{(h)}$*:*

$$PF^{(h)} = \frac{\boldsymbol{e}^{(h)} \cdot \boldsymbol{r}^{(h)}}{\|\boldsymbol{e}^{(h)}\| \cdot \|\boldsymbol{r}^{(h)}\|}, \ \forall\, h = 1, \ldots, H. \tag{6}$$

3.4 Multi-sided Fairness Re-ranking Recommendation Problem

We are now ready to define the re-ranking problem:

Definition 3 (Multi-sided Fairness Re-ranking Recommendation Problem). *Consider a recommendation system with a set of products* $\mathcal{I}$*, a set of users* $\mathcal{U}$*, and the users' interaction sequences* $\{s_u\}_{u \in \mathcal{U}}$*. For a user with an interaction sequence s, the initial recommendation algorithm generates a recommendation list L of length K, accompanied by the relevance score vector* $\boldsymbol{\gamma} = [\gamma_1, \ldots, \gamma_j, \ldots, \gamma_K]^T$*, with* γ_j *indicating the relevance of the j-th product for the user. The goal is to re-rank the initial recommendation list to form a refined list* L' *of length* K' $(K' \le K)$*, which simultaneously maximizes both user-side fairness (*$UF^{(h)}$*, Definition 1) and product-side fairness (*$PF^{(h)}$*, Definition 2) across all attributes* $h = 1, \ldots, H$*.*

4 Methodology

In this section, we develop a multi-objective optimization framework by 1) proposing a consistent and personalized optimization objective to balance the trade-off between $UF^{(h)}$ and $PF^{(h)}$ based on variety-seeking levels; 2) rigorously formalizing the problem as a constrained multi-objective integer programming (MOIP) model with accuracy guarantee; 3) proposing an alternating algorithm to solve the relaxed MOIP problem with theoretical guarantee, ensuring computational efficiency and tractability. Collectively, this method enables personalized fairness optimization while maintaining computational tractability and recommendation quality.

4.1 Consistent and Personalized Optimization Objective

To quantify users' variety-seeking levels as a moderating factor to adjust the fairness trade-offs, we introduce an entropy-based coefficient computed at the attribute level. Formally, for the attribute h, the entropy-based coefficient $\tau^{(h)}$ can be defined as

$$\tau^{(h)} = -\frac{\sum_{k=1}^{n^{(h)}} I^{(h)}_k \log I^{(h)}_k}{\log n^{(h)}}, \tag{7}$$

where $I^{(h)}_k = p^{(h)}_k / \sum_{k'=1}^{n^{(h)}} p^{(h)}_{k'}$ represents the empirical preference distribution over attribute values, $p^{(h)}_k$ is calculated using Equation (1), and $\log n^{(h)}$ is the normalization factor. This entropy formulation captures the dispersion of user preferences across different attribute values, with higher $\tau^{(h)}$ values indicating stronger variety-seeking tendencies.

To balance the trade-offs between multi-sided fairness, we integrate the personalized weights $\tau^{(h)}$ with a hyperparameter $\mu \in [0, 1]$ that represents a global and system-controllable weight for user-side fairness. By balancing $UF^{(h)}$ and $PF^{(h)}$ with the integrated weight, we obtain the following optimization objective:

$$\max \sum_{h=1}^{H} \mu \cdot UF^{(h)} + (1 - \mu) \cdot \tau^{(h)} PF^{(h)}. \tag{8}$$

Compared to low variety-seeking users, achieving user-side fairness for high variety-seeking users aligns more closely with improving the product-side fairness objective, so a larger weight $\tau^{(h)}$ is set for

$PF^{(h)}$. Notably, the optimization objective (8) possesses two key advantages. First, it integrates fairness objectives from both the user and product sides, both of which are consistently calculated using cosine similarity, ensuring unified and comparable measure across the two perspectives. Second, by leveraging personalized weights based on user entropy, the objective can adaptively adjust the trade-off between $UF^{(h)}$ and $PF^{(h)}$, enabling a personalized fairness-aware optimization.

Moreover, this multi-objective optimization approach is theoretically rigorous, and its optimal solution can be proven to be Pareto optimal, as demonstrated by the theorem below:

Theorem 1. *For any fixed set of weights $\mu \in [0,1]$, the optimal solution $\mathbf{y}^*$ to the optimization problem (8) is Pareto optimal with respect to the user-side fairness $UF^{(h)}$ and the weighted product-side fairness $\tau^{(h)} PF^{(h)}$.*

The set of Pareto optimal solutions represents the best possible trade-offs that can be achieved in a multi-objective optimization, where no objective (e.g., user-side fairness or product-side fairness) can be improved without worsening others. The detailed proof of Theorem 1 is shown in the Appendix B.1.

4.2 Multi-objective Integer Programming

In this section, we first define the decision variables for the MOIP and establish the corresponding constraints. To identify the products selected from the initial recommendation list L, we introduce the binary decision vector $\boldsymbol{y} = [y_1, \ldots, y_j, \ldots, y_K]^T \in \{0,1\}^K$, where each element is an indicator variable:

$$y_j = \begin{cases} 1, & \text{if the } j^{\text{th}} \text{ product in list } L \text{ is selected,} \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

To ensure that the length of the re-ranked list L' is K', the decision vector must satisfy the following constraint:

$$\mathbf{1}^T \boldsymbol{y} = K', \tag{10}$$

where $\mathbf{1}$ denotes a vector of ones.

To maintain recommendation quality, we introduce an accuracy constraint that bounds the relevance score of the re-ranked list. Let $q \in [0,1]$ be a hyperparameter controlling the minimum acceptable relevance threshold. The constraint is formally defined as:

$$\boldsymbol{\gamma}^T \boldsymbol{y} \geq q \cdot \sum_{j=1}^{K'} \gamma_j, \tag{11}$$

where $\boldsymbol{\gamma} = [\gamma_1, \ldots, \gamma_j, \ldots, \gamma_K]^T \in \mathbb{R}^K$ denotes the vector of relevance scores of the initial list L, sorted in descending order. The constraint (11) ensures that the re-ranked list's total relevance score (left-hand side) meets or exceeds a fraction q of the top-K' items' relevance score in the initial list (right-hand side).

Let $\tilde{\boldsymbol{r}}^{(h)}$ denote the recommendation list representation of the re-ranked list L' for the attribute h. Formally, $\tilde{\boldsymbol{r}}^{(h)}$ is constructed as:

$$\tilde{\boldsymbol{r}}^{(h)} = \boldsymbol{F}^{(h)^T} \boldsymbol{y}, \tag{12}$$

where $\boldsymbol{F}^{(h)} \in \{0,1\}^{K \times n^{(h)}}$ is the attribute value matrix corresponding to the initial recommendation list L. This matrix is derived by selecting the rows indexed by L from the complete attribute value matrix $\boldsymbol{A}^{(h)}$, i.e., $\boldsymbol{F}^{(h)} = \boldsymbol{A}^{(h)}[L,:]$.

By integrating the objective function from Equation (8) with the constraints (9), (10), (11), and (12), we formulate the MOIP problem as follows:

$$\begin{aligned} \max_{\boldsymbol{y}, \tilde{\boldsymbol{r}}^{(h)}} \quad & \sum_{h=1}^{H} \mu \cdot \frac{\boldsymbol{p}^{(h)} \cdot \tilde{\boldsymbol{r}}^{(h)}}{\|\boldsymbol{p}^{(h)}\| \|\tilde{\boldsymbol{r}}^{(h)}\|} + (1-\mu) \cdot \tau^{(h)} \frac{\boldsymbol{e}^{(h)} \cdot \tilde{\boldsymbol{r}}^{(h)}}{\|\boldsymbol{e}^{(h)}\| \|\tilde{\boldsymbol{r}}^{(h)}\|}, \\ \text{s.t.} \quad & \tilde{\boldsymbol{r}}^{(h)} = \boldsymbol{F}^{(h)^T} \boldsymbol{y}, \ \forall\, h = 1, \ldots, H, \\ & q \cdot \sum_{j=1}^{K'} \gamma_j - \boldsymbol{\gamma}^T \boldsymbol{y} \leq 0, \\ & \mathbf{1}^T \boldsymbol{y} = K', \\ & y_j \in \{0,1\}, \ \forall\, j = 1, \ldots, K. \end{aligned} \tag{13}$$

4.3 Solving Algorithm

To address the MOIP problem (13), we first relax the binary constraint $y_j \in \{0,1\}$ to a continuous interval $0 \leq y_j \leq 1$, $j = 1, \ldots, K$. Subsequently, we solve the relaxed problem and obtain the optimal solution $\boldsymbol{y}^*$. The final recommendations are derived by selecting top-K' elements in $\boldsymbol{y}^*$ with values closest to 1.

First, we denote the parameter $\boldsymbol{z}^{(h)} = \mu \frac{\boldsymbol{p}^{(h)}}{\|\boldsymbol{p}^{(h)}\|} + (1-\mu)\tau^{(h)} \frac{\boldsymbol{e}^{(h)}}{\|\boldsymbol{e}^{(h)}\|}$, $h = 1, \ldots, H$, as a weighted combination of the user preference distribution and expected product exposure distribution. Thus, the objective of problem (13) can be simplified as

$$\max_{\boldsymbol{y}} \sum_{h=1}^{H} \frac{\boldsymbol{z}^{(h)} \cdot \tilde{\boldsymbol{r}}^{(h)}}{\|\tilde{\boldsymbol{r}}^{(h)}\|} = \max_{\boldsymbol{y}} \sum_{h=1}^{H} \frac{\boldsymbol{z}^{(h)} \cdot \boldsymbol{F}^{(h)^T} \boldsymbol{y}}{\|\boldsymbol{F}^{(h)^T} \boldsymbol{y}\|}$$

by substituting $\tilde{\boldsymbol{r}}^{(h)} = \boldsymbol{F}^{(h)^T} \boldsymbol{y}$. In addition, to address the non-convex term introduced by the denominator (i.e., $\|\boldsymbol{F}^{(h)^T} \boldsymbol{y}\|$), we define $\beta^{(h)} = \frac{\boldsymbol{z}^{(h)} \cdot \boldsymbol{F}^{(h)^T} \boldsymbol{y}}{\|\boldsymbol{F}^{(h)^T} \boldsymbol{y}\|}$ with $\beta^{(h)} \geq 0$. To sum up, the relaxation of problem (13) is expressed as follows:

$$\begin{aligned} \max_{\boldsymbol{y}, \boldsymbol{\beta}} \quad & \sum_{h=1}^{H} \beta^{(h)} \\ \text{s.t.} \quad & \beta^{(h)} \|\boldsymbol{F}^{(h)^T} \boldsymbol{y}\| = \boldsymbol{z}^{(h)} \cdot \boldsymbol{F}^{(h)^T} \boldsymbol{y}, \ \forall\, h = 1, \ldots, H, \\ & q \cdot \sum_{j=1}^{K'} \gamma_j - \boldsymbol{\gamma}^T \boldsymbol{y} \leq 0, \\ & \mathbf{1}^T \boldsymbol{y} = K', \\ & 0 \leq y_j \leq 1, \ \forall\, j = 1, \ldots, K. \end{aligned} \tag{14}$$

Second, to effectively solve the relaxed problem (14), we introduce Lagrange multipliers $\boldsymbol{\xi}$ corresponding to the new constraint, leading to the Lagrangian formulation:

$$\begin{aligned} \max_{\boldsymbol{y}, \boldsymbol{\beta}, \boldsymbol{\xi}} \quad & \Psi(\boldsymbol{y}, \boldsymbol{\beta}, \boldsymbol{\xi}) = \sum_{h=1}^{H} \beta^{(h)} + \sum_{h=1}^{H} \xi^{(h)} \left(\boldsymbol{z}^{(h)} \boldsymbol{F}^{(h)^T} \boldsymbol{y} - \beta^{(h)} \|\boldsymbol{F}^{(h)^T} \boldsymbol{y}\| \right) \\ \text{s.t.} \quad & q \cdot \sum_{j=1}^{K'} \gamma_j - \boldsymbol{\gamma}^T \boldsymbol{y} \leq 0, \\ & \mathbf{1}^T \boldsymbol{y} = K', \\ & 0 \leq y_j \leq 1, \ \forall\, j = 1, \ldots, K. \end{aligned} \tag{15}$$

Inspired by the alternating optimization paradigm [33], we decompose problem (15) into two subproblems: (1) $\max_{\boldsymbol{\beta},\boldsymbol{\xi}} \Psi(\hat{\boldsymbol{y}}, \boldsymbol{\beta}, \boldsymbol{\xi})$ with fixed $\hat{\boldsymbol{y}}$; and (2) $\max_{\boldsymbol{y}} \Psi(\boldsymbol{y}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$ with fixed $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$.

The subproblem (1) $\max_{\boldsymbol{\beta},\boldsymbol{\xi}} \Psi(\hat{\boldsymbol{y}}, \boldsymbol{\beta}, \boldsymbol{\xi})$ becomes an unconstrained optimization problem. Setting the partial derivatives with respect to $\beta^{(h)}$ and $\xi^{(h)}$ for each h to be zero yields a stationary point solution $(\boldsymbol{\beta}_{sp}, \boldsymbol{\xi}_{sp})$ with

$$\beta_{sp}^{(h)} = \boldsymbol{z}^{(h)} \cdot \frac{\boldsymbol{F}^{(h)^T}\hat{\boldsymbol{y}}}{\|\boldsymbol{F}^{(h)^T}\hat{\boldsymbol{y}}\|},\ \xi_{sp}^{(h)} = \frac{1}{\|\boldsymbol{F}^{(h)^T}\hat{\boldsymbol{y}}\|},\ \forall\ h = 1, \ldots, H. \quad (16)$$

The subproblem (2) $\max_{\boldsymbol{y}} \Psi(\boldsymbol{y}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$ constitutes a Second-Order Cone Program (SOCP), which can be solved efficiently by standard commercial mathematical programming solvers (see Appendix B.2). This alternating scheme is formalized in Algorithm 1.

Algorithm 1: Alternating Algorithm for Solving CPFR-MOIP

Input: $\boldsymbol{A}^{(h)}, \boldsymbol{e}^{(h)}, \boldsymbol{p}^{(h)}, \tau^{(h)}, \boldsymbol{\gamma}, \mu, q$
Construct $\boldsymbol{F}^{(h)}$, $\boldsymbol{z}^{(h)}$ and initialize $\hat{\boldsymbol{y}}$;
repeat
  $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}}) \leftarrow (\boldsymbol{\beta}_{sp}, \boldsymbol{\xi}_{sp})$ in accordance with (16) given $\hat{\boldsymbol{y}}$;
  $\hat{\boldsymbol{y}} \leftarrow \arg\max_{\boldsymbol{y}} \Psi(\boldsymbol{y}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$ by SOCP solver;
until *termination condition satisfied*;
$\boldsymbol{y}^* \leftarrow$ Round the top-K' elements in $\hat{\boldsymbol{y}}$ to 1, others to 0.
Output: The re-ranked list L' corresponding to $\boldsymbol{y}^*$

This alternating optimization algorithm is supported by theoretical guarantee, as formalized in Theorem 2. Specifically, when the iterative loop terminates under the criterion that the difference between two consecutive iterates $\hat{\boldsymbol{y}}$ is sufficiently small, the algorithm converges to a stationary point of the problem (14). For non-convex problems such as (14), this result establishes practical guarantee under standard regularity conditions.

Theorem 2. *A solution triplet* $(\hat{\boldsymbol{y}}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$ *constitutes a stationary point of problem* (14) *if and only if the following conditions hold:*

$$\hat{\boldsymbol{y}} = \arg\max_{\boldsymbol{y}} \Psi(\boldsymbol{y}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}}), \quad (17a)$$

$$\hat{\beta}^{(h)} = \boldsymbol{z}^{(h)} \cdot \frac{\boldsymbol{F}^{(h)^T}\hat{\boldsymbol{y}}}{\|\boldsymbol{F}^{(h)^T}\hat{\boldsymbol{y}}\|},\ \ \hat{\xi}^{(h)} = \frac{1}{\|\boldsymbol{F}^{(h)^T}\hat{\boldsymbol{y}}\|},\ \forall\ h = 1, \ldots, H. \quad (17b)$$

The proof of Theorem 2 is grounded in the Karush-Kuhn-Tucker (KKT) conditions [3] for the associated optimization problems, with comprehensive details presented in Appendix B.3.

5 Empirical Evaluations

In this section, we present a comprehensive evaluation of the proposed fairness framework and aim to demonstrate its effectiveness and superiority by answering the following questions:

- **RQ1**: How does the proposed framework perform in comparison with the existing approaches in terms of multi-sided fairness and recommendation accuracy?
- **RQ2**: To what extent do the personalized weights in our framework contribute to a more favorable trade-off between user- and product-side fairness?
- **RQ3**: How does the choice of initial recommendation models affect the performance of our method?

5.1 Dataset and Experiment Setups

We conducted experiments on two public datasets, Tmall[1] and Dressipi[2]:

- **Tmall** dataset comes from IJCAI 2015 competition, containing anonymized users' shopping logs. Following previous work [14], we filtered out users with fewer than 3 interactions and items appearing less than 5 times. After preprocessing, the dataset contains 252,558 users and 36,886 items. For attribute selection, we define a commonly used "*popularity*" attribute by labeling the top 20% most interacted items in the training set as "popular" and the rest as "unpopular."
- **Dressipi** is a fashion recommendation dataset released in Recsys Challenge 2022. Each anonymized user session comprises a sequence of clicks and ends with a purchase. We applied the same k-filtering strategy as used for Tmall, resulting in 53,458 users and 4,477 items. In addition to the " *popularity*" attribute, we further incorporated an anonymous attribute that covers the most products and has 47 distinct attribute values. This setup allows us to evaluate our method's capability in handling diverse attribute values.

To evaluate recommendation performances, we introduce Recall@10 [21] and NDCG@10 [12] to assess recommendation accuracy. For fairness evaluation, we define User Fairness Matching Score (UFMS) and Product Fairness Matching Score (PFMS) aligned with Definitions 1 and 2, respectively. Formally, UFMS and PFMS are computed as follows:

$$\text{UFMS} = \frac{1}{H}\sum_{h=1}^{H} \frac{\boldsymbol{p}^{(h)} \cdot \boldsymbol{F}^{(h)^T}\boldsymbol{y}}{\|\boldsymbol{p}^{(h)}\|\|\boldsymbol{F}^{(h)^T}\boldsymbol{y}\|}, \quad (18)$$

$$\text{PFMS} = \frac{1}{H}\sum_{h=1}^{H} \frac{\boldsymbol{e}^{(h)} \cdot \boldsymbol{F}^{(h)^T}\boldsymbol{y}}{\|\boldsymbol{e}^{(h)}\|\|\boldsymbol{F}^{(h)^T}\boldsymbol{y}\|}. \quad (19)$$

We calculate PFMS using the expected product exposure distribution under demographic parity ($\boldsymbol{e}_{DP}^{(h)}$) and equal opportunity ($\boldsymbol{e}_{EO}^{(h)}$). Both UFMS and PFMS for a recommendation list fall in the range [0, 1], with higher values indicating a greater level of fairness.

The evaluation procedure is as follows: We first trained three initial recommendation models—FPMC (a matrix factorization-based method [22]), SASRec (a self-attention-based sequential model [13]), and MGS (a graph neural network-based model [14])—on the training set and tuned their hyperparameters using the validation set. Next, we retrieved the top-500 products sorted by the relevance score as the initial recommendation list for re-ranking. Subsequently, we executed re-ranking approaches to generate the final list consisting of top-10 products and evaluated their performances in terms of recommendation accuracy and multi-sided fairness.

5.2 Recommendation Performance (RQ1)

To answer RQ1, we compared our method with several benchmark models from the literature on fairness recommendation. For single-sided fairness optimization, we applied two re-ranking approaches:

[1] https://tianchi.aliyun.com/dataset/42
[2] http://www.recsyschallenge.com/2022/

PFAR [16] and OFAiR [24], which employ personalized weights to balance recommendation accuracy and product-side fairness. For multi-sided fairness optimization, we selected two re-ranking models: CPFair [18] and TFROM [31]. CPFair [18] proposes a multi-objective optimization problem and balances user- and product-side fairness using a global weight, while TFROM [31] designs a re-ranking algorithm that considers multi-sided fairness criteria, such as by constraining the group exposure of selected items to remain within a specified fairness threshold. In this experiment, we utilized MGS to generate the initial recommendation lists for re-ranking, for its superior performances in recommendation accuracy.

According to Table 1, our method (CPFR-MOIP) consistently ranks first or second in NDCG@10, UFMS, and PFMS across all experimental settings. Firstly, when compared to OFAiR and PFAR, two re-ranking methods that balance product-side group fairness with personalized weights, our CPFR-MOIP achieves higher UFMS while maintaining competitive PFMS. For example, under the DP principle on the Tmall dataset, CPFR-MOIP improves UFMS by 7.13%, compared to 0.11% reduction by OFAiR and 0.33% improvement by PFAR, and it also improves PFMS by 17.91%, demonstrating a higher enhancement over both benchmarks. This can be attributed to the design of unified similarity-based fairness metrics across both user- and product-sides, leading to a more balanced multi-sided fairness optimization. Secondly, in comparison to multi-sided fairness approaches, including CPFair and TFROM, our method improves UFMS and PFMS simultaneously with minimal loss in recommendation accuracy, demonstrating the advantage of our personalized weighting strategy over global weights for achieving multi-sided fairness optimization[3].

5.3 Ablation Study (RQ2)

To answer RQ2, we evaluated the effectiveness of the personalized weights through the following ablated models:

(1) CPFR-MOIP (mean tau), which replaces $\tau_u^{(h)}$ with the average value across all users, denoted as $\overline{\tau}^{(h)} = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \tau_u^{(h)}$.
(2) CPFR-MOIP (random tau), which replaces $\tau_u^{(h)}$ with a random $\tau_{u'}^{(h)}$ sampled from the original distribution $\boldsymbol{\tau}^h$ across users.
(3) CPFR-MOIP (w/o tau), which removes the personalized variety-seeking weight $\tau_u^{(h)}$ across users.

For our method and each ablated model, we computed their fairness metrics at different μ with step size 0.1 and plotted UFMS (vertical axis) against the weighted PFMS[4](horizontal axis). As shown in Figure 2, each method's curve represents a Pareto front, where each point reflects a different balance between user- and product-side fairness. To quantify the quality of these Pareto fronts, we introduce the *hypervolume* metric—the area enclosed between the Pareto curve and the origin. Notably, our method consistently achieves the largest hypervolume across all experimental settings, indicating superior trade-offs between the two fairness metrics [1]. The result demonstrates that our method dominates others in both user- and product-side fairness, regardless of the system designer's preferred multi-sided fairness trade-offs (i.e., the choice of μ).

5.4 Robustness Analysis (RQ3)

To investigate the impact of initial recommendation models, we applied our re-ranking approach to the recommendation lists generated by FPMC, SASRec, and MGS. As shown in Table 2, CPFR-MOIP substantially improves both user- and product-side fairness metrics (e.g., 24.53% lift in UFMS and 53.54% lift in PFMS under the DP principle on the Tmall dataset), with a maximum loss of 1.69% in NDCG@10. The results support the robustness of our method across different recommendation settings.

Additionally, we note that fairness and recommendation accuracy do not always conflict. On the Tmall dataset, all the metrics are improved when applying CPFR-MOIP on FPMC. This implies that our re-ranking method can enhance accuracy in certain circumstances, particularly when the initial recommendation model is inaccurate, primarily due to our user-side individual fairness metric that considers the alignment between user preference and recommendation list representations.

6 Conclusion

Fairness issues in recommendation systems are gaining increasing attention. To balance multi-sided fairness through re-ranking approaches, we propose a multi-objective integer program with consistent fairness metrics and personalized weights derived from users' variety-seeking levels. To effectively solve this problem, we develop an alternating algorithm with theoretical guarantee and demonstrate the Pareto optimality of its obtained solutions. Extensive experiments across two real-world datasets demonstrate that our approach can effectively balance the multi-sided fairness with minimal loss of recommendation accuracy, and it even yields simultaneous improvements in both accuracy and fairness when applied to initially inaccurate recommendation lists.

In future work, in addition to the multi-objective integer programming algorithm, we plan to develop alternative optimization algorithms, particularly reinforcement learning approaches, to explore the model-agnostic property of our framework. Furthermore, it would be worthwhile to apply our framework to address bilateral recommendation problems, such as talent recommendations, and investigate its efficacy in maintaining recommendation quality and achieving multi-sided fairness for both talents and enterprise. Finally, incorporating additional stakeholders—including platforms and governments—presents a promising direction for achieving comprehensive fairness in recommendation ecosystems.

A Appendix for Section 3

A.1 Proof of Proposition 1

Given $\boldsymbol{r} \in \mathbb{R}^n$ with normalized vector $\hat{\boldsymbol{r}} = \frac{\boldsymbol{r}}{\|\boldsymbol{r}\|}$, define function $f(\boldsymbol{p}) : \mathbb{R}^n \backslash \{0\} \to \mathbb{R}$ as $f(\boldsymbol{p}) = \cos(\boldsymbol{p}, \boldsymbol{r}) = \frac{\boldsymbol{r}^T \boldsymbol{p}}{\|\boldsymbol{r}\| \cdot \|\boldsymbol{p}\|} = \frac{\hat{\boldsymbol{r}}^T \boldsymbol{p}}{\|\boldsymbol{p}\|}$. This proposition is equivalently to show $f(\boldsymbol{p})$ is Lipschitz continuous.

Gradient Computation. The gradient of f is derived as:

$$\nabla f(\boldsymbol{p}) = \frac{\hat{\boldsymbol{r}}}{\|\boldsymbol{p}\|} - \frac{\hat{\boldsymbol{r}}^T \boldsymbol{p}}{\|\boldsymbol{p}\|^3} \boldsymbol{p} = \frac{\hat{\boldsymbol{r}}_\perp}{\|\boldsymbol{p}\|},$$

[3]Our method can be executed in parallel for each user. Empirical results show that, on average, it takes 0.21s per user on the Dressipi dataset with multiple attributes and 0.05s on the Tmall dataset with a single attribute.

[4]Aligned with Theorem 1, the weighted PFMS incorporates the personalized weights $\tau^{(h)}$ and is caluated by $\frac{1}{H} \sum_h \tau^{(h)} \frac{e^{(h)} \cdot F^{(h)T} y}{\|e^{(h)}\| \|F^{(h)T} y\|}$ under both DP and EO principles.

Table 1: Re-ranking Performance Comparison: CPFR-MOIP vs. Benchmark Models. Values in parentheses show improvements over initial recommendations; bold indicates best performance, and underline denotes second-best.

Fairness Principle	Benchmark Model	Tmall				Dressipi			
		R@10	NDCG@10	UFMS	PFMS	R@10	NDCG@10	UFMS	PFMS
DP	w/o reranking	**0.407**	**0.274**	0.897	0.748	**0.310**	**0.198**	0.757	0.558
	+CPFR-MOIP (Our method)	0.402 (-1.23%)	0.272 (-0.73%)	**0.961** (7.13%)	0.882 (17.91%)	0.300 (-3.23%)	0.195 (-1.52%)	**0.838** (10.7%)	0.678 (21.51%)
	+OFAiR	0.399 (-1.97%)	0.270 (-1.46%)	0.896 (-0.11%)	0.847 (13.24%)	0.298 (-3.87%)	0.193 (-2.53%)	0.731 (-3.43%)	0.532 (-4.66%)
	+PFAR	**0.407** (0.00%)	0.272 (-0.73%)	0.900 (0.33%)	0.754 (0.80%)	0.302 (-2.58%)	0.191 (-3.54%)	0.750 (-0.92%)	0.551 (-1.25%)
	+CPFair	0.397 (-2.46%)	0.271 (-1.09%)	0.927 (3.34%)	0.862 (15.24%)	0.296 (-4.52%)	0.193 (-2.53%)	0.769 (1.59%)	0.615 (10.22%)
	+TFROM	0.354 (-13.02%)	0.254 (-7.3%)	0.803 (-10.48%)	**0.988** (32.09%)	0.163 (-47.42%)	0.130 (-34.34%)	0.537 (-29.06%)	**0.845** (51.43%)
EO	w/o reranking	**0.407**	**0.274**	0.897	0.889	**0.310**	**0.198**	0.757	0.810
	+CPFR-MOIP (Our method)	0.405 (-0.49%)	0.273 (-0.36%)	**0.966** (7.69%)	0.910 (2.36%)	0.303 (-2.26%)	0.196 (-1.01%)	**0.821** (8.45%)	**0.892** (10.12%)
	+OFAiR	0.399 (-1.97%)	0.270 (-1.46%)	0.896 (-0.11%)	**0.983 (10.57%)**	0.298 (-3.87%)	0.193 (-2.53%)	0.731 (-3.43%)	0.808 (-0.25%)
	+PFAR	**0.407** (0.00%)	0.272 (-0.73%)	0.900 (0.33%)	0.892 (0.34%)	0.307 (-0.97%)	0.193 (-2.53%)	0.758 (0.13%)	0.808 (-0.25%)
	+CPFair	0.397 (-2.46%)	0.271 (-1.09%)	0.927 (3.34%)	0.852 (-4.16%)	0.296 (-4.52%)	0.193 (-2.53%)	0.769 (1.59%)	0.796 (-1.73%)
	+TFROM	0.405 (-0.49%)	0.273 (-0.36%)	0.887 (-1.11%)	0.902 (1.46%)	0.300 (-3.23%)	0.195 (-1.52%)	0.769 (1.59%)	0.817 (0.86%)

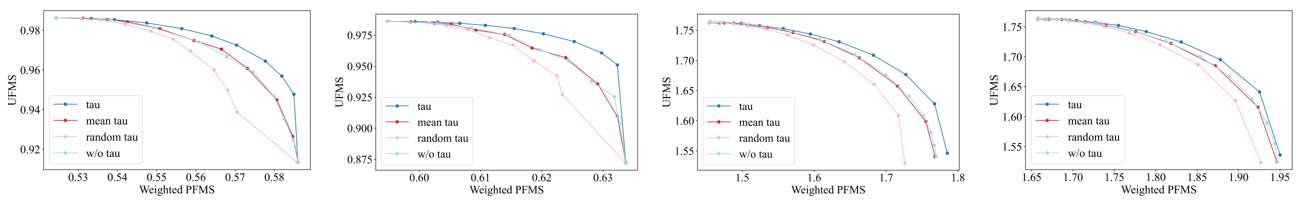

(a) Tmall (DP principle) (b) Tmall (EO principle) (c) Dressipi (DP principle) (d) Dressipi (EO principle)

Figure 2: Pareto Front Comparison: CPFR-MOIP vs. Ablated Models.

Table 2: CPFR-MOIP's Performances Given Different Initial Recommendation Models: FPMC, SASRec, and MGS.

Fairness Principle	Benchmark Model	Tmall				Dressipi			
		R@10	NDCG@10	UFMS	PFMS	R@10	NDCG@10	UFMS	PFMS
DP	FPMC	0.176	0.117	0.587	0.353	0.312	0.213	0.722	0.617
	+CPFR-MOIP	0.178 (1.14%)	0.118 (0.85%)	0.731 (24.53%)	0.542 (53.54%)	0.304 (-2.56%)	0.210 (-1.41%)	0.809 (12.05%)	0.731 (18.48%)
	SASRec	0.301	0.201	0.872	0.736	0.346	0.236	0.731	0.598
	+CPFR-MOIP	0.297 (-1.33%)	0.200 (-0.5%)	0.906 (3.9%)	0.931 (26.49%)	0.333 (-3.76%)	0.232 (-1.69%)	0.818 (11.9%)	0.716 (19.73%)
	MGS	0.407	0.274	0.897	0.748	0.310	0.198	0.757	0.558
	+CPFR-MOIP	0.402 (-1.23%)	0.272 (-0.73%)	0.961 (7.13%)	0.882 (17.91%)	0.300 (-3.23%)	0.195 (-1.52%)	0.838 (10.7%)	0.678 (21.51%)
EO	FPMC	0.176	0.117	0.587	0.840	0.312	0.213	0.722	0.811
	+CPFR-MOIP	0.178 (1.14%)	0.118 (0.85%)	0.731 (24.53%)	0.907 (7.98%)	0.308 (-1.28%)	0.211 (-0.94%)	0.840 (16.34%)	0.871 (7.4%)
	SASRec	0.301	0.201	0.872	0.920	0.346	0.236	0.731	0.814
	+CPFR-MOIP	0.301 (0%)	0.201 (0%)	0.964 (10.55%)	0.924 (0.43%)	0.337 (-2.6%)	0.233 (-1.27%)	0.846 (15.73%)	0.869 (6.76%)
	MGS	0.407	0.274	0.897	0.889	0.310	0.198	0.757	0.810
	+CPFR-MOIP	0.405 (-0.49%)	0.273 (-0.36%)	0.966 (7.69%)	0.910 (2.36%)	0.303 (-2.26%)	0.196 (-1.01%)	0.821 (8.45%)	0.892 (10.12%)

where $\hat{\boldsymbol{r}}_\perp = \hat{\boldsymbol{r}} - \frac{\hat{\boldsymbol{r}}^T \boldsymbol{p}}{\|\boldsymbol{p}\|^2}\boldsymbol{p}$ is the perpendicular component of $\hat{\boldsymbol{r}}$ in the direction of $\boldsymbol{p}$. Geometrically, let θ be the angle between $\hat{\boldsymbol{r}}$ and $\boldsymbol{p}$, and we have $\|\hat{\boldsymbol{r}}_\perp\| = \sqrt{1 - \left(\frac{\hat{\boldsymbol{r}}^T \boldsymbol{p}}{\|\boldsymbol{p}\|}\right)^2} = |\sin\theta|$.

Norm Bounding. In our context, $\mathbf{1}^T \boldsymbol{p}^{(h)} = \frac{1}{|s|}\sum_{i\in s}\sum_{k=1}^{n^{(h)}} a_{ik}^{(h)} \geq 1$, because every product must possess at least one nonzero value on the attribute h. Therefore, Cauchy-Schwarz inequality yields

$$1 \leq \mathbf{1}^T \boldsymbol{p} \leq \|\mathbf{1}\|\|\boldsymbol{p}\| = \sqrt{n}\|\boldsymbol{p}\| \implies \|\boldsymbol{p}\| \geq 1/\sqrt{n}.$$

Lipschitz Constant. Substituting $\|\boldsymbol{p}\| \geq 1/\sqrt{n}$ into the gradient norm yields

$$\|\nabla f(\boldsymbol{p})\| \leq \sqrt{n}|\sin\theta| \leq \sqrt{n}.$$

Thus, f is Lipschitz continuous with constant $\sqrt{n}$.

Remark: Under additional constraints $\theta \le \theta_{\max} < \frac{\pi}{2}$, the Lipschitz constant improves to $\sqrt{n}\sin\theta_{\max}$.

A.2 Proof of Proposition 2

As the distance metric $d(\boldsymbol{x}_1, \boldsymbol{x}_2) = \sqrt{1-\cos(\boldsymbol{x}_1,\boldsymbol{x}_2)}$ is defined based on cosine similarity, without loss of generality, we assume that the input vectors $\boldsymbol{p}_1, \boldsymbol{p}_2, \boldsymbol{r}_1, \boldsymbol{r}_2$ are normalized to 1. Therefore, supposing that for $\|\boldsymbol{x}_1\| = \|\boldsymbol{x}_2\| = 1$, we have the following property

$$\begin{aligned}\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|^2 &= \|\boldsymbol{x}_1\|^2 + \|\boldsymbol{x}_2\|^2 - 2\|\boldsymbol{x}_1\|\|\boldsymbol{x}_2\|\cos(\boldsymbol{x}_1,\boldsymbol{x}_2)\\ &= 2 - 2\cos(x_1, x_2) = 2d(\boldsymbol{x}_1,\boldsymbol{x}_2)^2.\end{aligned}$$

In other words, $d(\boldsymbol{x}_1,\boldsymbol{x}_2) = \frac{1}{\sqrt{2}}\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|$.

Now, we can show that

$$\begin{aligned} d(\boldsymbol{r}_1,\boldsymbol{r}_2) &= \frac{1}{\sqrt{2}}\|\boldsymbol{r}_1 - \boldsymbol{r}_2\|\\ &= \frac{1}{\sqrt{2}}\|(\boldsymbol{r}_1 - \boldsymbol{p}_1) + (\boldsymbol{p}_1 - \boldsymbol{p}_2) + (\boldsymbol{p}_2 - \boldsymbol{r}_2)\|\\ \text{(triangle inequality)} &\le \frac{1}{\sqrt{2}}\left(\|\boldsymbol{r}_1 - \boldsymbol{p}_1\| + \|\boldsymbol{p}_1 - \boldsymbol{p}_2\| + \|\boldsymbol{p}_2 - \boldsymbol{r}_2\|\right)\\ &= d(\boldsymbol{r}_1,\boldsymbol{p}_1) + d(\boldsymbol{p}_1,\boldsymbol{p}_2) + d(\boldsymbol{p}_2,\boldsymbol{r}_2).\end{aligned}$$

B Appendix for Section 4

B.1 Proof of Theorem 1

In this appendix, we prove that the solutions obtained through the personalized objective (8) in recommendation fairness optimization necessarily lie on the Pareto frontier.

We employ proof by contradiction. Suppose there exists a fixed set of weights μ such that the optimal solution $\mathbf{y}^*$ to the optimization problem is not Pareto optimal. According to the definition of Pareto optimality [35], this implies that there exists another feasible solution $\mathbf{y}'$ such that:

(1) $UF^{(h)}(\mathbf{y}') \ge UF^{(h)}(\mathbf{y}^*), \forall h = 1, 2, \ldots, H$;
(2) $\tau^{(h)}PF^{(h)}(\mathbf{y}') \ge \tau^{(h)}PF^{(h)}(\mathbf{y}^*), \forall h = 1, 2, \ldots, H$;
(3) There exists at least one index $h_0 \in 1, 2, \ldots, H$ such that $UF^{(h_0)}(\mathbf{y}') > UF^{(h_0)}(\mathbf{y}^*)$ or $\tau^{(h_0)}PF^{(h_0)}(\mathbf{y}') > \tau^{(h_0)}PF^{(h_0)}(\mathbf{y}^*)$.

Now we denote our objective function in Problem (8) as $J(\mathbf{y})$. For solutions $\mathbf{y}'$ and $\mathbf{y}^*$, we have:

$$\begin{aligned} J(\mathbf{y}') - J(\mathbf{y}^*) = &\sum_{h=1}^{H} \mu \cdot [UF^{(h)}(\mathbf{y}') - UF^{(h)}(\mathbf{y}^*)]\\ &+ (1-\mu)\cdot[\tau^{(h)}PF^{(h)}(\mathbf{y}') - \tau^{(h)}PF^{(h)}(\mathbf{y}^*)].\end{aligned}$$

Thus, we can safely conclude $J(\mathbf{y}') - J(\mathbf{y}^*) > 0$ by $\mu \in [0, 1]$ and properties (1), (2), (3).

In all, $J(\mathbf{y}') > J(\mathbf{y}^*)$ contradicts the assumption that $\mathbf{y}^*$ is the optimal solution. Hence, we can prove that $\mathbf{y}^*$ is the Pareto optimal.

B.2 The SOCP Reformulation

The equivalent SOCP formulation of $\max_{\boldsymbol{y}} \ \Psi(\boldsymbol{y}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$ is:

$$\begin{aligned} \max_{\boldsymbol{y}} \quad & \sum_{h=1}^{H} \hat{\beta}^{(h)} + \sum_{h=1}^{H} \hat{\xi}^{(h)} \boldsymbol{z}^{(h)} \boldsymbol{F}^{(h)^T} \boldsymbol{y} + \sum_{h=1}^{H} S_h\\ \text{s.t.} \quad & q \cdot \sum_{j=1}^{K'} \gamma_j - \boldsymbol{\gamma}^T \boldsymbol{y} \le 0,\\ & \mathbf{1}^T \boldsymbol{y} = K',\\ & \|\boldsymbol{F}^{(h)^T}\boldsymbol{y}\| \le -\frac{S_h}{\hat{\beta}^{(h)}\hat{\xi}^{(h)}}, \ \forall\, h = 1, \ldots, H,\\ & 0 \le y_j \le 1, \ \forall\, j = 1, \ldots, K.\end{aligned}$$

where S_h is an auxiliary variable introduced to linearize the norm term $\|\boldsymbol{F}^{(h)^T}\boldsymbol{y}\|$. In our experiment, we solved SOCP problems using the standard commercial solver COPT[5].

B.3 Proof of Theorem 2

For problem (14), let $\xi^{(h)}, h = 1, \ldots, H$, $\lambda \ge 0$, v, and $\alpha_j, \omega_j \ge 0$, $j = 1, \ldots, K$ be the dual variables associated with the respective constraints. Then the Lagrangian function can be formulated as:

$$\begin{aligned}\mathcal{L}(\boldsymbol{y},\boldsymbol{\beta},\xi,\lambda,v,\boldsymbol{\alpha},\boldsymbol{\omega}) =& \psi(\boldsymbol{y},\boldsymbol{\beta},\xi) + \lambda\left(q\sum_{j=1}^{K'}\gamma_j - \boldsymbol{\gamma}^T\boldsymbol{y}\right)\\ &+ v\left(\mathbf{1}^T\boldsymbol{y} - K'\right) + \sum_{j=1}^{K}\alpha_j\left(y_j - 1\right) - \sum_{j=1}^{K}\omega_j y_j,\end{aligned}$$

where

$$\psi(\boldsymbol{y},\boldsymbol{\beta},\xi) = \sum_{h=1}^{H}\beta^{(h)} + \sum_{h=1}^{H}\xi^{(h)}\left(\boldsymbol{z}^{(h)}\boldsymbol{F}^{(h)^T}\boldsymbol{y} - \beta^{(h)}\|\boldsymbol{F}^{(h)^T}\boldsymbol{y}\|\right).$$

The KKT conditions of problem (14) are derived as follows:

$$\nabla_{\boldsymbol{y}}\psi(\boldsymbol{y},\boldsymbol{\beta},\xi) - \lambda\boldsymbol{\gamma} + v\mathbf{1} + (\boldsymbol{\alpha} - \boldsymbol{\omega}) = 0. \tag{20a}$$

$$1 - \xi^{(h)}\|F^{(h)^T}\boldsymbol{y}\| = 0, \ \forall\, h = 1, \ldots, H. \tag{20b}$$

$$\boldsymbol{z}^{(h)}\boldsymbol{F}^{(h)^T}\boldsymbol{y} - \beta^{(h)}\|\boldsymbol{F}^{(h)^T}\boldsymbol{y}\| = 0, \ \forall\, h = 1, \ldots, H. \tag{20c}$$

$$\mathbf{1}^T\boldsymbol{y} - K' = 0. \tag{20d}$$

$$q\sum_{j=1}^{K'}\gamma_j - \boldsymbol{\gamma}^T\boldsymbol{y} \le 0. \tag{20e}$$

$$0 \le y_j \le 1, \ \forall\, j = 1, \cdots, K. \tag{20f}$$

$$\lambda \ge 0, \alpha_j, \omega_j \ge 0, \ \forall\, j = 1, \cdots, K. \tag{20g}$$

Similarly, we can derive the KKT conditions for the subproblem $\max_{\boldsymbol{y}} \Psi(\boldsymbol{y}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$, corresponding precisely to (20a) and (20d)-(20g). In addition, the equations in (17b) are exactly equivalent to conditions (20c) and (20b), respectively. Therefore, when the solution triple $(\hat{\boldsymbol{y}}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$ satisfies condition (17a) and (17b), the full set of KKT conditions for problem (14) is also fulfilled, indicating that $(\hat{\boldsymbol{y}}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\xi}})$ is a stationary point solution [3]. Hence, Theorem 2 holds.

Acknowledgments

This research is supported by the National Natural Science Foundation of China (NSFC) under grants No. 72001128, 72101139, and 72394360.

[5]https://guide.coap.online/copt/en-doc/index.html

GenAI Usage Disclosure

During the preparation of this work, GenAI tools were used solely to improve the spelling and grammar of the author-written text. The authors are fully accountable for the content.

References

[1] Anne Auger, Johannes Bader, Dimo Brockhoff, and Eckart Zitzler. 2009. Theory of the hypervolume indicator: optimal μ-distributions and the choice of the reference point. In *Proceedings of the Tenth ACM SIGEVO Workshop on Foundations of Genetic Algorithms* (Orlando, Florida, USA) *(FOGA '09)*. Association for Computing Machinery, New York, NY, USA, 87–102.

[2] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. ACM, New York, NY, USA, 405–414.

[3] Stephen P. Boyd and Lieven Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.

[4] Simon Caton and Christian Haas. 2024. Fairness in Machine Learning: A Survey. *Comput. Surveys* 56, 7 (April 2024), 166:1–166:38.

[5] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and Debias in Recommender System: A Survey and Future Directions. *ACM Transactions on Information Systems* 41, 3 (Feb. 2023), 67:1–67:39.

[6] Qinyi Chen, Jason Cheuk Nam Liang, Negin Golrezaei, and Djallel Bouneffouf. 2024. Interpolating Item and User Fairness in Multi-Sided Recommendations. In *Advances in Neural Information Processing Systems (NIPS'24)*. Curran Associates, Inc., 50189–50229.

[7] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness Through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS '12)*. ACM, New York, NY, USA, 214–226.

[8] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*. ACM, New York, NY, USA, 2221–2231.

[9] Elizabeth Gómez, Ludovico Boratto, and Maria Salamó. 2022. Provider Fairness across Continents in Collaborative Recommender Systems. *Information Processing & Management* 59, 1 (Jan. 2022), 102719.

[10] Sophie Greenwood, Sudalakshmee Chiniah, and Nikhil Garg. 2024. User-Item Fairness Tradeoffs in Recommendations. In *Advances in Neural Information Processing Systems (NIPS'24)*. Curran Associates, Inc., 114236–114288.

[11] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems* (Barcelona, Spain) *(NIPS'16)*. Curran Associates, Inc., Red Hook, NY, USA, 3323–3331.

[12] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated Gain-Based Evaluation of IR Techniques. *ACM Transactions on Information Systems* 20, 4 (Oct. 2002), 422–446.

[13] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.

[14] Siqi Lai, Erli Meng, Fan Zhang, Chenliang Li, Bin Wang, and Aixin Sun. 2022. An Attribute-Driven Mirror Graph Network for Session-based Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*. ACM, New York, NY, USA, 1674–1683.

[15] Yunqi Li, Yingqiang Ge, and Yongfeng Zhang. 2021. CIKM 2021 Tutorial on Fairness of Machine Learning in Recommender Systems. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (CIKM '21)*. ACM, Virtual Event Queensland Australia, 4857–4860.

[16] Weiwen Liu, Jun Guo, Nasim Sonboli, Robin Burke, and Shengyu Zhang. 2019. Personalized Fairness-Aware Re-Ranking for Microlending. In *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys '19)*. ACM, Copenhagen Denmark, 467–471.

[17] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a Fair Marketplace: Counterfactual Evaluation of the Trade-off between Relevance, Fairness & Satisfaction in Recommendation Systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. ACM, New York, NY, USA, 2243–2251.

[18] Mohammadmehdi Naghiaei, Hossein A. Rahmani, and Yashar Deldjoo. 2022. CPFair: Personalized Consumer and Producer Fairness Re-ranking for Recommender Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*. ACM, New York, NY, USA, 770–779.

[19] Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. 2020. FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms. In *Proceedings of The Web Conference 2020 (WWW '20)*. ACM, New York, NY, USA, 1194–1204.

[20] Felix Petersen, Debarghya Mukherjee, Yuekai Sun, and Mikhail Yurochkin. 2021. Post-Processing for Individual Fairness. In *Advances in Neural Information Processing Systems (NIPS '21)*. Curran Associates Inc., Red Hook, NY, USA, 25944–25955.

[21] David Powers. 2011. Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness & Correlation. *Journal of Machine Learning Technologies* 2, 1 (2011), 37–63.

[22] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing Personalized Markov Chains for Next-basket Recommendation. In *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*. ACM, New York, NY, USA, 811–820.

[23] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. Introduction to Recommender Systems Handbook. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). Springer US, Boston, MA, 1–35.

[24] Nasim Sonboli, Farzad Eskandanian, Robin Burke, Weiwen Liu, and Bamshad Mobasher. 2020. Opportunistic Multi-aspect Fairness through Personalized Re-ranking. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20)*. ACM, New York, NY, USA, 239–247.

[25] Weihao Song, Yushun Dong, Ninghao Liu, and Jundong Li. 2022. GUIDE: Group Equality Informed Individual Fairness in Graph Neural Networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22)*. ACM, New York, NY, USA, 1625–1634.

[26] Mengting Wan, Jianmo Ni, Rishabh Misra, and Julian McAuley. 2020. Addressing Marketing Bias in Product Recommendations. In *Proceedings of the 13th International Conference on Web Search and Data Mining (WSDM '20)*. ACM, New York, NY, USA, 618–626.

[27] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A Survey on the Fairness of Recommender Systems. *ACM Transactions on Information Systems* 41, 3 (July 2023), 1–43.

[28] Haolun Wu, Chen Ma, Bhaskar Mitra, Fernando Diaz, and Xue Liu. 2022. A Multi-Objective Optimization Framework for Multi-Stakeholder Fairness-Aware Recommendation. *ACM Transactions on Information Systems* 41, 2 (Dec. 2022), 47:1–47:29.

[29] Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. 2022. Joint Multisided Exposure Fairness for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*. ACM, New York, NY, USA, 703–714.

[30] Kun Wu, Jacob Erickson, Wendy Hui Wang, and Yue Ning. 2023. Equipping Recommender Systems with Individual Fairness via Second-Order Proximity Embedding. In *Proceedings of the 2022 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM '22)*. IEEE Press, Istanbul, Turkey, 171–175.

[31] Yao Wu, Jian Cao, Guandong Xu, and Yudong Tan. 2021. TFROM: A Two-sided Fairness-Aware Recommendation Model for Both Customers and Providers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*. ACM, New York, NY, USA, 1013–1022.

[32] Xiaopeng Ye, Chen Xu, Jun Xu, Xuyang Xie, Gang Wang, and Zhenhua Dong. 2024. Guaranteeing Accuracy and Fairness under Fluctuating User Traffic: A Bankruptcy-Inspired Re-ranking Approach. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*. ACM, Boise ID USA, 2991–3001.

[33] Kexin Yin, Xiao Fang, Bintong Chen, and Olivia R. Liu Sheng. 2023. Diversity Preference-Aware Link Recommendation for Online Social Networks. *Information Systems Research* 34, 4 (Dec. 2023), 1398–1414.

[34] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. FA*IR: A Fair Top-k Ranking Algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM '24)*. ACM, Singapore Singapore, 1569–1578.

[35] Eckart Zitzler and Lothar Thiele. 1999. Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach. *IEEE Transactions on Evolutionary Computation* 3, 4 (Nov. 1999), 257–271.