Invariant Debiasing Learning for Recommendation via Biased Imputation*

Ting Bai, Weijie Chen, Cheng Yang and Chuan Shi*

Beijing University of Posts and Telecommunications, Xitu Cheng Road, Beijing, 100876, China

ARTICLE INFO

Keywords: Recommender Systems Debiasing Learning Knowledge Distillation

ABSTRACT

Previous debiasing studies utilize unbiased data to make supervision of model training. They suffer from the high trial risks and experimental costs to obtain unbiased data. Recent research attempts to use invariant learning to detach the invariant preference of users for unbiased recommendations in an unsupervised way. However, it faces the drawbacks of low model accuracy and unstable prediction performance due to the losing cooperation with variant preference. In this paper, we experimentally demonstrate that invariant learning causes information loss by directly discarding the variant information, which reduces the generalization ability and results in the degradation of model performance in unbiased recommendations. Based on this consideration, we propose a novel lightweight knowledge distillation framework (KD-Debias) to automatically learn the unbiased preference of users from both invariant and variant information. Specifically, the variant information is imputed to the invariant user preference in the distanceaware knowledge distillation process. Extensive experiments on three public datasets, i.e., Yahoo!R3, Coat, and MIND, show that with the biased imputation from the variant preference of users, our proposed method achieves significant improvements with less than 50% learning parameters compared to the SOTA unsupervised debiasing model in recommender systems. Our code is publicly available at https://github.com/BAI-LAB/KD-Debias.

1. Introduction

As recommender systems have achieved great success in providing personalized service, numerous studies focus on improving the accuracy of model performance (Luo, Zou, Ai, Chen, Li, Yin, and Davison, 2024; Wei, Ren, Tang, Wang, Su, Cheng, Wang, Yin, and Huang, 2024; Song, Huang, Lian, and Jin, 2024; Hua, Ge, Xu, Ji, Li, and Zhang, 2024). Recently, increasing attention has been paid to the long-term impacts on user experience and retention in commercial platforms, especially in reducing the risks brought by multiple kinds of bias, including the exposure bias (Yang, Cui, Xuan, Wang, Belongie, and Estrin, 2018), position bias (Collins, Tkaczyk, Aizawa, and Beel, 2018), popularity bias (Chen, Xiao, Li, Ye, Sun, and Deng, 2020; Krishnan, Sharma, Sankar, and Sundaram, 2018) and so on. For example, popularity bias refers to popular items being overly recommended, which is at the expense of damages in recommendation ecology due to the under-recommendation of less popular items that users may be interested in. The studies of debiasing in recommender systems can be generally classified into two basic routes: one is to address a specific bias, like popularity bias (Chen et al., 2020; Krishnan et al., 2018) or position bias (Collins et al., 2018), the other is general debiasing with various biases (Li, Wu, Zheng, Xiao, Wang, Geng, Feng, He, and Wu, 2024; Li, Wang, Wang, and Xu, 2023; Zhang, Ma, Wei, Sheng, and Wang, 2024). In practice, in real recommendation scenarios, various types of biases are often mixed up, for example, a user's click behavior may be influenced by both popularity bias and exposure bias.

To deal with the mixed-up biases, general debiasing methods (Li et al., 2024, 2023; Zhang et al., 2024) are proposed, and they usually require the supervision of unbiased data in the training process. However, the acquisition of such unbiased data in real recommender systems often entails considerable trial risks and human costs. In response, recent studies attempt to apply invariant learning (Creager, Jacobsen, and Zemel, 2021; Du, Wu, Feng, He, and Tang, 2022) to detach the invariant preference of users for unbiased recommendations in an unsupervised way (i.e., training without unbiased data). The debiasing method (Wang, He, Liu, Zou, Yu, and Cui, 2022; Tang, Zou, Cui, Liang, and Jin, 2024) via invariant learning assumes that the observational user behaviors are determined by invariant preference (i.e.

^{*} This work was supported by the National Key Research and Development Program of China (No.2023YFC3303800). ^{*} Corresponding author

ORCID(s):



Figure 1: The drawbacks (i.e., low model accuracy and unstable prediction ability) of the invariant learning method InvPref which solely relying on invariant information to make unbiased recommendations. (a) The model performances of our model KD-Debias and InvPref on the Yahoo!R3 dataset with random seeds in the 20 repeated experiments, the horizontal axis represents different experimental groups; (b) The distributions of invariant and variant preference in InvPref; (c) The information entropies of invariant and variant preference; (d) The distance between the unbiased true label and the predicted label from invariant and variant preference: the smaller the distance is, the more accurate the prediction is. The cases in which the prediction from variant preference is more accurate to the unbiased true label are circled with the dotted line, implying their potential utility to enhance the model performance, the horizontal axis refers to different data samples.

a user's true preference) and the variant preference (affected by some unobserved confounders)¹. They believe that the unobserved confounders may cause various biases, and only the invariant preference should be used for unbiased recommendations. Conducted in an unsupervised way, such a debiasing model (Wang et al., 2022) avoids the high risks in collecting unbiased data.

However, we find it faces the drawbacks of low model accuracy and unstable prediction ability. Fig.1 (a) illustrates the utilization of solely the invariant preference for generating unbiased recommendations across various random initialized seeds. Notably, the performance of the invariant learning model (Wang et al., 2022) (marked in blue line) exhibits marked instability, indicative of its limited generalization capability. To figure out the reasons, we characterize the distributions of invariant preference and the variant preference via t-SNE algorithm (Van der Maaten and Hinton, 2008), as well as their information entropy in Fig.1 (b-c). We find out that the distribution of the invariant and variant preference are distinct different (subfigure b), and both of them contain abundant information (subfigure c). Besides, to verify whether the variant preferences are useless for debiasing learning, we evaluate their capability to make unbiased recommendations, as shown in Fig.1 (d): for a certain number of cases (circled with the dotted line), the prediction from variant preference is more accurate to the true unbiased label. This indicates that solely relying on invariant information to make unbiased recommendations will result in information loss, which degrades model performance and reduces its generalization ability. The presence of abundant informations. Based on the above consideration, our aim is to take full advantage of the variant information to enhance unbiased recommendations.

To tackle this challenge, we propose a novel distance-aware knowledge distillation framework (**KD-Debias**) to adaptively fuse the invariant and variant information for general debiasing. The knowledge distillation process provides a way to distill useful information from the variant information, which is automatically imputed to the unbiased data distribution following the rules: comparing to the distribution of invariant information, the more different the distribution of the variant information is (i.e., greater distance), the more attention will be paid to the variant preference in the knowledge distillation process. Then a simple lightweight matrix factorization (MF) based student model is designed to capture the collaborative signals. With the soft predicted label from the teacher model, the MF based student model not only improves the model efficiency by reducing the number of parameters, but also enhance its generalization ability for the final unbiased predictions. Our contributions are as follows:

- We use invariant learning to conduct general debiasing with biased observational data only and demonstrate the necessity to incorporate the biased imputation from variant information in unbiased recommendations.
- We propose a lightweight distance-aware knowledge distillation framework, termed as KD-Debias, to automatically fuse the invariant and variant information, the distillated student model achieves the best model

¹In the debiasing model (Wang et al., 2022): the invariant preference refers to the unbiased preference, and the variant preference refers to the biased preference affected by unobserved confounders.

performance with less than 50% learning parameters compared to the SOTA unsupervised debiasing method (i.e., InvPref).

• We conduct extensive experiments on three public datasets, i.e., Yahoo!R3, Coat, and Mind, to show that with the imputation from biased information, our method achieves the best model performance in unbiased recommendations.

2. Related Work

In this section, we introduce the concepts related to our study briefly, including debiasing learning, invariant learning, and knowledge distillation methods in recommender systems.

2.1. Debiasing in recommender systems

There are various kinds of biases in recommender systems, such as selection bias (Li et al., 2023; Saito, 2020; Li, Zheng, Zhou, and Wu, 2022), exposure bias (Mansoury, 2022; Ning, Sangineto, Porrello, Calderara, and Cucchiara, 2023), popularity bias (Zhao, Chen, Zhou, He, Cao, Zhang, and Wu, 2022; Chen, Dong, Wang, Feng, Wang, and He, 2023), and other undefined biases caused by latent confounders (Nam, Cha, Ahn, Lee, and Shin, 2020; Wang, Shen, Wang, Chen, Chen, and Wen, 2022; Zhu, Zhang, Yang, Wang, and Feng, 2024). The biases would hurt the experience of users and the revenues of platforms. As increasing attention has been paid to debiasing in recommender systems, more and more debiasing methods have been proposed (Zhao, Cui, Li, Bao, Li, Zhou, Liu, and Feng, 2023; Wei, Feng, Chen, Wu, Yi, and He, 2021; Qiu, Wang, Chen, Yin, and Huang, 2021). The studies of debiasing in recommender systems can be generally classified into two basic routes: one is to address a specific bias, like the exposure bias (Du, Yuan, Huang, Zhao, and Zhou, 2023; Yang et al., 2018), position bias (Collins et al., 2018; Wu, Qiu, Zheng, Zhu, and Chen, 2024), popularity bias (Chen et al., 2020; Krishnan et al., 2018) and so on. For example, KDRec (Yang, Huang, Xia, Huang, Luo, and Lin, 2023) introduces contrastive learning to capture item transition patterns and address popularity bias. Some researchers design inverse propensity score estimators to eliminate biases across different domains (Li, Yao, Mu, Zhao, Li, Guo, Ding, and Wen, 2021). The other is general debiasing with various biases mixed up (Guo, Cai, Zheng, Zhang, and Chen, 2024; Zhang et al., 2024; Yang, Li, Liu, Wang, Lu, and Liu, 2024). For example, the IPS-based debiasing methods (Schnabel, Swaminathan, Singh, Chandak, and Joachims, 2016; Zhu, He, Zhang, and Caverlee, 2020) address biases by re-weighting the samples to make the biased distribution of training data approximate the distribution of unbiased data, iDCF (Zhang, Zhang, Liu, Wang, Gao, Zhang, and Guo, 2023) applies proximal causal inference to infer the unmeasured confounders and identify the counterfactual feedback. DIB (Liu, Cheng, Zhu, Dong, He, Pan, and Ming, 2023) alleviates the confounding bias via information bottleneck. AdvDrop (Zhang et al., 2024) employs adversarial learning to split the neighborhood into bias-mitigated interactions and bias-aware interactions for graph-based recommendations. Despite their effectiveness, most of these methods train models under the supervision of unbiased data which is usually generated with high experimental costs, besides, the imbalance in the quantity of unbiased and normal data also raises concerns about overfitting issues.

Recently, some researchers have solved the debiasing issue in unsupervised ways (Wang, Chen, Wen, Huang, Kuruoglu, and Zheng, 2020; Wang et al., 2022), i.e., training without unbiased data. For example, CVIB (Wang et al., 2020) builds a counterfactual variational information bottleneck, mitigating bias by balancing learning between the factual and counterfactual domains; InvPref (Wang et al., 2022) captures the biases through estimating pseudo environments labels and generate unbiased user preferences with invariant learning. Despite their ability to eliminate the dependence on unbiased data, these methods still face the drawbacks of low and unstable model performance according to our findings in Section 1.

2.2. Invariant Learning

Invariant learning (Arjovsky, Bottou, Gulrajani, and Lopez-Paz, 2019; Creager et al., 2021) assumes that the observed data originate from multiple different environments, which caused the heterogeneity of the distribution of observed data. Relying on the principle of invariance inherent in causality, invariant learning captures representations with invariant predictive ability across environments. By capitalizing on the invariant relationship that exists between features and labels across diverse distributions, invariant learning can ensure generalization and disregard pseudo-correlations even when there are deviations in distribution. Drawing from these strengths, the concept of invariant learning has been extensively applied in recommendation tasks, extending its utility into the domain of debiasing recommendations as well. For example, InvPref (Wang et al., 2022) attributes the heterogeneity of environments



Figure 2: An Overview of KD-Debias framework. It contains two components: the Disentangled Preference Modeling component and the Distance-aware Knowledge Distillation component. The disentangled preference modeling is conducted on the left causal graph with an environment classifier to identify the biases from different environments. Then the knowledge distillation are conducted on the disentangled invariant preference and variant preference with a distance-aware fusion strategy.

will cause biases. They assume that the observational user behaviors are determined by invariant preference (i.e. a user's true preference) and the variant preference (affected by some unobserved environment confounders). They capture the invariant user preferences across environments for unbiased recommendations. Compared with the existing debiasing methods with invariant learning, we not only focus on the generalizability of invariant preference but also the information loss exhibited by invariant learning when it comes to personalized recommendations.

2.3. Knowledge Distillation

The knowledge distillation framework (Hinton, Vinyals, and Dean, 2015) can distill the knowledge from teacher models (an ensemble of cumbersome models) to the student model (a small model) by using the "soft targets or label" produced by the teacher model for training the student model. In addition to the information of positive samples, aligning with "soft targets" could also carry a lot of information from negative ones, which improves the generalization ability of the student model. Benefiting from these strengths, many researchers employ knowledge distillation into recommendation systems to lightweight the model (Gou, Yu, Maybank, and Tao, 2021; Liu, Cheng, Dong, He, Pan, and Ming, 2020) and improve the generalization capabilities (Kang, Hwang, Kweon, and Yu, 2020; Tang and Wang, 2018; Zhu, Si, Wang, Yang, and Xiao, 2022) by ensemble models at a fine-grained level. Recently, some studies (Ding, Feng, He, Jin, Wang, Liao, and Zhang, 2022; Liu, Cheng, Lin, Luo, Dong, He, Pan, and Ming, 2022; Chen, Chen, Feng, Zhou, and He, 2023) attempt to use knowledge distillation to improve the debiasing ability in recommender systems. KDCRec (Liu et al., 2022) proposes a flexible knowledge distillation method with four types of distillation strategies from the teacher model to the student model. InterD (Ding et al., 2022) achieves competitive performance in both the norm test and debiasing test by distilling knowledge from a biased model and an unbiased model. However, both of them are trained under the supervision of unbiased data, which brings the high experimental costs and concerns of the overfitting problem. Different from the existing knowledge distillation approaches on debiasing learning, we aim to achieve unbiased recommendations without the participation of unbiased data.

3. Our Proposed Method

In order to make unbiased recommendations with the biased observational data alone, we use invariant learning to obtain the invariant and variant preferences of users in an unsupervised way. To fully make utilization of variant information, we propose a distance-aware knowledge distillation debiasing framework (KD-Debias) to adaptively fuse the invariant and variant information for debiasing in recommender systems. KD-Debias consists of two components: i.e., the disentangled preference modeling component and the distance-aware knowledge distillation component. The illustration of our proposed framework is shown in Fig. 2.

3.1. Disentangled Preference Modeling

To learn the unbiased preference of users, we disentangle user preference into user invariant preference and variant preference. Following the method in invariant learning (Wang et al., 2022), which assumes that the variant information is influenced by different latent confounders from environments and is regarded as biased information. Specifically, the disentangled preference modeling is conducted on a causal graph (see in the left part of Fig. 2) that describes the generation of observational user behaviors. Let U and I represent the set of users and items respectively. The preference of users U towards items I can be disentangled into invariant preference M and variant preference N. Different from the invariant preference M, the variant preference N is also directly influenced by the latent confounders in the environment E. The observed user feedback Y is inevitably biased since it is affected by both invariant and variant preference.

3.1.1. Basic Summarization.

To disentangle the invariant and variant preferences, invariant debiasing learning is conducted by an adversarial network. It designs an environment classifier to identify the biases from different environments. The invariant preference tries to confuse the environment classifier, while the classifier aims to identify the environment label of the invariant preference. A well-trained classifier can easily predict the environment label of the biases mixed in invariant preference. A well-optimized invariant user preference contains no bias that increases the difficulty of the classifier in identifying which environment the bias comes from. Through the joint optimizations of the invariant preference and environment classifier, we can obtain the disentangled invariant and variant preferences of users. Specifically, invariant learning can be divided into two steps, i.e., the invariant and variant preference modeling and the adversarial learning with the environment classifier.

3.1.2. The Invariant and Variant Preference Modeling.

For observed data user u and item i, we use \mathbf{u}^{inv} , \mathbf{i}^{var} , \mathbf{i}^{var} to denote the invariant embedding and variant embedding of user u and item i. \mathbf{q}_e is the corresponding environment embedding. The invariant preference $\mathbf{m}_{u,i}$ and variant preference $\mathbf{n}_{u,i}$ of the user toward item i can be computed as:

$$\mathbf{m}_{u,i} = \mathbf{u}^{inv} \odot \mathbf{i}^{inv}, \tag{1}$$
$$\mathbf{n}_{u,i} = \mathbf{u}^{var} \odot \mathbf{i}^{var} \odot \mathbf{q}_{e}, \tag{2}$$

where \odot is the Hadamard product of two embeddings.

In order to enhance the recommendation abilities of invariant preference, we use the observed feedback label *y* as the supervision information, and the invariant representations are optimized by:

$$L_{inv} = \frac{1}{|D|} \sum_{(u,i) \in D} l_{rec}(\varphi(\mathbf{m}_{u,i}), y_{u,i}),$$
(3)

where $y_{u,i}$ is the user feedback in observed data, $l_{rec}(\cdot)$ is the binary cross entropy loss and $\varphi(\cdot)$ is a predictive function as follow:

$$\varphi(\mathbf{x}) = Sigmoid(\mathbf{x}^T \mathbf{1}_L),\tag{4}$$

where $\mathbf{1}_L$ is an all-ones vector with the same dimension as \mathbf{x} .

3.1.3. The Adversarial Learning with the Environment Classifier.

The invariant preference of a user towards an item $\mathbf{m}_{u,i}$ tries to confuse the environment classifier ϕ_c , while the classifier tries to identify the environment label of this pair according to its invariant preference. The ground truth environment label $e_{u,i}$ can be assigned according to the prediction loss of the observed user feedback (see in Eq.6), since the smaller difference between observed biased feedback and prediction is, the more possibility the pair is from this environment. The optimization of the environment classifier ϕ_c can be defined as:

$$L_{env} = \frac{1}{|D|} \sum_{(u,i)\in D} l_c(\phi_c(\mathbf{m}_{u,i}), e_{u,i}),$$
(5)

where ϕ_c is a single-layer MLP network whose output dimensions are equal to the number of environments, and l_c is the cross entropy loss for multiple classifications.

The environment label $e_{u,i}$ for user-item pair (u, i) is reallocated according to the prediction loss L_{var} of the observed feedback label, formulated as:

$$e_{u,i} \leftarrow \underset{e \in \mathcal{E}}{\operatorname{arg\,min}}(L_{var}), \tag{6}$$

$$L_{var} = l_{rec}(f(\varphi(\mathbf{n}_{u,i}), \varphi(\mathbf{m}_{u,i})), y_{u,i}),$$
(7)

where $y_{u,i}$ is the biased label in observed data, φ is the predictive function (defined in Eq.4), $f(\cdot, \cdot)$ is a product function to fuse the invariant and variant predictions, and $l_{rec}(\cdot)$ is the binary cross entropy loss.

3.1.4. Overall Objective Functions.

The goal of invariant preference $\mathbf{m}_{u,i} = \mathbf{u}^{inv} \odot \mathbf{i}^{inv}$ is to confuse the environment classifier ϕ_c , i.e., maximize L_{env} , while the aim of the classifier is to identify the label of the environment, i.e., minimize the parameters θ_c in L_{env} . The adversarial process can be formulated as:

$$\min_{\theta_c} \max_{\mathbf{u}^{inv}, \mathbf{i}^{inv}} L_{env}.$$
(8)

The overall loss function of disentangled preference modeling can be formulated as:

$$L_{major} = L_{inv} - \alpha \cdot L_{env} + \beta \cdot L_{var},\tag{9}$$

where α , β are the hyper-parameters to balance the proportions of the variant and invariant preferences.

3.2. Distance-aware Unbiased Knowledge Distillation

Utilizing the Disentangled Preference Modeling module, we systematically decompose user preferences into two distinct categories: invariant and variant preferences. To achieve unbiased recommendations, we propose a distance-aware knowledge distillation approach. By conducting the distance-aware knowledge distillation, we distill useful information from both invariant and variant preference into a lightweight student model, which improves the model efficiency by reducing the number of parameters. Besides, the use of soft predicted labels from the teacher model further enhances the model's generalization capability.

3.2.1. Basic Summarization.

Invariant preferences represent stable aspects of users' tastes unaffected by external factors, reflecting enduring interests or values. In contrast, variant preferences fluctuate in response to external influences such as time and context. By adaptively integrating these two preferences, recommender systems can better understand and anticipate user needs, ultimately facilitating unbiased recommendations. Our method is grounded in a fundamental principle: when there is a significant difference in the invariant and variant preferences, it indicates that the user's current needs diverge markedly from their long-term preferences. By incorporating variant preferences, the system can effectively capture this contextual information, allowing it to deliver recommendations that are not only more relevant but also unbiased. In practice, this means that the greater the discrepancy between the variant preference and invariant preference, the more attention should be paid to the information in the variant preference. Conversely, when the discrepancy is minimal, such imputation information may be less required. Ultimately, we derive users' invariant and variant-dominated feedback, measuring their differences through distance calculations. This difference informs the adaptive fusion of preferences, which serves as soft labels to guide the training of the student model.

3.2.2. Teacher Model.

Given a user-item pair (u, i), we can calculate the invariant-dominated feedback $\hat{y}_{u,i}^{inv}$ and the variant-dominated feedback $\hat{y}_{u,i}^{var}$ respectively:

$$\hat{y}_{u,i}^{inv} = \varphi(\mathbf{m}_{u,i}), \tag{10}$$
$$\hat{y}_{u,i}^{var} = \varphi(\mathbf{n}_{u,i}). \tag{11}$$

where φ is the predictive function (defined in Eq.4).

Then the distance between the prediction labels of invariant and variant preferences is calculated as an indicator to fuse the prediction results. The fusion process in the teacher model can be formulated as follows:

$$d_{u,i} = |\hat{y}_{u,i}^{inv} - \hat{y}_{u,i}^{var}|, \tag{12}$$

$$w_{u,i}^{inv} = (1 - d_{u,i})^{\gamma},\tag{13}$$

$$w_{u,i}^{var} = 1 - w_{u,i}^{inv},$$
 (14)

$$\hat{y}_{u,i}^* = w_{u,i}^{inv} * \hat{y}_{u,i}^{inv} + w_{u,i}^{var} * \hat{y}_{u,i}^{var}, \tag{15}$$

where γ is a hyper-parameter and acts as the temperature coefficient in the fusion process. With the increasing of γ , more information on variant preferences be used to predict the final unbiased user feedback.

3.2.3. Student Model.

In order to enhance model performance and facilitate downstream recommendation tasks, we introduce a distanceaware knowledge distillation strategy to distill the unbiased information from the teacher model to a lightweight student model. Specifically, a simple matrix factorization-based student model is designed to capture the collaborative signals. The prediction label $\hat{y}_{u,i}^{stu}$ of the student model is defined as:

$$\hat{y}_{u,i}^{stu} = \mathbf{s}_{u,i} \odot \mathbf{t}_{u,i},\tag{16}$$

where $\mathbf{s}_{u,i}$ and $\mathbf{t}_{u,i}$ are the learning representations of the user and item in the student model.

For each user-item pair (u, i) in observed data, we treat the predicted unbiased user feedback $\hat{y}_{u,v}^*$ in the teacher model (see in Eq.15) as the ground truth label. The loss function in knowledge distillation process is defined as:

$$L_{KD} = \frac{1}{|D|} \sum_{(u,i)\in D} l_{rec}(\hat{y}_{u,i}^{stu}, \hat{y}_{u,i}^{*}),$$
(17)

where $l_{rec}(\cdot)$ is the binary cross-entropy loss.

In summary, we first disentangle the variant and invariant preference of users, then propose a student model to make the distillation of fused information, which not only lightweight the model parameters but also increases the generalization ability of the debiasing student model. Without the requirement of unbiased data, our debiasing model is trained in an unsupervised way, which expands the applicability of KD-Debias.

4. Experiments

We conduct experiments on three public datasets and compare our model with representative debiasing models to show the effectiveness of KD-Debias.

4.1. Experimental Settings

4.1.1. Datasets.

To validate the effectiveness of KD-Debias, we utilize three real-world datasets: Yahoo!R3, Coat, and Mind datasets. The statistics of three datasets are shown in Table 1.

• Yahoo!R3²: The Yahoo! R3 dataset is a collection of user ratings for songs, collected within the context of music recommendation. The rating scale range from 1 to 5. It consists of two subsets: a set of biased data collected from norm recommendation scenarios, and a set of unbiased data collected by a random logging strategy where items are posted randomly.

²https://webscope.sandbox.yahoo.com/catalog.php?datatype=r&did=3

Table 1		
Statistics	of	datasets.

Dataset	#Users	#Item	Biased Data	Unbiased Data
Yahoo!R3	15.4k	1.0k	250k	54k
Coat	290	300	6.9k	4.2k
Mind	100k	161k	24M	-

- **Coat**³: The dataset is sourced from a coat recommendation platform, capturing user ratings for items on a scale of 1 to 5. It consists of two subsets: a set of biased data collected from norm recommendation scenarios, and a set of unbiased data collected by a random logging strategy where items are posted randomly.
- **Mind**⁴: is a widely used news dataset and it contains user clicks for news and which news is exposed to users. It is widely used to study the debiasing problem of exposure bias.

Following the experimental settings in (Ding et al., 2022; Wang et al., 2022), the biased data is used in the training process. For the methods under the supervision of unbiased data, we randomly sample 5% of unbiased data as the supervised labels in the training set (this part is not used in unsupervised methods), 5% of unbiased data as the validation set for choosing parameters, the remaining 90% of the unbiased data is used for testing. The rating scores (1-5) larger than 3 are treated as positive feedback, while the other rating scores are treated as negative feedback.

4.1.2. Evaluation metrics.

We adopt the NDCG@k and Recall@k to evaluate the performance of our model.

• NDCG@k: measures the quality of recommendations by combining relevance and position information.

$$DCG@k = \sum_{(u,i)} \frac{\mathbb{I}(\hat{z}_{u,i} \le k)}{log(\hat{z}_{u,i} + 1)},$$
(18)

$$NDCG@k = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{DCG@k}{IDCG@k},$$
(19)

where $\hat{z}_{u,i}$ is the rank position of item *i* in the recommended list for user *u*, IDCG@*k* is the ideal DCG@*k*. $\mathbb{I}(\hat{z}_{u,i} \leq k)$ returns 1 if $\hat{z}_{u,i} \leq k$, and 0 otherwise.

• Recall@k: measures the number of items appear in the top-K clicked list of users.

$$Recall_u@k = \frac{\sum_{(u,i)} \mathbb{I}(\hat{z}_{u,i} \le k)}{|I_u|},\tag{20}$$

$$Recall@k = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} Recall_u@k,$$
(21)

where I_u is the item set of all interactions of user u.

4.1.3. Baselines.

We compare our proposed framework KD-Debias with the mainstream debiasing methods in recommender systems.

- MF (Koren, Bell, and Volinsky, 2009): Matrix Factorization is a widely used model in the recommendation, which is a biased model without debiasing capability.
- MF-IPS (Schnabel et al., 2016): is a recommendation model that incorporates inverse propensity scores into the MF framework to address bias issues.

³https://www.cs.cornell.edu/ schnabts/mnar/ ⁴https://paperswithcode.com/dataset/mind

Table 2

	Debiasing	Unsupervised	General	KD
MF	×		Х	×
MF-IPS		×		×
AutoDebias		×		×
CauseE	, V	×	v	
InterD	, V	×	v	v
InvPref	, V		v	×
WMF	, V	, V	×	×
EXMF	, V	, V	×	×
KD-Debias				

Properties of debiasing methods. Debiasing: a debiasing model?; Unsupervised: training with unbiased data? General: a general debiasing model? KD: knowledge distillation model?

- **CauseE** (Bonner and Vasile, 2018): is a popular embedding method that is able to predict unbiased user feedback by using biased data and unbiased data.
- AutoDebias (Chen, Dong, Qiu, He, Xin, Chen, Lin, and Yang, 2021): is a debiasing model trained with normal biased and unbiased data by utilizing meta-learning.
- InterD (Ding et al., 2022): is a SOTA debiasing knowledge distillation method in recommender systems, interD takes MF and AutoDebias as biased-teacher and unbiased teacher to distill a student model.
- **InvPref** (Wang et al., 2022): is a SOTA unsupervised debiasing method in recommender systems. InvPref treats environments as the agent of confounders and applies invariant learning to learn the user's unbiased preference.

To further verify the debiasing ability of our model on a specific bias, we select one of the most common biases in recommender systems, i.e., the exposure bias. In addition to comparing with the general debiasing methods, we also compare our method with the representative methods that specially designed for exposure bias, including:

- WMF (Hu, Koren, and Volinsky, 2008): is the most representative debiasing model for exposure bias. WMF reduces exposure bias by imputing and assigning lower weights to unobserved samples.
- **EXMF** (Liang, Charlin, McInerney, and Blei, 2016): it models the probability of exposure and converted it as confidence degree weight to reduce exposure bias.

The above methods cover different kinds of debiasing approaches in recommender systems: MF-IPS and AutoDebias leverage IPS and meta-learning techniques for debiasing. CauseE and InterD are knowledge distillation debiasing approaches with lightweight unbiased models. All of them need the supervision of unbiased data. The most similar work to our proposed framework is InvPref, which conducts unbiased recommendations without relying on unbiased data. Different from InvPref, our KD-Debias framework is a general debiasing framework, which automatically incorporates both invariant and variant information via a distance-aware knowledge distillation process. Table 2 summarizes the properties of different methods.

4.1.4. Parameter Settings.

For each baseline method, a grid search is applied to find the optimal settings. These include embedding dimensions from {10, 20, 30, 40, 50}, and the learning rate from {0.1, 0.01, 0.001, 0.0001, 0.00001}. We report the result of each method with its optimal hyperparameter settings on the validation data. In our KD-Debias model, we set the dimension of invariant and variant embedding to 40, the number of environments is 2, the learning rate in disentangled preference modeling is 0.003, the learning rate in distance-aware distillation is 0.005, the α , β , γ are set to 1.9, 9.9, 0.17. The code will be publicly available after the review process.

4.2. Main Results

We compare the performances of methods in the general debiasing and specific debiasing (i.e., exposure bias) settings. As shown in Table 3, Table 4 and Table 5, we can observe that:

Table 3

Performace comparison on Yahoo!R3 dataset. The improvements achieved by KD-Debias are statistically significant (i.e., two-sided t-test with p-value < 0.05) over the best baseline (marked by underline).

Mathad	Yahoo!R3				
wiethou	NDCG@5	Imp.↑	Recall@5	Imp.↑	# Parameters
MF	0.6007	-	0.7639	-	492k
MF-IPS	0.6113	+1.76%	0.7754	+1.51%	492k
CauseE	0.6108	+1.68%	0.7694	+0.72%	492k
AutoDebias	0.6134	+2.11%	0.7762	+1.61%	180k
InterD	0.6263	+4.26%	0.7853	+2.80%	180k
InvPref	0.6578	+9.51%	0.8065	+5.58%	1.31M
KD-Debias	0.6739*	+12.19 %	0.8217*	+7.57%	656k

Table 4

Performace comparison on Coat dataset. The improvements achieved by KD-Debias are statistically significant (i.e., twosided t-test with p-value < 0.05) over the best baseline (marked by underline).

Method	Coat					
Wethou	NDCG@5	lmp. ↑	Recall@5	lmp. ↑	# Parameters	
MF	0.4128	-	0.4716	-	17.7k	
MF-IPS	0.4131	+0.07%	0.4775	+1.25%	17.7k	
CauseE	0.4166	+0.92%	0.4734	+0.38%	17.7k	
AutoDebias	0.4380	+6.10%	0.4899	+3.88%	6.5k	
InterD	0.4475	+8.41%	0.5121	+8.59%	6.5k	
InvPref	0.4241	+2.74%	0.4721	+0.11%	47.3k	
KD-Debias	0.4687*	+13.54%	0.5180*	+9.83%	23.6k	

Table 5

Performace comparisons in dealing with exposure bias on MIND dataset.

Method	NDCG			Recall		
Method	Top10	Top20	Top40	Top10	Top20	Top40
MF	0.222	0.258	0.283	0.362	0.499	0.591
WMF	0.234	0.267	0.301	0.379	0.513	0.607
EXMF	0.228	0.264	0.289	0.362	0.499	0.588
InvPref	0.234	0.270	0.298	0.386	0.532	0.624
KD-Debias	0.238	0.278	0.301	0.386	0.545	0.661

(1) MF performs the worst in general debiasing test due to it is not a debiased model. Compared to MF, MF-IPS addressing specific biases without uniform data can bring benefits.

(2) CauseE, AutoDebias, and InterD steadily outperform MF and MF-IPS. Under the supervision of unbiased data, these methods gain improvements in resisting affect from biases.

(3) The unsupervised model InvPerf outperforms other baseline methods in Yahoo!R3 dataset, suggesting that the invariant preference can reveal the unbiased preference of users to some extent. However, due to the information loss problem, InvPerf loses the advantages of the supervised method InterD in the Coat dataset.

(4) Our proposed method KD-Debias achieves the best performance on all datasets, showing the effectiveness of our model by incorporating the biased information in the distance-aware knowledge distillation framework. (5) For the model efficiency of the training methods with observed biased data only, our surrogate student model in KD-Debias is a lightweight model, which achieves the best performance with less than 50% learning parameters compared to the SOTA baseline method InvPerf.

(6) For the specific exposure bias, we report the NDCG and Recall at top K = 10, 20, 40 in the rank lists. As shown in Table 5, our model achieves the best performance on both NDCG and Recall in terms of addressing exposure bias. This observation indicates that with the imputation of biased information, our knowledge distillation method learns more useful information for unbiased recommendations.

Dataset	Yahoo!R3		Coat	
Metric	NDCG@5 Recall@5		NDCG@5	Recall@5
Ours	0.6739	0.8217	0.4687	0.5180
Ours-Variant	0.6693	0.8125	0.4278	0.4695
Ours-Distance	0.6596	0.8089	0.4643	0.5142
Ours-KD	0.6326	0.8046	0.4291	0.4723



Figure 3: Stability comparison on Yahoo!R3 dataset. The standard deviation is computed based on 20 times repeated experiments with different random seeds.

4.3. Experimental Analysis

We conduct in-depth analyses of our model to quantitatively show the effectiveness and stability of KD-Debias.

4.3.1. Ablation Study.

To demonstrate the effectiveness of each component (including the variant information, the distance-aware fusion strategy and the knowledge distillation process) in our proposed KD-Debias, we conduct the ablation study on Yahoo!R3 and Coat datasets. The degeneration versions of KD-Debias are as follows:

- Ours-Variant: it only utilizes invariant user preference to guide the learning of student model.
- Ours-Distance: it replaces the distance-aware combination with a simple combination strategy (i.e., information are fused with the equal weight 0.5) of the invariant and variant preference in the knowledge distillation process.
- Ours-KD: the invariant and variant preference are fused with distance-aware mode for unbiased recommendations without knowledge distillation process. The teacher model is used to make prediction for unbiased recommendations.

As shown in Table 6, we can observed that: (1) The performance of Ours-Variant is worse than KD-Debias on all metrics, which verifies the effectiveness of incorporating biased information. (2) Replacing the fusion function (Ours-Distance) will result in the degeneration of model performance, showing the usefulness of our proposed distance-aware module. This emphasizing the critical role of effective information fusion in enhancing overall model performance. (3) By removing the knowledge distillation module (Ours-KD), although the variant and invariant preferences are intuitively fused with weighted coefficients, the model performance decreases a lot, showing the necessity to design a knowledge distillation framework with the surrogate student model to fuse information. This emphasizes the importance of knowledge distillation in effectively blending diverse information for better performance.



Figure 4: The performances comparison of NDCG@5 with different hyper-parameter γ .



Figure 5: The distance between the prediction results of invariant preferences (marked by the red circles), variant preferences (marked by the blue triangles), and the fused preference (marked by the green stars). The dots under the red dashed line (y = 0.5) can be regarded as correct predictions.

4.3.2. The Stability Analysis.

We conducted 20 repeated experiments with different random seeds in the same settings to evaluate the stability of models. We compute the standard deviation of the performance on each model on Yahoo!R3 dataset, and choose the base model (i.e., MF (Koren et al., 2009)) in these methods as the baseline to compare with. As shown in Fig. 3, there is a certain gap between InvPref and the base model MF, while the standard deviations of the model performance in InterD and AutoDebias are 13 times larger than MF, showing the instability of these methods. The stability of our KD-Debias is comparable to the base model MF. It may benefit from the distance-aware knowledge distillation module, which combines invariant preference and variant preference by considering the inconsistency distribution between them and meanwhile increases the generalization ability via learning the soft predicted labels from the teacher model.

4.3.3. Hyperparameters Analysis.

We conduct analysis of the hyperparameters (e.g. the temperature coefficient γ) that may influence the model performance. As mentioned in Eq.13, γ is a hyper-parameter acting as the temperature co-efficient in knowledge distillation to smooth the fusion of unbiased information and biased information. As shown in Fig. 4, as the increasing of γ , *NDCG*@5 has a clear trend of increasing and then decreasing. The increasing part indicated that incorporating biased information enhances the performance of unbiased recommendations, while incorporating too much variant information may confound invariant preference and substantially decrease the performance of unbiased recommendations.

4.3.4. Case Study.

We have demonstrated the effectiveness of incorporating variant information in knowledge distillation. We intuitively display its effect on the model performance in Fig. 5. The value on the vertical axis represents the distance to the true unbiased label. The smaller value is, the more precise of prediction results. We can see that variant information (marked by blue triangles) plays a complementary role to invariant information (marked by the red circle) in debiasing learning, and their fusion makes the prediction results (marked by green stars) closer to the true unbiased label. It shows the the effectiveness of variant preferences in our model in enhancing the debiasing performance in recommender systems.

5. Conclusion

In this paper, we investigate the problem of general debiasing with only the biased observation data. We point out the instability and low accuracy of invariant learning in unbiased recommendations, and propose a novel distanceaware knowledge distillation framework (KD-Debias) to make unbiased predictions with both invariant and variant information. Extensive experiments show that with the imputation from variant information, our knowledge distillation method achieves the best model performance in unbiased recommendation task. In our framework, a simple MF method is used as the student model to improve the model efficiency and generalization ability in unbiased recommendations. In the future, we will expand the applications of our framework with more complex student models in different recommendation scenarios.

References

- D. Luo, L. Zou, Q. Ai, Z. Chen, C. Li, D. Yin, B. D. Davison, Unbiased learning-to-rank needs unconfounded propensity estimation, in: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2024, pp. 1535–1545.
- W. Wei, X. Ren, J. Tang, Q. Wang, L. Su, S. Cheng, J. Wang, D. Yin, C. Huang, Llmrec: Large language models with graph augmentation for recommendation, in: Proceedings of the 17th ACM International Conference on Web Search and Data Mining, 2024, pp. 806–815.
- X. Song, H. Huang, J. Lian, H. Jin, Xgcn: a library for large-scale graph neural network recommendations, Frontiers of Computer Science 18 (2024) 183343.
- W. Hua, Y. Ge, S. Xu, J. Ji, Z. Li, Y. Zhang, Up5: Unbiased foundation model for fairness-aware recommendation, in: Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers), 2024, pp. 1899–1912.
- L. Yang, Y. Cui, Y. Xuan, C. Wang, S. Belongie, D. Estrin, Unbiased offline recommender evaluation for missing-not-at-random implicit feedback, in: Proceedings of the 12th ACM conference on recommender systems, 2018, pp. 279–287.
- A. Collins, D. Tkaczyk, A. Aizawa, J. Beel, A study of position bias in digital library recommender systems, arXiv preprint arXiv:1802.06565 (2018).
- Z. Chen, R. Xiao, C. Li, G. Ye, H. Sun, H. Deng, Esam: Discriminative domain adaptation with non-displayed items to improve long-tail performance, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2020, pp. 579–588.
- A. Krishnan, A. Sharma, A. Sankar, H. Sundaram, An adversarial approach to improve long-tail performance in neural collaborative filtering, in: Proceedings of the 27th ACM International Conference on information and knowledge management, 2018, pp. 1491–1494.
- H. Li, K. Wu, C. Zheng, Y. Xiao, H. Wang, Z. Geng, F. Feng, X. He, P. Wu, Removing hidden confounding in recommendation: a unified multi-task learning approach, Advances in Neural Information Processing Systems 36 (2024).
- Q. Li, X. Wang, Z. Wang, G. Xu, Be causal: De-biasing social network confounding in recommendation, ACM Transactions on Knowledge Discovery from Data 17 (2023) 1–23.
- A. Zhang, W. Ma, P. Wei, L. Sheng, X. Wang, General debiasing for graph-based collaborative filtering via adversarial graph dropout, arXiv preprint arXiv:2402.13769 (2024).
- E. Creager, J.-H. Jacobsen, R. Zemel, Environment inference for invariant learning, in: International Conference on Machine Learning, PMLR, 2021, pp. 2189–2200.
- X. Du, Z. Wu, F. Feng, X. He, J. Tang, Invariant representation learning for multimedia recommendation, in: Proceedings of the 30th ACM International Conference on Multimedia, 2022, pp. 619–628.
- Z. Wang, Y. He, J. Liu, W. Zou, P. S. Yu, P. Cui, Invariant preference learning for general debiasing in recommendation, in: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2022, pp. 1969–1978.
- M. Tang, L. Zou, S. Cui, S.-n. Liang, Z. Jin, Unbiased recommendation through invariant representation learning, in: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer, 2024, pp. 280–296.
- L. Van der Maaten, G. Hinton, Visualizing data using t-sne., Journal of machine learning research 9 (2008).
- Y. Saito, Asymmetric tri-training for debiasing missing-not-at-random explicit feedback, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2020, pp. 309–318.
- H. Li, C. Zheng, X.-H. Zhou, P. Wu, Stabilized doubly robust learning for recommendation on data missing not at random, arXiv preprint arXiv:2205.04701 (2022).
- M. Mansoury, Understanding and mitigating multi-sided exposure bias in recommender systems, ACM SIGWEB Newsletter (2022) 1-4.
- M. Ning, E. Sangineto, A. Porrello, S. Calderara, R. Cucchiara, Input perturbation reduces exposure bias in diffusion models, arXiv preprint arXiv:2301.11706 (2023).

- Z. Zhao, J. Chen, S. Zhou, X. He, X. Cao, F. Zhang, W. Wu, Popularity bias is not always evil: Disentangling benign and harmful bias for recommendation, IEEE Transactions on Knowledge and Data Engineering (2022).
- J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, X. He, Bias and debias in recommender system: A survey and future directions, ACM Transactions on Information Systems 41 (2023) 1–39.
- J. Nam, H. Cha, S. Ahn, J. Lee, J. Shin, Learning from failure: De-biasing classifier from biased classifier, Advances in Neural Information Processing Systems 33 (2020) 20673–20684.
- Z. Wang, S. Shen, Z. Wang, B. Chen, X. Chen, J.-R. Wen, Unbiased sequential recommendation with latent confounders, in: Proceedings of the ACM Web Conference 2022, 2022, pp. 2195–2204.
- X. Zhu, Y. Zhang, X. Yang, D. Wang, F. Feng, Mitigating hidden confounding effects for causal recommendation, IEEE Transactions on Knowledge and Data Engineering (2024).
- H. Zhao, Q. Cui, X. Li, R. Bao, L. Li, J. Zhou, Z. Liu, J. Feng, Mdi: A debiasing method combining unbiased and biased data, in: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2023, pp. 3280–3284.
- T. Wei, F. Feng, J. Chen, Z. Wu, J. Yi, X. He, Model-agnostic counterfactual reasoning for eliminating popularity bias in recommender system, in: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, 2021, pp. 1791–1800.
- R. Qiu, S. Wang, Z. Chen, H. Yin, Z. Huang, Causalrec: Causal inference for visual debiasing in visually-aware recommendation, in: Proceedings of the 29th ACM International Conference on Multimedia, 2021, pp. 3844–3852.
- H. Du, H. Yuan, Z. Huang, P. Zhao, X. Zhou, Sequential recommendation with diffusion models, arXiv preprint arXiv:2304.04541 (2023).
- L. Wu, Z. Qiu, Z. Zheng, H. Zhu, E. Chen, Exploring large language model for graph data understanding in online job recommendations, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, 2024, pp. 9178–9186.
- Y. Yang, C. Huang, L. Xia, C. Huang, D. Luo, K. Lin, Debiased contrastive learning for sequential recommendation, in: Proceedings of the ACM web conference 2023, 2023, pp. 1063–1073.
- S. Li, L. Yao, S. Mu, W. X. Zhao, Y. Li, T. Guo, B. Ding, J.-R. Wen, Debiasing learning based cross-domain recommendation, in: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, 2021, pp. 3190–3199.
- Y. Guo, F. Cai, J. Zheng, X. Zhang, H. Chen, Disentangled variational auto-encoder enhanced by counterfactual data for debiasing recommendation, Complex & Intelligent Systems 10 (2024) 3119–3132.
- X. Yang, X. Li, Z. Liu, Y. Wang, S. Lu, F. Liu, Disentangled causal representation learning for debiasing recommendation with uniform data, Applied Intelligence (2024) 1–16.
- T. Schnabel, A. Swaminathan, A. Singh, N. Chandak, T. Joachims, Recommendations as treatments: Debiasing learning and evaluation, in: international conference on machine learning, PMLR, 2016, pp. 1670–1679.
- Z. Zhu, Y. He, Y. Zhang, J. Caverlee, Unbiased implicit recommendation and propensity estimation via combinational joint learning, in: Proceedings of the 14th ACM Conference on Recommender Systems, 2020, pp. 551–556.
- Q. Zhang, X. Zhang, Y. Liu, H. Wang, M. Gao, J. Zhang, R. Guo, Debiasing recommendation by learning identifiable latent confounders, in: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2023, pp. 3353–3363.
- D. Liu, P. Cheng, H. Zhu, Z. Dong, X. He, W. Pan, Z. Ming, Debiased representation learning in recommendation via information bottleneck, ACM Transactions on Recommender Systems 1 (2023) 1–27.
- Z. Wang, X. Chen, R. Wen, S.-L. Huang, E. Kuruoglu, Y. Zheng, Information theoretic counterfactual learning from missing-not-at-random feedback, Advances in Neural Information Processing Systems 33 (2020) 1854–1864.
- M. Arjovsky, L. Bottou, I. Gulrajani, D. Lopez-Paz, Invariant risk minimization, arXiv preprint arXiv:1907.02893 (2019).
- G. Hinton, O. Vinyals, J. Dean, Distilling the knowledge in a neural network, arXiv preprint arXiv:1503.02531 (2015).
- J. Gou, B. Yu, S. J. Maybank, D. Tao, Knowledge distillation: A survey, International Journal of Computer Vision 129 (2021) 1789-1819.
- D. Liu, P. Cheng, Z. Dong, X. He, W. Pan, Z. Ming, A general knowledge distillation framework for counterfactual recommendation via uniform data, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2020, pp. 831–840.
- S. Kang, J. Hwang, W. Kweon, H. Yu, De-rrd: A knowledge distillation framework for recommender system, in: Proceedings of the 29th ACM International Conference on Information & Knowledge Management, 2020, pp. 605–614.
- J. Tang, K. Wang, Ranking distillation: Learning compact ranking models with high performance for recommender system, in: Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining, 2018, pp. 2289–2298.
- Z. Zhu, S. Si, J. Wang, Y. Yang, J. Xiao, Debias the black-box: A fair ranking framework via knowledge distillation, in: Web Information Systems Engineering–WISE 2022: 23rd International Conference, Biarritz, France, November 1–3, 2022, Proceedings, Springer, 2022, pp. 395–405.
- S. Ding, F. Feng, X. He, J. Jin, W. Wang, Y. Liao, Y. Zhang, Interpolative distillation for unifying biased and debiased recommendation, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2022, pp. 40–49.
- D. Liu, P. Cheng, Z. Lin, J. Luo, Z. Dong, X. He, W. Pan, Z. Ming, Kdcrec: Knowledge distillation for counterfactual recommendation via uniform data, IEEE Transactions on Knowledge and Data Engineering (2022).
- G. Chen, J. Chen, F. Feng, S. Zhou, X. He, Unbiased knowledge distillation for recommendation, in: Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, 2023, pp. 976–984.
- Y. Koren, R. Bell, C. Volinsky, Matrix factorization techniques for recommender systems, Computer 42 (2009) 30-37.
- S. Bonner, F. Vasile, Causal embeddings for recommendation, in: Proceedings of the 12th ACM conference on recommender systems, 2018, pp. 104–112.
- J. Chen, H. Dong, Y. Qiu, X. He, X. Xin, L. Chen, G. Lin, K. Yang, Autodebias: Learning to debias for recommendation, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2021, pp. 21–30.
- Y. Hu, Y. Koren, C. Volinsky, Collaborative filtering for implicit feedback datasets, in: 2008 Eighth IEEE international conference on data mining, Ieee, 2008, pp. 263–272.

D. Liang, L. Charlin, J. McInerney, D. M. Blei, Modeling user exposure in recommendation, in: Proceedings of the 25th international conference on World Wide Web, 2016, pp. 951–961.