FRS4CPP: A Fair Recommendation Strategy Considering Interests of Users, Providers and Platform

Haiyan Zhao[1], Ping Zhou[1], Jian Cao[2(✉)], and Nengjun Zhu[3]

[1] University of Shanghai for Science and Technology, Shanghai, China
[2] Shanghai Jiaotong University, Shanghai, China
cao-jian@sjtu.edu.cn
[3] Shanghai University, Shanghai, China
zhu_nj@shu.edu.cn

Abstract. Currently, the research on the fairness of recommender systems has expanded beyond considering only the interests of users and product providers. However, the stakeholders of recommender systems go beyond just users and product providers; the platform that provides the recommendation service is also an important player whose interests are currently overlooked in the recommendation algorithm. In this study, we address this gap by considering the potential gains of the platform in addition to recommendation quality and fairness among providers. We analyze the theoretical relationships between recommendation quality, fairness among providers, and the platform's potential gains. Subsequently, we propose a fair recommendation strategy that takes into account the interests of all three parties. Through experiments conducted on a real-world dataset, we demonstrate that our models successfully achieve the desired design goals.

Keywords: three-sided interests · fairness of recommendation · fairness among providers · platform's potential gains

1 Introduction

The objective of a personalized recommender system is to provide recommendations to users (customers) by understanding their preferences. Studies have shown that recommendations significantly impact user decision-making and contribute to a positive user experience [20]. However, it is important to note that recommender systems involve multiple stakeholders [1–3]. Apart from users, there are also product providers, the platforms that operate the recommender system, and other participants.

Users aim to have personalized recommendations that cater to their individual needs, which is a fundamental objective of existing recommendation algorithms [20]. Simultaneously, the recommended products originate from various providers. Typically, products ranked higher in the recommendation list tend

Y. Sun et al. (Eds.): ChineseCSCW 2023, CCIS 2012, pp. 363–377, 2024.
https://doi.org/10.1007/978-981-99-9637-7_27

to gather more attention. Consequently, if provider A consistently receives lower rankings compared to provider B, it leads to a situation where provider A obtains fewer sales due to reduced exposure of their products in the recommendations [23]. Thus, it is crucial to ensure fairness in the exposure of products among providers to prevent dissatisfaction and potential attrition of providers from the platform.

The platform serves as a mediator by offering recommendation services that facilitate transactions between providers and users [18]. The platform derives benefits from charging for product exposure and clicks [6,14,23]. However, if the platform prioritizes its own revenue generation too heavily, it may lean towards granting superior positions to products whose providers can afford to pay more. Unfortunately, such a approach not only diminishes fairness among providers but also negatively impacts user satisfaction. Ultimately, this results in the loss of users and the departure of providers from the platform.

It is evident that the objective of recommender systems must expand beyond solely ensuring user satisfaction to encompass the interests of multiple stakeholders. Currently, there are limited studies focusing on enhancing fairness among users or providers within recommender systems. Some research, such as [4,9], addresses fairness among users, while others, such as [17,21,22], tackle fairness among product providers. Furthermore, a few studies simultaneously consider fairness among both parties, i.e., users and product providers, in the recommender system. For example, [15,16] examine the fairness among both product providers and users.

We introduce a novel fair recommendation strategy that considers the interests of multiple stakeholders simultaneously, including users, providers, and the platform. This approach represents the first attempt to address the concerns of all parties involved. Specifically, users' interests are assessed based on recommendation quality, providers' interests are captured through fairness in the exposure of their products, and the platform's interests are reflected by the revenue generated from users' clicks on the recommended products from providers.

Our contributions can be summarized as follows:

1) We analyze the interrelationships among recommendation quality, fairness among providers, and the platform's potential gains. Additionally, we propose the fair recommendation problem, which, for the first time, considers the interests of users, providers, and the platform.
2) We propose the Fair Recommendation Strategy considering the interests of users, providers, and the platform (FRS4CPP). FRS4CPP dynamically adjusts the user's recommended list while ensuring recommendation quality, thus promoting individual fairness among providers and serving the interests of the platform. To implement FRS4CPP, we design three algorithms.
3) We conduct experiments using a real-world dataset of user behaviors on the Taobao platform, incorporating simulated pay-per-click information. The experimental results demonstrate that our strategy effectively upholds the interests of all stakeholders involved.

2 Related Work

The fairness of recommender systems can be classified from various perspectives [23]. Individual fairness emphasizes equal treatment for each individual, while group fairness focuses on equal treatment among groups with specific attributes. In terms of the number of recommendations, fairness can be divided into cumulative recommendation fairness and single round recommendation fairness. Regarding the stage at which fairness is ensured within the recommender system, algorithms can be classified into three categories: pre-processing [7], in-processing [4], and post-processing [12,13,19]. Depending on the stakeholders within the system, the fairness of recommender systems can be further divided into three main categories: user-side fairness, provider-side fairness, and fairness for other stakeholders [2,3].

Currently, most research on recommendation fairness primarily focuses on a single stakeholder's perspective, with only a few studies examining fairness from the viewpoints of both users and product providers.

In terms of user-side fairness, researchers primarily focus on investigating whether there are systematic differences in the quality or types of recommended products for different users [1,23]. For instance, [9] addresses the issue of varying recommendation quality for users with different levels of activity, both at the individual and group levels, within the context of knowledge graph recommendations. They aim to eliminate or mitigate this unfair phenomenon. Another study, [4], mitigates unfairness to users by learning a set of adversarial filters and removing user-sensitive attribute information from graph embeddings. Additionally, [10] considers user preferences from an Envy-Free perspective [5,8], ensuring that users with similar product preferences are presented with similar recommendations.

Regarding provider-side fairness, the focus is mainly on exploring how providers can obtain fair opportunities for selection and whether their products meet fairness requirements within user recommendation lists. For example, [21] proposes explaining the unfair treatment of product providers using non-sensitive attributes and employs causal graphs to assess unfair treatment. In [22], provider-side fairness in top-k ranking is quantified by analyzing the distribution of products in user recommendation lists. The fairness of recommendation based on provider exposure is considered in [17].

A few studies take into account both user-side fairness and provider-side fairness. In [15], the FairRec algorithm is introduced to maximize the exposure of most product providers while ensuring envy-free fairness among users. Regarding the impact of platform updating algorithms on the fairness of product exposure, [16] proposes an integer linear programming (ILP) based algorithm that guarantees stable exposure rates for product providers after algorithm updates, while ensuring each user's minimum utility.

The aforementioned studies primarily focus on user-side fairness or provider-side fairness within recommender systems. However, there is limited research that considers the interests of other stakeholders. For the first time, we consider the simultaneous inclusion of recommendation quality, provider-side fairness, and the platform's interests in our study.

3 Fair Recommendation for the Interests of Users, Providers, and Platform

We consider the application scenario of top-k recommendation, where recommendations are generated based on users' historical behavior. To ensure a balance among recommendation quality, fairness among product providers, and the potential gains of the platform, we propose to reorder the original recommendation list. Specifically, we employ nDCG [11] as a measure of recommendation quality. Additionally, we quantify the benefits of product providers using Click-Through Rate (CTR), and Pay-Per-Click (PPC) [14] serves as the basis for providers to pay the platform.

3.1 Notations

We introduce the following notations:

- $U = \{u_1, u_2, \ldots, u_m\}$ denotes the set of users.
- $P = \{p_1, p_2, \ldots, p_n\}$ denotes the set of providers.
- $I = \{i_1, i_2, \ldots, i_s\}$ denotes the set of recommended items.
- $L^{ori} = \{l^{ori}_{u_1}, l^{ori}_{u_2}, \ldots, l^{ori}_{u_m}\}$ represents the set of original recommendation lists.
- $L^{rec} = \{l^{rec}_{u_1}, l^{rec}_{u_2}, \ldots, l^{rec}_{u_m}\}$ represents the set of recommendation lists outputted to customers.
- $i^p = \{i^p_1, i^p_2, \ldots, i^p_l\}$ represents the set of items provided by provider p.
- $i^{p,u} = \{i^{p,u}_1, i^{p,u}_2, \ldots, i^{p,u}_t\}$ represents the set of recommended items provided by provider p in the recommendation list of customer u.
- $v_{u,i}$ denotes the relevance between user u and item i.

3.2 Recommendation Quality

Ideally, in a recommendation algorithm, a user's recommendation list should be sorted based on their preferences, reflecting the relevance between the user and the recommended products. However, when other objectives are introduced, such as fairness considerations, the item rankings in the list may be adjusted. As a result, some items with lower relevance may be ranked higher, leading to a potential decline in the recommendation quality. Therefore, we use the original recommendation list l^{ori}_u as a baseline to evaluate the recommendation quality after adjustments.

To quantify the recommendation quality, we employ two classic metrics from information retrieval: Discount Cumulative Gain (DCG) and Normalized Discount Cumulative Gain (nDCG).

3.3 Fairness Among Product Providers

Previous studies on provider fairness mainly focused on measuring the fairness based on the exposure of products. However, exposure alone does not fully capture the economic benefits that product providers can potentially gain. In e-commerce, Click-Through Rate (CTR) is commonly used to measure the economic gains of product providers [17]. Therefore, we utilize CTR as a metric to quantify the possible economic benefits of product providers.

CTRs of Products and Providers. The position of a product in the recommendation list significantly influences its CTR. We assume that the higher a product is ranked in the list, the higher its CTR will be. Although there are other factors that can affect CTR, such as product images displayed on the platform, we consider the position as the sole factor influencing CTR in this research.

CTR represents the click probability of a product at each position in the recommendation list and is calculated as follows:

$$CTR\,(i_x^{p,u}) = \frac{1}{\log_2 x + 1}. \tag{1}$$

Since each provider typically has multiple products on the platform, we introduce the concept of CTRs of providers. Let T_R denote the total number of recommendations, and T_R^p represent the total number of times provider p appears in the recommendation lists. The CTRs of provider p are calculated by summing the CTRs of its products and dividing by T_R^p:

$$CTR^p = \frac{\sum_{j=1}^{|U|} \sum_{x=1}^{|i^{p,u_j}|} CTR(i_x^{p,u_j})}{T_R^p} \tag{2}$$

Fairness of CTRs Among Providers. A fair recommendation system aims to achieve consistency in the CTRs among product providers. In other words, the recommendation is fair if the CTRs of all providers are equal:

$$CTR^{p_i} = CTR^{p_j}, \forall p_i, p_j \in P \tag{3}$$

To measure the level of fairness among providers, we calculate the average of all providers' CTRs as the baseline:

$$CTR_{avg} = \frac{\sum_{j=1}^{|P|} \sum_{s=1}^{|U|} CTR(i^{p_j,u_s})}{T_R}. \tag{4}$$

The difference between the baseline and a provider's CTR is defined as:

$$PCF^p = CTR_{avg} - CTR^p. \tag{5}$$

When $PCF^p > 0$, it indicates that the CTR of product provider p is lower than the average level, which is unfair to product provider p. Conversely, when $PCF^p < 0$, it implies that the CTR of product provider p is higher than the average, which is unfair to other product providers.

Furthermore, we use the variance of all providers' CTRs to evaluate the overall fairness level among product providers.

3.4 Potential Gains of the Platform

Pay-Per-Click (PPC) is a common pricing model used by platforms to charge product providers, and it is widely adopted by companies such as Taobao, Google, and Baidu [14]. Therefore, we utilize PPC as a metric to estimate the total potential gains of the platform.

The PPC can vary for different providers as it is negotiated between the providers and the platform. Let $value^p$ denote the PPC for provider p. The expected gains from a position can be calculated as follows:

$$TC(i_x^p) = CTR(i_x^p) * value^p \tag{6}$$

When the platform prioritizes its own potential gains, it tends to promote products with higher PPC rankings to increase its revenue. The total potential gains of the platform can be estimated using the following equation:

$$Gains = \sum_{y=1}^{|P|} \sum_{x=1}^{|i^{p_y}|} CTR(i_x^{p_y}) * value^{p_y}. \tag{7}$$

3.5 Relations Among Recommendation Quality, Provider Fairness, and Platform Gains

Experimental studies [20] have shown that when a recommender system perfectly aligns with users' preferences, it often leads to unfairness among product providers. Additionally, the potential gains of the platform are closely related to the PPC. When we make adjustments to the rankings to achieve fairness among product providers or maximize platform gains, the following three observations can be made:

Observation 1: Any recommendation that prioritizes objectives other than user preferences will result in lower recommendation quality compared to the original recommendation list l_u^{ori}, which is designed to best match users' preferences.

When the system modifies the original recommendation lists to enhance fairness among product providers' CTRs, the $nDCG\,(l_u^{rec})$ value becomes less than 1. However, the PPC of a provider remains independent of their CTRs, and the platform's interests do not exhibit any systematic changes. Hence, we observe:

Observation 2: Strategic adjustments made to the recommendation list l_u^{ori} to improve fairness among providers lead to a decrease in recommendation quality and random changes in the platform's interests.

When the system modifies the original recommendation lists to increase the platform's potential gains, the rankings of products from providers with higher PPCs are promoted, causing $nDCG\,(l_u^{rec})$ to become less than 1. However, the fairness among providers remains unchanged since the CTRs of each position are independent of the PPCs of individual providers. Therefore, we have:

Observation 3: As the potential gains of the platform increase, recommendation quality decreases, while fairness among providers' CTRs remains unaffected.

Based on the above observations, achieving fairness in recommendations for the interests of all three stakeholders requires sacrificing recommendation quality in order to enhance fairness among providers and maximize platform gains.

4 FRS4CPP: The Fair Recommendation Strategy Considering Interests of Customers, Providers, and Platform

Based on the analysis in Sect. 3.5, we have devised the strategy for fair recommendation that takes into account the interests of users, product providers, and platforms.

4.1 Strategy Design

To achieve fairness among product providers or maximize platform gains, it is necessary to lower the recommendation quality. However, we aim to minimize the impact on recommendation quality. Therefore, we combine all these objectives into a single value:

$$Value_{FRS4CPP} = w_1 * v_{rq} + w_2 * v_{fp} + w_3 * v_{PG} \tag{8}$$

Here, v_{rq}, v_{fp}, and v_{PG} represent utilities related to recommendation quality, fairness among providers, and potential gains of the platform, respectively.

As analyzed in Sect. 3, this problem can be reduced to a knapsack problem, which is known to be NP-complete even when considering the interests of only two sides. For instance, when considering recommendation quality and provider fairness, the length of the recommendation list can be considered as the knapsack capacity, the products to be recommended as the items to be placed in the knapsack, and maximizing provider fairness as the objective. The problem becomes more complex when the potential gains of the platform are taken into consideration. Hence, we propose heuristic fairness recommendation strategies to address this problem.

4.2 Algorithms for Implementing the Strategy

We employ three algorithms to implement FRS4CPP.

Overall Score-Based Algorithm (OSA). OSA calculates a synthetic score for each product in each position of the list and selects the products with the highest score for each position, starting from the highest position and moving downward. The score of a product for position x is computed using Eq. 7:

$$score_x^i = \frac{v_{i,u}}{\log_2(x+1)} + \alpha * PCF^p + \beta * value^p \tag{9}$$

The factors influencing recommendation quality are the relevance between users and products. The fairness among providers PCF^p represents the cumulative fairness of an individual provider over multiple rounds of recommendations, thus affecting the system's dynamic adjustment of the current recommendation list. The main factor influencing the potential gains of the platform is the pc^p

of product provider p. We assign weights to PCF^p and $value^p$ according to the following rules: 1) to account for magnitude differences among the three interest measurements and 2) to adjust the relative importance of the three interest measurements. If the system emphasizes recommendation quality, the values of α or β should be reduced. If the fairness among product providers or the potential gains of the platform are given more attention, the values of α or β can be correspondingly increased.

OSA reorders the original list l_u^{ori} only if $score_{x-1}^{i_{x-1}} < score_x^{i_x}$. In this case, the original list l_u^{ori} is modified, and PCF_x^p or the platform's potential gains increase.

Two-Round Reordering Algorithm (2RA). The 2RA strategy consists of two steps:

Step 1: We calculate the benefit score for both recommendation quality and fairness among providers by weighting and evaluating the factors. Each product's score for a given position in the top-k ranking is determined, and the product with the highest score is selected for each position, from highest to lowest, as the first round of recommendation results. The calculation method for both sides is presented in Eq. 10:

$$\text{score}_x^i = \frac{v_{i,u}}{\log_2(x+1)} + a * PCF^p. \tag{10}$$

Step 2: The recommendation list obtained in the first step is fine-tuned to maximize the platform's potential gains. We compare the price per click of adjacent products in the recommendation list obtained from the first step. If the difference between the products at positions x and $x-1$ exceeds a threshold value ($\triangle price$), we exchange the products at these positions.

In the first step, we adjust the list based on the interests of both sides. In the original list l_u^{ori}, for products ranked at positions x and $x-1$, where $v_{x-1} \geq v_x$, we change the list l_u^{ori} only if the fairness among providers' click-through rates (CTRs) satisfies $PCF_x^p > PCF_{x-1}^p$. As a result, the position of the provider with a lower CTR rises, increasing the fairness of provider CTRs.In the second step, the products at positions x and $x-1$ are exchanged when the following condition is met:

$$value_x^p - value_{x-1}^p > \triangle value. \tag{11}$$

This leads to an increase in the platform's potential gains. The value of $\triangle price$ determines the extent to which the second step influences the first step.

Tree-Round Reordering Algorithm (3RA). 3RA determines the final list of recommendations through three rounds of selection and sorting. We select the products in the order of the original list to guarantee the user's recommendation quality:

The 1st Round: The products whose providers' CTRs are lower than fair CTR ($PCF > 0$), and whose PPC are lower than the average of PPC ($value_i^p > \overline{value}$) are listed in the recommended list l_u^{rec}.

The 2nd Round: The products whose product providers' CTRs are lower than fair CTR ($PCF > 0$), or whose PPCs are lower than the average CTR ($value_i^p > \overline{value}$) are still listed in the recommended list l_u^{rec}.

The 3rd Round: The remaining products are sequentially added to the list based on their relevance, taking into account the recommendation quality.

In the 3RA strategy, the three rounds of selection are conducted in a sequential manner, following the original order, with each round considering the recommendation quality during reordering. In the first round, the emphasis is on prioritizing products that enhance fairness among providers and maximize the platform's potential gains. The ranking of products that align with the interests of all three parties is placed at the forefront of the list, ensuring the maximum alignment of interests. In the second round, the focus is on selecting and sorting products that further improve fairness among providers or increase the platform's potential gains. Products that cater to the interests of the user and either the provider or the platform are given secondary priority. This ensures a certain degree of alignment between the interests of the two sides. The first two rounds of reordering already include products that significantly enhance fairness among providers or maximize the platform's potential gains. Therefore, in the third round of reordering, only the recommendation quality is considered, as the primary objective.

4.3 Time Complexity Analysis

In OSA, the algorithm calculates the scores of each product in each position, selects the product for each position based on the scores, and updates the CTRs of providers. In the worst-case scenario, the time complexity is $O(k^2 + kn)$. However, in practice, the number of providers (n) is typically greater than the length of the recommendation list (k), resulting in a time complexity of $O(kn)$ for OSA. Similarly, in the first step of 2RA, the algorithm calculates the scores of each product in each position, selects the product for each position based on the scores, and then fine-adjusts the product's position in the second step using click values. Finally, it updates the providers' CTRs. Therefore, the time complexity of 2RA is also $O(kn)$. In 3RA, the algorithm updates the CTRs of providers after three rounds of selection. Thus, the time complexity of 3RA is also $O(kn)$.

5 Experiments

5.1 Datasets and Metrics

To conduct our experiments, we utilize a desensitized dataset from Taobao. The original dataset comprises behavior data from 1.14 million users over a span of 8 days. For our experiment, we extract behavior data from 10,782 users to create a new dataset.

To generate the original recommendation lists, we employ the ALS algorithm to match users with products. However, obtaining the PPC (Pay-Per-Click) information for each provider is challenging, if not impossible, due to its confidential nature. To simulate the PPC information, we generate random numbers between 0 and 1, following a normal distribution, to represent the PPC values.

For evaluating the recommendation quality, fairness among product providers, and potential gains of the platform, we utilize the metrics introduced in Sect. 3.

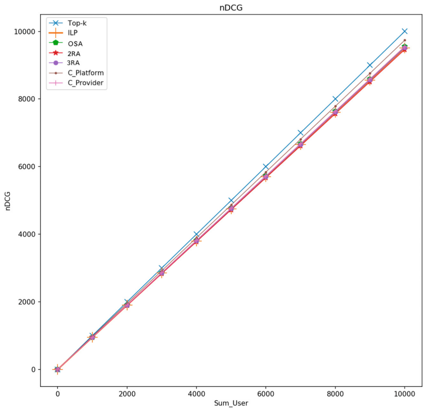

Fig. 1. Cumulative Value of nDCG

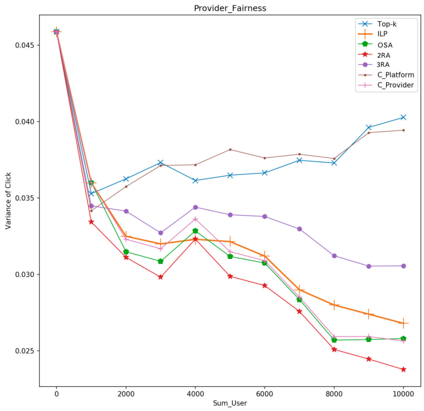

Fig. 2. Variance of CTRs of Providers

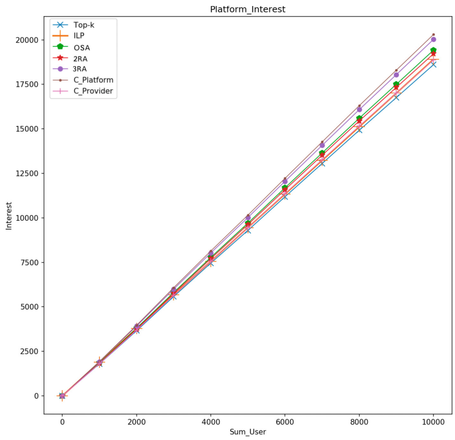

Fig. 3. Cumulative Potential Gains of Platform

5.2 Compared Approaches

In our comparative analysis, we evaluate our proposed approach against the following algorithms:

1) *Top-k.* This algorithm is based on item-based collaborative filtering and directly recommends the top-k items, focusing solely on maximizing recommendation quality.
2) *An ILP-based fair ranking mechanism.* This algorithm utilizes Integer Linear Programming (ILP) to ensure Quality Weighted Fairness in provider exposure. It aims to minimize the absolute difference between the two cumulative values while considering recommendation quality as a constraint.
3) *C-Provider.* This algorithm considers only recommendation quality and provider fairness, without any deliberate bias towards specific providers. We utilize the result of the first step of the 2RA algorithm to generate a recommendation list.
4) *C-Platform.* This algorithm considers only recommendation quality and the platform's potential gains, without favoring any specific providers. We utilize the second step of the 2RA algorithm to generate a recommendation list.

5.3 Experimental Results

Figures 1, 2 and 3 present the experimental results, demonstrating that the three algorithms trade off a portion of the recommendation quality to improve fairness among providers and increase the potential gains of the platform.

Table 1. The Result of the Experiment Considering Two side Interests

Algorithms	nDCG	Variance of CTRs	Platform Gains
Top-k	10782	0.0404	20000
C-Platform	10503	0.0398	21808
C-Provider	10300	0.0234	20180

It can be observed that the cumulative nDCGs of the post-processed recommendation results, denoted as the sum of cumulative nDCGs, are lower than the sum of cumulative nDCGs in $l_u^o ri$ (the original list) when considering the objectives of platform potential gains and fairness among providers. This confirms the validity of Observation 1.

Table 1 presents the experimental results of the C-Provider algorithm. It can be seen that when the system considers the goals of recommendation quality and fairness among providers, the platform's potential gains in the C-Provider algorithm (20180) only slightly deviate from the original list $l_u^o ri$ of the Top-k algorithm (20000). Therefore, Observation 2 is supported. Similarly, the experimental results of the C-Platform algorithm demonstrate that when the system considers recommendation quality and platform's potential gains, the variance of providers' CTRs (0.0398) is slightly lower compared to the original list $l_u^o ri$ (0.0404). This supports Observation 3.

Furthermore, the results of the ILP-based fair ranking mechanism highlight that while the algorithm ensures recommendation quality and fairness among providers, it performs poorly in terms of platform potential gains.

Experiment Results of OSA. In the OSA algorithm, this experiment involves adjusting the weight coefficient α of PCF (Provider Cumulative Fairness) and the weight coefficient β of the potential gains of the platform. The results of adjusting α are presented in Fig. 4:

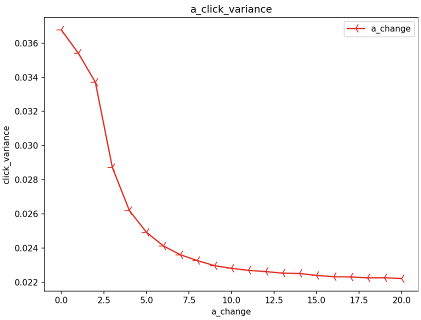

Fig. 4. The Change of CTRs of Providers

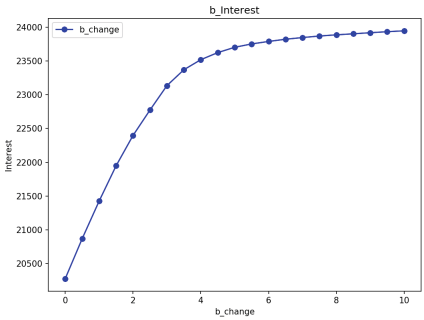

Fig. 5. The Change of Potential Gains of Platform

As the weight coefficient α gradually increases from 0 to 20, the variance of the provider's CTRs shows a decreasing trend and eventually stabilizes. This indicates that when the system prioritizes the fairness among providers, the level of unfairness among them gradually diminishes but cannot be completely eliminated.

The experimental results of adjusting β are depicted in Fig. 5. As β increases from 0 to 10, the platform's potential gains demonstrate an upward trend and eventually stabilize. Assuming that the recommended products are sorted in descending order of PPCs (Pay-Per-Click), the platform's potential gains reach their maximum. However, in practice, the achieved potential gains may not reach the absolute maximum due to other constraints and factors.

It is evident that while the OSA algorithm experiences a decrease in recommendation quality to some extent, both the fairness among providers and the platform's potential gains are improved.

Experiment Results of 2RA. To examine the impact of $\triangle price$ on the first step and its influence on the platform's potential gains, this study adjusts the value of $\triangle price$.

To investigate the effect of the algorithm on recommendation quality, we introduce a binary weight Q to represent the consideration of recommendation quality in the first step of scoring. When $Q = 1$, the algorithm takes recommendation quality into account. As presented in Table 2, as $\triangle price$ gradually decreases, the platform's potential gains increase, the fairness among providers improves, and the recommendation quality decreases. When $Q = 0$, indicating that recommendation quality is not considered, the nDCG (Normalized Discounted Cumulative Gain) for recommendation quality is 9845. In the case of $Q = 1$, as $\triangle price \rightarrow 0$, the second step has the most significant impact on the outcome of the first step, resulting in cumulative nDCGs of 10100, which surpass the result when recommendation quality is not considered. Hence, even when $\triangle price \rightarrow 0$, the influence of the second step on the recommendation quality of the first step remains limited, and the 2RA algorithm can ensure recommendation quality.

Table 2. The Result of the Experiment on 2RA

$\triangle price$	Q	a	nDCG	Variance of CTRs	Platform Gains
1	1	5	10300	0.0234	20277
0.5	1	5	10286	0.0223	20313
0.25	1	5	10205	0.0214	20636
0.1	1	5	10160	0.0211	20736
0.01	1	5	10106	0.0211	20891
0.001	1	5	10100	0.0212	20900
0.0001	1	5	10100	0.0212	20900
0.00001	1	5	10100	0.0212	20900
0.25	0	5	9845	0.0208	20501
0.25	1	0	10503	0.0398	21808

Regarding the 2RA algorithm, as depicted in Fig. 1, compared to the Top-k algorithm, there is a certain decrease in recommendation quality. However, there are noticeable improvements in both the fairness among providers and the platform's potential gains.

Experiment Results of 3RA. When the recommendation quality is not considered, as in the case of $Q = 0$, an experiment called CFPRM is conducted. The results of the CFPRM algorithm and 3RA are presented in Table 3. The result indicates that 3RA ensures a certain level of recommendation quality.

Table 3. The Result of the Experiment on 3RA

Algorithms	nDCG	Variance of CTRs	Platform Gains
Top-k	10782	0.0404	20000
CPFRM	9845	0.0208	20501
3RA	10266	0.0291	21520

It can be found from the results that the 3RA manages to control the decrease in recommendation quality to a certain extent, while simultaneously improving both the fairness among providers and the potential gains of the platform.

6 Conclusions

This paper aims to address the interests of users, providers, and the platform by improving fairness among providers, maximizing the potential gains of the platform, and ensuring recommendation quality. To achieve this goal, we propose a fair recommendation strategy and design three algorithms to implement it.

In this paper, we process the original top-k recommendation list through post-processing methods (reordering) to obtain a recommendation list that satisfies all three objectives. In the future, we can explore in-processing approaches, where we directly revise the recommendation algorithm itself to achieve these objectives.

References

1. Abdollahpouri, H., et al.: Multistakeholder recommendation: survey and research directions. User Model. User-Adap. Inter. **30**(1), 127–158 (2020)
2. Abdollahpouri, H., Burke, R.: Multi-stakeholder recommendation and its connection to multi-sided fairness (2019)
3. Abdollahpouri, H., Burke, R., Mobasher, B.: Recommender systems as multistakeholder environments. In: Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization, pp. 347–348. ACM, New York (2017)
4. Bose, A.J., Hamilton, W.L.: Compositional fairness constraints for graph embeddings (2019)
5. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. J. Polit. Econ. **119**(6), 1061–1061 (2010)
6. Burke, R.: Multisided fairness for recommendation (2017)
7. Calmon, F.P., Wei, D., Vinzamuri, B., Ramamurthy, K.N., Varshney, K.R.: Optimized pre-processing for discrimination prevention. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, pp. 3995–4004. Curran Associates Inc., Red Hook (2017)
8. Do, V., Corbett-Davies, S., Atif, J., Usunier, N.: Online certification of preference-based fairness for personalized recommender systems. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, no. 6, pp. 6532–6540 (2022)

9. Fu, Z., et al.: Fairness-aware explainable recommendation over knowledge graphs. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 69–78. Association for Computing Machinery, New York (2020)
10. Ilvento, C., Jagadeesan, M., Chawla, S.: Multi-category fairness in sponsored search auctions. In: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* 2020, pp. 348–358. Association for Computing Machinery, New York (2020). https://doi.org/10.1145/3351095.3372848
11. Järvelin, K., Kekäläinen, J.: Cumulated gain-based evaluation of IR techniques. ACM Trans. Inf. Syst. **20**(4), 422–446 (2002)
12. Karako, C., Manggala, P.: Using image fairness representations in diversity-based re-ranking for recommendations. In: Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization, pp. 23–28. Association for Computing Machinery, New York (2018). https://doi.org/10.1145/3213586.3226206
13. Liu, W., Burke, R.: Personalizing fairness-aware re-ranking (2018)
14. Liu Peng, W.C.: Advertising Computing: Market and Technology of Realizing Internet Business. Posts & Telecom, Beijing (2019)
15. Patro, G.K., Biswas, A., Ganguly, N., Gummadi, K.P., Chakraborty, A.: FairRec: two-sided fairness for personalized recommendations in two-sided platforms. In: Proceedings of The Web Conference 2020, pp. 1194–1204. ACM, New York (2020). https://doi.org/10.1145/3366423.3380196
16. Patro, G.K., Chakraborty, A., Ganguly, N., Gummadi, K.: Incremental fairness in two-sided market platforms: on smoothly updating recommendations. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, no. 01, pp. 181–188 (2020)
17. Singh, A., Joachims, T.: Fairness of exposure in rankings. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 2219–228. ACM, New York (2018). https://doi.org/10.1145/3219819.3220088
18. Sühr, T., Biega, A.J., Zehlike, M., Gummadi, K.P., Chakraborty, A.: Two-sided fairness for repeated matchings in two-sided markets: a case study of a ride-hailing platform. In: the 25th ACM SIGKDD International Conference, pp. 3082–3092 (2019)
19. Wu, C., Wu, F., Wang, X., Huang, Y., Xie, X.: FairRec: fairness-aware news recommendation with decomposed adversarial learning (2020)
20. Wu, Y., Cao, J., Xu, G., Tan, Y.: 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, pp. 1013–1022. Association for Computing Machinery, New York (2021)
21. Wu, Y., Zhang, L., Wu, X.: On discrimination discovery and removal in ranked data using causal graph. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 2536–2544. ACM, New York (2018). https://doi.org/10.1145/3219819.3220087
22. Yang, K., Stoyanovich, J.: Measuring fairness in ranked outputs. In: Proceedings of the 29th International Conference on Scientific and Statistical Database Management (2017). https://doi.org/10.1145/3085504.3085526
23. Zhao, H.Y., Zhou, P., Chen, Q.K., Cao, J.: Fairness evaluation methods in recommender system: current research and prospects. J. Chin. Comput. Syst. **43**(3), 456–465 (2022)