

CounterCLR: Counterfactual Contrastive Learning with Non-random Missing Data in Recommendation

Jun Wang¹, Haoxuan Li², Chi Zhang³, Dongxu Liang³, Enyun Yu³, Wenwu Ou³, Wenjia Wang^{1,4}

¹Hong Kong University of Science and Technology

²Peking University

³Kuaishou Technology Co., Ltd.

⁴Hong Kong University of Science and Technology (Guangzhou)

Abstract—Recommender systems are designed to learn user preferences from observed feedback and comprise many fundamental tasks, such as rating prediction and post-click conversion rate (pCVR) prediction. However, the observed feedback usually suffer from two issues: *selection bias* and *data sparsity*, where biased and insufficient feedback seriously degrade the performance of recommender systems in terms of accuracy and ranking. Existing solutions for handling the issues, such as data imputation and inverse propensity score, are highly susceptible to additional trained imputation or propensity models. In this work, we propose a novel counterfactual contrastive learning framework for recommendation, named CounterCLR, to tackle the problem of non-random missing data by exploiting the advances in contrast learning. Specifically, the proposed CounterCLR employs a deep representation network, called CauNet, to infer non-random missing data in recommendations and perform user preference modeling by further introducing a self-supervised contrastive learning task. Our CounterCLR mitigates the selection bias problem without the need for additional models or estimators, while also enhancing the generalization ability in cases of sparse data. Experiments on real-world datasets demonstrate the effectiveness and superiority of our method.

Index Terms—recommendation system, non-random missing data, causal inference, contrastive learning.

I. INTRODUCTION

In real-world recommender systems, users' interactive feedback on items such as rating and purchase is used to represent and develop user preferences. By using the observed sparse feedback matrix to inference the potential preferences or relevance of non-interactive user-item pairs, recommendation systems can be facilitated to provide diverse and personalized recommendations to users. Thus, the task of rating or conversion rate prediction has become the core problem and has attracted increasingly attention [1]–[6].

Although many techniques have been proposed to tackle the missing data problem in recommendation, recent studies showed that the missing data prediction task usually suffers from two major issues: *selection bias* and *data sparsity* [2], [7]–[9]. The selection bias, due to the users' self-selection behaviors or the policies of the recommender systems, results that the observed ratings are highly biased, which is widely known as data *Missing Not At Random* (MNAR) [9]–[12].

This work was done when Jun Wang (jwangfx@connect.ust.hk) was a research intern at Kuaishou Technology Co., Ltd. Wenjia Wang (wenjiawang@ust.hk) is the corresponding author.

Specifically, since users tend to choose the items they like to rate or the recommender systems are more likely to recommend popular items to users, the higher (or the more positive) ratings are more likely to be observed. The data sparsity means that only a small portion of ratings are observed while the most of the ratings are missing. It occurs since the interactions between the users and items are rare, compared with the number of entire user-item pairs. In summary, the selection bias and data sparsity issues severely prevent recommender systems from learning users' true preferences and thus degrade the performance of the recommender systems.

There are growing literature focusing on addressing the selection bias and data sparsity issues [8], [9]. Among them, causality-based methods have become increasingly popular in recommender systems [1], [13]–[19], in which the missingness of feedback are modeled under the widely-used “*Potential Outcome Framework*” (POF) in causal inference literature. One class of causality-based methods is reweighting-based methods, including inverse propensity score (IPS) [13], [19], self-normalized inverse propensity score (SNIPS) [13], and doubly robust (DR) estimation [1], [14]–[20]. Despite showing superior performance on debiasing, they highly rely on the accuracy of the propensity estimation, in which an extra model and/or unbiased datasets are required, and the performance degenerates with mis-specified model or sparse interacted data.

As for data sparsity, many studies have proposed data augmentation-based methods, by eliciting more pseudo user ratings or merging the ratings of reliable similar users, so as to provide sufficient data for training the prediction model [21]. However, these methods only exploit the observed feedback for data augmentation and ignore the distributional difference between the observed and potential feedback on all user-item pairs, which leads to biased predictions. Recent studies have proposed the use of multi-task learning [6], [22], [23] to address both selection bias and data sparsity simultaneously. Different from these studies, we propose a novel contrastive learning approach to mitigate the data sparsity problem while addressing non-random missing data.

In this work, we propose a **Counterfactual Contrastive Learning** framework for **Recommendations**, named **CounterCLR**, to address the non-random missing data and data sparsity issues simultaneously. The proposed CounterCLR is composed of a causality-based network named CauNet and a

contrastive learning auxiliary task. Specifically, the CauNet is designed under the potential outcome framework to model the non-random missing feedback from the observed interactions. The contrastive learning objective in CauNet conducts an auxiliary task for learning user and item representations, which enhances the generalization ability in cases of sparse data.

The contributions of this paper are summarized as follows. **i)** It not only exhibits a substantial enhancement of generalization ability under sparse interactions, but also alleviates the problem of non-random missing data, thus can address selection bias and data sparsity issues simultaneously in recommender systems. **ii)** Our framework does not require any additional models or unbiased datasets, and thus it is practically preferred. **iii)** A novel casually contrastive paradigm is proposed for debiased recommendation and user preference modeling. Extensive experiments on real-world datasets illustrate the above merits of CounterCLR, and show that CounterCLR substantially outperforms state-of-the-art methods.

II. RELATED WORK

In this section, we review many previous methods designed for debiased recommendation and recent advances on self-supervised contrastive learning (SSCL).

Debiased Recommendation. The inconsistency between training and test set distributions has been widely studied [4], [24]–[31]. Existing methods tackling the non-random missing data can be divided into three categories: generative modeling-based, relabeling-based, and reweighting-based methods. By assuming a data generation process and adjusting for biases accordingly, generative modeling-based methods [32], [33] leverage heuristic human prior knowledge and provide an explainable solution for debiasing. Relabeling methods [1], [34] typically mitigate the selection bias by data imputation, which adaptively downweights the contribution of imputed ratings for unobserved user-item pairs to the loss function. Reweighting methods involve assigning weights induced from the propensity scores to each instance, in order to rescale their contributions during model training. Typical reweighting methods include IPS [13], [19], self-normalized inverse propensity score (SNIPS) [13], and DR [1], [3], [14]–[16], [35]–[37], where the propensity scores are obtained through naive Bayes or logistic regression. Nevertheless, the above methods highly depend on additional imputation models or propensity estimators, which need to be trained separately.

On the other hand, there have also been several works that avoid the introduction of extra models or estimators. Specifically, [38] developed a model-agnostic meta learning method by deploying the asymmetric tri-training framework for unsupervised domain adaptation. [39], [40] built information-theoretic frameworks, where novel non-exposure variational information bottlenecks are derived for addressing the non-random missing data problem. Alternatively, [21] proposed a self-supervised learning-based method, in which an extra collected unbiased *Missing At Random* (MAR) rating data is required to calibrate the rating distribution. However, the collection of unbiased MAR ratings can be costly or even

unavailable in practice. In contrast, in this work, we proposed a causally contrastive learning-based framework that essentially introduces neither additional models and estimators nor extra unbiased MAR data, which is more practically preferred.

Self-supervised Contrastive Learning. Self-supervised contrastive learning (SSCL) aims to boost the model generalization ability by representation learning, where the embeddings of the augmented versions of the same sample are trained to be close to each other, while those of different samples are required to be pushed away. This method has been widely utilized in CV and NLP areas [41]–[43]. [43] empirically and theoretically shows that self-supervised contrastive learning can learn more representative features for helping classification task in long-tailed labeled image datasets. Recently, [44]–[46] studied the contrastive learning methods in sequential recommendations to learn better item representations in the presence of long-tail items. [47] exploited a SSCL-based method, named CLRec, to alleviate the popularity bias in deep candidate generation. However, these methods failed to tackle the problem of non-random missing data in recommendation.

III. PRELIMINARIES

In this section, we formulate the problem of non-random missing data in recommendation using the widely-adopted *Potential Outcome Framework* (POF) in causal inference literature, and introduce the motivation of the proposed framework (CounterCLR) from a causal inference view.

A. Problem Setup

Let the user set and item set be $\mathcal{U} = \{u_1, u_2, \dots, u_N\}$ and $\mathcal{I} = \{i_1, i_2, \dots, i_M\}$, respectively, and the set of all user-item pairs be $\mathcal{D} = \mathcal{U} \times \mathcal{I}$. Let $\mathbf{R} \in \mathbb{R}^{N \times M}$ be the matrix of true feedback (e.g., rating, conversion, etc.), with elements $r_{u,i}$ be the true feedback of user u to item i . For modeling the missing mechanism of the observed feedback, we introduce an observing indicator matrix $\mathbf{O} = (o_{u,i}) \in \mathbb{R}^{N \times M}$, where $o_{u,i} = 1$ represents that $r_{u,i}$ is observed or missing $o_{u,i} = 0$. For the rating or the post-click conversion rate prediction task, we aim to train a prediction model minimizing the training loss

$$\text{Error}(\hat{\mathbf{R}}, \mathbf{R}) = \sum_{(u,i) \in \mathcal{D}} \ell(\hat{r}_{u,i}, r_{u,i}), \quad (1)$$

where $\hat{\mathbf{R}} = (\hat{r}_{u,i}) \in \mathbb{R}^{N \times M}$ is the predicted rating matrix, and $\ell(\cdot, \cdot)$ is a loss function, e.g., squared loss.

In practice, recommender systems conceptually model the user u and item i by K -dimensional embeddings, i.e., $\mathbf{e}_u, \mathbf{e}_i \in \mathbb{R}^K$, respectively. Then, each user-item pair (u, i) is represented by the embedding concatenation as $\mathbf{x}_{u,i} = (\mathbf{e}_u, \mathbf{e}_i)$. Two common rating prediction models in recommender systems are matrix factorization (MF) [48] and neural collaborative filtering (NCF) [49]. The MF model directly processes $\mathbf{x}_{u,i}$ to conduct the rating prediction by $\hat{r}_{u,i} = \mathbf{e}_u^\top \mathbf{e}_i$, while the NCF model applies a multi-layer feedforward neural network to obtain the feedback prediction $\hat{r}_{u,i}$.

B. Problem Formulation under POF

Conceptually, the POF is formulated by three components: the covariate set \mathcal{X} , the treatment set \mathcal{T} , and the potential outcome set \mathcal{Y} . For example, for a diabetic patient with covariate $x \in \mathcal{X}$, doctors are usually interested in whether a new medical treatment affects the blood sugar level $y \in \mathcal{Y}$. The potential outcomes in this example can be modeled as $y(0), y(1) \in \mathcal{Y}$, and $\mathcal{T} = \{0, 1\}$ with 0 and 1 denote the patient not take and take this new medication, respectively.

Similarly, the missing mechanism in recommendation can be modeled by an treatment assignment mechanism in the POF. To see this, note that $o_{u,i} = 1$ can be treated as assigning an exposure treatment to the user-item pair (u, i) , i.e., the item i is recommended to the user u . Consequently, we can follow the POF to redefine the rating prediction task as follows.

Example 1 (Rating Prediction Task under POF). For each user-item pair (u, i) , $o_{u,i}$ is the treatment indicator and there are two potential outcomes $r_{u,i}(1)$ and $r_{u,i}(0)$, named *exposure rating* and *non-exposure rating*, respectively. The rating prediction task is to estimate the exposure ratings of users to all items, i.e., estimate $r_{u,i}(1)$ for all $(u, i) \in \mathcal{D}$.

For simplicity, hereafter we denote $r_{u,i}(1)$ and $r_{u,i}(0)$ as $r_{u,i}^1$ and $r_{u,i}^0$, respectively. In the above example, the observed ratings are the exposure ratings, i.e., the potential outcomes of the units with the exposure treatment. Meanwhile, the ratings of the units without the exposure treatment are missing. This motivates us to introduce a causal rating prediction model CauNet into the rating prediction task. Notably, the rating prediction task under POF is different with the traditional POF problems like investigating the therapeutic effects of different treatments as in the previous diabetes example, and two challenges are raised from the differences in consequence. First, since users tend to rate the items they like, the higher ratings become more likely to be observed. Therefore, the collected user-item pairs with observed ratings could not be used as a representative of the target population, i.e., the entirety of the user-item pairs, which leads to the sample selection bias. Second, only the potential outcomes $r_{u,i}^1$ are observable, while $r_{u,i}^0$ are unavailable if the item i is not recommended to the user u , which arises the data sparsity issue. These challenges call for new methods for modeling the causalities in the rating prediction task.

IV. METHODOLOGY

In this section, we introduce the proposed causality-based contrastive learning framework for addressing non-random missing data in recommendation, named CounterCLR, which also greatly alleviates the data sparsity issue simultaneously. The CounterCLR consists of two parts: (1) a causality-based prediction model and (2) a contrastive learning objective.

A. Causality-Based Prediction

The CauNet is a three-headed architecture, which can predict the exposure and non-exposure ratings, i.e., predict $r_{u,i}^1$ and $r_{u,i}^0$, and also estimate the propensity score with a simple

and efficient procedure at the same time. The predicted exposure and non-exposure ratings are passed to the contrastive learning objective to construct positive and negative pairs. In addition, the propensity score estimation can help to learn a better representation for predicting the treatment and reduce the selection bias.

The processing steps in CauNet are presented in Figure 1. Given an embedding concatenation $\mathbf{x}_{u,i} = (\mathbf{e}_u, \mathbf{e}_i)$, we first use a neural network for encoding $\mathbf{x}_{u,i}$ to $\mathbf{z}_{u,i} \in \mathbb{R}^p$, i.e., $\mathbf{z}_{u,i} = h_{\mathbf{W}_1}(\mathbf{x}_{u,i})$. With the feature embedding $\mathbf{z}_{u,i}$, we use another two neural networks to predict the exposure and non-exposure ratings $\hat{r}_{u,i}^1$ and $\hat{r}_{u,i}^0$ when $o_{u,i} = 1$ and $o_{u,i} = 0$, respectively, i.e., $\hat{r}_{u,i}^1 = h_{\mathbf{W}_2}(\mathbf{z}_{u,i}, 1)$ and $\hat{r}_{u,i}^0 = h_{\mathbf{W}_3}(\mathbf{z}_{u,i}, 0)$. Here, \mathbf{W}_1 , \mathbf{W}_2 , and \mathbf{W}_3 are the parameters inside the neural networks $h_{\mathbf{W}_1}$, $h_{\mathbf{W}_2}$, and $h_{\mathbf{W}_3}$, respectively. Since the true ratings are determined by the user preference and should not be influenced by the exposure treatment in the rating prediction task, we update \mathbf{W}_3 in a momentum manner, i.e., $\mathbf{W}_3 \leftarrow m\mathbf{W}_3 + (1-m)\mathbf{W}_2$, to ensure that $h_{\mathbf{W}_2}$ and $h_{\mathbf{W}_3}$ are close, where m is a hyper-parameter controlling the weights of \mathbf{W}_2 in the momentum update. Additionally, in Section IV-B, we will see this momentum update mechanism cooperates with the contrastive learning objective to help minimize the discrepancy between the distributions of $r_{u,i}^1$ and $r_{u,i}^0$.

Since the propensity score estimates are not passed to the contrastive learning objective and do not explicitly influence the quality of the CounterCLR, we can construct a rather simple model for the propensity score estimation. Motivated by the random feature regression models [50], [51], we use a linear transformation with a random vector $\mathbf{h} \in \mathbb{R}^p$ followed by a sigmoid function, i.e., $\hat{o}_{u,i} = \text{Sigmoid}(\mathbf{h}^\top \mathbf{z}_{u,i})$. Note that \mathbf{h} is fixed after initialization.

In summary, the model parameter of CauNet is $\theta = (\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3)$, which is trained to minimize

$$\mathcal{L}_{cau} = \mathcal{L}_{base} + \alpha \mathcal{L}_{pro}, \quad (2)$$

where $\mathcal{L}_{base} = \sum_{(u,i) \in \mathcal{O}^1} (\hat{r}_{u,i}^1 - r_{u,i})^2 / \hat{o}_{u,i}$, and $\mathcal{L}_{pro} = \sum_{(u,i) \in \mathcal{D}} \text{CrossEntropy}(\hat{o}_{u,i}, o_{u,i})$. Hereby $\mathcal{O}^1 = \{(u, i) \in \mathcal{D} : o_{u,i} = 1\}$, and α is a hyper-parameter.

B. Contrastive Learning Objective

In order to better learn the causality in rating prediction task, we deploy a contrastive learning objective cascading to CauNet. The contrastive learning objective leverages SSCL on user preference embeddings extracted from the ratings. First, let $\hat{\mathbf{r}}_u = [\hat{r}_{u,i_1}, \dots, \hat{r}_{u,i_M}]$ stores the predicted ratings of a user u to all items. An ideal user preference embedding should not only be informative to the ratings but also reserve the statistical properties of the rating distribution. Hence, we adopt the aggregation function $f(\cdot)$ proposed by [21] as a user preference extractor, for mapping the predicted ratings $\hat{\mathbf{r}}_u$ to the user preference embedding $f_{\hat{\mathbf{r}}_u}$ as

$$f_{\hat{\mathbf{r}}_u} = f(\hat{\mathbf{r}}_u) = [f^{(1)}(\hat{\mathbf{r}}_u), \dots, f^{(K)}(\hat{\mathbf{r}}_u)], \quad (3)$$

where K is the dimension of the user embedding \mathbf{e}_u , $f^{(k)}(\hat{\mathbf{r}}_u) = \frac{1}{M} \sum_{i=1}^M \sigma\left(\tau \left(\frac{k}{K} r_{\max} + \frac{K-k}{K} r_{\min} - \hat{r}_{u,i}\right)\right)$, r_{\max}

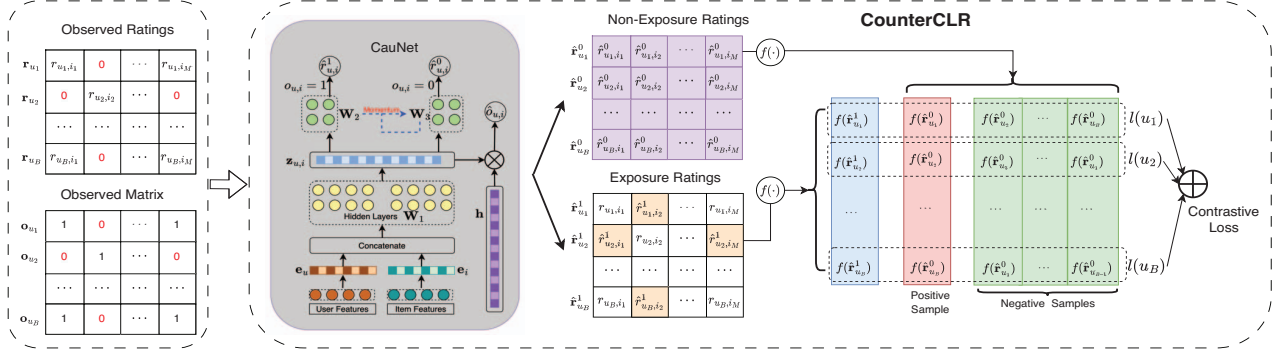


Fig. 1. The architecture of CounterCLR, including a causality-based prediction model with an auxiliary contrastive learning objective.

and r_{\min} are the maximum and minimum possible ratings respectively, $\sigma(\cdot)$ is the sigmoid function, and τ is the scale parameter. It has been proved by [21] that the user preference extractor $f(\cdot)$ is indeed an approximation of the empirical cumulative distribution function of \hat{r}_u . In addition, Eq. (3) is differentiable and adapted to gradient-based optimization methods, and theoretically guarantees the user preference embeddings and rating vectors share the same distributions between observed and unobserved user-item pairs.

Next, we define the exposure rating vector of the user u as $\hat{r}_u^1 = [o_{u,i_1} r_{u,i_1} + (1 - o_{u,i_1}) \hat{r}_{u,i_1}^1, \dots, o_{u,i_M} r_{u,i_M} + (1 - o_{u,i_M}) \hat{r}_{u,i_M}^1]$, and his/her non-exposure rating vector as $\hat{r}_u^0 = [\hat{r}_{u,i_1}^0, \dots, \hat{r}_{u,i_M}^0]$. Through the user preference extractor, we can obtain the exposure user preference embedding $f_{\hat{r}_u^1} = f(\hat{r}_u^1)$ and the non-exposure one $f_{\hat{r}_u^0} = f(\hat{r}_u^0)$. Then, we define our causally contrastive loss as

$$\mathcal{L}_{con} = \sum_{u \in \mathcal{U}} \ell(u) = \sum_{u \in \mathcal{U}} -\log \left(\frac{\exp(f_{\hat{r}_u^1}^\top f_{\hat{r}_u^0} / t)}{\sum_{u' \in \mathcal{U}} \exp(f_{\hat{r}_u^1}^\top f_{\hat{r}_{u'}^0} / t)} \right), \quad (4)$$

where t is called temperature hyper-parameter.

An illustration of the contrastive learning objective is presented in Figure 1. In practice, \mathcal{L}_{con} is optimized in a batch-wise manner. Notice that, as previously highlighted in Section IV-A, the distributions of $\hat{r}_{u,i}^1$ and $\hat{r}_{u,i}^0$ should be close, indicating that the exposure and non-exposure user preference embeddings, i.e. $f_{\hat{r}_u^1}$ and $f_{\hat{r}_u^0}$, should be similar. To this end, by minimizing \mathcal{L}_{con} , we can pull closer $f_{\hat{r}_u^1}$ and $f_{\hat{r}_u^0}$ for the user u , and push away $f_{\hat{r}_u^1}$ and $f_{\hat{r}_{u'}^0}$ for $u' \neq u$.

C. Overall Loss

Taking the above two components together, we define the overall loss \mathcal{L} as

$$\mathcal{L} = \mathcal{L}_{cau} + \beta \mathcal{L}_{con} = \mathcal{L}_{base} + \alpha \mathcal{L}_{pro} + \beta \mathcal{L}_{con}, \quad (5)$$

where β is a hyper-parameter to control the contribution of the contrastive learning objective.

V. EXPERIMENT

In this section, we empirically validate the performance of CounterCLR by answering the research questions (RQs):

RQ1. Does the proposed CounterCLR achieve the state-of-the-art capability in mitigating selection bias in rating prediction task? **RQ2.** Does the proposed CounterCLR outperform the baselines with varying data sparsity level?

A. Experiment Setup

We consider three real world datasets: **Coat**¹, **Yahoo! R3**², and **KuaiRec**³. All datasets contain biased training data acquired through traditional data collection process and unbiased test data acquired through randomized controlled trials (**Coat**, **Yahoo! R3**), and full exposure of test items (**KuaiRec**). For methods requiring unbiased data for training or propensity estimation, we randomly split out 5% test data and use Naive Bayes propensity estimator (NB) and User-Item propensity estimator (UI). We adopt three widely-used metrics, namely MSE, MAE, and NDCG@5, for performance evaluation and report the mean results over five runs.

We choose representative debiasing methods as baselines, including HEI [34], Naive [48], DR [14], AT [38], CVIB [39], ESCM2 [23], RDC [21], and SDR [17]. All methods are taking Matrix Factorization (MF) model and Neural Collaborative Filtering (NCF) model as base model on Pytorch with Adam as the optimizer. We use cross validation to select the learning rate in $\{5e-3, \dots, 1e-1\}$, weight decay in $\{1e-7, \dots, 1e-1\}$, and batch size in $\{128, \dots, 4096\}$. For the proposed CounterCLR, we also use cross-validation to select K among $\{5, 10, 20, 30, 40\}$, α among $\{0.1, 1, 5, 10, 50, 100, 200\}$, and β among $\{0.01, 0.1, 1, 5, 10\}$. We set the momentum number m as 0.999, the hyper-parameter τ as 1, and temperature hyper-parameter t as 0.07 by following [21] and [52].

B. Experiments Results (RQ1)

The debiasing performance is shown in Table I. First, the proposed CounterCLR outperforms the baseline methods both with unbiased data and without unbiased data on all three datasets, which suggests that the integration of the CauNet and the contrastive learning objective does facilitate reducing the influence of the selection bias. Second, we note that SNIPS-NB and DR-NB methods perform better than SNIPS-UI and

¹<https://www.cs.cornell.edu/~schnabts/mnar/>

²<http://webscope.sandbox.yahoo.com/>

³<https://kuaiREC.com/>

TABLE I

THE MSE, MAE AND NDCG@5 ON THREE REAL-WORLD DATASETS. THE BEST RESULTS AMONG THE METHODS WITHOUT UNBIASED MAR DATA ARE BOLD. THE BEST RESULTS AMONG ALL METHODS ARE UNDERLINED. NRU REPRESENTS “NOT REQUIRE UNBIASED MAR DATA”.

Model	Method	NRU	Coat			Yahoo! R3			KuaiRec		
			MSE	MAE	nDCG@5	MSE	MAE	nDCG@5	MSE	MAE	nDCG@5
MF	HEI	✓	1.3614	0.8927	0.7785	2.3874	1.2127	0.8013	2.3227	1.3705	0.3448
	Naive	✓	1.2851	0.8704	0.7850	2.3887	1.2131	0.8015	2.3817	1.3935	0.3326
	SNIPS-UI	✓	1.2342	0.8564	0.7846	1.9772	1.0519	0.7774	1.2084	0.8663	0.3255
	DR-UI	✓	1.2330	0.8486	0.7704	2.5912	1.2543	0.7813	2.0233	1.0457	0.3466
	AT	✓	1.1603	0.8294	0.7912	1.8389	0.9613	0.7973	1.2274	0.8752	0.3622
	CVIB	✓	1.2000	0.9025	0.7232	1.1270	0.8537	0.7321	1.1940	0.9733	0.3327
	ESCM2	✓	1.1368	0.8173	0.7976	1.1675	0.8154	0.8157	2.0414	1.2712	0.5403
	CounterCLR	✓	1.0956	0.8008	0.8002	1.1137	0.8117	0.8049	1.1586	0.8523	0.5893
	RDC	×	<u>1.0946</u>	0.8097	0.7930	1.3260	0.8521	0.8040	1.2733	0.9297	0.5941
	SNIPS-NB	×	1.1770	0.8367	0.7720	1.3314	0.8524	0.8039	1.3599	0.9692	0.4920
DR-NB	×	1.2254	0.8673	0.7617	1.1875	0.8278	0.8058	1.4308	0.9872	0.5484	
SDR-NB	×	1.1923	0.8645	0.7749	1.1599	<u>0.8103</u>	0.8125	1.2793	0.9053	<u>0.6102</u>	
NCF	HEI	✓	1.3557	0.9388	0.7321	2.1483	1.2070	0.7995	0.4239	0.5009	0.3576
	Naive	✓	1.4343	0.9583	0.7320	2.1967	1.2521	0.7963	0.4268	0.5279	0.4252
	SNIPS-UI	✓	1.3052	0.9228	0.7341	2.3317	1.2412	0.7930	0.5735	0.5978	0.3439
	DR-UI	✓	1.3998	0.9453	0.7347	2.7590	1.3169	0.7769	0.3517	0.4243	0.6093
	AT	✓	1.2641	0.9102	0.7402	2.2167	1.2346	0.7961	0.5030	0.5958	0.3183
	CVIB	✓	1.2197	0.8893	0.7265	1.2505	0.9795	0.7895	0.3978	0.4774	0.5298
	ESCM2	✓	1.2505	0.9105	0.7410	2.1919	1.2279	0.7999	0.4780	0.4973	0.5916
	CounterCLR	✓	1.1743	0.8733	0.7421	1.2203	0.8180	0.8040	0.3425	0.4129	0.6198
	RDC	×	1.2857	0.9147	0.7415	2.0037	1.2411	0.8006	0.3590	0.4389	0.6193
	SNIPS-NB	×	1.2353	0.8788	0.7376	1.0838	0.8311	0.7948	0.3949	0.4820	0.6050
DR-NB	×	1.1760	<u>0.8593</u>	0.7362	1.3213	0.8815	0.7913	<u>0.3302</u>	<u>0.3874</u>	<u>0.6257</u>	
SDR-NB	×	1.2330	0.8824	0.7367	<u>1.0821</u>	0.8195	0.8027	0.3865	0.4323	0.6180	

DR-UI methods, which is because that NB can provide a more accurate propensity score estimation since the additional 5% unbiased MAR test data. It verifies that the propensity-based methods rely on an accurate propensity score estimator to achieve desirable performance. The proposed CounterCLR, on the other hand, does not require additional unbiased MAR data, and thus is more robust and practical.

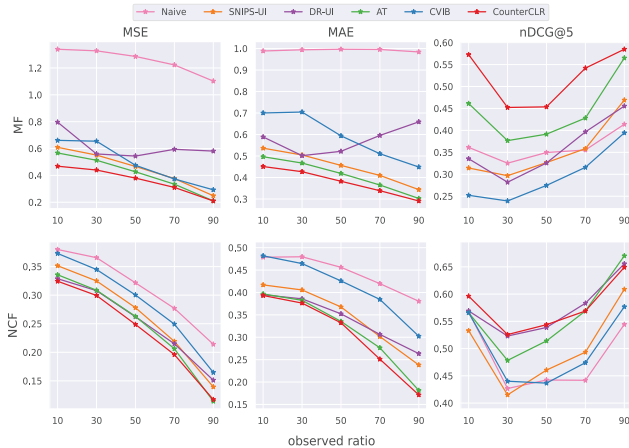


Fig. 2. Rating prediction accuracy and recommendation quality with varying observed ratio of the Small matrix in KuaiRec Dataset.

C. Influence of Data Sparsity (RQ2)

In order to evaluate the recommendation performance under data sparsity issue, we exploit the “full observed” *small*

matrix in **KuaiRec** to synthesize partially-observed data, with the observed ratio varying in $\{10\%, 30\%, 50\%, 70\%, 90\%\}$. Meanwhile, we follow the positive-oriented exposure strategy in [53] to simulate MNAR training data to model different selection bias level. The unobserved parts in *small matrix* are used as the test sets to evaluate the imputed missing data.

We compare our CounterCLR with the baselines without unbiased MAR data for a fair comparison. From Figure 2, we first can observe that our CounterCLR stably outperforms the baselines in both rating prediction accuracy and recommendation quality. This shows the superior generalization ability of CounterCLR under data sparsity issues. Besides, with the observed ratios increasing, the curves of our CounterCLR and all the baselines show a downward trend in terms of MSE and MAE, and an upward trend in terms of nDCG@5. This is because the higher observed ratio, the more training data will be accessed to train the recommender systems for providing more accurate recommendations.

VI. CONCLUSION

In this work, we propose a novel causality-based contrastive learning framework CounterCLR for debiased rating prediction, which consists of a causal network CauNet and a contrastive learning objective. The proposed method CounterCLR can effectively address the selection bias and data sparsity issues simultaneously without introducing separate imputation and propensity estimators, or unbiased MAR data. Extensive experiments on real-world datasets show that CounterCLR outperforms other state-of-the-art rating prediction methods.

REFERENCES

- [1] X. Wang, R. Zhang, Y. Sun, and J. Qi, "Doubly robust joint learning for recommendation on data missing not at random," in *ICML*, 2019.
- [2] H. Wang, Z. Chen, J. Fan, Y. Huang, W. Liu, and X. Liu, "Entire space counterfactual learning: Tuning, analytical properties and industrial applications," *arXiv preprint arXiv:2210.11039*, 2022.
- [3] J. Chen, H. Dong, Y. Qiu, X. He, X. Xin, L. Chen, G. Lin, and K. Yang, "AutoDebias: Learning to debias for recommendation," in *SIGIR*, 2021.
- [4] D. Liu, P. Cheng, Z. Lin, J. Luo, Z. Dong, X. He, W. Pan, and Z. Ming, "KDCRec: Knowledge distillation for counterfactual recommendation via uniform data," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [5] Y. Saito and M. Nomura, "Towards resolving propensity contradiction in offline recommender learning," in *IJCAI*, 2022.
- [6] H. Li, K. Wu, C. Zheng, Y. Xiao, H. Wang, F. Feng, X. He, Z. Geng, and P. Wu, "Removing hidden confounding in recommendation: a unified multi-task learning approach," in *NeurIPS*, 2023.
- [7] X. Ma, L. Zhao, G. Huang, Z. Wang, Z. Hu, X. Zhu, and K. Gai, "Entire space multi-task model: An effective approach for estimating post-click conversion rate," in *SIGIR*, 2018.
- [8] J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, and X. He, "Bias and debias in recommender system: A survey and future directions," *ACM Transactions on Information Systems*, 2022.
- [9] P. Wu, H. Li, Y. Deng, W. Hu, Q. Dai, Z. Dong, J. Sun, R. Zhang, and X.-H. Zhou, "On the opportunity of causal learning in recommendation systems: Foundation, estimation, prediction and challenges," in *IJCAI*, 2022.
- [10] D. B. Rubin, "Inference and missing data," *Biometrika*, vol. 63, no. 3, pp. 581–592, 1976.
- [11] C. Gao, Y. Zheng, W. Wang, F. Feng, X. He, and Y. Li, "Causal inference in recommender systems: A survey and future directions," *arXiv preprint arXiv:2208.12397*, 2022.
- [12] W. Wang, Y. Zhang, H. Li, P. Wu, F. Feng, and X. He, "Causal recommendation: Progresses and future directions," in *Tutorial on SIGIR*, 2023.
- [13] T. Schnabel, A. Swaminathan, A. Singh, N. Chandak, and T. Joachims, "Recommendations as treatments: Debiasing learning and evaluation," in *ICML*, 2016.
- [14] Y. Saito, H. Sakata, and K. Nakata, "Doubly robust prediction and evaluation methods improve uplift modeling for observational data," in *ICDM*, 2019.
- [15] S. Guo, L. Zou, Y. Liu, W. Ye, S. Cheng, S. Wang, H. Chen, D. Yin, and Y. Chang, "Enhanced doubly robust learning for debiasing post-click conversion rate estimation," in *SIGIR*, 2021.
- [16] Q. Dai, H. Li, P. Wu, Z. Dong, X.-H. Zhou, R. Zhang, X. He, R. Zhang, and J. Sun, "A generalized doubly robust learning framework for debiasing post-click conversion rate prediction," in *KDD*, 2022.
- [17] H. Li, C. Zheng, and P. Wu, "StableDR: Stabilized doubly robust learning for recommendation on data missing not at random," in *ICLR*, 2023.
- [18] H. Li, Y. Lyu, C. Zheng, and P. Wu, "TDR-CL: Targeted doubly robust collaborative learning for debiased recommendations," in *ICLR*, 2023.
- [19] H. Li, Y. Xiao, C. Zheng, P. Wu, and P. Cui, "Propensity matters: Measuring and enhancing balancing for recommendation," in *ICML*, 2023.
- [20] H. Li, Y. Xiao, C. Zheng, and P. Wu, "Balancing unobserved confounding with a few unbiased ratings in debiased recommendations," in *Proceedings of the ACM Web Conference 2023*, 2023, pp. 1305–1313.
- [21] H. Liu, D. Tang, J. Yang, X. Zhao, H. Liu, J. Tang, and Y. Cheng, "Rating distribution calibration for selection bias mitigation in recommendations," in *Proceedings of the ACM Web Conference*, 2022.
- [22] X. Ma, L. Zhao, G. Huang, Z. Wang, Z. Hu, X. Zhu, and K. Gai, "Entire space multi-task model: An effective approach for estimating post-click conversion rate," in *SIGIR*, 2018.
- [23] H. Wang, T.-W. Chang, T. Liu, J. Huang, Z. Chen, C. Yu, R. Li, and W. Chu, "ESCM2: Entire space counterfactual multi-task model for post-click conversion rate estimation," in *SIGIR*, 2022.
- [24] Z. Wang, Y. He, J. Liu, W. Zou, P. S. Yu, and P. Cui, "Invariant preference learning for general debiasing in recommendation," in *KDD*, 2022.
- [25] X. Du, Z. Wu, F. Feng, X. He, and J. Tang, "Invariant representation learning for multimedia recommendation," in *Proceedings of the 30th ACM International Conference on Multimedia*, 2022.
- [26] S. Huang, H. Li, Q. Li, C. Zheng, and L. Liu, "Pareto invariant representation learning for multimedia recommendation," in *Proceedings of the 31th ACM International Conference on Multimedia*, 2023.
- [27] H. Wang, Q. Dai, J. Fan, Z. Chen, H. Li, W. Liu, T. Liu, Y. Wang, Z. Dong, and R. Tang, "Optimal transport for causal effect estimation," in *NeurIPS*, 2023.
- [28] H. Wang, K. Kuang, L. Lan, Z. Wang, W. Huang, F. Wu, and W. Yang, "Out-of-distribution generalization with causal feature separation," *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [29] H. Wang, W. Yang, L. Yang, A. Wu, L. Xu, J. Ren, F. Wu, and K. Kuang, "Estimating individualized causal effect with confounded instruments," in *KDD*, 2022.
- [30] H. Wang, K. Kuang, H. Chi, L. Yang, M. Geng, W. Huang, and W. Yang, "Treatment effect estimation with adjustment feature selection," in *KDD*, 2023.
- [31] H. Zou, H. Wang, R. Xu, B. Li, J. Pei, Y. J. Jian, and P. Cui, "Factual observation based heterogeneity learning for counterfactual prediction," in *2nd Conference on Causal Learning and Reasoning*, 2023.
- [32] J. M. Hernández-Lobato, N. Houlsby, and Z. Ghahramani, "Probabilistic matrix factorization with non-random missing data," in *ICML*, 2014.
- [33] W. Wang, X. Lin, F. Feng, X. He, and T.-S. Chua, "Generative recommendation: Towards next-generation recommender paradigm," *arXiv preprint arXiv:2304.03516*, 2023.
- [34] H. Steck, "Training and testing of recommender systems on data missing not at random," in *KDD*, 2010.
- [35] H. Li, Q. Dai, Y. Li, Y. Lyu, Z. Dong, X.-H. Zhou, and P. Wu, "Multiple robust learning for recommendation," in *AAAI*, 2023.
- [36] H. Li, C. Zheng, P. Wu, K. Kuang, Y. Liu, and P. Cui, "Who should be given incentives? counterfactual optimal treatment regimes learning for recommendation," in *KDD*, 2023.
- [37] Z. Song, J. Chen, S. Zhou, Q. Shi, Y. Feng, C. Chen, and C. Wang, "CDR: Conservative doubly robust learning for debiased recommendation," in *CIKM*, 2023.
- [38] Y. Saito, "Asymmetric tri-training for debiasing missing-not-at-random explicit feedback," in *SIGIR*, 2020.
- [39] Z. Wang, X. Chen, R. Wen, S.-L. Huang, E. Kuruoglu, and Y. Zheng, "Information theoretic counterfactual learning from missing-not-at-random feedback," *NeurIPS*, 2020.
- [40] D. Liu, P. Cheng, H. Zhu, Z. Dong, X. He, W. Pan, and Z. Ming, "Mitigating confounding bias in recommendation via information bottleneck," in *RecSys*, 2021.
- [41] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*, 2020.
- [42] L. Kong, C. d. M. d'Autume, W. Ling, L. Yu, Z. Dai, and D. Yogatama, "A mutual information maximization perspective of language representation learning," *arXiv preprint arXiv:1910.08350*, 2019.
- [43] H. Liu, J. Z. HaoChen, A. Gaidon, and T. Ma, "Self-supervised learning is more robust to dataset imbalance," in *ICLR*, 2022.
- [44] X. Xie, F. Sun, Z. Liu, S. Wu, J. Gao, J. Zhang, B. Ding, and B. Cui, "Contrastive learning for sequential recommendation," in *ICDE*, 2022.
- [45] K. Zhou, H. Wang, W. X. Zhao, Y. Zhu, S. Wang, F. Zhang, Z. Wang, and J.-R. Wen, "S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization," in *CIKM*, 2020.
- [46] X. Cai, C. Huang, L. Xia, and X. Ren, "LightGCL: Simple yet effective graph contrastive learning for recommendation," in *ICLR*, 2023.
- [47] C. Zhou, J. Ma, J. Zhang, J. Zhou, and H. Yang, "Contrastive learning for debiased candidate generation in large-scale recommender systems," in *KDD*, 2021.
- [48] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.
- [49] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proceedings of the ACM Web Conference*, 2017.
- [50] Z. Yang, Y. Bai, and S. Mei, "Exact gap between generalization error and uniform convergence in random feature models," in *ICML*, 2021.
- [51] S. Mei and A. Montanari, "The generalization error of random features regression: Precise asymptotics and the double descent curve," *Communications on Pure and Applied Mathematics*, 2022.
- [52] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *CVPR*, 2020.
- [53] C. Gao, S. Li, W. Lei, J. Chen, B. Li, P. Jiang, X. He, J. Mao, and T.-S. Chua, "Kuairc: A fully-observed dataset and insights for evaluating recommender systems," in *CIKM*, 2022.