

Enhanced Doubly Robust Learning for Debiasing Post-Click Conversion Rate Estimation

Siyuan Guo^{1,*}, Lixin Zou^{2,†}, Yiding Liu², Wenwen Ye², Suqi Cheng², Shuaiqiang Wang²,
Hechang Chen^{1,4,†}, Dawei Yin² and Yi Chang^{1,3,4,†}

¹School of Artificial Intelligence, Jilin University, ²Baidu Inc., ³International Center of Future Science, Jilin University

⁴Key Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, Jilin University

{guosyjl, zoulixin15, liuyiding.tanh, yewenwen.hi, chengsuqi, shqiang.wang}@gmail.com

yindawei@acm.org, {chenhc, yichang}@jlu.edu.cn

ABSTRACT

Post-click conversion, as a strong signal indicating the user preference, is salutary for building recommender systems. However, accurately estimating the post-click conversion rate (CVR) is challenging due to the selection bias, i.e., the observed clicked events usually happen on users' preferred items. Currently, most existing methods utilize counterfactual learning to debias recommender systems. Among them, the doubly robust (DR) estimator has achieved competitive performance by combining the error imputation based (EIB) estimator and the inverse propensity score (IPS) estimator in a doubly robust way. However, inaccurate error imputation may result in its higher variance than the IPS estimator. Worse still, existing methods typically use simple model-agnostic methods to estimate the imputation error, which are not sufficient to approximate the dynamically changing model-correlated target (i.e., the gradient direction of the prediction model). To solve these problems, we first derive the bias and variance of the DR estimator. Based on it, a more robust doubly robust (MRDR) estimator has been proposed to further reduce its variance while retaining its double robustness. Moreover, we propose a novel double learning approach for the MRDR estimator, which can convert the error imputation into the general CVR estimation. Besides, we empirically verify that the proposed learning scheme can further eliminate the high variance problem of the imputation learning. To evaluate its effectiveness, extensive experiments are conducted on a semi-synthetic dataset and two real-world datasets. The results demonstrate the superiority of the proposed approach over the state-of-the-art methods. The code is available at <https://github.com/guosyjl/MRDR-DL>.

CCS CONCEPTS

• Information systems → Recommender systems.

*Work performed during an internship at Baidu Inc.

†Joint corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '21, July 11–15, 2021, Virtual Event, Canada.

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8037-9/21/07...\$15.00

<https://doi.org/10.1145/3404835.3462917>

KEYWORDS

Selection Bias; Missing-Not-At-Random Data; Doubly Robust; Post-click Conversion Rate Estimation; Recommender System

ACM Reference Format:

Siyuan Guo, Lixin Zou, Yiding Liu, Wenwen Ye, Suqi Cheng, Shuaiqiang Wang, Hechang Chen, Dawei Yin, Yi Chang. 2021. Enhanced Doubly Robust Learning for Debiasing Post-Click Conversion Rate Estimation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3404835.3462917>

1 INTRODUCTION

E-commerce recommender systems aim at not only helping users explore the items of their interests, but also increasing revenues for the platform. Therefore, estimating the post-click conversion rate (CVR), i.e., the probability of an item being purchased after it is clicked, is a crucial task for building such systems in practice. Moreover, post-click conversion feedbacks have been widely recognized as strong signals for the learning of the recommender systems, as they explicitly express the user preference and directly contribute to the gross merchandise volume (GMV) of the platform [23, 33]. However, it is very challenging to model such signals, which are extremely sparse in real-world applications. In particular, the post-click conversion feedbacks can only be observed in clicked events, which make up a tiny fraction of all possible user behaviors, while the other conversion feedbacks for unclicked events are missing. As such, a fundamental problem of CVR estimation is to study the missing mechanism in the post-click conversion feedbacks.

For simplification, conventional CVR models usually assume that the missing conversion feedbacks are *missing-at-random* (MAR). Such assumption can barely hold under the selection bias and recent studies [24, 25, 37] have shown that a recommendation model with MAR assumption often leads to sub-optimal results. On real-world e-commerce platforms, as users are free to click the items that they are likely to want to purchase (i.e., user self-selection), the observed clicked events are not representative samples of all the events, which makes the missing conversions *missing-not-at-random* (MNAR). In other words, the fundamental reason behind the selection bias is that the users' propensities vary from item to item. Here, the propensity is defined as the probability of an item being clicked by a user, i.e., the click-through rate (CTR). Hence, in this paper, we adopt the MNAR assumption when estimating the

post-click conversion rate, and focus on addressing the selection bias problem.

In recent years, three unbiased estimators in counterfactual learning have been applied to debiasing the CVR estimation. **(1) The error imputation based (EIB) estimator** [10, 26] computes an imputed error, i.e., the estimated value of the prediction error, for each unclicked event, and then uses it to estimate the true prediction error of all the events. However, this estimator often has a large bias due to the inaccurate error imputation, which will easily mislead the CVR estimation. **(2) The inverse propensity score (IPS) estimator** [1, 18, 25] inversely weights the prediction error of each clicked event with its estimated CTR to correct the mismatch between the distributions of the clicked events and unclicked events. Although this estimator is unbiased given the ground-truth CTRs, it typically suffers from a high variance problem, which would lead to sub-optimal results. **(3) The doubly robust (DR) estimator** [5, 30, 37] combines the EIB estimator and IPS estimator to ensure both the low variance and low bias. Its unbiasedness is guaranteed if either the imputed errors or the CTRs are accurate. This property is called the double robustness.

Among the aforementioned estimators, the DR estimator has achieved initial success for debiasing recommender systems [23, 30, 37]. However, there are still two inherent challenges to be solved. Despite the double robustness, the DR estimator may increase the variance of the IPS estimator under inaccurate error imputation, which makes the learning process even complicated and leads to sub-optimal results. Hence, further variance reduction for the DR approaches deserves to investigate. Furthermore, although the DR estimator is more robust than the EIB and IPS estimator, it still requires relatively accurate CTR estimation and error imputation. In terms of the two tasks, the former has been extensively investigated by a lot of works [7, 38], whereas the latter rarely investigated. To estimate the imputed errors, previous DR based approaches typically introduce an extra imputation model that is agnostic of the prediction model, such as linear regression [5], matrix factorization [30], multilayer perceptron (MLP) [37], etc. Here, the imputed errors, utilized as the gradient directions of the prediction model, should be dynamically changing during its learning process. However, simply using model-agnostic methods are not sufficient to approximate such a model-correlated target. Thus, it still calls for a better solution on how to model the error imputation.

To address the above-mentioned challenges, we propose the enhanced doubly robust learning approach for debiasing post-click conversion rate estimation. To tackle the first challenge, we propose to reduce the variance of the DR estimator, by redesigning the goal of the imputation learning (i.e., the learning process of the imputation model) as the minimization of its variance [3, 6]. Specifically, we derive the bias and variance of the DR estimator, based on which we propose the more robust doubly robust (MRDR) estimator as a variant of the DR estimator to derive lower variance while retaining the double robustness. Moreover, inspired by Double DQN [28] in reinforcement learning, we propose a novel double learning approach for the MRDR estimator to tackle the second limitation. In particular, we adopt two CVR models with same structure but different parameters. The first one serves as the prediction model to learn from both the imputed errors and the true prediction errors

for final CVR estimation. The second one serves as the imputation model to generate the pseudo label using its predicted CVR for each event. During the learning of the prediction model, the imputed error can be directly computed with the pseudo label and the predicted CVR. As such, we convert the error imputation into the general CVR estimation, and further, the imputed errors can be dynamically estimated in a model-correlated way. For the learning of both models, we alternate their learning process to enable them to be mutually regularized. In addition, we periodically update the parameters of the imputation model with the parameters of the prediction model, which is empirically beneficial for eliminating the high variance problem of the imputation learning. Extensive experiments are conducted on both semi-synthetic and real-world datasets to verify the effectiveness of both the proposed MRDR estimator and double learning approach.

The main contributions of this work are summarized as follows.

- We conduct theoretical analysis on the bias and variance of the DR estimator, based on which we propose the more robust doubly robust (MRDR) estimator. It can achieve further variance reduction while retaining the double robustness.
- To dynamically utilize the information of the prediction model for error imputation, we propose a novel double learning approach for the MRDR estimator, which is also empirically beneficial for addressing the high variance problem of the imputation learning.
- Experimental results on the semi-synthetic dataset empirically verify the effectiveness of the proposed MRDR estimator. Furthermore, we conduct extensive experiments on two real-world datasets. The results show that the proposed enhanced doubly robust learning approach MRDR-DL outperforms the state-of-the-art methods.

2 PRELIMINARIES

In this section, we detail the problem formulation, and introduce some existing unbiased estimators in the post-click conversion setting.

2.1 Problem Formulation

Let $\mathcal{U} = \{u_1, u_2, \dots, u_m\}$ be the set of m users, $\mathcal{I} = \{i_1, i_2, \dots, i_n\}$ be the set of n items, and $\mathcal{D} = \mathcal{U} \times \mathcal{I}$ be the set of all user-item pairs. We denote $\mathbf{R} \in \{0, 1\}^{m \times n}$ as the conversion label matrix where each entry $r_{u,i} \in \{0, 1\}$ indicates whether a conversion action occurs after user u clicks item i . We use $\hat{\mathbf{R}} \in \mathbb{R}^{m \times n}$ to represent the predicted conversion rate matrix, where $\hat{r}_{u,i} \in [0, 1]$ represents the conversion rate predicted by a model. If we have a fully observed conversion label matrix \mathbf{R} , the ideal loss function for minimization can be formulated as

$$\mathcal{L}_{ideal}(\hat{\mathbf{R}}) = \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} e_{u,i}, \quad (1)$$

where $e_{u,i}$ is the prediction error. We usually adopt the cross entropy, $e_{u,i} = CE(r_{u,i}, \hat{r}_{u,i}) = -r_{u,i} \log \hat{r}_{u,i} - (1 - r_{u,i}) \log(1 - \hat{r}_{u,i})$ as the optimization goal for binary classification. Let $\mathbf{O} \in \{0, 1\}^{m \times n}$ be the click indicator matrix with each entry $o_{u,i} = 1$ if user u clicks item i , and 0 otherwise. Since only post-click conversions for clicked events can be observed, the naive estimator estimates the ideal loss

function by averaging the prediction error for clicked events as

$$\begin{aligned}\mathcal{L}_{naive}(\hat{\mathbf{R}}) &= \frac{1}{|\mathcal{O}|} \sum_{(u,i) \in \mathcal{O}} e_{u,i} \\ &= \frac{1}{|\mathcal{O}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} e_{u,i},\end{aligned}\quad (2)$$

where $\mathcal{O} = \{(u,i) | (u,i) \in \mathcal{D}, o_{u,i} = 1\}$ denotes the clicked events. The naive estimator is intuitive and widely adopted by many existing methods. However, due to the selection bias, the conversions for unclicked events are MNAR, which leads to a biased estimation, i.e., $\mathbb{E}_{\mathcal{O}}[\mathcal{L}_{naive}(\hat{\mathbf{R}})] \neq \mathcal{L}_{ideal}(\hat{\mathbf{R}})$. Previous works [24, 25, 30] have proved that the learning process based on a biased estimator often leads to a sub-optimal prediction model. Hence, it is essential to develop an unbiased estimator to address the MNAR problem. In the following, we will introduce three existing unbiased estimators.

2.2 Error Imputation Based Estimator

The error imputation based (EIB) estimator introduces an imputation model to compute the imputed error $\hat{e}_{u,i}$, i.e., the estimated value of the prediction error [10, 26]. Leveraging the imputed errors for unclicked events and the prediction errors for clicked events, we estimate the ideal loss function with the EIB estimator as

$$\mathcal{L}_{EIB}(\hat{\mathbf{R}}) = \frac{1}{|\mathcal{O}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} e_{u,i} + (1 - o_{u,i}) \hat{e}_{u,i}. \quad (3)$$

When the imputed error $\hat{e}_{u,i}$ is accurate for any given unclicked event, the EIB estimator is unbiased, i.e., $\mathbb{E}_{\mathcal{O}}[\mathcal{L}_{EIB}(\hat{\mathbf{R}})] = \mathcal{L}_{ideal}(\hat{\mathbf{R}})$. However, the EIB estimator can hardly achieve accurate error imputation, and thus often has a large bias in practice, which would easily mislead the learning of the prediction model.

2.3 Inverse Propensity Score Estimator

The inverse propensity score (IPS) estimator [1, 24, 25] weights each clicked event with $1/p_{u,i}$, where the propensity $p_{u,i} = \mathbb{P}(o_{u,i} = 1) = \mathbb{E}[o_{u,i}]$ refers to the probability of the item being clicked by the user, i.e., the click-through rate (CTR) in the post-click conversion setting. By introducing an auxiliary CTR task to estimate the propensity $\hat{p}_{u,i}$, the IPS estimator can be formulated as

$$\mathcal{L}_{IPS}(\hat{\mathbf{R}}) = \frac{1}{|\mathcal{O}|} \sum_{(u,i) \in \mathcal{D}} \frac{o_{u,i} e_{u,i}}{\hat{p}_{u,i}}. \quad (4)$$

The IPS estimator derives an unbiased estimate of the ideal loss function, i.e., $\mathbb{E}_{\mathcal{O}}[\mathcal{L}_{IPS}(\hat{\mathbf{R}})] = \mathcal{L}_{ideal}(\hat{\mathbf{R}})$, when the estimated propensity $\hat{p}_{u,i}$ is accurate for any given clicked event. However, as the clicked events merely account for a small part of \mathcal{D} , the CTR is typically assigned with a small value. Hence, the IPS estimator suffers from an especially severe high variance problem.

2.4 Doubly Robust Estimator

To address the large bias problem of the EIB estimator and the high variance problem of the IPS estimator, the doubly robust (DR) estimator is adopted by many previous works [5, 30, 37]. It combines the EIB estimator and the IPS estimator in a doubly robust way. Particularly, this estimator uses the imputed errors $\hat{e}_{u,i}$ to estimate the prediction errors for all the events, and correct the error deviation $\delta_{u,i} = e_{u,i} - \hat{e}_{u,i}$ for the unclicked events. The propensity

$\hat{p}_{u,i}$ is inversely weighted to the error deviation for eliminating the MNAR effect. The loss function of the DR estimator can be defined as

$$\mathcal{L}_{DR}(\hat{\mathbf{R}}) = \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \hat{e}_{u,i} + \frac{o_{u,i}(e_{u,i} - \hat{e}_{u,i})}{\hat{p}_{u,i}}. \quad (5)$$

The DR estimator is unbiased, i.e., $\mathbb{E}_{\mathcal{O}}[\mathcal{L}_{DR}(\hat{\mathbf{R}})] = \mathcal{L}_{ideal}(\hat{\mathbf{R}})$, if either the imputed error $\hat{e}_{u,i}$ of any event or the propensity $\hat{p}_{u,i}$ of any clicked event is accurate. This property is recognized as double robustness. To compute the imputed errors, previous works typically introduce a separate imputation model. Since the imputation learning is actually a regression problem, DR uses the squared loss,

$$\mathcal{L}_e^{DR} = \sum_{(u,i) \in \mathcal{O}} \frac{(\hat{e}_{u,i} - e_{u,i})^2}{\hat{p}_{u,i}}, \quad (6)$$

to train the imputation model. The inverse propensity score is weighted to consider the MNAR effect, which also leads to the high variance problem of the imputation learning.

3 ENHANCED DOUBLY ROBUST LEARNING APPROACH

In this section, we elaborate the proposed enhanced doubly robust learning approach. We first analyze the bias and variance of the doubly robust estimator, based on which we propose the more robust doubly robust estimator for further variance reduction. Then, we detail the proposed novel double learning approach for the MRDR estimator.

3.1 Bias and Variance Analysis of DR Estimator

Initially, we formulate the bias of the DR estimator to prove its double robustness.

THEOREM 3.1. *Let $\delta_{u,i} = e_{u,i} - \hat{e}_{u,i}$ denote the additive error deviation, and $\Delta_{u,i} = 1 - \frac{p_{u,i}}{\hat{p}_{u,i}}$ the multiplicative propensity deviation. Then, the bias of the DR estimator is*

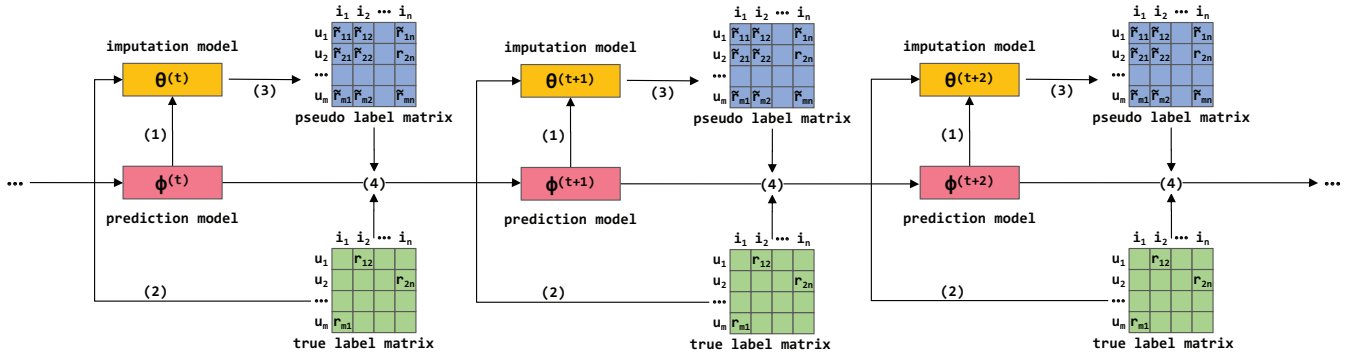
$$\text{Bias}[\mathcal{L}_{DR}(\hat{\mathbf{R}})] = \frac{1}{|\mathcal{D}|} \left| \sum_{(u,i) \in \mathcal{D}} \Delta_{u,i} \delta_{u,i} \right|. \quad (7)$$

PROOF. See Theorem 3.2 in [37] for the proof. \square

As shown in Theorem 3.1, the DR estimator is close to the ideal loss function, i.e., $\text{Bias}[\mathcal{L}_{DR}(\hat{\mathbf{R}})] \approx 0$, if either $\delta_{u,i} \approx 0$ or $\Delta_{u,i} \approx 0$, whereas the EIB estimator requires $\delta_{u,i} \approx 0$ and the IPS estimator requires $\Delta_{u,i} \approx 0$. This property is called double robustness. Then, we derive the variance of the DR estimator.

THEOREM 3.2. *The variance of the DR estimator is*

$$\mathbb{V}_{\mathcal{O}}[\mathcal{L}_{DR}(\hat{\mathbf{R}})] = \frac{1}{|\mathcal{D}|^2} \sum_{(u,i) \in \mathcal{D}} \frac{p_{u,i}(1-p_{u,i})}{\hat{p}_{u,i}^2} (e_{u,i} - \hat{e}_{u,i})^2. \quad (8)$$



- (1) Copy the parameters of the prediction model to update those of the imputation model;
- (2) Learn the imputation model based on the observed true conversion label matrix;
- (3) Use the imputation model to generate the pseudo label matrix;
- (4) Learn the prediction model with true conversion label matrix and the generated pseudo label matrix.

Figure 1: Workflow of the double learning approach.

PROOF. For a single term of the DR estimator, its variance on the click indicator $o_{u,i}$ is

$$\begin{aligned}
& \mathbb{V}_{o_{u,i}} \left[\hat{e}_{u,i} + \frac{o_{u,i}(e_{u,i} - \hat{e}_{u,i})}{\hat{p}_{u,i}} \right] \\
&= \mathbb{V}_{o_{u,i}} [o_{u,i}] \cdot \left(\frac{e_{u,i} - \hat{e}_{u,i}}{\hat{p}_{u,i}} \right)^2 \\
&= \left(\mathbb{E}_{o_{u,i}} [o_{u,i}^2] - (\mathbb{E}_{o_{u,i}} [o_{u,i}])^2 \right) \cdot \left(\frac{e_{u,i} - \hat{e}_{u,i}}{\hat{p}_{u,i}} \right)^2 \\
&= \mathbb{E}_{o_{u,i}} [o_{u,i}] (1 - \mathbb{E}_{o_{u,i}} [o_{u,i}]) \cdot \left(\frac{e_{u,i} - \hat{e}_{u,i}}{\hat{p}_{u,i}} \right)^2 \\
&= \mathbb{E}_{o_{u,i}} \left[o_{u,i} \frac{1 - p_{u,i}}{\hat{p}_{u,i}^2} (e_{u,i} - \hat{e}_{u,i})^2 \right]. \quad (9)
\end{aligned}$$

Then, summing across all terms of the DR estimator, we can derive the variance:

$$\begin{aligned}
\mathbb{V}_{\mathcal{O}}[\mathcal{L}_{DR}(\hat{\mathbf{R}})] &= \mathbb{V}_{\mathcal{O}} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \hat{e}_{u,i} + \frac{o_{u,i}(e_{u,i} - \hat{e}_{u,i})}{\hat{p}_{u,i}} \right] \\
&= \frac{1}{|\mathcal{D}|^2} \sum_{(u,i) \in \mathcal{D}} \mathbb{V}_{o_{u,i}} \left[\hat{e}_{u,i} + \frac{o_{u,i}(e_{u,i} - \hat{e}_{u,i})}{\hat{p}_{u,i}} \right] \quad (10) \\
&= \frac{1}{|\mathcal{D}|^2} \sum_{(u,i) \in \mathcal{D}} \frac{p_{u,i}(1 - p_{u,i})}{\hat{p}_{u,i}^2} (e_{u,i} - \hat{e}_{u,i})^2.
\end{aligned}$$

□

Similarly, we can derive the variance of the IPS estimator as

$$\mathbb{V}_{\mathcal{O}}[\mathcal{L}_{IPS}(\hat{\mathbf{R}})] = \frac{1}{|\mathcal{D}|^2} \sum_{(u,i) \in \mathcal{D}} \frac{p_{u,i}(1 - p_{u,i})}{\hat{p}_{u,i}^2} e_{u,i}^2. \quad (11)$$

Theorem 3.2 and Equation 11 illustrate that the variance of both estimators depends on the estimated propensity, i.e., the predicted CTR $\hat{p}_{u,i}$, which may lead to a high variance problem. However, it is worth noting that the DR estimator still reduces the variance of the IPS estimator, if any given event satisfies $0 \leq \hat{e}_{u,i} \leq 2e_{u,i}$.

3.2 More Robust Doubly Robust Estimator

The theoretical analysis demonstrates that despite the double robustness, the DR estimator suffers from the risk of increasing the variance of the IPS estimator under inaccurate error imputation. Hence, we propose a more robust doubly robust (MRDR) estimator for further variance reduction. Specifically, we propose to learn the imputation model of the DR estimator by minimizing its variance. In other words, it is a variation of the DR estimator, and the only difference is that its loss function for imputation learning is derived from minimizing DR's variance. This means that the proposed MRDR estimator not only retains the double robustness, but also derives a lower variance than the original DR estimator. Based on Equation 9, we take the expectation on \mathcal{D} and estimate $p_{u,i}$ with $\hat{p}_{u,i}$ to derive the loss function of the imputation learning in the MRDR estimator as

$$\begin{aligned}
\mathcal{L}_e^{MRDR} &= \sum_{(u,i) \in \mathcal{D}} o_{u,i} \cdot \frac{1 - \hat{p}_{u,i}}{\hat{p}_{u,i}^2} \cdot (\hat{e}_{u,i} - e_{u,i})^2 \\
&= \sum_{(u,i) \in \mathcal{O}} \frac{1 - \hat{p}_{u,i}}{\hat{p}_{u,i}} \cdot \frac{(\hat{e}_{u,i} - e_{u,i})^2}{\hat{p}_{u,i}}. \quad (12)
\end{aligned}$$

Comparing the loss function of imputation learning in MRDR with that in DR, we note that MRDR changes the weights from $1/\hat{p}_{u,i}$ to $(1 - \hat{p}_{u,i})/\hat{p}_{u,i}^2$, which has the property

$$\begin{cases} \frac{1}{\hat{p}_{u,i}} < \frac{1 - \hat{p}_{u,i}}{\hat{p}_{u,i}^2}, & \text{if } \hat{p}_{u,i} < \frac{1}{2} \\ \frac{1}{\hat{p}_{u,i}} > \frac{1 - \hat{p}_{u,i}}{\hat{p}_{u,i}^2}, & \text{if } \hat{p}_{u,i} > \frac{1}{2} \end{cases}. \quad (13)$$

As such, the MRDR estimator increases the penalty of the clicked events with low propensity, and decreases the penalty of the rest of the clicked events. In this way, the imputation model is learned better, which further enables MRDR to reduce the variance of the DR estimator.

3.3 Double Learning Approach

In this subsection, we detail the proposed double learning approach for the MRDR estimator. Given a vector $\mathbf{x}_{u,i}$ encoding all the features of user u and item i , previous works typically introduce two separate models: an imputation model $\hat{e}_{u,i} = f_{\theta}(\mathbf{x}_{u,i})$ estimates the imputed errors, and a prediction model $\hat{r}_{u,i} = g_{\phi}(\mathbf{x}_{u,i})$ learns from the imputed errors and true conversion labels to predict the CVR. Here, the imputation model is agnostic of the prediction model, and merely takes the user-item features $\mathbf{x}_{u,i}$ for error imputation. In other words, during the learning process of the prediction model, the imputed error cannot be dynamically estimated. From an optimization perspective, the imputation model plays the role of estimating the gradients for the learning of the prediction model. However, we argue that simply utilizing model-agnostic methods is not sufficient to approximate such a model-correlated target. To this end, we propose a novel double learning approach, which utilizes the pseudo-labelling technique to provide dynamically changing imputed errors for the prediction model. As such, the complicated error imputation is simplified as a general CVR estimation task. We show the workflow of the double learning approach in Figure 1.

Specifically, we introduce two models with the same structure but different parameters: the prediction model $g_{\phi}(\mathbf{x}_{u,i})$ and the imputation model $g_{\theta}(\mathbf{x}_{u,i})$. When we need to learn the prediction model, we first generate the pseudo label $\tilde{r}_{u,i}$ for each event based on the imputation model. Then, we estimate the imputed error by computing the cross entropy between the predicted conversion rate $\hat{r}_{u,i}$ and the pseudo label $\tilde{r}_{u,i}$, i.e., $\hat{e}_{u,i} = CE(\tilde{r}_{u,i}, \hat{r}_{u,i})$. In this way, the imputation model $g_{\theta}(\mathbf{x}_{u,i})$ is converted to a CVR estimation task; further to this, the original regression problem is converted to a binary classification problem. Therefore, we replace the squared loss term $(\hat{e}_{u,i} - e_{u,i})^2$ in Equation 12 with a cross-entropy term. The imputation learning process of the MRDR estimator is thereby redesigned as

$$\mathcal{L}_e^{MRDR}(\theta) = \sum_{(u,i) \in \mathcal{O}} \frac{1 - \hat{p}_{u,i}}{\hat{p}_{u,i}^2} \cdot CE(r_{u,i}, g_{\theta}(\mathbf{x}_{u,i})) + \lambda \|\theta\|_F^2, \quad (14)$$

where θ denotes all the parameters of the imputation model $g_{\theta}(\mathbf{x}_{u,i})$ and λ controls the L_2 regularization strength to prevent overfitting. Note that, although we change the original formulation of the loss function of the imputation model in MRDR, the idea of increasing the penalty of the low-propensity clicked events and decreasing the penalty of the rest is kept. Meanwhile, we formulate the learning of the prediction model as

$$\mathcal{L}_r^{MRDR}(\phi) = \sum_{(u,i) \in \mathcal{D}} \hat{e}_{u,i} + \frac{o_{u,i}(e_{u,i} - \hat{e}_{u,i})}{\hat{p}_{u,i}} + \mu \|\phi\|_F^2, \quad (15)$$

where ϕ , $e_{u,i} = CE(r_{u,i}, g_{\phi}(\mathbf{x}_{u,i}))$ and $\hat{e}_{u,i} = CE(g_{\theta}(\mathbf{x}_{u,i}), g_{\phi}(\mathbf{x}_{u,i}))$ denote all the parameters of the prediction model $g_{\phi}(\mathbf{x}_{u,i})$, the prediction error, and the imputed error, and μ controls the L_2 regularization strength to prevent overfitting.

Inspired by Double DQN [28], we redesign the learning approach of the both models. Generally, we alternate the learning process between the imputation model and the prediction model via mini-batch stochastic gradient descent. As such, two models regularize each other and jointly reach convergence. Since that the MRDR estimator merely enhances the inverse propensity weight of the

Algorithm 1: The Proposed Enhanced Doubly Robust Learning Approach, MRDR-DL

Input: $\mathcal{O}, \mathcal{D}, \hat{p}$
Output: ϕ

- 1 Initialize the parameters θ, ϕ
- 2 **while** stopping criteria is not satisfied **do**
- 3 $\theta := \phi$
- 4 **for** number of steps for training the imputation model **do**
- 5 Sample a batch of clicked events from \mathcal{O}
- 6 Update θ by descending along the gradient
 $\nabla_{\theta} \mathcal{L}_e^{MRDR}(\theta)$
- 7 **end**
- 8 Generate pseudo label $\tilde{r}_{u,i}$ for any event $\forall (u, i) \in \mathcal{D}$
- 9 **for** number of steps for training the prediction model **do**
- 10 Sample a batch of events from \mathcal{D}^1
- 11 Update ϕ by descending along the gradient
 $\nabla_{\phi} \mathcal{L}_r^{MRDR}(\phi)$
- 12 **end**
- 13 **end**

	ML 100K	Coat Shopping	Yahoo! R3
#users	943	290	15400
#items	1682	300	1000
#MNAR ratings	100000	6960	311704
#MAR ratings	0	4640	54000

Table 1: Statistic details of the datasets.

imputation learning into $(1 - \hat{p}_{u,i})/\hat{p}_{u,i}^2$, it suffers from the high variance of the imputation learning, which also happens in the DR estimator as mentioned in Section 2.4. Therefore, each time before learning the imputation model, we update its parameters with those of the prediction model, i.e., $\theta := \phi$. In this way, the imputation model will be periodically corrected, and the information that the enhanced inverse propensity weight brings is kept. We empirically demonstrate that such learning scheme is beneficial for eliminating the high variance problem of the imputation learning. We summarize the proposed enhanced doubly robust learning approach, named MRDR-DL, in Algorithm 1.

4 SEMI-SYNTHETIC EXPERIMENTS

Following previous works [23, 25, 30], we conduct semi-synthetic experiments to investigate the following research question (RQ).

RQ1. Does the MRDR estimator lead to more accurate loss estimation than other estimators?

4.1 Experimental Setup

4.1.1 Dataset and preprocessing. To compute the accuracy of the estimated loss, we need a fully observed conversion label matrix,

¹Due to the sparsity of the clicked events, we decrease the sample probability of the unclicked events in practice.

which is unavailable in real-world datasets. Thus, we create a semi-synthetic evaluation dataset using the MovieLens (ML) 100K² [9] dataset in order to allow us to conduct the semi-synthetic experiment. The statistical details of the dataset are presented in Table 1. We employ the following preprocessing procedures [23] to convert the explicit feedback setting to the post-click conversion setting, and derive a fully observed conversion label matrix and a click indicator matrix.

(1) Use matrix factorization [12] to complete the rating matrix, but the predicted ratings are unrealistically high for all user-item pairs. To match a more realistic rating distribution $[p_1, p_2, p_3, p_4, p_5]$ given in [19], we sort all the ratings in ascending order, assign a value of 1 to the bottom p_1 fraction of the matrix entries, assign a value of 2 to the next p_2 fraction, and so on.

(2) Transform the predicted ratings $R_{u,i} \in \{1, 2, 3, 4, 5\}$ into CTR $p_{u,i} \in (0, 1)$ with $p_{u,i} = p\alpha^{\min(4, 6-R_{u,i})}$, where p is set to 1 and α is set to 0.5 in our experiments.

(3) Transform the predicted ratings $R_{u,i}$ into true CVR $r_{u,i}^{true}$ by correspondingly replacing the rating $\{1, 2, 3, 4, 5\}$ with the conversion rate $\{0.1, 0.3, 0.5, 0.7, 0.9\}$. Note that we can only observe the binary conversion labels rather than the true values of the CVR in practice. Thus, we simply assign fixed values to them based on different predicted ratings.

(4) Sample the binary click indicator and conversion label with the Bernoulli sampling; that is,

$$o_{u,i} \sim \text{Bern}(p_{u,i}), r_{u,i} \sim \text{Bern}(r_{u,i}^{true}), \forall (u, i) \in \mathcal{D}, \quad (16)$$

where $\text{Bern}(\cdot)$ denotes the Bernoulli distribution. Thereafter, we can derive a fully-observed conversion label matrix \mathbf{R} and a click indicator matrix \mathbf{O} .

4.1.2 Experimental details. Given a predicted CVR matrix $\hat{\mathbf{R}}$, we can directly compute the ideal loss by averaging the prediction error of each entry between \mathbf{R} and $\hat{\mathbf{R}}$. In contrast, the estimators derive the estimated loss with partial entries in \mathbf{R} whose corresponding click indicators $o_{u,i}$ equal 1. To evaluate the performance of loss estimation, we use the following five predicted CVR matrices [25, 30] for comparison.

- **ONE:** The predicted conversion rate $\hat{r}_{u,i}$ is identical to the true CVR $r_{u,i}^{true}$, except that $|\{(u, i) | r_{u,i}^{true} = 0.9\}|$ randomly selected true CVR of 0.1 are flipped to 0.9.
- **THREE:** Same as ONE, but flipping true CVR of 0.3 instead.
- **FIVE:** Same as ONE, but flipping true CVR of 0.5 instead.
- **SKEW:** The predicted conversion rate is sampled from the Gaussian distribution $\mathcal{N}(\mu = r_{u,i}^{true}, \sigma = \frac{1-r_{u,i}^{true}}{2})$, and clipped to $[0.1, 0.9]$.
- **CRS:** The predicted conversion rate $\hat{r}_{u,i} = 0.1$ if the true CVR $r_{u,i}^{true} \leq 0.7$. Otherwise, $\hat{r}_{u,i} = 0.5$.

We compare the MRDR estimator with the naive, EIB, IPS, and DR estimators. We estimate the propensity as $\frac{1}{\hat{p}_{u,i}} = \frac{1-\beta}{p_{u,i}} + \frac{\beta}{p_e}$, where $p_e = \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i}$, and β is set to 0.5 to introduce noises. For EIB and DR, the imputed error is computed as $\hat{e}_{u,i} =$

	naive	EIB	IPS	DR	MRDR
ONE	0.0686	0.5427	0.0346	0.0131	0.0073
THREE	0.0792	0.5869	0.0401	0.0172	0.0047
FIVE	0.1023	0.6152	0.0515	0.0138	0.0119
SKEW	0.0255	0.3574	0.0124	0.0081	0.0013
CRS	0.1773	0.0610	0.0888	0.0551	0.0503

Table 2: RE of the five estimators compared to the ideal loss.

$CE(\frac{\sum_{(u,i) \in \mathcal{D}} r_{u,i}/\hat{p}_{u,i}}{\sum_{(u,i) \in \mathcal{D}} 1/\hat{p}_{u,i}}, \hat{r}_{u,i})$. For MRDR, we compute the imputed errors as $\hat{e}_{u,i} = CE(\frac{\sum_{(u,i) \in \mathcal{D}} (1-\hat{p}_{u,i})r_{u,i}/\hat{p}_{u,i}^2}{\sum_{(u,i) \in \mathcal{D}} 1/\hat{p}_{u,i}}, \hat{r}_{u,i})$.

4.1.3 Evaluation metric. We compare the performance of the five estimators by the relative error (RE) as

$$\text{RE}(\mathcal{L}_{estimator}) = \frac{|\mathcal{L}_{ideal}(\hat{\mathbf{R}}) - \mathcal{L}_{estimator}(\hat{\mathbf{R}})|}{\mathcal{L}_{ideal}(\hat{\mathbf{R}})}, \quad (17)$$

where $\mathcal{L}_{estimator}$ denotes the estimator to be compared. RE evaluates the accuracy of the estimated loss, and a smaller value of MRE means a higher accuracy.

4.2 Experiment Results (RQ1)

In Table 2, we report the averaged RE of the five estimators over 20 times of sampling with Equation 16. We can see that IPS, DR and MRDR estimators outperform the naive estimator in every setting. This is caused by the selection bias that we introduce by controlling the propensity $p_{u,i}$. In contrast, the EIB estimator derives the worst RE in four settings; this is mainly due to the large bias of the heuristic error imputation. Additionally, the DR estimator improves the performance of the IPS estimator by jointly considering the imputed errors and the estimated propensities. Over all the settings, the MRDR estimator achieves the best performance, which can be attributed to both the double robustness and the reduced variance. Overall, the results conclude that our proposed method can achieve more accurate loss estimation. Next, we further evaluate our method on the task of CVR estimation on real-world datasets.

5 REAL-WORLD EXPERIMENTS

In this section, we compare the proposed learning approach with other existing debiasing approaches using real-world datasets. We anticipate the experimental results to answer the following RQs.

RQ2. Does the proposed approach MRDR-DL lead to higher debiasing performance than existing approaches?

RQ3. What influence do the various designs have on the proposed approach MRDR-DL?

RQ4. How does the sample ratio of unclicked events to clicked events influence the performance of MRDR-DL?

RQ5. How does the proposed double learning approach work for both the imputation model and the prediction model?

5.1 Experimental Setup

5.1.1 Datasets and preprocessing. To evaluate the performance of the unbiased CVR estimation, we need an MAR test set. However, as stated in [37], we cannot force users to randomly click items in

²<https://grouplens.org/datasets/movielens/>

Datasets	Methods	DCG@K			Recall@K		
		K=2	K=4	K=6	K=2	K=4	K=6
Coat Shopping	naive	0.6694±0.0136	0.9432±0.0138	1.1321±0.0126	0.8054±0.0159	1.3903±0.0225	1.8991±0.0233
	IPS	0.7093±0.0232	0.9552±0.0223	1.1248±0.0217	0.8249±0.0298	1.3520±0.0353	1.8078±0.0399
	DR-JL	0.6771±0.0273	0.9266±0.0282	1.0962±0.0272	0.7949±0.0337	1.3286±0.0420	1.7849±0.0456
	MRDR-DL	0.7219±0.0211	0.9905±0.0204	1.1696±0.0217	0.8499±0.0265	1.4249±0.0321	1.9060±0.0430
Yahoo! R3	Naive	0.5469±0.0058	0.7466±0.0049	0.8714±0.0040	0.6479±0.0066	1.0745±0.0074	1.4098±0.0062
	IPS	0.5502±0.0018	0.7520±0.0018	0.8751±0.0014	0.6545±0.0021	1.0797±0.0025	1.4168±0.0025
	DR-JL	0.5310±0.0045	0.7273±0.0053	0.8512±0.0045	0.6292±0.0049	1.0495±0.0082	1.3822±0.0081
	MRDR-DL	0.5561±0.0058	0.7549±0.0023	0.8811±0.0036	0.6595±0.0074	1.0846±0.0054	1.4237±0.0059

Table 3: A comparison of the overall performance of MRDR-DL and competing methods on two real-world datasets.

order to generate unbiased data for CVR estimations. Previous work [23] simulates the unbiased CVR estimation setting by using the datasets with specific properties. First, the datasets need to contain explicit feedback, which can reveal ground-truth user preference information. Next, the datasets need to contain an additional MAR test set, where users are asked to rate randomly selected sets of items. This enables us to evaluate the performance of the unbiased CVR estimation. To the best of our knowledge, there are only two publicly available datasets that satisfy these requirements, i.e., Coat Shopping³ and Yahoo! R3⁴. The statistical details for both datasets are presented in Table 1.

For both the MNAR data and the MAR data of both datasets, we follow [23] and employ the following preprocessing procedure.

(1) We define the binary click indicator as $o_{u,i} = 1$ if the item i is rated by user u , and $o_{u,i} = 0$ otherwise.

(2) We define the binary conversion label as $r_{u,i} = 1$ if the item i is rated greater than or equal to 4 by user u , and $r_{u,i} = 0$ otherwise.

(3) We derive the post-click conversion dataset as $\{(u, i, r_{u,i}) | o_{u,i} = 1, \forall (u, i) \in \mathcal{D}\}$.

For both datasets, we randomly split the MNAR datasets into training (90%) and validation (10%) sets, while the MAR datasets are kept as the test sets. Following the previous works [23, 35], we filter out users who have no conversion records in the test set.

5.1.2 Baselines. We compare the proposed method with the following baselines:

- **Naive:** It simply uses the naive estimator as the loss function to estimate CVR.
- **IPS** [25]: It derives the IPS estimator as the loss function by estimating the CTR as the propensity score.
- **DR-JL** [30]: It utilizes the DR estimator by jointly learning the imputation model and prediction model.

Due to the high bias problem, the EIB estimator is widely recognized as a weak baseline [23, 25, 30], and thus is not included in our comparison. In our experiments, both the CTR and the CVR are estimated by factorization machine (FM) [20].

5.1.3 Experimental Protocols. We adopt the mini-batch Adam to optimize all the methods, with the default learning rate set at 0.001. We fix the mini-batch size to 1024 for both datasets. In terms of FM, the embedding size is fixed as 64. We tune the L_2 regularization

coefficient λ in the range of $\{1e^{-5}, 1e^{-4}, \dots, 1\}$. Note that, for DR based methods, we apply a grid search when tuning the L_2 regularization coefficient of the imputation model and the prediction model; also, the sample ratio for unclicked events to clicked events is tuned in the range of $\{2, 4, 6, 8\}$. For CTR estimation, we fix the negative sampling ratio to 4.

For all the methods, we first choose the best hyper-parameters based on the validation set. Then, we perform the early stopping strategy (which applies if the model performance does not improve for five epochs) and report the best test result from the best-performing model on the validation set.

We use recall and discounted cumulative gain (DCG) to evaluate the debiasing performance of all the methods. We calculate both metrics for each user in the test set and report the average score.

5.2 Overall Performance (RQ2)

Table 3 shows the overall performance in terms of DCG@K and Recall@K ($K \in \{2, 4, 6\}$) on two real-world datasets. To reduce the effect of randomness, we repeat the experiments 100 times for Coat Shopping and 20 times for Yahoo! R3, and then report the mean and standard deviation for each. The best results are highlighted in boldface. From the table, we can see that the debiasing methods, IPS and MRDR-DL, outperform Naive for both datasets, demonstrating the necessity of handling the selection bias in the CVR estimation. Meanwhile, we find that although DR-JL utilizes the unbiased DR estimator, it still gives the worst performance on both datasets. One possible explanation for this is that DR-JL is originally designed for debiasing explicit MNAR feedbacks; as such, its joint learning approach may not be applicable to CVR estimation. Overall, the proposed method MRDR-DL consistently outperforms other methods on both datasets, which verifies the effectiveness of both the proposed MRDR estimator and the double learning approach.

5.3 Ablation Study (RQ3)

To apply the DR estimator to the post-click conversion setting, the proposed method, MRDR-DL, has specific design features. In this section, we will analyze their respective impacts on the method’s performance via an ablation study. The experimental results for MRDR-DL and its three variants on two datasets are summarized in Table 4. The results that are better than MRDR-DL are highlighted in boldface. We detail the variants and analyze their respective effects as follows.

³<https://www.cs.cornell.edu/~schnabts/mnar>

⁴<http://webscope.sandbox.yahoo.com/>

Datasets	Methods	DCG@K					Recall@K				
		K=2	K=3	K=4	K=5	K=6	K=2	K=3	K=4	K=5	K=6
Coat Shopping	MRDR-DL	0.7219	0.8728	0.9905	1.0878	1.1695	0.8499	1.1518	1.4249	1.6765	1.9060
	DR-DL	0.7205	0.8670	0.9806	1.0778	1.1601	0.8438	1.1368	1.4004	1.6517	1.8827
	MRDR-JL	0.6948	0.8442	0.9613	1.0582	1.1442	0.8227	1.1215	1.3935	1.6439	1.8852
	MRDR-DL with SL	0.7255	0.8720	0.9871	1.0827	1.1651	0.8504	1.1434	1.4107	1.6580	1.8892
Yahoo! R3	MRDR-DL	0.5561	0.6694	0.7549	0.8234	0.8811	0.6595	0.8860	1.0846	1.2616	1.4237
	DR-DL	0.5463	0.6602	0.7459	0.8145	0.8714	0.6484	0.8762	1.0752	1.2525	1.4123
	MRDR-JL	0.5546	0.6668	0.7544	0.8221	0.8786	0.6584	0.8828	1.0862	1.2612	1.4199
	MRDR-DL with SL	0.5321	0.6439	0.7287	0.7963	0.8538	0.6298	0.8535	1.0503	1.2251	1.3863

Table 4: Ablation study of MRDR-DL on two real-world datasets.

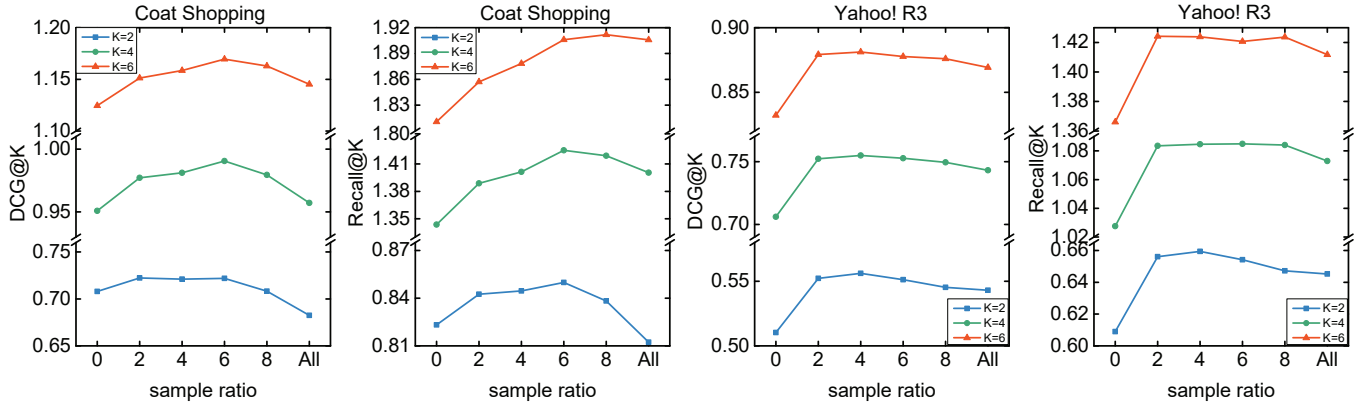


Figure 2: Effect of the sample ratio of unclicked events to clicked events. "All" means that we sample from all the events.

(1) DR-DL: We replace the MRDR estimator with the DR estimator, i.e., we change the weights of the imputation learning from $(1 - \hat{p}_{u,i})/\hat{p}_{u,i}^2$ to $1/\hat{p}_{u,i}$. The results imply that enhancing the weights to adjust the penalty for the clicked events based on varying propensities is conducive to the variance reduction of the DR estimator, and further improves the performance of the prediction model.

(2) MRDR-JL: We alternate the training of the imputation model and the prediction model without sharing the parameters periodically (i.e., we skip Step 3 of the Algorithm 1). The experiment results verify the necessity of periodically correcting the imputation model based on the prediction model, which is empirically beneficial for eliminating the high variance problem of the imputation learning.

(3) MRDR-DL with SL: We replace the cross-entropy term of the imputation learning with the squared loss term, which is theoretically derived from the variance of the DR estimator. The results of the variant are consistent with MRDR-DL on Coat Shopping and significantly better than MRDR-DL on Yahoo! R3. One reason is that the squared loss aims at minimizing the deviation between imputed errors and true prediction errors, whereas the pseudo-label generation is essentially a binary classification problem. Hence, it is more intuitive to have cross entropy as the optimization goal.

5.4 Parameter Sensitivity Study (RQ4)

By jointly considering both clicked and unclicked events, DR based estimators can enjoy the double robustness. To investigate the

impact of the unclicked events for the proposed MRDR-DL method, we vary the sample ratio of unclicked events to clicked events in the range of $\{0, 2, 4, 6, 8, \text{All}\}$. Here, "All" means that the sample ratio is set to the maximum possible value, which is 12.5 for Coat Shopping and 49.4 for Yahoo! R3. Figure 2 shows DCG@K and Recall@K for MRDR-DL with respect to different sample ratios on both datasets. As we can see, MRDR-DL with a sample ratio of 0 (i.e., we merely sample from the clicked events) derives the worst performance in most settings. This shows that the well-learned imputation model enables the unclicked events to provide the prediction model with useful information. Furthermore, we find that sampling from all the events adversely hurts the performance of the prediction model, even though we should in theory. One reason for this might be that clicked events are typically sparse in the real-world datasets, meaning that we cannot ensure that the prediction model obtains sufficient information. For both datasets, the optimal sample ratio is around 4 to 8. Setting the sample ratio too conservatively or too aggressively may adversely affect the prediction performance.

5.5 Analysis of the Double Learning Approach (RQ5)

In this subsection, we will further investigate the proposed double learning approach. We plot the training curves of the prediction model and the imputation model of MRDR-DL on Coat Shopping in Figure 3. In the proposed method MRDR-DL, the prediction

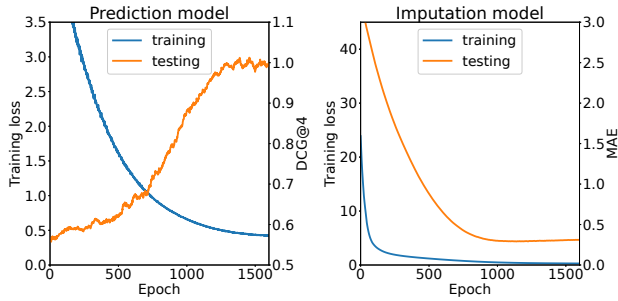


Figure 3: Training curves for the prediction model and the imputation model of MRDR-DL on Coat Shopping.

model aims at estimating CVR, while the imputation model aims at computing the imputed errors. Hence, we adopt DCG@4, and mean absolute error (MAE) for imputed errors and true prediction errors, respectively, to evaluate their testing performance. As shown in Figure 2, the training loss of the prediction model slightly fluctuates in the first 300 epochs before gradually reaching convergence. In contrast, the training curve of the imputation model is more stable. The reason for this is that the imputation model is not well-trained enough to provide the prediction model with sufficiently accurate information at the very beginning, whereas the imputation model is trained on clicked events with ground-truth labels. Further epochs of the double learning approach enable both models to exchange their information periodically. In this way, both models are jointly well-learned, reaching convergence together after about 900 epochs. Note that the testing curve of the prediction model fluctuates in the training process. This is reasonable because we train it with point-wise loss (i.e., cross entropy), whereas we evaluate it with a list-wise metric (i.e., DCG@4) when checking its debiasing performance.

6 RELATED WORK

6.1 General Approaches to CVR Estimation

CVR estimation is a key component of the recommender system because it directly contributes to the final revenue. Due to the inherent similarity, CVR estimation typically refers to the advances made by the CTR prediction task and implicit recommendation in practice, such as traditional models [11, 15], deep learning based models [7, 27, 36, 38] and reinforcement learning based models [39–41]. However, few studies directly investigate the CVR estimation tasks. Previous works often employ traditional models such as logistic regression [13, 21] and gradient boosting decision tree [16], while deep learning techniques like neural network [18, 32] and graph convolution network [1, 14] are also adopted for CVR estimation. However, the selection bias issue is still underexplored, which has a significant influence on improving performance in practice.

6.2 Counterfactual Learning from MNAR Data

Most data for learning the recommender systems are MNAR, which is caused by various biases, including selection bias, conformity bias, exposure bias, etc [4]. Previous works typically adopt counterfactual learning methods to address these issues. Specifically,

EIB methods [10, 26] rely on a missing data model to model the missing mechanism. IPS methods employ logistic regression [25], expectation-maximization algorithm [34], and matrix completion [17] to estimate the propensity for correcting the mismatch between observed and unobserved data. DR methods [30, 37] utilize an imputation model and a prediction model to jointly learn from the MNAR data. Other methods based on information bottleneck [31], meta learning [22], and causal embedding [2, 29] have been also explored to address these biases. Among above, IPS and DR has been widely applied to the recommender systems. However, how to specify appropriate error imputation and propensity estimation is a critical issue affecting their unbiasedness, which needs to be resolved in the post-click conversion setting.

6.3 Selection Bias in CVR Estimation

Selection bias is ubiquitous in recommender systems, especially in the CVR estimation task. A few studies have investigated it, achieving effective results. ESMM [18] models both the CTR and CVR tasks, using multi-task learning to eliminate the selection bias issue in a heuristic way. Similarly, ESM^2 [32], which is also essentially biased, extends ESMM by introducing additional auxiliary tasks. In contrast, GMCM [1] uses the IPS estimator to derive unbiased error evaluation when learning the CVR estimation task. In addition, Multi-IPW and Multi-DR [37] enjoy the unbiasedness of the IPS and DR estimator by learning both CTR and CVR tasks through multi-task learning. Although they consider the selection bias, the above approaches are evaluated using biased datasets; thus, their experimental results cannot be used to verify their debiasing performance, which is a widely-recognized limitation in practice. Furthermore, a recent work [23] proposes to utilize the DR estimator for debiasing ranking metric with post-click conversions, which mainly concerns the evaluation of the recommender systems. In contrast, we focus on debiasing the learning of the CVR estimation, and we use two real-world datasets containing unbiased data to evaluate the debiasing performance.

7 CONCLUSION AND FUTURE WORK

In this paper, we explore the problem of the selection bias in post-click conversion rate estimation. First, we analyze the bias and the variance of the DR estimator. Then, based on the theoretical analysis, we propose the more robust doubly robust estimator, which reduces the variance of the DR estimator while retaining the double robustness. Finally, we propose a novel double learning approach for MRDR estimator. It can dynamically utilize the information of the prediction model for the imputation model and empirically eliminate the high variance problem of the imputation learning. In the experiments, we verify the effectiveness of the proposed MRDR estimator on semi-synthetic datasets. In addition, we conduct extensive experiments on two real-world datasets to demonstrate the superiority of the proposed debiasing approach. For future work, we believe that the explainability [8] of the debiasing approach warrants further investigation.

ACKNOWLEDGMENTS

This work is supported by National Natural Science Foundation of China (No.61976102, No.U19A2065 and No.61902145).

REFERENCES

- [1] Wentian Bao, Hong Wen, Sha Li, Xiao-Yang Liu, Quan Lin, and Keping Yang. Gmcm: Graph-based micro-behavior conversion model for post-click conversion rate estimation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2201–2210, 2020.
- [2] Stephen Bonner and Flavian Vasile. Causal embeddings for recommendation. In *Proceedings of the 12th ACM conference on recommender systems*, pages 104–112, 2018.
- [3] Weihua Cao, Anastasios A Tsiatis, and Marie Davidian. Improving efficiency and robustness of the doubly robust estimator for a population mean with incomplete data. *Biometrika*, 96(3):723–734, 2009.
- [4] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. Bias and debias in recommender system: A survey and future directions. *arXiv preprint arXiv:2010.03240*, 2020.
- [5] Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, page 1097–1104, Madison, WI, USA, 2011.
- [6] Mehrdad Farajtabar, Yinlam Chow, and Mohammad Ghavamzadeh. More robust doubly robust off-policy evaluation. In *International Conference on Machine Learning*, pages 1447–1456, 2018.
- [7] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. Deepfm: a factorization-machine based neural network for ctr prediction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 1725–1731, 2017.
- [8] Siyuan Guo, Ying Wang, Hao Yuan, Zeyu Huang, Jianwei Chen, and Xin Wang. Taert: Triple-attentional explainable recommendation with temporal convolutional network. *Information Sciences*, 567:185–200, 2021.
- [9] F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
- [10] José Miguel Hernández-Lobato, Neil Houlsby, and Zoubin Ghahramani. Probabilistic matrix factorization with non-random missing data. In *International Conference on Machine Learning*, pages 1512–1520. PMLR, 2014.
- [11] Yuchin Juan, Yong Zhuang, Wei-Sheng Chin, and Chih-Jen Lin. Field-aware factorization machines for ctr prediction. In *Proceedings of the 10th ACM conference on recommender systems*, pages 43–50, 2016.
- [12] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [13] Kuang-chieh Lee, Burkay Orten, Ali Dasdan, and Wentong Li. Estimating conversion rate in display advertising from past performance data. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 768–776, 2012.
- [14] Yiding Liu, Yulong Gu, Zhuoye Ding, Junchao Gao, Ziyi Guo, Yongjun Bao, and Weipeng Yan. Decoupled graph convolution network for inferring substitutable and complementary items. *CIKM '20*, page 2621–2628, New York, NY, USA, 2020.
- [15] Yiding Liu, T. Pham, G. Cong, and Quan Yuan. An experimental evaluation of point-of-interest recommendation in location-based social networks. *Proc. VLDB Endow.*, 10:1010–1021, 2017.
- [16] Quan Lu, Shengjun Pan, Liang Wang, Junwei Pan, Fengdan Wan, and Hongxia Yang. A practical framework of conversion rate prediction for online display advertising. In *Proceedings of the ADKDD'17*, pages 1–9, 2017.
- [17] Wei Ma and George H Chen. Missing not at random in matrix completion: The effectiveness of estimating missingness probabilities under a low nuclear norm assumption. *Advances in Neural Information Processing Systems*, 32, 2019.
- [18] Xiao Ma, Liqin Zhao, Guan Huang, Zhi Wang, Zelin Hu, Xiaoqiang Zhu, and Kun Gai. Entire space multi-task model: An effective approach for estimating post-click conversion rate. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1137–1140, 2018.
- [19] Benjamin M Marlin and Richard S Zemel. Collaborative prediction and ranking with non-random missing data. In *Proceedings of the third ACM conference on Recommender systems*, pages 5–12, 2009.
- [20] Steffen Rendle. Factorization machines. In *2010 IEEE International Conference on Data Mining*, pages 995–1000. IEEE, 2010.
- [21] Rómer Rosales, Haibin Cheng, and Eren Manavoglu. Post-click conversion modeling and analysis for non-guaranteed delivery display advertising. In *Proceedings of the fifth ACM International Conference on Web Search and Data Mining*, pages 293–302, 2012.
- [22] Yuta 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*, pages 309–318, 2020.
- [23] Yuta Saito. Doubly robust estimator for ranking metrics with post-click conversions. In *Fourteenth ACM Conference on Recommender Systems*, pages 92–100, 2020.
- [24] Yuta Saito, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. Unbiased recommender learning from missing-not-at-random implicit feedback. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, pages 501–509, 2020.
- [25] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *International Conference on Machine Learning*, pages 1670–1679, 2016.
- [26] Harald Steck. Training and testing of recommender systems on data missing not at random. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 713–722, 2010.
- [27] Yi Tay, Shuai Zhang, Anh Tuan Luu, Siu Cheung Hui, Lina Yao, and Tran Dang Quang Vinh. Holographic factorization machines for recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5143–5150, 2019.
- [28] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- [29] Xiang Wang, Tinglin Huang, Dingxian Wang, Yancheng Yuan, Zhenguang Liu, Xiangnan He, and Tat-Seng Chua. Learning intents behind interactions with knowledge graph for recommendation. In *WWW*, 2021.
- [30] Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. Doubly robust joint learning for recommendation on data missing not at random. In *International Conference on Machine Learning*, pages 6638–6647, 2019.
- [31] Zifeng Wang, Xi Chen, Rui Wen, Shao-Lun Huang, Ercan E Kuruoglu, and Yefeng Zheng. Information theoretic counterfactual learning from missing-not-at-random feedback. In *Neural Information Processing Systems*, 2020.
- [32] Hong Wen, Jing Zhang, Yuan Wang, Fuyu Lv, Wentian Bao, Quan Lin, and Keping Yang. Entire space multi-task modeling via post-click behavior decomposition for conversion rate prediction. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2377–2386, 2020.
- [33] Liang Wu, Diane Hu, Liangjie Hong, and Huan Liu. Turning clicks into purchases: Revenue optimization for product search in e-commerce. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 365–374, 2018.
- [34] Xinwei Wu, Hechang Chen, Jiashu Zhao, Li He, Dawei Yin, and Yi Chang. Unbiased learning to rank in feeds recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 490–498, 2021.
- [35] Longqi Yang, Yin Cui, Yuan Xuan, Chenyang Wang, Serge Belongie, and Deborah Estrin. Unbiased offline recommender evaluation for missing-not-at-random implicit feedback. *RecSys '18*, page 279–287, New York, NY, USA, 2018.
- [36] Wenwen Ye, Shuaiqiang Wang, Xu Chen, Xuepeng Wang, Zheng Qin, and Dawei Yin. Time matters: Sequential recommendation with complex temporal information. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '20, page 1459–1468, 2020.
- [37] Wenhao Zhang, Wentian Bao, Xiao-Yang Liu, Keping Yang, Quan Lin, Hong Wen, and Ramin Ramezani. Large-scale causal approaches to debiasing post-click conversion rate estimation with multi-task learning. In *Proceedings of The Web Conference 2020*, pages 2775–2781, 2020.
- [38] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1059–1068, 2018.
- [39] Lixin Zou, Long Xia, Zhuoye Ding, Jiaying Song, Weidong Liu, and Dawei Yin. Reinforcement learning to optimize long-term user engagement in recommender systems. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '19, page 2810–2818, 2019.
- [40] Lixin Zou, Long Xia, Pan Du, Zhuo Zhang, Ting Bai, Weidong Liu, Jian-Yun Nie, and Dawei Yin. Pseudo dyna-q: A reinforcement learning framework for interactive recommendation. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, WSDM '20, page 816–824, 2020.
- [41] Lixin Zou, Long Xia, Pan Du, Zhuo Zhang, Ting Bai, Weidong Liu, Jian-Yun Nie, and Dawei Yin. Pseudo dyna-q: A reinforcement learning framework for interactive recommendation. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, WSDM '20, page 816–824, 2020.