



Balancing Between Accuracy and Fairness for Interactive Recommendation with Reinforcement Learning

Weiwen Liu¹, Feng Liu², Ruiming Tang^{1,2,3}(✉), Ben Liao^{1,2,3},
Guangyong Chen³(✉), and Pheng Ann Heng^{1,3}

¹ The Chinese University of Hong Kong, Sha Tin, Hong Kong
tangruiming2015@163.com

² Harbin Institute of Technology, Harbin, China

³ Guangdong Provincial Key Laboratory of Computer Vision and Virtual Reality
Technology, Shenzhen Institutes of Advanced Technology,
Chinese Academy of Sciences, Shenzhen, China
gychen@link.cuhk.edu.hk

Abstract. Fairness in recommendation has attracted increasing attention due to bias and discrimination possibly caused by traditional recommenders. In Interactive Recommender Systems (IRS), user preferences and the system's fairness status are constantly changing over time. Existing fairness-aware recommenders mainly consider fairness in static settings. Directly applying existing methods to IRS will result in poor recommendation. To resolve this problem, we propose a reinforcement learning based framework, *FairRec*, to dynamically maintain a long-term balance between accuracy and fairness in IRS. User preferences and the system's fairness status are jointly compressed into the state representation to generate recommendations. FairRec aims at maximizing our designed cumulative reward that combines accuracy and fairness. Extensive experiments validate that FairRec can improve fairness, while preserving good recommendation quality.

1 Introduction

Interactive Recommender Systems (IRS) have been widely implemented in various fields, *e.g.*, news, movies, and finance [20]. Different from the conventional recommendation settings [11], IRS consecutively recommend items to individual users and receive their feedback in interactive processes. IRS gradually refine the recommendation policy according to the obtained user feedback in an online manner. The goal of such a system is to maximize the total utility over the whole interaction period. A typical utility of IRS is user acceptance of recommendations. Conversion Rate (CVR) is one of the most commonly used measures of recommendation acceptance, computing the ratio of users *performing a system's desired activity* to users *having viewed recommended items*. A desired activity could be downloading from App stores, or making loans for microlending.

However, optimizing CVR solely may result in fairness issues, one of which is the unfair allocation of desired activities, like clicks or downloads, over different demographic groups. Under such unfair circumstances, majority (over-representing) groups may dominate recommendations, thereby holding a higher proportion of opportunities and resources, while minority groups are largely under-represented or even totally ignored. A fair allocation is a critical objective in recommendation due to the following benefits:

Legal. Recommendation in particular settings are explicitly mandated to guarantee fairness. In the setting of employment, education, housing, or public accommodation, a fair treatment with respect to race, color, religion, *etc.*, is required by the anti-discrimination laws [8]. For job recommendation, it is expected that jobs at minority-owned businesses are being recommended and applied at the same rate as jobs at white-owned businesses. In microlending, loan recommender systems must ensure borrowers of different races or regions have an equal chance of being recommended and funded.

Financial. Under-representing for some groups leads to the abandonment of the system. For instance, video sharing platforms like YouTube involve viewers and creators. It is desirable to ensure each creator has a fair chance of being recommended and promoted. Otherwise, if the new creators do not get adequate exposure and appreciation, they tend to leave the platform, resulting in less user-generated content. Consequently, users' satisfaction from both viewers and creators, as well as the platform's total income are affected in the long run.

The fairness concern in recommender systems is quite challenging, as accuracy and fairness are usually conflicting goals to be achieved to some extent. On the one hand, to obtain the ideal fairness, one could simply divide the recommendation opportunities equally to each item group, but users' satisfaction will be affected by being persistently presented with unattractive items. On the other hand, existing recommender systems have been demonstrated to favor popular items [5], resulting in extremely unbalanced recommendation results. Thus, our work aims to answer this question: *Can we achieve a fairer recommendation while preserving or just sacrificing a little recommendation accuracy?*

Most prior works consider fairness for the conventional recommender systems [1, 2], where the recommendation is regarded as a static process at a certain time instant. A general framework that formulates fairness constraints on rankings in terms of exposure allocation is proposed in [19]. Individual attention fairness is discussed in [3]. [21] models re-ranking with fairness constraints in Multi-sided Recommender Systems (MRS) as an integer linear programming. The balanced neighborhoods method [4] balances protected and unprotected groups by reformulating the Sparse Linear Method (SLIM) with a new regularizer.

However, it is hard to directly apply those methods to IRS due to:

- (i) It is infeasible to impose fairness constraints at every time instant. Forcing the system to be fair at any time and increasing fairness uniformly for all users will result in poor recommendations. In fact, IRS focus on the long-term cumulative utility over the whole interaction session, where the system

could focus on improving accuracy for users with particular favor, and the lack of fairness at the time can later be compensated when recommending items to users with diversified interests. As such, we can achieve long-term system’s fairness while preserving satisfying recommendation quality.

- (ii) Existing work only considers the distribution of the number of recommendations (exposure) an item group received. Actually, the distribution of the desired activities that take place after an exposure like clicks or downloads has much larger commercial value and can be directly converted to revenue.

To resolve the problem, we design a **Fairness-aware Recommendation** framework with reinforcement learning (FairRec) for IRS. FairRec jointly compresses the user preferences and the system’s fairness status into the current state representation. A two-fold reward is designed to measure the system gain regarding accuracy and fairness. FairRec is trained to maximize the long-term cumulative reward to maintain an accuracy-fairness balance. The major contributions of this paper are as follows:

- We formulate a fairness objective for IRS. To the best of our knowledge, this is the first work that balances between accuracy and fairness in IRS.
- We propose a reinforcement learning based framework, FairRec, to dynamically maintain a balance between accuracy and fairness in IRS. In FairRec, user preferences and the system’s fairness status are jointly compressed into the state representation to generate recommendations. We also design a two-fold reward to combine accuracy and fairness.
- We evaluate our proposed FairRec algorithm on both synthetic and real-world data. We show that FairRec can achieve a better balance between accuracy and fairness, compared to the state-of-the-art methods.

2 Problem Formulation

2.1 Markov Decision Process for IRS

In this paper, we model the fairness-aware recommendation for IRS as a finite time Markov Decision Process (MDP), with an action space \mathcal{A} , a state space \mathcal{S} , and a reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$. When a user u arrives at time step $t = 1, \dots, T$, the system observes the current state $s_t \in \mathcal{S}$ of the user u and takes an action $a_t \in \mathcal{A}$ (*e.g.*, recommending an item to the user).

The user views the item and provides feedback y_{a_t} , *e.g.*, clicking or downloading on the recommended item, if she feels interested. Let $y_{a_t} \in \{0, 1\}$ denote the user’s feedback, with $y_{a_t} = 1$ meaning the user performs desired activities, and 0 otherwise. The system then receives a reward r_t (a function of y_{a_t}), and updates the model. The problem formulation is formally presented as follows:

States \mathcal{S} : The state s_t is described by user preferences and the system’s fairness status. We jointly embed them into the current state representation. The detailed design of the state representation is given in Sect. 3.2.

Transitions \mathcal{P} : The transition of states models the dynamic change of user preferences and the system’s fairness. The successor state s_{t+1} is obtained once the user’s feedback at time t is collected.

Action \mathcal{A} : An action a_t is recommending an item chosen from the available candidate item set \mathcal{A} . Our framework can be easily extended to the case of recommending a list of items. To simplify our presentations, we focus on recommending an item at a time in this paper.

Reward \mathcal{R} : The reward r_t is a scalar measuring the system’s gain regarding accuracy and fairness after taking action a_t , elaborated in Sect. 3.3.

We aim to learn a policy π , mapping from states to actions $a_t = \pi(s_t)$, to generate recommendations that are both accurate and fair. The goal is to maximize the sum of discounted rewards (return) from time t onward, which is defined by $R_t^\gamma = \sum_{k=t}^T \gamma^{k-t} r_k$, and γ is the discount factor.

2.2 Weighted Proportional Fairness for IRS

Each item is associated with a categorical protected attribute $C \in \{c_1, \dots, c_l\}$. Let $\mathcal{A}_c = \{a | C = c, a \in \mathcal{A}\}$ denote the group of items with an attribute value c . Take loan recommendation for instance, if the protected attribute is the geographical region, then \mathcal{A}_c with $c = \text{“Oceania”}$ contains all the loans applied from Oceania. Denote by $x_t \in \mathbb{R}_+^l$ the *allocation vector*, where x_t^i represents the allocation proportion of group i up to time t ,

$$x_t^i = \frac{\sum_{k=1}^t y_{a_k} \mathbb{1}_{\mathcal{A}_{c_i}}(a_k)}{\sum_{i'=1}^l \sum_{k=1}^t y_{a_k} \mathbb{1}_{\mathcal{A}_{c_{i'}}}(a_k)}. \quad (1)$$

where $\mathbb{1}_A(x)$ equals to 1 if $x \in A$, and 0 otherwise. Recall that y_{a_k} is the user’s feedback on recommended item a_k . In loan recommendation, x_t^i denotes the rate of funded loans from the region i over all funded ones up to time t .

In this work, we focus on a well-accepted and axiomatically justified metric of fairness, the weighted proportional fairness [9]. Weighted proportional fairness is a generalized Nash solution for multiple groups.

Definition 1 (Weighted Proportional Fairness). *An allocation of desired activities x_t is weighted proportionally fair if it is the solution of the following optimization problem,*

$$\max_{x_t} \sum_{i=1}^l w_i \log(x_t^i), \quad s.t. \quad \sum_{i=1}^l x_t^i = 1, \quad x_t^i \geq 0, \quad i = 1, \dots, l. \quad (2)$$

The coefficient $w_i \in \mathbb{R}_+$ is a pre-defined parameter weighing the importance of each group. The optimal solution can be easily solved by standard Lagrangian multiplier methods, namely

$$x_*^i = \frac{w_i}{\sum_{i'=1}^l w_{i'}}. \quad (3)$$

We aim to improve the weighted proportional fairness $\sum_{i=1}^l w_i \log(x_T^i)$ while preserving high conversions $\sum_{t=1}^T y_{a_t}$ up to time T .

3 Proposed Model

This section begins with a brief overview of our proposed FairRec. After that, we introduce the components of FairRec and the learning algorithm in detail.

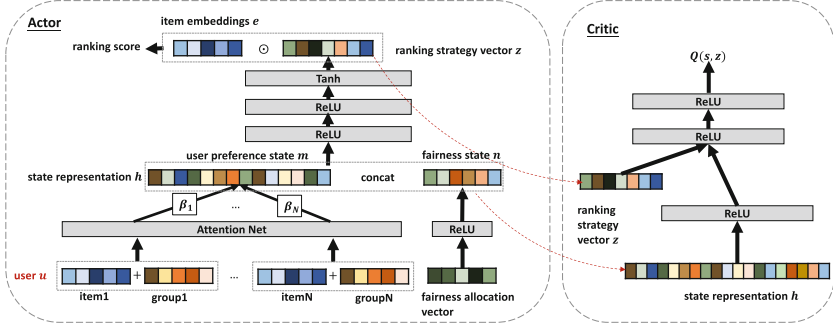


Fig. 1. The architecture of FairRec.

3.1 Overview

To balance between accuracy and fairness in the long run, we formulate IRS recommendation as an MDP, which is then solved by reinforcement learning.

The previously studied reinforcement learning models can be categorized as follows: *Value-based* methods approximate the value function, then the action with the largest value is selected [26, 27]. Value-based methods are more sample-efficient and steady, but the computational cost is high when the action space is large. *Policy-based* methods directly learn a policy that takes as input of the current user state and outputs an action [6, 24], which generally have a faster convergence. *Actor-critic* architectures take advantage of both value-based and policy-based methods [15, 25]. Therefore, we design our model following the actor-critic framework.

The overall architecture of FairRec is illustrated in Fig. 1, which consists of an *actor* network and a *critic* network. The actor network performs time-varying recommendations according to the dynamic user preferences and the fairness status. The critic network estimates the value of the outputs associated with the actor network to encourage or discourage the recommended items.

3.2 Personalized Fairness-Aware State Representation

We propose a personalized fairness-aware state representation to jointly consider accuracy and fairness, which is composed of the the User Preference State (UPS)

and the Fairness State (FS). State representation learns a non-linear transformation $h_t = f_s(s_t)$ that maps the current state s_t to a continuous vector h_t .

User Preference State (UPS). UPS represents personalized user preferences. We propose a two-level granularity representation: the item-level and the group-level. The item-level representation indicates the user’s fine-grained preferences to each item, while the group-level representation shows the user’s coarse-grained interests in each item group. Such two-level granularity representation provides more information on the propensity of different users towards diverse recommendation. Therefore, the agent could focus on accuracy for the users with particular favor, and the lack of fairness at a point in time can later be compensated when recommending items to users with diverse interests.

The input of UPS is the sequence of the user u ’s N most recent positively interacted items, as well as the corresponding group IDs that the items belong to at t . Items belonging to the same group share the same protected attribute value c . Each item a is mapped to a continuous embedding vector $e_a \in \mathbb{R}^d$. The embedding vector of each group ID e_g is the average of the embedding vectors of all items belonging to the group g . Then each item is represented by

$$\epsilon_a = e_a + e_g, \quad (4)$$

where $\epsilon_a \in \mathbb{R}^d$, and item a belongs to group g . The group embedding e_g is added to serve as a global bias (or a regularizer), allowing items belonging to the same group to share the same group information.

As for a specific user u , the affects of different historical interactions on her future interest may vary significantly. To capture this sequential dependencies among the historical interacted items, we apply an attention mechanism [23] to weigh each item in the interacted item sequence. The attention net learns a weight vector β of size N , $\beta = \text{Softmax}(\omega^1 \sigma(\omega^2 [\epsilon_{a_1}, \dots, \epsilon_{a_N}] + b^2) + b^1)$, where $\omega^1, b^1, \omega^2, b^2$ are the network parameters and $\sigma(\cdot)$ is the ReLU activation function. The user preference state representation m_t is obtained by multiplying the attention weights with the corresponding item representations as $m_t = [\beta_1 \epsilon_{a_1}, \dots, \beta_N \epsilon_{a_N}]$, where m_t is of dimension $N \times d$ and β_i denotes the i -th entry in the weight vector β . Therefore, the items currently contributing more to the outcome are assigned with higher weights.

Fairness State (FS). The input of FS is the current allocation distribution of the desired activities at time t . As a complementary for UPS, FS provides evidence of the current fairness status and helps the agent to promote items belonging to under-represented groups. In particular, we deploy a Multi-Layer Perceptron (MLP) to map the allocation vector x_t to a latent space, $n_t = \text{MLP}(x_t)$. Then we concatenate m_t and n_t to obtain the final state representation,

$$h_t = [m_t || n_t], \quad (5)$$

with $||$ denotes concatenation operation.

3.3 Reward Function Design

The reward is designed to measure the system’s gain regarding accuracy and fairness. Existing reinforcement learning frameworks for recommendation only consider the recommendation accuracy, and one commonly used definition of reward is $r = 1$ if the user performs desired activities and -1 otherwise [15, 25]. To incorporate the fairness measure into IRS, we propose a two-fold reward by first examining whether the user performs the desired activities on the recommended item, and then evaluating the fairness gain of performing such a desired activity.

As discussed in Sect. 2, to achieve the weighted proportional fairness, the optimal allocation vector is $x_*^i = \frac{w_i}{\sum_{i'=1}^I w_{i'}}$, with w_i the pre-defined target allocation proportion of group i . Therefore, we incorporate the deviation from the optimal solution $x_*^i - x_t^i$ into the reward as the fairness indicator:

$$r_t = \begin{cases} \sum_{i=1}^I \mathbb{1}_{\mathcal{A}_{c_i}}(a_t) (x_*^i - x_t^i + 1), & \text{if } y_{a_t} = 1 \\ -\lambda, & \text{if } y_{a_t} = 0 \end{cases}, \quad (6)$$

where $\mathbb{1}_A(x)$ is the indicator function and is 1 when $x \in A$, 0 otherwise, x_t^i is the allocation proportion of group i at time t . The constant $\lambda > 1$ is the penalty value for inaccurate recommendations and manages the accuracy-fairness tradeoff. A larger λ means that the agent focuses more on accuracy.

Since the fairness metric (Eq. (1) and Eq. (2)) is computed according to the number of the desired activities, only positive y_{a_t} influences fairness. Therefore, we simply give a negative reward $-\lambda$ for $y_{a_t} = 0$ to punish the undesired activities. When $y_{a_t} = 1$, we compute the fairness score $x_*^i - x_t^i$, which is the difference between the optimal distribution and current allocation. Suppose the user performs a desired activity on the item $a_t \in \mathcal{A}_{c_i}$. Then the fairness score $x_*^i - x_t^i$ is negative if the i -th group is over-representing ($x_t^i > x_*^i$), and is more negative if \mathcal{A}_{c_i} already has a higher rate of the desired activity, indicating that the system should focus more on other groups. Similarly, the fairness score $x_*^i - x_t^i$ is positive if the i -th group is currently under-representing ($x_t^i < x_*^i$), and is more positive if \mathcal{A}_{c_i} is more lacking in the desired activity. We add 1 to the fairness score to ensure the reward is positive if $y_{a_t} = 1$.

To sum up, the agent receives a large positive reward if the user performs a desired activity on the item and the item belongs to an under-representing group. Whereas the reward is a smaller positive number if the activity is desired, but the item belongs to an over-representing (majority) group. We punish the most severely with $y_{a_t} = 0$, as it neither contributes to accuracy nor fairness.

3.4 Model Update

Actor Network. The actor network extracts latent features from s_t and outputs a ranking strategy vector z_t . The recommendation is performed according to the ranking vector by $a_t = \arg \max_{a \in \mathcal{A}} e_a^\top z_t$. In particular, we first embed s_t to h_t following the architecture described in Sect. 3.2, then we stack fully-connected

layers on top of h_t to learn the nonlinear transformation and generate z_t , as presented in Fig. 1.

Suppose the policy $\pi_\theta(s)$ learned by the actor is parameterized by θ . The actor is trained according to $Q_\eta(s_t, z_t)$ from the critic, and updated by the sampled policy gradient [18] with α_θ as the learning rate, B as the batch size,

$$\theta \leftarrow \theta + \alpha_\theta \frac{1}{B} \sum_t \nabla_z Q_\eta(s, z)|_{s=s_t, z=\pi_\theta(s_t)} \nabla_\theta \pi_\theta(s)|_{s=s_t}, \quad (7)$$

Critic Network. The critic adopts a deep neural network $Q_\eta(s_t, z_t)$, parameterized by η , to estimate the expected total discounted reward $\mathbb{E}[R_t^\gamma | s_t, z_t; \pi]$, given the state s_t and the ranking strategy vector z_t under the policy π . Specifically for this problem, the network structure is designed as follows

$$Q_\eta(s_t, z_t) = \text{MLP}([\sigma(W_h h_t + b_h) || z_t]), \quad (8)$$

by first mapping h_t to the same space as z_t with a fully-connected layer and then concatenating it with z_t , while $\text{MLP}(\cdot)$ denotes a multi-layer perceptron, and $h_t = f_s(s_t)$ is the state representation as presented in Sect. 3.2.

We use the temporal-difference (TD) learning [22] to update the critic. The loss function is the mean square error $L = \sum_t (\nu_t - Q_\eta(s_t, z_t))^2$, where $\nu_t = r_t + \gamma Q_{\eta'}(s_{t+1}, \pi_{\theta'}(s_{t+1}))$. The term $\nu_t - Q_\eta(s_t, z_t)$ is called time difference (TD), η' and θ' are the parameters of the target critic and actor network that are periodically copied from η, θ and kept constant for a number of iterations to ensure the stability of the training [14]. The parameter θ is updated by gradient descent, with α_η the learning rate and B the batch size:

$$\eta \leftarrow \eta + \alpha_\eta \frac{1}{B} \sum_t (\nu_t - Q_\eta(s_t, z_t)) \nabla_\eta Q_\eta(s_t, z_t). \quad (9)$$

4 Experiments

4.1 Experimental Settings

We evaluate the proposed FairRec algorithm on both synthetic and real-world data, comparing with the state-of-the-art recommendation methods in terms of fairness and accuracy.

Datasets. We use MovieLens¹ and Kiva.org datasets for evaluation. **MovieLens** is a public benchmark dataset for recommender systems, with 943 users, 1,602 items and 100,000 user-item interactions. Since the MovieLens data do not have protected attributes, we created 10 groups to represent differences among group inventories, and randomly assigned movies to each of such groups following a geometric distribution. An interaction with the rating (ranging from 1 to

¹ <https://grouplens.org/datasets/movielens>.

5) larger than 3 is defined as a desired activity in calculating CVR. **Kiva.org** is a proprietary dataset obtained from Kiva.org, consisting of lending transactions over a 6-month period. We followed the pre-processing technique used in [16] to densify the dataset. The retained dataset has 1,589 loans, 589 lenders and 43,976 ratings. The geographical region of loans is selected as the protected attribute, as Kiva.org has a stated mission of equalizing access to capital across different regions so that loans from each region have a fair chance to be funded. We define a transaction amount greater than USD25 as the desired activity for Kiva.

Evaluation Metrics. We evaluate the recommendation accuracy by the Conversion Rate (CVR):

$$\text{CVR} = \frac{\sum_{k=1}^T y_{a_k}}{T}, \quad (10)$$

and measure the fairness by Weighted Proportional Fairness (PropFair)²:

$$\text{PropFair} = \sum_{i=1}^l w_i \log(1 + x_T^i). \quad (11)$$

Table 1. Experimental results on MovieLens and Kiva.

	MovieLens			Kiva		
	CVR	PropFair	UFG	CVR	PropFair	UFG
NMF	0.7972	0.8592	4.2362	0.4211	0.8473	1.4635
SVD	0.8478	0.8337	5.4795	0.4870	0.8686	1.6931
DeepFM	<u>0.8612</u>	0.8098	5.8323	0.6349	0.8671	2.3752
LinUCB	0.8577	0.8464	5.9476	0.6517	0.8697	2.4970
DRR	0.8592	0.8470	<u>6.0177</u>	<u>0.6567</u>	0.8645	<u>2.5183</u>
MRPC	0.8361	<u>0.8608</u>	5.2508	0.4286	<u>0.8761</u>	1.5332
FairRec	0.8702*	0.8666*	6.6776*	0.6905*	0.8838*	2.8555*

We conduct a two-sided significant test [17] between FairRec and the strongest baseline DRR, where * means the p-value is smaller than 0.05.

Moreover, we propose a Unit Fairness Gain (UFG) to jointly consider accuracy and fairness,

$$\text{UFG} = \frac{\text{PropFair}}{\text{CVR}_{\max} - \text{CVR}} = \frac{\text{PropFair}}{1 - \text{CVR}}. \quad (12)$$

UFG indicates the fairness of the system under unit accuracy budget. For any recommendation system, the ideal maximum CVR, namely CVR_{\max} , equals to 1. Thus UFG can be interpreted as the slope of fairness versus accuracy—the fairness gain if we decrease a unit accuracy from CVR_{\max} . A larger UFG means a higher value of PropFair can be achieved with unit deviation from CVR_{\max} , namely, the larger, the better.

² The input of PropFair is shifted by one to avoid infinity results.

Reproducibility. We randomly sample 80% of the user with associated rating sequences for training, and 10% for validation, 10% for testing, so that the item dependencies within each session can be learned. We use grid search to select the hyper-parameters for all the methods to maximize the hybrid metric UFG: the embedding dimension in $\{10, 30, 50, 100\}$, the learning rate in $\{0.0001, 0.001, 0.01\}$. Embedding vectors are pre-trained using standard matrix factorization [11] following the traditional processing as in [15, 25]. For the proposed FairRec, we set the number of recent interacted items $N = 5$, discount factor $\gamma = 0.9$, the width of each hidden layer of the actor-critic network is 1000. The batch size is set to 1024, and the optimization method is Adam. Without loss of generality, we set $w_i = 1, i = 1, \dots, l$. All results are averaged from multiple independent runs.

4.2 Results and Analysis

Comparison with Existing Methods. We compare our proposed FairRec with six representative recommendation algorithms: **(i) NMF.** Non-negative Matrix Factorization (NMF) [12] estimates the rating matrix with positive user and item factors; **(ii) SVD.** Singular Value Decomposition (SVD) [10] is the classic matrix factorization based method that decomposes the rating matrix via a singular value decomposition; **(iii) DeepFM.** DeepFM [7] is the state-of-the-art deep learning model in recommendation that combines the factorization machines and deep neural networks; **(iv) LinUCB.** LinUCB [13] is the state-of-the-art contextual bandits algorithm that sequentially selects items and balances between exploitation and exploration in IRS; **(v) DRR.** DRR [15] is a deep reinforcement learning framework designed for IRS to maximize the long-term reward; **(vi) MRPC.** Multi-sided Recommendation with Provider Constraints (MRPC) [21] is the state-of-the-art fairness-aware method by formulating the fairness problem as an integer programming.

Table 1 shows the results. Bold numbers are the best results and underlined numbers are the strongest baselines. We have the following observations:

First, the deep learning based method (DeepFM) outperforms matrix factorization based methods (NMF and SVD) in CVR, while PropFair of DeepFM is lower. This is consistent with our expectation that DeepFM combines low-order and high-order feature interactions and has great fitting capability, yet it solely maximizes the accuracy, with fairness issues overlooked.

Second, LinUCB and DRR generally achieve better CVR than matrix factorization and deep learning methods. It is because LinUCB and DRR consider the IRS setting, and aims to maximize the long-term reward. Compared LinUCB to DRR, LinUCB underperforms DRR since LinUCB assumes states of the system remain unchanged and fails to tailor the recommendation to match the dynamic user preferences. DRR is the strongest baseline as it achieves the best tradeoff between accuracy and fairness, with UFG = 6.0177 on MovieLens and UFG = 2.5183 on Kiva, respectively.

Third, MRPC considers fairness by adding fairness constraints for static recommendation. Therefore, MRPC generates the fairest recommendation on both

datasets, but the CVR significantly decreases as MRPC ignores the dynamic change of user preferences and the fairness status.

Fourth, FairRec consistently yields the best performance in terms of CVR, PropFair, and UFG on both datasets, demonstrating FairRec is effective in maintaining the accuracy-fairness tradeoff over time. FairRec outperforms the strongest baselines, DRR, by 1.3%, 2.3%, and 11% in CVR, PropFair, and UFG on MovieLens, and 5.1%, 2.2%, and 13.4% on Kiva. Considering UFG, with unit accuracy loss, FairRec achieves the most fairness gain. FairRec observes the current user preferences and the fairness status, and estimates the long-term discounted cumulative reward. Therefore, FairRec is capable of long-term planning to manage the balance between accuracy and fairness.

Influence of Embedding Dimension. Embedding dimension d is an important factor for FairRec. We study how the embedding dimension d influences the performance of FairRec. We vary d in $\{10, 30, 50\}$, and run 2500 epochs. The cumulative reward and the test performance are plotted in Fig. 2.

We observe that when d is large ($d = 30$ and $d = 50$), the algorithm benefits from sufficient expressive power and the reward converges at a high level. As for $d = 10$, the cumulative reward converges fast at a relatively low value, indicating that the model suffers from the limited fitting capability. In terms of UFG value, UFG = 6.68 when $d = 50$, which is slightly better than 6.6 as $d = 30$. Similar results can be found on Kiva, which is omitted for limited space. Therefore, we select $d = 50$ in FairRec for all the experiments.




Fig. 2. Experimental results with embedding dimension d on MovieLens: cumulative reward (left) and CVR, PropFair, and UFG (right).

Table 2. Ablation study on MovieLens and Kiva.

	MovieLens			Kiva		
	CVR	PropFair	UFG	CVR	PropFair	UFG
FairRec(reward-)	0.8561	0.8053	5.5957	0.6935	0.8670	2.8290
FairRec(state-)	0.8194	0.8758	4.8494	0.6723	0.8746	2.6688
FairRec	0.8702	0.8666	6.6776	0.6905	0.8838	2.8555

Ablation Study. To evaluate the effectiveness of different components (*i.e.*, the state representation and the reward function) in FairRec, we replace a component of FairRec with the standard setting in RL at each time, and compare the performance with the full-fledged FairRec. Experimental results are presented in Table 2. We design two variants: **FairRec(reward-)** with standard reward as in [15, 25]; and **FairRec(state-)** with simple concatenation of item embeddings as the state representation as in [15].

Results show that FairRec(reward-) generally has high CVR, as no punishment on unfair recommendation. Moreover, the model simply optimizes accuracy, failing to balance between accuracy and fairness. As for FairRec(state-), CVR is downgraded significantly, validating the importance of our designed state representation. Overall, UFG of FairRec is the largest, confirming that all the components of FairRec work together yield the best results.

5 Conclusion

In this work, we propose a fairness-aware recommendation framework in IRS to dynamically balance between accuracy and fairness in the long run with reinforcement learning. In the proposed state representation component, UPS models both personalized preference and propensity to diversity; FS is utilized to describe the current fairness status of IRS. A two-fold reward is designed to combine accuracy and fairness. Experimental results demonstrate the effectiveness in the balance of accuracy and fairness of our proposed framework over the state-of-the-art models.

Acknowledgements. This work is supported by National Natural Science Foundation of China (No. U1813204).

References

1. Abdollahpouri, H., et al.: Beyond personalization: Research directions in multi-stakeholder recommendation. *User Model. User Adap. Inter.* (2020)
2. Abdollahpouri, H., Burke, R.: Multi-stakeholder recommendation and its connection to multi-sided fairness. *arXiv preprint arXiv:1907.13158* (2019)
3. Biega, A.J., et al.: Equity of attention: amortizing individual fairness in rankings. In: *SIGIR* (2018)
4. Burke, R., et al.: Balanced neighborhoods for multi-sided fairness in recommendation. In: *FAT** (2018)
5. Celma, Ö., Cano, P.: From hits to niches? Or how popular artists can bias music recommendation and discovery. In: *Proceedings of the 2nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition* (2008)
6. Chen, H., et al.: Large-scale interactive recommendation with tree-structured policy gradient. In: *AAAI* (2019)
7. Guo, H., et al.: DeepFM: a factorization-machine based neural network for CTR prediction. In: *IJCAI* (2017)
8. Holmes, E.: Antidiscrimination rights without equality. *Mod. Law Rev.* (2) (2005)

9. Kelly, F.P., et al.: Rate control for communication networks: shadow prices, proportional fairness and stability. *JORS* **49**, 237–252 (1998). <https://doi.org/10.1057/palgrave.jors.2600523>
10. Koren, Y.: Factor in the neighbors: Scalable and accurate collaborative filtering. *TKDD* **4**(1), 1:1–1:24 (2010)
11. Koren, Y., et al.: Matrix factorization techniques for recommender systems. *Computer* (2009)
12. Lee, D.D., et al.: Algorithms for non-negative matrix factorization. In: *NIPS* (2000)
13. Li, L., et al.: A contextual-bandit approach to personalized news article recommendation. In: *WWW* (2010)
14. Lillicrap, T.P., et al.: Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* (2015)
15. Liu, F., et al.: Deep reinforcement learning based recommendation with explicit user-item interactions modeling. *arXiv preprint arXiv:1810.12027* (2018)
16. Liu, W., Burke, R.: Personalizing fairness-aware re-ranking. *arXiv preprint arXiv:1809.02921* (2018)
17. Ruxton, G.D.: The unequal variance t-test is an underused alternative to student’s t-test and the mann–whitney u test. *Behav. Ecol.* (2006)
18. Silver, D., et al.: Deterministic policy gradient algorithms. In: *ICML* (2014)
19. Singh, A., Joachims, T.: Fairness of exposure in rankings. In: *KDD* (2018)
20. Steck, H., et al.: Interactive recommender systems: tutorial. In: *RecSys* (2015)
21. Sürer, Ö., et al.: Multistakeholder recommendation with provider constraints. In: *RecSys* (2018)
22. Sutton, R.S., et al.: *Introduction to Reinforcement Learning* (1998)
23. Vaswani, A., et al.: Attention is all you need. In: *NeurIPS* (2017)
24. Wang, X., et al.: A reinforcement learning framework for explainable recommendation. In: *ICDM* (2018)
25. Zhao, X., et al.: Deep reinforcement learning for page-wise recommendations. In: *RecSys* (2018)
26. Zhao, X., et al.: Recommendations with negative feedback via pairwise deep reinforcement learning. In: *KDD* (2018)
27. Zheng, G., et al.: DRN: a deep reinforcement learning framework for news recommendation. In: *WWW* (2018)