



ProFairRec: Provider Fairness-aware News Recommendation

Tao Qi

Department of Electronic
Engineering, Tsinghua University
Beijing, China
taoqi.qt@gmail.com

Fangzhao Wu*

Microsoft Research Asia
Beijing, China
wufangzhao@gmail.com

Chuhan Wu

Department of Electronic
Engineering, Tsinghua University
Beijing, China
wuchuhan15@gmail.com

Peijie Sun

Hefei University of Technology
Beijing, China
sun.hfut@gmail.com

Le Wu

Hefei University of Technology
Anhui, China
lewu.ustc@gmail.com

Xiting Wang

Microsoft Research Asia
Beijing, China
xitwan@microsoft.com

Yongfeng Huang

Department of Electronic
Engineering, Tsinghua University
Beijing, China
yfhuang@tsinghua.edu.cn

Xing Xie

Microsoft Research Asia
Beijing, China
xingx@microsoft.com

ABSTRACT

News recommendation aims to help online news platform users find their preferred news articles. Existing news recommendation methods usually learn models from historical user behaviors on news. However, these behaviors are usually biased on news providers. Models trained on biased user data may capture and even amplify the biases on news providers, and are unfair for some minority news providers. In this paper, we propose a provider fairness-aware news recommendation framework (named *ProFairRec*), which can learn news recommendation models fair for different news providers from biased user data. The core idea of *ProFairRec* is to learn provider-fair news representations and provider-fair user representations to achieve provider fairness. To learn provider-fair representations from biased data, we employ provider-biased representations to inherit provider bias from data. Provider-fair and -biased news representations are learned from news content and provider IDs respectively, which are further aggregated to build fair and biased user representations based on user click history. All of these representations are used in model training while only fair representations are used for user-news matching to achieve fair news recommendation. Besides, we propose an adversarial learning task on news provider discrimination to prevent provider-fair news representation from encoding provider bias. We also propose an orthogonal regularization on provider-fair and -biased representations to better reduce provider bias in provider-fair representations. Moreover, *ProFairRec* is a general framework and can be applied to different

news recommendation methods. Extensive experiments on a public dataset verify that our *ProFairRec* approach can effectively improve the provider fairness of many existing methods and meanwhile maintain their recommendation accuracy.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

News Recommendation, Provider Fairness, Adversarial Learning

ACM Reference Format:

Tao Qi, Fangzhao Wu, Chuhan Wu, Peijie Sun, Le Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. 2022. ProFairRec: Provider Fairness-aware News Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*, July 11–15, 2022, Madrid, Spain. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3477495.3532046>

1 INTRODUCTION

In recent years, online news platforms such as Apple News have become more and more popular for people to read news [41, 49, 55]. Since massive news articles are collected by online news platforms from different news providers everyday [20, 45], it is very difficult for users to manually find their interested news [42, 58]. Thus, news recommendation, which aims to recommend news according to users' personal interest [28, 30, 57], has become very critical for online news platforms to improve user experience [2, 11, 13, 31].

Most of the existing news recommendation methods first learn news and user representations from news content and user reading history respectively, and then match them for personalized news recommendation [24, 56]. Besides, these methods usually rely on historical users' behavior data on news to train recommendation models [37, 47]. For example, An et al. [1] employed an attentive convolutional network to learn news representations from news titles and categories. They proposed to capture short-term user interest from user's clicked news via a GRU network and long-term

*The corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '22, July 11–15, 2022, Madrid, Spain.

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-8732-3/22/07...\$15.00
<https://doi.org/10.1145/3477495.3532046>




Figure 1: Average exposure number per news of providers on the *MIND* dataset [50]. We randomly select 50 providers.

user interest via ID embeddings and combined them to form user interest representations. Besides, they further model user interest in news based on the inner product of their representations and exploit user historical click and non-click data for model training. Wu et al. [44] proposed to learn news representations from news titles, entities, and tags via transformers and proposed a heterogeneous graph pooling network to learn user representations from personalized user graphs. Besides, they also performed interest matching via inner product and trained models on click data.

In general, users' reading behaviors on different news providers are usually biased [21]. Some popular news providers may take up most clicks while clicks on some minority news providers are usually rare. Thus, news recommendation models trained on biased user data may capture and even amplify the biases on news providers, which may tend to recommend more news from popular providers and become unfair for some minority news providers [27, 33, 52]. For example, Fig. 1 shows the average exposure number of news from different news providers in the real-world news recommendation dataset *MIND* [50], we can see that news exposures are biased on different news providers. Fig. 2 shows average numbers of news from popular and unpopular providers recommended by several SOTA news recommendation models trained on *MIND*. Fig. 2 verifies that models learned from biased data can capture provider bias and become unfair for minority providers. However, most of the existing news recommendation methods do not consider provider fairness, which may hurt the diversity of news sources and perspectives, as well as the revenue of minority providers [34].

In this paper, we propose a provider fairness-aware news recommendation framework named *ProFairRec*, which can improve provider fairness of recommendation models learned from biased data. The core idea of *ProFairRec* is to learn provider-fair news and user representations. In order to learn provider-fair representations from biased data, we utilize provider-biased representations to inherit bias from data. Specifically, for news modeling, we learn provider-fair news representations from news content and learn provider-biased news representations from provider IDs. For user modeling, we learn provider-fair and -biased user representations by aggregating provider-fair and -biased representations of users' clicked news, respectively. We aggregate provider-fair and -biased representations of news and users for training the recommendation model to achieve decomposition of bias-aware and bias-independent information and only match provider-fair news and user representations to achieve fair news recommendation. In addition, we propose an adversarial learning task on news




Figure 2: Average exposure opportunity of popular and unpopular providers in the top recommended results of different news recommendation models on *MIND*. Popular providers are defined as top 50% providers ranked by average news click number per news article in historical data, and the remaining are regarded as unpopular providers.

provider discrimination and apply it to provider-fair news representations to prevent them from encoding provider bias. We also design an orthogonal regularization that enforces provider-fair and -biased representations to be orthogonal for better reducing bias in provider-fair representations. Moreover, *ProFairRec* is a general framework and can be applied to different news recommendation models. We conduct extensive experiments on a public news recommendation dataset. Results show that *ProFairRec* can effectively improve the provider fairness of many news recommendation methods and meanwhile maintain their recommendation accuracy.

The contribution of this paper is three-fold:

- To our best knowledge, we are the first to study the provider fairness problem in news recommendation.
- We propose a unified framework to learn provider-fair news representations and user representations from biased data.
- Extensive experiments verify that *ProFairRec* can effectively improve provider fairness of many news recommendation methods without significant accuracy drop.

2 RELATED WORK

2.1 News Recommendation

News recommendation technique is important for users to improve their experience on online news platforms [4, 18]. Mainstream news recommendation methods usually first model news from its content and model user interest from user's clicked news, and further match user interest with candidate news for news recommendation [17, 36, 41, 42, 46]. Besides, these methods usually train recommendation models on historical user behavior data. For example, Wu et al. [43] proposed to apply self-attention networks to learn news representations from news texts and learn user representations from user behavior contexts. They further measured user interest in candidate news based on the inner product of their representations. Qi et al. [32] proposed to learn news representations from news titles and entities via transformer networks. They further proposed to learn hierarchical user representations to model diverse and multi-grained user interest for matching candidate news. Moreover, both of these two methods trained news recommendation models on click and non-click data. In general, historical click data

on different news providers are usually biased. Popular providers may accumulate most clicks while click data on some minority news providers are usually rare. However, existing news recommendation models trained on historical data may inherit provider bias and become unfair for some minority news providers. Different from these methods, we propose a decomposed framework with adversarial learning, which can effectively reduce provider bias in news recommendation models learned from biased data.

2.2 Fair Recommender System

In recent years, fair recommender systems have attracted more and more attentions [6, 15, 23, 52]. Existing works usually consider recommendation fairness from two perspectives, i.e., user and provider [3]. Many existing fair recommender systems were proposed to improve user-side fairness [16, 48]. Some of these methods were designed to provide unbiased recommendations for users with different sensitive attributes (e.g., genders and ages) [48, 51] while some of them were proposed to provide recommendation services with equal qualities for users in different groups (e.g., user groups partitioned by gender) [7, 15, 54]. Different from these user-side fair recommendation methods, our *ProFairRec* framework focuses on provider-side recommendation fairness. There are also some works that have been proposed to improve provider-side recommendation fairness [3, 33]. These methods are usually designed to provide equal recommendation opportunities for different news providers [23, 25, 52]. For example, Liu et al. [21] proposed to re-rank candidate items based on their relevance with user interest and proportion of protected providers (e.g., minority providers) in recommendation lists. Sonboli et al. [33] proposed to prompt provider fairness by enhancing recommendation diversity. They proposed to re-rank candidate items based on their matching with user interest as well as intra-list item similarities. Wu et al. [52] proposed to use the difference between the proportion of protected providers in the recommendation list and item set as a constraint to re-rank candidate items. In short, these methods are usually based on manually designed re-ranking rules to re-balance the recommendation opportunities of different news providers to improve provider fairness. However, the manually designed re-ranking rules may be sub-optimal for achieving an effective trade-off between recommendation fairness and accuracy. Different from these methods, we propose a unified provider fairness-aware news recommendation framework to reduce provider bias in recommendation models with adversarial learning, which can jointly optimize recommendation fairness and accuracy during model training.

3 METHODOLOGY

We will first give a formal definition on news recommendation and provider fairness in it. Then, we will introduce our provider fairness-aware news recommendation method (named *ProFairRec*).

3.1 Problem Formulation

In news recommendation, we assume that there are H news providers in total: $\mathcal{P} = \{P_i | i = 1, \dots, H\}$, where P_i is the ID of the i -th provider and \mathcal{P} is the set of news providers. The set of news articles belonging to the provider $p \in \mathcal{P}$ is denoted as \mathcal{D}_p and the set of all news articles is denoted as \mathcal{D} . For each news article $d \in \mathcal{D}$, we assume

that the article d is composed of the textual content c and the ID of its provider p , i.e., $d = \{c, p\}$. Besides, we assume that the textual content c is composed of T words: $c = [t_1, \dots, t_T]$, where t_i is the i -th word. For each target user $u \in \mathcal{U}$, we assume that the target user has previously clicked m news articles $[d_1^u, \dots, d_i^u, \dots, d_m^u]$, where d_i^u is the i -th news clicked by the target user and \mathcal{U} is the set of users. Given a list of candidate news $[d_1^n, \dots, d_i^n, \dots, d_N^n]$ and a target user u , the recommendation model needs to predict the relevance scores $[\hat{r}_1^c, \dots, \hat{r}_i^c, \dots, \hat{r}_N^c]$ for matching user interest and candidate news, where \hat{r}_i^c is the predicted relevance score of the i -th candidate news and user interest. These candidate news are further ranked based on their relevance scores and the news recommender system will only display part of candidate news with the highest relevance scores. The goal of provider fairness-aware news recommendation is to recommend news according to user interest and meanwhile provide fair exposure opportunities for different news providers.

Following existing works [3, 33], we employ group-level fairness concept to define provider fairness to reduce the influence of the quality differences of news from different providers. News providers are first partitioned into multiple groups $\{\mathcal{P}_i | i = 1, 2, \dots, g\}$ based on concerned provider features (e.g., provider popularity [33] or provider revenue [3]), where \mathcal{P}_i is the i -th group and g is the number of provider groups. The group-level provider fairness requires that the recommendation model should provide equal recommendation opportunities for news of providers in different groups:

$$\frac{\sum_{u \in \mathcal{U}} \sum_{p \in \mathcal{P}_i} \sum_{d_i \in \mathcal{D}_p} \mathbb{1}(d_i, u)}{|\mathcal{U}| \sum_{p \in \mathcal{P}_i} |\mathcal{D}_p|} = \frac{\sum_{u \in \mathcal{U}} \sum_{p \in \mathcal{P}_j} \sum_{d_j \in \mathcal{D}_p} \mathbb{1}(d_j, u)}{|\mathcal{U}| \sum_{p \in \mathcal{P}_j} |\mathcal{D}_p|}, \quad (1)$$

where $\mathbb{1}(d, u)$ is an indicator function representing the strategy of the news recommender system. The indicator function $\mathbb{1}(d, u)$ will output 1 if the news d is recommended to the user u and will output 0 if the news d is not recommended to the user u .

3.2 Overall Framework

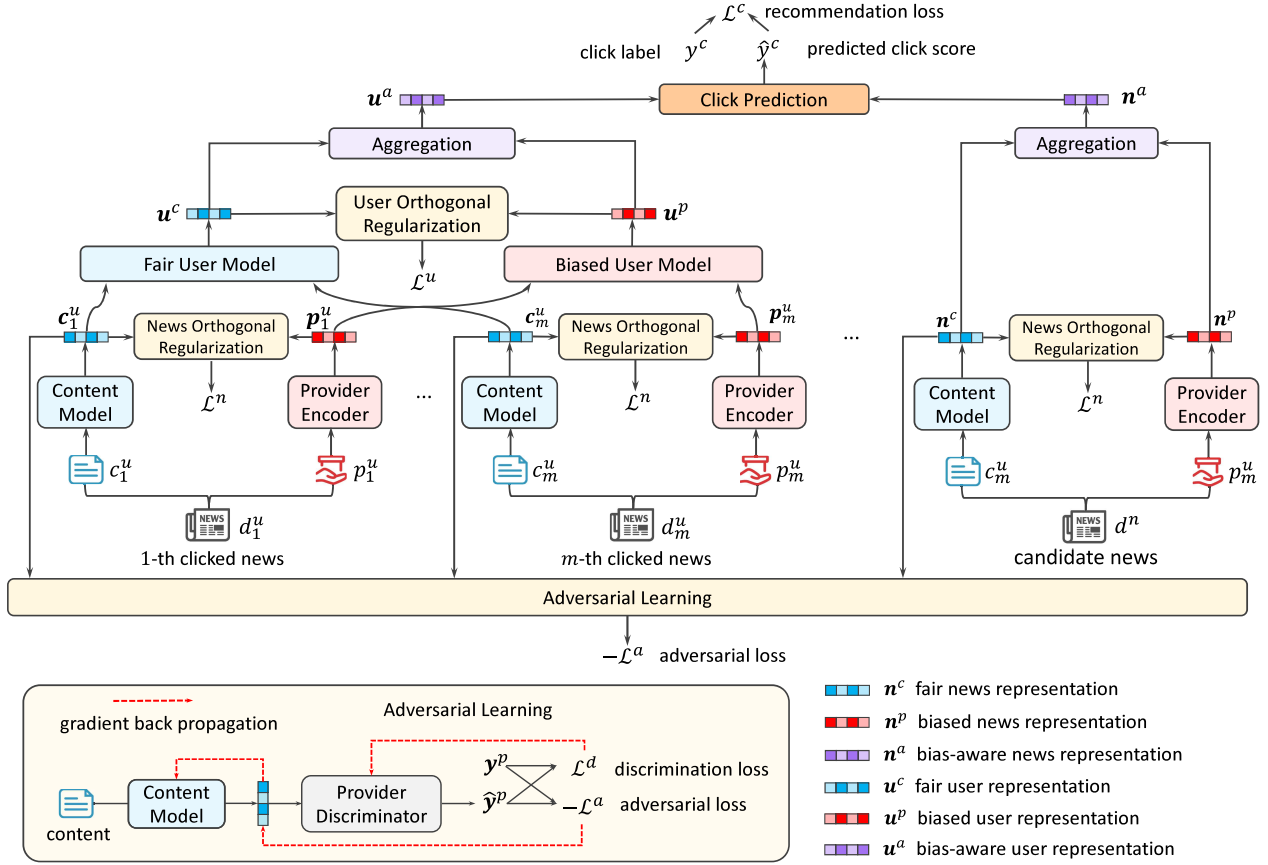
Next, we will introduce the overall framework of our provider fairness-aware news recommendation framework (*ProFairRec*), which can learn fair news recommendation models from biased data.

As shown in Fig. 3, the core idea of *ProFairRec* is to learn provider-fair news and user representations to achieve provider fairness. Note that training data usually contains both provider-biased signals (e.g., provider popularity) and provider-independent signals (e.g., the matching of user interest and news content), it may be difficult to employ a single news or user representation to purely encode provider-independent signals without encoding provider-biased signals. To tackle this challenge, motivated by Wu et al. [48], we employ provider-biased representations to inherit provider bias from biased data to further learn provider-fair representations. Specifically, given a news d with its content c and provider ID p , *ProFairRec* employs a content model Φ_c^n to learn provider-fair news representation \mathbf{c} from its content c , and a provider encoder Φ_p^n to learn provider-biased news representation \mathbf{p} from its provider ID p :

$$\mathbf{c} = \Phi_c^n(c), \quad \mathbf{p} = \Phi_p^n(p). \quad (2)$$

The content model can be implemented by news modeling techniques proposed in existing methods [29, 39, 43]¹ and the provider

¹We will introduce the architecture of the content model in Section 3.5.

Figure 3: The overall framework of *ProFairRec*.

encoder is implemented by the stack of a provider embedding layer and an MLP network. Based on this decomposition, the provider-biased news representations have the potential to encode provider bias in training data and provider-fair news representations have the potential to purely encode bias-independent information. Next, given a target user u , we apply the content model and the provider encoder to her clicked news to obtain the corresponding provider-fair news representations $[c_1^u, \dots, c_m^u]$ and provider-biased news representations $[p_1^u, \dots, p_m^u]$, where c_i^u and p_i^u denote provider-fair and -biased representations of the i -th clicked news d_i^u , respectively. Given a candidate news d^n , we also apply the content model and the provider encoder to learn its provider-fair news representation n^c and provider-biased news representation n^p .

Next, for user modeling, we also employ provider-biased representations to inherit provider-biased signals from training data to learn provider-fair representations to capture user interest. We employ a fair user model Φ_c^u to learn provider-fair user representation u^c from provider-fair representations of user's clicked news:

$$u^c = \Phi_c^u([c_1^u, \dots, c_i^u, \dots, c_m^u]), \quad (3)$$

where the fair user model can be implemented by user modeling methods proposed in existing works [1, 32, 44]². We also employ a

²We will introduce the framework of the user model in Section 3.5.

biased user model Φ_p^u to learn provider-biased user representation u^p from provider-biased representations of user's clicked news:

$$u^p = \Phi_p^u([p_1^u, \dots, p_i^u, \dots, p_m^u]), \quad (4)$$

where Φ_p^u is also implemented by existing user modeling methods. Based on this decomposition, provider-biased user representations can encode biased signals in training data and provider-fair user representations could purely encode bias-independent information.

Since training data usually contains provider bias signals, provider-fair representations may inherit bias information if we only employ them for model training. Thus, we use both provider-fair and -biased representations of users and news for training to fit the distribution of biased data. Provider-biased representations are used to inherit bias from training data and provider-fair representations are used to capture bias-independent information. In addition, to better prevent provider-fair news representations from encoding bias, we propose an adversarial learning task on news provider discrimination and apply it to provider-fair news representations. We also propose an orthogonal regularization and apply it to the provider-fair and -biased representations of both user and news to better eliminate bias in provider-fair representations. Thus, *ProFairRec* has the ability to learn provider-fair representations from that purely encode bias-independent information from biased data. Finally, we only use

provider-fair news and user representations for interest matching to achieve provider fairness-aware news recommendation.

3.3 Model Training

Next, we will introduce the training framework of *ProFairRec* in detail. Following existing works [1, 43], we also train news recommendation models on historical user click data. As mentioned above, user click behaviors on news providers are usually biased and training data may encode bias on news providers. News recommendation models trained on biased data may inherit provider bias and become unfair for minority providers. To tackle this challenge, during model training of *ProFairRec*, both provider-fair and provider-biased representations are used for click prediction to fit the distribution of biased training data. The provider-biased representations are used to inherit provider bias from training data and provider-fair representations are used to encode bias-independent information. Specifically, we first aggregate provider-fair and -biased representation of user and news to learn bias-aware user representation \mathbf{u}^a and bias-aware news representation \mathbf{n}^a : $\mathbf{u}^a = \mathbf{u}^c + \mathbf{u}^p$, $\mathbf{n}^a = \mathbf{n}^c + \mathbf{n}^p$. Next, we predict click score \hat{y}^c of the target user u on the candidate news d^n based on the matching of their bias-aware representations: $\hat{y}^c = \mathbf{u}^a \cdot \mathbf{n}^a$. Following existing works [39, 40], we apply the noise contrastive estimation (NCE) technique to formulate the recommendation loss \mathcal{L}^c . We treat each news d^+ clicked by the user as a positive sample, and randomly select F negative samples $[d_1^-, \dots, d_F^-]$ for each positive sample, where negative samples are non-clicked news in the same impression with the positive sample and d_j^- denotes the j -th negative sample. Next we calculate the click score \hat{y}_+^c of the positive sample d^+ and click scores $[\hat{y}_1^-, \dots, \hat{y}_F^-]$ of the negative samples. Finally, we obtain the recommendation loss \mathcal{L}^c based on the NCE loss:

$$\mathcal{L}^c = \mathbb{E}_{x \sim \Omega} \left[-\log \left(\frac{\exp(\hat{y}_+^c)}{\exp(\hat{y}_+^c) + \sum_{j=1}^F \exp(\hat{y}_j^-)} \right) \right], \quad (5)$$

where Ω is the training dataset, x is a training sample that is composed of a positive sample and F negative samples.

Although in *ProFairRec* we employ provider-biased representations to inherit provider bias in training data, provider-fair representations may also encode some bias information. To protect provider-fair representations from encoding provider bias, we propose an adversarial learning task to enforce provider-fair news representations to be indistinguishable to different news providers. We first employ a provider discriminator $\Psi(\cdot)$ to predict provider of a news d from its provider-fair representation \mathbf{c} , i.e., $\hat{\mathbf{y}}^p = \Psi(\mathbf{c})$, where $\hat{\mathbf{y}}^p \in \mathbb{R}^{|\mathcal{P}|}$ denotes the predicted probability vector and its j -th element is the predicted probability of the news d belonging to the j -th provider P_j . We use an MLP network to implement the provider discriminator $\Psi(\cdot)$:

$$\hat{\mathbf{y}}^p = \text{softmax}(\mathbf{W}_d \tilde{\mathbf{c}} + \mathbf{b}_d), \quad \tilde{\mathbf{c}} = \text{MLP}(\mathbf{c}), \quad (6)$$

where \mathbf{W}_d and \mathbf{b}_d are trainable parameters, and $\text{MLP}(\cdot)$ is an MLP network. Next, we formulate the provider discrimination loss \mathcal{L}^d and minimize \mathcal{L}^d to learn the optimal provider discriminator Ψ^* :

$$\Psi^* = \arg \min_{\Psi} \mathcal{L}^d, \quad \mathcal{L}^d = \mathbb{E}_{d \in \mathcal{D}} [-\mathbf{y}^p \cdot \log \hat{\mathbf{y}}^p], \quad (7)$$

where \mathbf{y}^p is the label of the provider of news d . After obtaining the optimal provider discriminator Ψ^* , we recalculate $\hat{\mathbf{y}}^p$ based on Ψ^* and further calculate the adversarial loss \mathcal{L}^a :

$$\mathcal{L}^a = \mathbb{E}_{d \in \mathcal{D}} [-\mathbf{y}^p \cdot \log \hat{\mathbf{y}}^p | \Psi^*]. \quad (8)$$

Based on this equation, we can calculate and employ the negative adversarial gradients on the recommendation model as a penalty to reduce provider bias encoded in provider-fair news representations.

Besides, to achieve a better decomposition of provider bias and bias-independent information, we also apply an orthogonal regularization to provider-fair and -biased representations of both news and users. For each news or user, orthogonal regularization constrains its provider-fair representation to be orthogonal with the corresponding provider-biased representation. We first calculate cosine similarities between provider-fair and -biased representations, and then calculate the orthogonal regularization:

$$\mathcal{L}^u = \mathbb{E}_{u \in \mathcal{U}} [|s^u|], \quad s^u = \frac{\mathbf{u}^c \cdot \mathbf{u}^p}{\|\mathbf{u}^c\|_2 \|\mathbf{u}^p\|_2}, \quad (9)$$

$$\mathcal{L}^n = \mathbb{E}_{d \in \mathcal{D}} [|s^n|], \quad s^n = \frac{\mathbf{n}^c \cdot \mathbf{n}^p}{\|\mathbf{n}^c\|_2 \|\mathbf{n}^p\|_2}, \quad (10)$$

where s^u is the cosine similarity between provider-fair and -biased representation of the user u , s^n is the cosine similarity between provider-fair and -biased representation of the news d , \mathcal{L}^u and \mathcal{L}^n denote the user and news orthogonal regularization, respectively.

Finally, we combine different loss functions into an overall optimization objective \mathcal{L} for model training:

$$\mathcal{L} = \lambda_c \mathcal{L}^c + \lambda_u \mathcal{L}^u + \lambda_n \mathcal{L}^n - \lambda_a \mathcal{L}^a, \quad (11)$$

where $\lambda_c, \lambda_u, \lambda_n, \lambda_a$ are hyper-parameters controlling the weights of loss $\mathcal{L}^c, \mathcal{L}^u, \mathcal{L}^n$ and \mathcal{L}^a , respectively.

3.4 Fair News Recommendation

Next, we will introduce the fair recommendation framework of *ProFairRec*. After training the recommendation model on user data, provider bias is enforced to be encoded in provider-biased representations and be reduced from provider-fair representations. Since provider-fair representations only encode bias-independent information, we can match provider-fair news and user representations to achieve fair news recommendation. Thus, given a target user u and a candidate news d^n , based on the content model and fair user model, we first learn provider-fair news representation \mathbf{n}^c for the candidate news d^n and provider-fair user representation \mathbf{u}^c for the target user u . We further match them to model user interest in candidate news without being affected by provider bias: $\hat{r}^c = \sigma(\mathbf{u}^c \cdot \mathbf{n}^c)$, where \hat{r}^c is the predicted unbiased click probability and $\sigma(\cdot)$ is the sigmoid activation function. Finally, given a target user u and a list of candidate news $[d_1^n, \dots, d_j^n, \dots, d_N^n]$, we can predict unbiased click probability for each candidate news and further rank these candidate news for recommendation.

3.5 Recommendation Model

As mentioned in Section 3.2, the content model and the user model of *ProFairRec* can be implemented by news modeling and user modeling techniques proposed in existing methods. Thus, *ProFairRec* is model-agnostic and can be combined with existing news recommendation models to improve their provider fairness. Next, we will

briefly introduce the architecture of mainstream news recommendation models. These models usually contain a content model to learn news representations from textual news content and a user model to learn user representations from users' reading history.

The content model contains three major components, i.e., a word embedding layer, a textual context modeling network, and a text pooling network. Given words in news content, they are first converted into a word embedding via the word embedding layer. Next, a textual context modeling network is applied to the word embedding sequence to capture relatedness among words and learn contextual representations for them. This text context modeling network can be implemented by various NLP techniques, such as CNN network [1, 14], LSTM network [9, 50], and transformer network [5, 35]. Finally, contextual word representations are aggregated by the text pooling network to obtain a news representation \mathbf{d} . The text pooling network can be implemented by various pooling operations, such as average pooling [38] and attention network [39].

The user model usually contains two major components, i.e., a behavior context modeling network and a behavior pooling network. User's clicked news are first converted to news embeddings via the content model. Then we apply a behavior context modeling network to understand user interest from user behavior contexts. It takes embeddings of user's clicked news as input and generates contextual news representations. This layer is usually implemented by some effective sequence modeling techniques, such as GRU network [1, 24] and multi-head self-attention network [43]. Next, we apply the behavior pooling network to build user interest representation \mathbf{u} . It is usually implemented by attention network [39, 43].

4 EXPERIMENT

4.1 Datasets and Experimental Settings

We conduct extensive experiments on a public news recommendation dataset [50] (*MIND*)³ to evaluate model performance and fairness. *MIND* is constructed by click data of 1 million users on the Microsoft News⁴ from October 12 to November 22, 2019 (6 weeks). User data in the first four weeks are used to construct click history, user data in the last week are used to construct the test set and other user data are used to construct the training and validation set. Besides, since *MIND* does not contain provider information, we extract them from the website via a web crawler. More details of the *MIND* dataset are in Table 1.

Table 1: Details of the *MIND* dataset.

# News	130,379	# Providers	1,705
# Users	1,000,000	# Impressions	4,979,946
# Clicks	7,583,733	# Non-clicks	183,124,199
Avg. # words in news titles		11.78	

Next, we will introduce hyper-parameter settings of *ProFairRec*. Following Wu et al. [40, 43], we utilize the first 30 words in news titles for news content modeling and users' recent 50 clicked news for user modeling. We select a famous news recommendation method,

³<https://msnews.github.io/index.html>.

⁴<https://www.msn.com/en-us/news/us>

Table 2: Comparisons of *ProFairRec* with baseline methods on recommendation performance.

	AUC	MRR	nDCG@10
GRU [24]	66.50±0.04	32.06±0.04	40.45±0.02
DKN [38]	66.70±0.15	32.41±0.11	40.88±0.12
HiFiArk [22]	67.49±0.19	33.04±0.15	41.57±0.15
NAML [39]	67.22±0.20	33.01±0.10	41.54±0.12
NPA [40]	67.13±0.07	32.90±0.07	41.45±0.07
KRED [19]	67.55±0.11	33.27±0.05	41.83±0.05
GNewsRec [10]	68.41±0.10	33.59±0.10	42.24±0.11
LSTUR [1]	68.35±0.10	33.48±0.10	42.16±0.09
NRMS [43]	68.08±0.13	33.43±0.09	42.07±0.10
OFAiR [33]	67.46±0.17	33.08±0.14	41.62±0.17
TFORM [52]	67.53±0.16	33.12±0.13	41.63±0.16
ProFairRec	67.64±0.10	33.08±0.05	41.67±0.07

i.e., *NRMS* [43], to implement the content and user model of *ProFairRec*. The content model is based on the stack of a word embedding layer, a multi-head self-attention (MHSA) network, and an attention pooling network. Specifically, the word embedding layer is initialized by pre-trained 300-dimensional Glove word embeddings [26] and fine-tuned in experiments. The MHSA network is set to contain 20 attention heads, and the output dimension of each head is set to 20. The attention network is implemented by a two-layer MLP network. The user model is based on the stack of an MHSA network and an attention pooling network. Similarly, the MHSA generates 400-dimensional output vectors, and the attention pooling network is a two-layer MLP network. Besides, for the provider encoder, the 400-dimensional provider embeddings are randomly initialized and fine-tuned in experiments. We randomly select 4 negative samples for each positive sample to formulate the recommendation loss \mathcal{L}^c and set its weight λ_c to 1. Both the weights λ_u and λ_n of news and user orthogonal regularization \mathcal{L}^u and \mathcal{L}^n are set to 1. In addition, since there are many providers in *MIND* which brings difficulties to the convergence of the adversarial learning, we retain 50 providers with the most clicks and merge other providers into a single class category for discrimination. Besides, following Goodfellow et al. [8], we iteratively train discrimination and adversarial task for a single step rather than finding the optimal discriminator in practice. The weight λ_a of the adversarial loss \mathcal{L}^a is set to 0.004. We utilize the Adam algorithm [12] with 0.0001 learning rate for model optimization. We select hyper-parameters based on the validation dataset. Codes are released at <https://github.com/taoqi98/ProFairRec>.

Following existing works [1, 5, 50], we use *AUC*, *nDCG@10* and *MRR* for recommendation evaluation. Besides, following existing works [3, 53], group-level fairness metrics are used for fairness evaluation and provider popularity is used for group partition. We evaluate the provider fairness based on the difference of exposure number of providers in a protected group and an unprotected group. Providers are first ranked based on their average click number, where unpopular providers ranked in the bottom $r\%$ are partitioned into the protected group \mathcal{P}^+ while others are partitioned into the unprotected group \mathcal{P}^- . Fairness is evaluated under different r : 10%, 30% and 50%. Following existing works [3, 53], we employ two

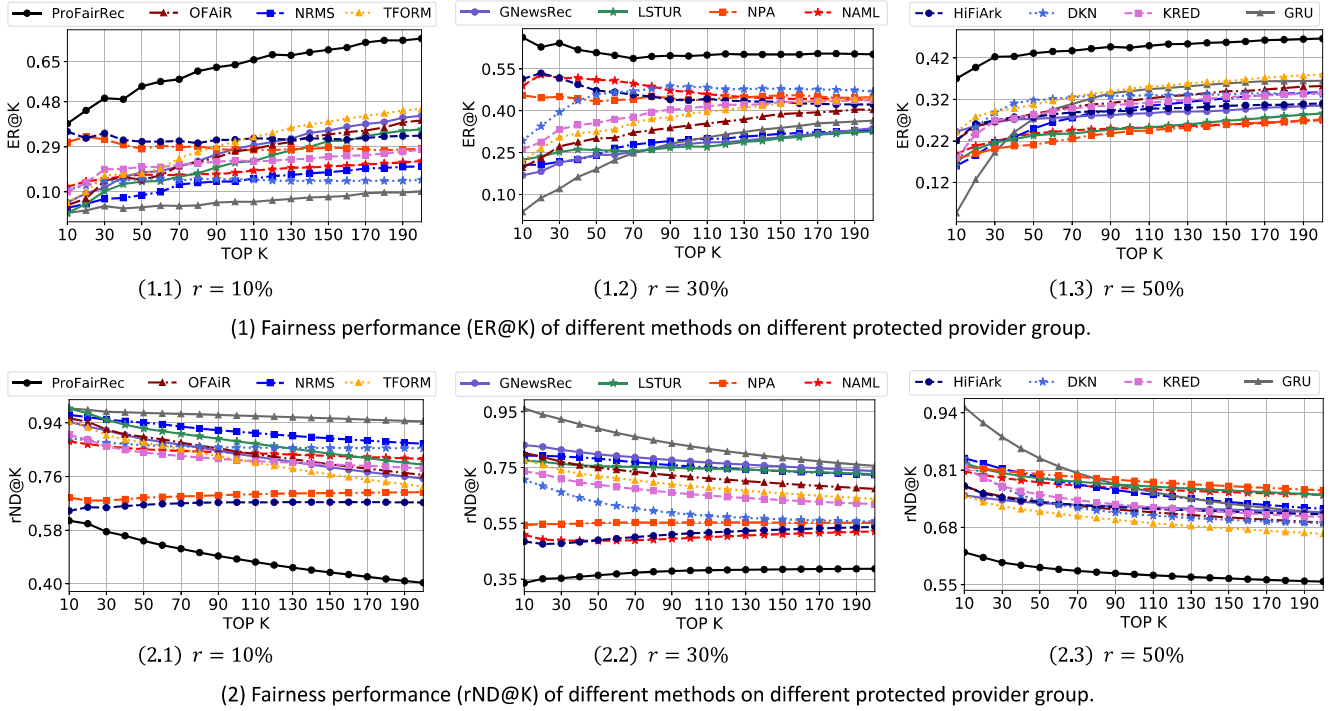


Figure 4: Provider fairness of *ProFairRec* and baseline methods based on different partition of protected provider group. Fairness is better if $ER@K$ is closer to 1 and $rND@K$ is closer to 0.

metrics to measure the exposure difference of providers in different groups based on the top K recommended news, i.e., relative exposure ratio $ER@K$ and normalized discounted difference $rND@K$:

$$ER@K = \frac{\mathbb{E}_{u \in \mathcal{U}}[|\mathcal{R}_K^u \cap \mathcal{D}^+|/|\mathcal{D}^+|]}{\mathbb{E}_{u \in \mathcal{U}}[|\mathcal{R}_K^u \cap \mathcal{D}^-|/|\mathcal{D}^-|]}, \quad (12)$$

$$rND@K = \frac{1}{Z} \mathbb{E}_{u \in \mathcal{U}} \left[\sum_{n=10,20,\dots}^K \frac{1}{\log_2 n} \left| \frac{|\mathcal{R}_n^u \cap \mathcal{D}^+|}{n} - \frac{|\mathcal{D}^+|}{|\mathcal{D}|} \right| \right], \quad (13)$$

$$\mathcal{D}^+ = \bigcup_{p \in \mathcal{P}^+} \mathcal{D}_p, \quad \mathcal{D}^- = \bigcup_{p \in \mathcal{P}^-} \mathcal{D}_p, \quad (14)$$

where \mathcal{R}_K^u is the top K news recommended to the user u and Z is the normalization factor. Eq. 12 and Eq. 13 show that the model is fairer if $ER@K$ is closer to 1 or $rND@K$ is closer to 0. Besides, we rank all news in the dataset to compare the fairness of different methods, since original news impression may be biased.

4.2 Performance Evaluation

In this section, we compare *ProFairRec* with baseline methods on both recommendation performance and provider fairness. The first group of baseline methods includes several mainstream news recommendation methods: (1) *GRU* [24]: using a text auto-encoder to learn news representations and a GRU network to learn user representations. (2) *DKN* [38]: proposing a knowledge-aware CNN network to learn news representations from news texts and entities. (3) *HiFi-Ark* [22]: proposing to learn multiple achieved user representations to model diverse user interest. (4) *NAML* [39]: learning news representations from news titles, topics, and

bodies via an attentive CNN network. (5) *NPA* [40]: proposing personalized attention networks to learn news representations from news titles and user representations by aggregating user’s clicked news. (6) *KRED* [19]: proposing a knowledge-aware graph attention network to learn news representations from news texts, entities, and neighbors of entities on knowledge graphs. (7) *GNewsRec* [10]: modeling long-term user interest from news-user graph via a GNN network and short-term user interest from users’ recent clicked news via a GRU network. (8) *LSTUR* [1]: learning user representations by combining short-term user interest inferred from reading history via a GRU network and long-term user interest modeled by user ID embeddings. (9) *NRMS* [43]: proposing to use multi-head self-attention networks to learn news and user representations.

The second group of baseline methods includes several provider fairness-aware recommendation methods. (1) *OFaiR* [33]: re-ranking items based on personalized relevance and provider unfairness of recommendation results. (2) *TFORM* [52]: dynamically re-balancing the exposure opportunities of providers based on their historical exposures. Note that these methods apply re-ranking techniques to improve the provider fairness in item recommendation and cannot be directly applied to the news recommendation task, thus we apply the core re-ranking techniques in them to re-rank recommendation results of *NRMS* as baseline provider fairness-aware news recommendation methods for comparisons. Besides, since the trade-off between recommendation performance and fairness of these two baseline methods and *ProFairRec* can be adjusted via

Table 3: Provider fairness of different news recommendation methods and their combination with *ProFairRec*. $r = 50\%$.

	ER@10	ER@30	ER@50	rND@10	rND@30	rND@50
NAML	0.1646±0.0080	0.2099±0.0107	0.2236±0.0102	0.8316±0.0082	0.8130±0.0077	0.8022±0.0077
+ProFairRec	0.2644±0.0329	0.2721±0.0229	0.2822±0.0255	0.7303±0.0332	0.7271±0.0261	0.7232±0.0230
KRED	0.1518±0.0399	0.2271±0.0349	0.2572±0.0321	0.8447±0.0407	0.8138±0.0379	0.7940±0.0356
+ProFairRec	0.2223±0.0777	0.2627±0.0649	0.2918±0.0596	0.7730±0.0788	0.7565±0.0734	0.7425±0.0699
NPA	0.1503±0.0358	0.1894±0.0203	0.2086±0.0220	0.8462±0.0364	0.8301±0.0298	0.8189±0.0276
+ProFairRec	0.2444±0.0382	0.2499±0.0361	0.2584±0.0332	0.7506±0.0388	0.7483±0.0373	0.7452±0.0358
NRMS	0.1237±0.0243	0.1991±0.0213	0.2378±0.0195	0.8734±0.0248	0.8423±0.0225	0.8202±0.0212
+ProFairRec	0.3644±0.0847	0.3941±0.0689	0.4156±0.0658	0.6295±0.0854	0.6174±0.0779	0.6071±0.0738
LSTUR	0.1583±0.0206	0.2086±0.0292	0.2254±0.0326	0.8381±0.0210	0.8174±0.0232	0.8050±0.0257
+ProFairRec	0.3765±0.1067	0.4491±0.0791	0.4865±0.0663	0.6174±0.1070	0.5879±0.0952	0.5668±0.0869

Table 4: Recommendation performance of baseline news recommendation methods and their combinations with our *ProFairRec* approach.

	AUC	MRR	nDCG@10
NAML	67.22±0.20	33.01±0.10	41.54±0.12
+ProFairRec	67.13±0.08	32.86±0.09	41.38±0.09
KRED	67.55±0.11	33.27±0.05	41.83±0.05
+ProFairRec	67.51±0.19	33.11±0.16	41.71±0.16
NPA	67.13±0.07	32.90±0.07	41.45±0.07
+ProFairRec	67.13±0.03	32.86±0.05	41.39±0.05
NRMS	68.08±0.13	33.43±0.09	42.07±0.10
+ProFairRec	67.64±0.10	33.08±0.05	41.67±0.07
LSTUR	68.35±0.10	33.48±0.10	42.16±0.09
+ProFairRec	67.46±0.11	32.83±0.13	41.35±0.13

hyper-parameters, we compare the provider fairness of them under similar recommendation performance for fair comparisons. We repeat each experiment five times and report average results.

Table 2 shows the recommendation performance of different methods and Fig. 4 shows the provider fairness of different methods. From these results, we have several findings. First, baseline news recommendation methods tend to recommend news of popular providers. This is because user click data on news providers are usually biased. Recommendation models trained on biased data may learn some shortcuts on news providers for recommendation and may be unfair for some minority news providers. Second, our *ProFairRec* method outperforms baseline news recommendation methods on provider fairness. This is because we propose to decompose provider bias and bias-independent information of user data into provider-biased and -fair representations and only employ provider-fair representations for news ranking. Besides, we also propose an adversarial learning task on provider discrimination and orthogonal regularization to achieve better bias decomposition. Third, our *ProFairRec* method effectively outperforms baseline fair recommendation methods on provider fairness when their recommendation performance are similar. This is because existing provider fairness-aware recommendation methods are usually based on re-ranking

techniques. These methods independently optimize model performance and fairness and may only achieve a sub-optimal trade-off between performance and fairness. Different from these methods, our *ProFairRec* method jointly optimizes recommendation performance and provider fairness based on adversarial decomposition technique, which can more effectively improve provider fairness of news recommendation models. Forth, compared with baseline news recommendation methods, our *ProFairRec* method can achieve comparable or even better recommendation performance. These results validate that our *ProFairRec* method can effectively improve recommendation fairness and meanwhile accurately recommend users' interested news as existing news recommendation methods.

4.3 Generalization of *ProFairRec*

As mentioned above, *ProFairRec* is a general framework and we can combine *ProFairRec* with models of existing news recommendation methods to improve their fairness. Thus, to verify the generalization of *ProFairRec*, we apply *ProFairRec* to several news recommendation models that achieve good recommendation performance in Table 2, i.e., *NAML*, *KRED*, *NPA*, *NRMS* and *LSTUR* and evaluate their fairness and performance. Specifically, we first employ the adversarial training framework of *ProFairRec* to learn fair user and news models for these methods. Then we employ the fair recommendation framework of *ProFairRec* for recommendation based on the fair user and news models in these methods. Results on provider fairness are summarized in Table 3 and results on recommendation performance are summarized in Table 4. Due to space limitations, we only show results under the protected provider group with $r = 50\%$. From these results, we find that our *ProFairRec* approach can effectively and consistently improve the provider fairness of existing news recommendation models with minor performance decline. This is because our *ProFairRec* approach employs provider-biased representations to inherit provider-bias signals from the biased training data and thereby can learn provider-fair representations via adversarial learning for both user and news to achieve provider fairness in news recommendation. In addition, our *ProFairRec* approach is a model-agnostics framework, which can be combined with existing news recommendation methods to improve their fairness. These results further validate the effectiveness and generalization of *ProFairRec* on improving provider fairness.




Figure 5: Ablation analysis of our *ProFairRec* approach.

4.4 Ablation Study

Next, we conduct an ablation study to evaluate the effectiveness of several important modules in *ProFairRec*, i.e., provider-biased representations, adversarial learning, and orthogonal regularization, on the provider fairness by removing them individually. Results are summarized in Fig. 5, from which we have three findings. First, after removing the provider-biased representations, provider fairness of *ProFairRec* significantly declines. This is because training data usually includes both provider bias and bias-independent signals, and it is difficult for a single news or user representation to purely encode bias-independent information without encoding provider bias. To tackle this challenge, *ProFairRec* decomposes provider bias and bias-independent information of biased data into provider-biased and -fair representations respectively, and only employ provider-fair representations for recommendation to improve provider fairness. Second, adversarial learning can also improve the provider fairness of *ProFairRec*. This is because provider-fair representations may also inherit some bias information from biased training data. The adversarial learning enforces provider-fair news representations to be indistinguishable with respect to different news providers, which can further prevent provider-fair news representations from encoding provider bias. Third, removing the orthogonal regularization also hurts the provider fairness of *ProFairRec*. This may be because the orthogonal regularization constrains the orthogonality of provider-biased and -fair representations, which can better decompose provider bias and bias-independent information.

4.5 Influence of Hyper-Parameters

In this section, we will explore how the adversarial learning in *ProFairRec* affects the recommendation performance and provider fairness. We analyze its influence by evaluating *ProFairRec* under different weights of the adversarial learning loss λ_a . Results are summarized in Fig. 6, from which we have two major findings. First, with the increase of λ_a , the recommendation fairness of *ProFairRec* consistently increases. This is intuitive since larger λ_a makes it




Figure 6: Influence of adversarial learning loss λ_a on recommendation performance and fairness.

more important for the model to reduce news provider information in provider-fair news representations, which can effectively improve the provider fairness. Second, large λ_a hurts the recommendation performance of *ProFairRec*. This is because large λ_a makes the recommendation model overly emphasize the adversarial learning task while the recommendation task cannot get enough respect. Thus, by adjusting λ_a , we can achieve a trade-off between recommendation performance and provider fairness.

5 CONCLUSION

In this paper, we propose a provider fairness-aware news recommendation framework (*ProFairRec*), which can learn fair news recommendation models from biased user data. The core of *ProFairRec* is to learn provider-fair representations for both news and users to encode bias-independent information for recommendation. Specifically, we propose to learn provider-fair and -biased representations for news from their content and provider IDs respectively. These representations are further aggregated to build provider-fair and -biased user representations based on user reading history. Both provider-fair and -biased representations are used for model training while only provide-fair representations are used for recommendation to improve provider fairness. In addition, we apply adversarial learning to provider-fair news representations to prevent them from encoding bias. We also apply orthogonal regularization to provider-fair and -biased representations to better reduce bias in provider-fair representations. Besides, *ProFairRec* is a model-agnostic framework and can be applied to existing news recommendation models to improve their fairness. Experiments on a public dataset demonstrate that *ProFairRec* can effectively improve the fairness of mainstream news recommendation methods with minor performance declines. In our future work, we plan to apply *ProFairRec* to improve provider fairness in more recommendation scenarios (e.g. item recommendation).

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China under Grant numbers 2021ZD0113902, U1936208, U1936126 and Tsinghua-Toyota Joint Research Fund 20213930033.

REFERENCES

- [1] Mingxiao An, Fangzhao Wu, Chuhan Wu, Kun Zhang, Zheng Liu, and Xing Xie. 2019. Neural news recommendation with long-and short-term user representations. In *ACL*. 336–345.
- [2] Trapit Bansal, Mrinal Das, and Chiranjib Bhattacharyya. 2015. Content driven user profiling for comment-worthy recommendations of news and blog articles. In *RecSys*. 195–202.
- [3] Robin Burke, Nasim Sonboli, and Aldo Ordóñez-Gauger. 2018. Balanced neighborhoods for multi-sided fairness in recommendation. In *ACM FAccT*. 202–214.
- [4] Abhinandan S Das, Mayur Datar, Ashutosh Garg, and Shyam Rajaram. 2007. Google news personalization: scalable online collaborative filtering. In *WWW*. 271–280.
- [5] Suyu Ge, Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020. Graph enhanced representation learning for news recommendation. In *WWW*. 2863–2869.
- [6] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, et al. 2021. Towards Long-term Fairness in Recommendation. In *WSDM*. 445–453.
- [7] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *KDD*. 2221–2231.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* (2020), 139–144.
- [9] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* (1997), 1735–1780.
- [10] Linmei Hu, Chen Li, Chuan Shi, Cheng Yang, and Chao Shao. 2020. Graph neural news recommendation with long-term and short-term interest modeling. *IP&M* (2020), 102142.
- [11] Dhruv Khattar, Vaibhav Kumar, Vasudeva Varma, and Manish Gupta. 2018. Weave&Rec: A word embedding based 3-D Convolutional network for news recommendation. In *CIKM*. 1855–1858.
- [12] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [13] Michal Kompan and Mária Bielíková. 2010. Content-based news recommendation. In *EC-Web*. 61–72.
- [14] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* (1998), 2278–2324.
- [15] Yunqi Li, Hanxiong Chen, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2021. User-oriented Fairness in Recommendation. In *WWW*. 624–632.
- [16] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. 2021. Towards Personalized Fairness based on Causal Notion. In *SIGIR*. 1054–1063.
- [17] Jianxun Lian, Fuzheng Zhang, Xing Xie, and Guangzhong Sun. 2018. Towards better representation learning for personalized news recommendation: a multi-channel deep fusion approach. In *IJCAI*. 3805–3811.
- [18] Chen Lin, Runquan Xie, Xinjun Guan, Lei Li, and Tao Li. 2014. Personalized news recommendation via implicit social experts. *Information Sciences* (2014), 1–18.
- [19] Danyang Liu, Jianxun Lian, Shiyin Wang, Ying Qiao, Jiun-Hung Chen, Guangzhong Sun, and Xing Xie. 2020. KRED: Knowledge-aware document representation for news recommendations. In *RecSys*. 200–209.
- [20] Jiahui Liu, Peter Dolan, and Elin Rønby Pedersen. 2010. Personalized news recommendation based on click behavior. In *IUI*. 31–40.
- [21] Weiwen Liu, Jun Guo, Nasim Sonboli, Robin Burke, and Shengyu Zhang. 2019. Personalized fairness-aware re-ranking for microlending. In *RecSys*. 467–471.
- [22] Zheng Liu, Yu Xing, Fangzhao Wu, Mingxiao An, and Xing Xie. 2019. Hi-Fi ark: deep user representation via high-fidelity archive network. In *IJCAI*. 3059–3065.
- [23] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling fairness and bias in dynamic learning-to-rank. In *SIGIR*. 429–438.
- [24] Shumpei Okura, Yukihiko Tagami, Shingo Ono, and Akira Tajima. 2017. Embedding-based news recommendation for millions of users. In *KDD*. 1933–1942.
- [25] 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 *WWW*. 1194–1204.
- [26] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*. 1532–1543.
- [27] Evaggelia Pitoura, Kostas Stefanidis, and Georgia Koutrika. 2021. Fairness in rankings and recommendations: An overview. *The VLDB Journal* (2021), 1–28.
- [28] Tao Qi, Fangzhao Wu, Chuhan Wu, and Yongfeng Huang. 2021. Personalized News Recommendation with Knowledge-aware Interactive Matching. In *SIGIR*. 61–70.
- [29] Tao Qi, Fangzhao Wu, Chuhan Wu, and Yongfeng Huang. 2021. PP-Rec: News Recommendation with Personalized User Interest and Time-aware News Popularity. In *ACL*. 5457–5467.
- [30] Tao Qi, Fangzhao Wu, Chuhan Wu, Yongfeng Huang, and Xing Xie. 2020. Privacy-Preserving News Recommendation Model Learning. In *EMNLP: Findings*. 1423–1432.
- [31] Tao Qi, Fangzhao Wu, Chuhan Wu, Yongfeng Huang, and Xing Xie. 2021. UniFedRec: A Unified Privacy-Preserving News Recommendation Framework for Model Training and Online Serving. In *EMNLP: Findings*. 1438–1448.
- [32] Tao Qi, Fangzhao Wu, Chuhan Wu, Peiru Yang, Yang Yu, Xing Xie, and Yongfeng Huang. 2021. HieRec: Hierarchical User Interest Modeling for Personalized News Recommendation. In *ACL*. 5446–5456.
- [33] Nasim Sonboli, Farzad Eskandarian, Robin Burke, Weiwen Liu, and Bamshad Mobasher. 2020. Opportunistic Multi-aspect Fairness through Personalized Re-ranking. In *UMAP*. 239–247.
- [34] Daniel Trielli and Nicholas Diakopoulos. 2019. Search as news curator: The role of Google in shaping attention to news information. In *CHI*. 1–15.
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*. 6000–6010.
- [36] Chong Wang and David M Blei. 2011. Collaborative topic modeling for recommending scientific articles. In *KDD*. 448–456.
- [37] Heyuan Wang, Fangzhao Wu, Zheng Liu, and Xing Xie. 2020. Fine-grained interest matching for neural news recommendation. In *ACL*. 836–845.
- [38] Hongwei Wang, Fuzheng Zhang, Xing Xie, and Minyi Guo. 2018. DKN: Deep knowledge-aware network for news recommendation. In *WWW*. 1835–1844.
- [39] Chuhan Wu, Fangzhao Wu, Mingxiao An, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019. Neural news recommendation with attentive multi-view learning. In *IJCAI*. 3863–3869.
- [40] Chuhan Wu, Fangzhao Wu, Mingxiao An, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019. Npa: Neural news recommendation with personalized attention. In *KDD*. 2576–2584.
- [41] Chuhan Wu, Fangzhao Wu, Mingxiao An, Yongfeng Huang, and Xing Xie. 2019. Neural News Recommendation with Topic-Aware News Representation. In *ACL*. 1154–1159.
- [42] Chuhan Wu, Fangzhao Wu, Mingxiao An, Tao Qi, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019. Neural news recommendation with heterogeneous user behavior. In *EMNLP*. 4876–4885.
- [43] Chuhan Wu, Fangzhao Wu, Suyu Ge, Tao Qi, Yongfeng Huang, and Xing Xie. 2019. Neural news recommendation with multi-head self-attention. In *EMNLP*. 6390–6395.
- [44] Chuhan Wu, Fangzhao Wu, Yongfeng Huang, and Xing Xie. 2021. User-as-Graph: User Modeling with Heterogeneous Graph Pooling for News Recommendation. In *IJCAI*. 1624–1630.
- [45] Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020. User modeling with click preference and reading satisfaction for news recommendation. In *IJCAI*. 3023–3029.
- [46] Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2021. Empowering News Recommendation with Pre-trained Language Models. In *SIGIR*. 1652–1656.
- [47] Chuhan Wu, Fangzhao Wu, Tao Qi, Jianxun Lian, Yongfeng Huang, and Xing Xie. 2020. PTUM: Pre-training User Model from Unlabeled User Behaviors via Self-supervision. In *EMNLP: Findings*. 1939–1944.
- [48] Chuhan Wu, Fangzhao Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. 2021. Fairness-aware News Recommendation with Decomposed Adversarial Learning. In *AAAI*. 4462–4469.
- [49] Chuhan Wu, Fangzhao Wu, Yang Yu, Tao Qi, Yongfeng Huang, and Qi Liu. 2021. NewsBERT: Distilling Pre-trained Language Model for Intelligent News Application. In *EMNLP: Findings*. 3285–3295.
- [50] Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, et al. 2020. MIND: A large-scale dataset for news recommendation. In *ACL*. 3597–3606.
- [51] Le Wu, Lei Chen, Pengyang Shao, Richang Hong, Xiting Wang, and Meng Wang. 2021. Learning Fair Representations for Recommendation: A Graph-based Perspective. In *WWW*. 2198–2208.
- [52] Yao Wu, Jian Cao, Guandong Xu, and Yudong Tan. 2021. TFROM: A Two-sided Fairness-Aware Recommendation Model for Both Customers and Providers. In *SIGIR*. 1013–1022.
- [53] Ke Yang and Julia Stoyanovich. 2017. Measuring fairness in ranked outputs. In *SSDBM*. 1–6.
- [54] Sirui Yao and Bert Huang. 2017. Beyond Parity: Fairness Objectives for Collaborative Filtering. In *NIPS*. 2921–2930.
- [55] Jingwei Yi, Fangzhao Wu, Chuhan Wu, Ruixuan Liu, Guangzhong Sun, and Xing Xie. 2021. Efficient-FedRec: Efficient Federated Learning Framework for Privacy-Preserving News Recommendation. In *EMNLP*. 2814–2824.
- [56] Qi Zhang, Jingjie Li, Qinglin Jia, Chuyuan Wang, Jieming Zhu, Zhaowei Wang, and Xiuqiang He. 2021. UNBERT: User-News Matching BERT for News Recommendation. In *IJCAI*. 3356–3362.
- [57] Guanjie Zheng, Fuzheng Zhang, Zihan Zheng, Yang Xiang, Nicholas Jing Yuan, Xing Xie, and Zhenhui Li. 2018. DRN: A deep reinforcement learning framework for news recommendation. In *WWW*. 167–176.
- [58] Qianman Zhu, Xiaofei Zhou, Zeliang Song, Jianlong Tan, and Guo Li. 2019. DAN: Deep attention neural network for news recommendation. In *AAAI*. 5973–5980.