

Towards Resolving Propensity Contradiction in Offline Recommender Learning

Yuta Saito¹, Masahiro Nomura²

¹Cornell University

²CyberAgent, Inc.

ys552@cornell.edu, nomura_masahiro@cyberagent.co.jp

Abstract

We study offline recommender learning from explicit rating feedback in the presence of *selection bias*. A current promising solution for the bias is the *inverse propensity score (IPS)* estimation. However, the performance of existing propensity-based methods can suffer significantly from the propensity estimation bias. In fact, most of the previous IPS-based methods require some amount of *missing-completely-at-random (MCAR)* data to accurately estimate the propensity. This leads to a critical **self-contradiction**: IPS is ineffective without MCAR data, even though it originally aims to learn recommenders from only *missing-not-at-random* feedback. To resolve this **propensity contradiction**, we derive a **propensity-independent generalization error bound** and propose a novel algorithm to minimize the theoretical bound via adversarial learning. Our theory and algorithm do not require a propensity estimation procedure, thereby leading to a well-performing rating predictor without the true propensity information. Extensive experiments demonstrate that the proposed approach is superior to a range of existing methods both in rating prediction and ranking metrics in practical settings without MCAR data.

1 Introduction

In industrial recommender systems, it is essential to obtain a well-performing rating predictor from sparse rating feedback to recommend items relevant to users. An important challenge is that most of the missing mechanisms of real-world explicit rating data are *missing-not-at-random (MNAR)* [Marlin and Zemel, 2009; Schnabel *et al.*, 2016; Wang *et al.*, 2019; Chen *et al.*, 2020]. There are two main causes behind the MNAR mechanism. The first is the past recommendation policy. Suppose that we rely on a policy that recommends popular items with high probabilities. Then, the observed ratings under that policy include more data of popular items [Bonner and Vasile, 2018]. The second is self-selection by users. It is well-known that the users tend to rate items for which they exhibit positive preferences, and the ratings for items with negative preferences

are more likely to be missing [Marlin and Zemel, 2009; Schnabel *et al.*, 2016].

1.1 Open Problem: Propensity Contradiction

Selection bias makes it difficult to learn rating predictors, as naive methods typically result in suboptimal and biased recommendations with MNAR data [Schnabel *et al.*, 2016; Steck, 2010; Wang *et al.*, 2019]. One of the most established solutions to this problem is the propensity-based approach. It defines the probability of each feedback being observed as a *propensity score* and obtains an unbiased estimator for the true metric via inverse propensity weighting [Liang *et al.*, 2016; Schnabel *et al.*, 2016; Wang *et al.*, 2019]. In general, unbiasedness of IPS is desirable, however, it is valid only when the true propensity score is available. Previous studies utilize some amount of *missing-completely-at-random (MCAR)* test data to estimate the propensity score and ensure their empirical performance [Schnabel *et al.*, 2016; Wang *et al.*, 2019]. In most real-world recommender systems, however, the true propensity score is unknown, and MCAR data are unavailable as well, resulting in a severe bias in the estimation of the loss function of interest. Thus, there remains a critical **contradiction**: the performance of the IPS variants relies heavily on mostly inaccessible MCAR data, although they originally aim at training recommenders using only MNAR data.

Two previous methods tackle the challenge of those propensity-based methods. The first one is *causal embeddings (CausE)* [Bonner and Vasile, 2018]. It introduces a domain regularization term to address the bias. However, CausE requires some amount of MCAR data by its design; it cannot be generalized to a realistic setting with only MNAR training data. Moreover, its domain regularization technique is a heuristic approach and lacks a theoretical guarantee. The other method is to use *1-bit matrix completion (1BitMC)* [Ma and Chen, 2019] to estimate propensity scores using only MNAR training data, along with a theoretical guarantee. However, the problem is that these methods presuppose the debiasing procedure with **inverse** propensity weighting and thus cannot be used when there is a user-item pair with zero observed probability. Furthermore, the experiments on 1BitMC were conducted using only extremely small datasets (Coat and MovieLens 100k) and prediction accuracy measures (MSE); accordingly, its performance on moderate size benchmark datasets (e.g., Yahoo! R3

[Mnih and Salakhutdinov, 2008]) and on a ranking task remain unknown.

1.2 Contributions

To overcome the limitations of the existing methods, we establish a new theory for offline recommender learning inspired by the theoretical framework of *unsupervised domain adaptation* (UDA) [Ben-David *et al.*, 2010]. UDA aims to obtain a good predictor in settings where the feature distributions between the training and test sets are different. To this end, UDA utilizes distance metrics that measure the dissimilarity between probability distributions and does not depend on the propensity score [Ganin *et al.*, 2016; Saito *et al.*, 2017; Saito, 2020a]. Thus, the framework is usable even when the true propensity score is unknown and is expected to alleviate the issues caused by the propensity estimation bias in the absence of MCAR data. Moreover, the method is valid even when there is a user-item pair with zero observed probability.

To figure out a solution to the propensity contradiction, we first define a novel discrepancy metric to quantify the similarity between two missing mechanisms of rating feedback. Building on our discrepancy measure, we derive a **propensity independent generalization error bound** for the loss function of interest. We further propose **domain adversarial matrix factorization** (DAMF), which minimizes the derived theoretical bound in an adversarial manner. Our theoretical bound and algorithm are independent of the propensity score and thus address the contradiction of the previous propensity-based solutions. Finally, we conduct extensive experiments using public real-world datasets. In particular, we demonstrate that the proposed approach outperforms existing propensity-based methods in terms of rating prediction and ranking performance under a realistic situation where the true propensity score is inaccessible.

2 Preliminaries

Let $u \in [m]$ be a user and $i \in [n]$ be an item in a recommender system. We also define $\mathcal{D} := [m] \times [n]$ as the set of all user-item pairs. Let $\mathbf{R} \in \mathbb{R}^{m \times n}$ denote a (deterministic) true rating matrix, where each entry $R_{u,i}$ represents the real-valued true rating of user u for item i .

Our goal is to develop an algorithm to obtain a better predicted rating matrix (or a *hypothesis*) $\hat{\mathbf{R}}$, where each entry $\hat{R}_{u,i}$ denotes a predicted rating value for (u, i) . To this end, we formally define **the ideal loss function of interest**, which should ideally be minimized to obtain a rating predictor as follows:

$$\mathcal{L}_{ideal}^\ell(\hat{\mathbf{R}}) := \frac{1}{mn} \sum_{(u,i) \in \mathcal{D}} \ell(R_{u,i}, \hat{R}_{u,i}), \quad (1)$$

where $\ell(\cdot, \cdot) : \mathbb{R} \times \mathbb{R} \rightarrow [0, \Delta]$ is a loss function bounded by Δ . For example, when $\ell(x, y) = (x - y)^2$, Eq. (1) represents the *mean-squared-error* (MSE).

In real-life recommender systems, one cannot directly calculate the ideal loss function, as most rating data are missing in nature. To precisely formulate this missing mechanism, we introduce two additional matrices. The first one is the

propensity matrix, denoted as $\mathbf{P} \in \mathcal{P}$, where \mathcal{P} represents the space of probability distributions over \mathcal{D} . Each of its entries $P_{u,i} \in [0, 1]$ is the *propensity score* of (u, i) , and it represents the probability of the rating feedback being observed. Next, let $\mathbf{O} \in \{0, 1\}^{m \times n}$ be an *observation matrix* where each entry $O_{u,i} \in \{0, 1\}$ is a Bernoulli random variable with its expectation $\mathbb{E}[O_{u,i}] = P_{u,i}$. If $O_{u,i} = 1$, the rating of the pair is observed, otherwise, it is unobserved. Finally, we use $\mathbf{O} \sim \mathbf{P}$ when the entries of \mathbf{O} are the realizations of Bernoulli distributions defined by the entries of \mathbf{P} . For simplicity and without loss of generality, we assume that $M := \sum_{(u,i) \in \mathcal{D}} O_{u,i}$ hereinafter.

2.1 Existing Methods

In our formulation, it is essential to approximate the ideal loss function using only observable feedback. Here, we summarize the existing methods to estimate the ideal loss function and discuss their limitations.

Naive Estimator. Given a set of observed rating feedback, the most basic estimator for the ideal loss is the naive estimator, which is defined as follows:

$$\hat{\mathcal{L}}_{naive}^\ell(\hat{\mathbf{R}}; \mathbf{O}) := \frac{1}{M} \sum_{(u,i) \in \mathcal{D}} O_{u,i} \cdot \ell(R_{u,i}, \hat{R}_{u,i}). \quad (2)$$

The naive estimator simply averages the loss of user-item pairs with the observed rating feedback (i.e., $O_{u,i} = 1$). This is valid when the missing mechanism is MCAR¹, as it is unbiased against the ideal loss function with MCAR data [Schnabel *et al.*, 2016]. However, several previous studies indicate that this estimator exhibits a bias under general MNAR settings (i.e., $\mathbb{E}[\hat{\mathcal{L}}_{naive}^\ell] \neq \mathcal{L}_{ideal}^\ell$ for some $\hat{\mathbf{R}}$). This means that the naive estimator does not converge to the ideal loss, even with infinite data. Therefore, one should use an estimator that addresses the bias as an alternative to the naive one [Schnabel *et al.*, 2016].

Inverse Propensity Score (IPS) Estimator. To improve the naive estimator, several previous studies apply the IPS estimation to alleviate the bias of MNAR rating feedback [Liang *et al.*, 2016; Schnabel *et al.*, 2016; Saito *et al.*, 2020c; Saito, 2020c; Saito *et al.*, 2021]. In causal inference, propensity scoring estimators are widely used to estimate the causal effects of a treatment from observational data [Imbens and Rubin, 2015]. In our formulation, we can derive an unbiased estimator for the loss function of interest by using the true propensity score as follows:

$$\hat{\mathcal{L}}_{IPS}^\ell(\hat{\mathbf{R}}; \mathbf{O}) := \frac{1}{mn} \sum_{(u,i) \in \mathcal{D}} O_{u,i} \cdot \frac{\ell(R_{u,i}, \hat{R}_{u,i})}{P_{u,i}}. \quad (3)$$

This estimator is unbiased for the ideal loss (i.e., $\mathbb{E}[\hat{\mathcal{L}}_{IPS}^\ell] = \mathcal{L}_{ideal}^\ell$ for any $\hat{\mathbf{R}}$) and is more desirable than the naive estimator in terms of bias. However, its unbiasedness is ensured only when the true propensity score is available, and it can have

¹A missing mechanism is said to be MCAR if the propensity score is constant, i.e., $P_{u,i} = C, \forall (u, i) \in \mathcal{D}$. Ensuring that the data is MCAR is extremely difficult, because the users' self-selection as to which items they provide rating feedback is out of our control.

a bias with an inaccurate propensity estimator (see Lemma 5.1 of Schnabel et al. [Schnabel et al., 2016]). The bias of IPS can arise in most real-world recommender systems, because the missing mechanism of rating feedback can depend on user self-selection, which cannot be controlled by analysts and is difficult to estimate [Marlin and Zemel, 2009; Schnabel et al., 2016]. Indeed, most previous studies estimate the propensity score by using some amount of MCAR test data to ensure empirical performance [Schnabel et al., 2016; Wang et al., 2019]. However, this is infeasible owing to the costly annotation process [Joachims et al., 2017]. In addition, this method is valid only when the ratings for some user-item pairs are observed with non-zero probabilities (i.e., $P_{u,i} \in (0, 1], \forall (u, i) \in \mathcal{D}$), which is a criterion difficult to verify in practice [Ma and Chen, 2019]. Therefore, we explore a theory and algorithm independent of the propensity score and MCAR data, aiming to alleviate the concurrent issues of the propensity-based methods.

Doubly Robust (DR) Estimator. The DR estimator utilizes the error imputation model and propensity score to decrease the variance of the IPS approach. MF-DR [Wang et al., 2019] optimizes the following DR estimator of the ideal loss.

$$\begin{aligned} & \hat{\mathcal{L}}_{DR}^{\ell}(\hat{\mathbf{R}}; \mathbf{O}) \\ & := \frac{1}{mn} \sum_{(u,i) \in \mathcal{D}} \left\{ \hat{\ell}_{u,i} + O_{u,i} \frac{\ell(R_{u,i}, \hat{R}_{u,i}) - \hat{\ell}_{u,i}}{P_{u,i}} \right\}, \quad (4) \end{aligned}$$

where $\hat{\ell}_{u,i}$ is the *imputation model*, which estimates $\ell(R_{u,i}, \hat{R}_{u,i})$. As discussed in [Wang et al., 2019], this estimator satisfies the unbiasedness with the true propensity (i.e., $\mathbb{E}[\hat{\mathcal{L}}_{DR}^{\ell}] = \mathcal{L}_{ideal}^{\ell}$) and often achieves a tighter estimation error bound compared to IPS. MF-DR obtains the final prediction ($\hat{\mathbf{R}}$) and the imputation model simultaneously via a joint learning procedure.

However, the proposed joint learning algorithm still requires a pre-estimated propensity score [Wang et al., 2019]. Furthermore, the estimation performance of the DR estimator is significantly degraded when the error imputation model and the propensity model are both misspecified [Dudík et al., 2011]. In fact, in the empirical evaluations of [Wang et al., 2019], MCAR test data are used to estimate the propensity score. In the experiments, we validate the performance of the DR estimator in a more practical setting where only MNAR data are available (without any MCAR data).

Causal Embeddings (CausE) The work most closely related to ours is [Bonner and Vasile, 2018], which proposed CausE, a domain adaptation-inspired method to address the selection bias in rating feedback. The loss function of CausE is given as follows:

$$\begin{aligned} & \hat{\mathcal{L}}_{CausE}^{\ell}(\hat{\mathbf{R}}_{MCAR}, \hat{\mathbf{R}}_{MNAR}; \mathbf{O}_{MCAR}, \mathbf{O}_{MNAR}) \\ & := \hat{\mathcal{L}}_{naive}^{\ell}(\hat{\mathbf{R}}_{MCAR}; \mathbf{O}_{MCAR}) \\ & \quad + \hat{\mathcal{L}}_{naive}^{\ell}(\hat{\mathbf{R}}_{MNAR}; \mathbf{O}_{MNAR}) \\ & \quad + \beta \cdot \Omega_{domain}(\hat{\mathbf{R}}_{MCAR}, \hat{\mathbf{R}}_{MNAR}), \quad (5) \end{aligned}$$

where $\beta \geq 0$ is a trade-off hyper-parameter. Two prediction matrices ($\hat{\mathbf{R}}_{MCAR}$ and $\hat{\mathbf{R}}_{MNAR}$) are for the MCAR and

MNAR data, respectively, where $\hat{\mathbf{R}}_{MNAR}$ is used to make the final predictions. The last term represents the *regularizer between tasks* and penalizes the divergence between the predictions for the MCAR and MNAR data.

This method empirically outperforms propensity-based methods in a binary rating feedback setting [Bonner and Vasile, 2018]. A problem is that this algorithm requires some MCAR data, which are generally unavailable. Moreover, it uses the idea of domain adaptation in a heuristic manner by myopically adding $\Omega_{domain}(\cdot)$; there is no theoretical guarantee for its loss function. Therefore, we explore an algorithm that is preferable to CausE in the following ways: (i) Our algorithm does not use any MCAR data during training, and thus is feasible in a realistic situation having no MCAR training data, (ii) Our proposed method is theoretically refined in the sense that it minimizes the propensity-independent upper bound of the ideal loss function.

3 Proposed Framework and Algorithm

This section derives a propensity-independent generalization error bound for the ideal loss and an algorithm to minimize the bound using only MNAR data.

3.1 Theoretical Bound

First, we define a discrepancy measure to quantify the similarity between two different propensity matrices.

Definition 1. (*Propensity Matrix Divergence (PMD)*) Let \mathcal{H} be a set of real-valued predicted rating matrices and $\hat{\mathbf{R}} \in \mathcal{H}$ be a specific prediction. The PMD between any two given propensity matrices \mathbf{P} and \mathbf{P}' is defined as follows:

$$\psi_{\hat{\mathbf{R}}, \mathcal{H}}(\mathbf{P}, \mathbf{P}') := \sup_{\hat{\mathbf{R}}' \in \mathcal{H}} \{ \mathcal{L}^{\ell}(\hat{\mathbf{R}}, \hat{\mathbf{R}}'; \mathbf{P}) - \mathcal{L}^{\ell}(\hat{\mathbf{R}}, \hat{\mathbf{R}}'; \mathbf{P}') \},$$

where $\mathcal{L}^{\ell}(\hat{\mathbf{R}}, \hat{\mathbf{R}}'; \mathbf{P}) := \mathbb{E}_{\mathbf{O} \sim \mathbf{P}}[\hat{\mathcal{L}}_{naive}^{\ell}(\hat{\mathbf{R}}, \hat{\mathbf{R}}'; \mathbf{O})] = \frac{1}{M} \sum_{(u,i) \in \mathcal{D}} P_{u,i} \cdot \ell(\hat{R}_{u,i}, \hat{R}'_{u,i})$.

Note that the PMD is well defined because $\psi_{\hat{\mathbf{R}}, \mathcal{H}}(\mathbf{P}, \mathbf{P}) = 0, \forall \mathbf{P} \in \mathcal{P}$, and it satisfies nonnegativity and subadditivity. Moreover, it is independent of the true rating matrices and can be calculated for any given pair of propensity matrices without true rating information.

We now use $\psi_{\hat{\mathbf{R}}, \mathcal{H}}$ and derive a propensity-independent generalization error bound for the ideal loss function.²

Theorem 1. (*Propensity Independent Generalization Error Bound*) Suppose two observation matrices with MCAR and MNAR mechanisms ($\mathbf{O}_{MCAR} \sim \mathbf{P}_{MCAR}$ and $\mathbf{O}_{MNAR} \sim \mathbf{P}_{MNAR}$) are given. For any prediction matrix $\hat{\mathbf{R}} \in \mathcal{H}$ and

²Note that the theoretical bounds of all existing studies depend on the cardinality of the hypothesis space $|\mathcal{H}|$ [Schnabel et al., 2016; Wang et al., 2019]. This is unrealistic, as $|\mathcal{H}|$ is infinite in almost all cases (as the parameter space is continuous). Under a realistic setting with infinite \mathcal{H} , therefore, the existing analysis becomes invalid. Thus, we use the Rademacher complexity to consider a realistic situation with infinite hypothesis spaces.

Algorithm 1 Domain Adversarial Matrix Factorization

Input: MNAR observation matrix \mathbf{O}_{MNAR} ; set of user-item pairs \mathcal{D} ; trade-off parameter β ; regularization parameter λ ; batch size B ; number of steps k

Output: learnt prediction matrix $\hat{\mathbf{R}} = \mathbf{U}\mathbf{V}^\top$

- 1: Randomly initialize \mathbf{U} , \mathbf{V} , \mathbf{U}' , and \mathbf{V}'
 - 2: **repeat**
 - 3: Sample mini-batch data of size B from \mathbf{O}_{MNAR}
 - 4: **for** $1, \dots, k$ **do**
 - 5: Update \mathbf{U} and \mathbf{V} by gradient descent according to Eq. (9) with fixed $\hat{\mathbf{R}}^* = \mathbf{U}'(\mathbf{V}')^\top$
 - 6: **end for**
 - 7: Uniformly sample user-item pairs of size T from \mathcal{D} to construct \mathbf{O}_{MCAR}
 - 8: **for** $1, \dots, k$ **do**
 - 9: Update \mathbf{U}' and \mathbf{V}' by gradient ascent according to Eq. (8) with fixed $\hat{\mathbf{R}} = \mathbf{U}\mathbf{V}^\top$
 - 10: **end for**
 - 11: **until** convergence;
-

for any $\delta \in (0, 1)$, the following inequality holds with a probability of at least $1 - \delta$:

$$\begin{aligned} & \mathcal{L}_{ideal}^\ell(\hat{\mathbf{R}}) \\ & \leq \underbrace{\hat{\mathcal{L}}_{naive}^\ell(\hat{\mathbf{R}}; \mathbf{O}_{MNAR})}_{\text{naive loss on MNAR training data}} + \underbrace{\psi_{\hat{\mathbf{R}}, \mathcal{H}}(\mathbf{O}_{MCAR}, \mathbf{O}_{MNAR})}_{\text{discrepancy}} \\ & \quad + 2L(3\mathfrak{R}_{\mathcal{P}, M}(\mathcal{H}) + 2\mathfrak{R}_{\mathcal{P}', M}(\mathcal{H})) + 3\Delta \sqrt{\frac{\log(6/\delta)}{2M}}, \end{aligned} \quad (6)$$

where L is the Lipschitz constant of the loss function and $\mathfrak{R}_{\mathcal{P}, M}(\mathcal{H})$ is the Rademacher complexity of \mathcal{H} , which is defined in Appendix A. See also Appendix A for the proof.

The theoretical bound consists of the following four factors: (i) Naive loss on MNAR data, (ii) Empirical PMD, (iii) Complexity of the hypothesis class, (iv) Confidence term that depends on δ . Note that (i) and (ii) can be optimized with only observable data, and (iii) can be controlled by adding a regularization term, as described in the next section. Moreover (iv) converges to zero as M increases. We will empirically show that the bound is informative in the sense that optimizing (i) and (ii) results in a desired value of the ideal loss function.

3.2 Algorithm

Here, we describe the proposed algorithm. Building on the theoretical bound derived in Theorem 1, we aim at minimizing the following objective:

$$\begin{aligned} & \min_{\hat{\mathbf{R}} \in \mathcal{H}} \underbrace{\hat{\mathcal{L}}_{naive}^\ell(\hat{\mathbf{R}}; \mathbf{O}_{MNAR})}_{\text{naive loss on MNAR training data}} \\ & \quad + \underbrace{\beta \cdot \psi_{\hat{\mathbf{R}}, \mathcal{H}}(\mathbf{O}_{MCAR}, \mathbf{O}_{MNAR})}_{\text{discrepancy}} + \underbrace{\lambda \cdot \Omega(\hat{\mathbf{R}})}_{\text{regularization}}, \end{aligned} \quad (7)$$

where $\beta \geq 0$ is the trade-off hyper-parameter between the naive loss and discrepancy measure. $\Omega(\cdot)$ is a regularization function that penalizes the complexity of $\hat{\mathbf{R}}$, and $\lambda \geq 0$ is

the regularization hyper-parameter. In addition, we can obtain \mathbf{O}_{MCAR} by uniformly sampling unlabeled user-item pairs from \mathcal{D} . The objective in Eq. (7) builds on the two controllable terms of the theoretical bound in Eq. (6). Thus, it is independent of the propensity score, and we need *not* estimate the propensity score to optimize this objective. Note that we add the regularization term to our objective to control the complexity of the hypothesis class \mathcal{H} , as it is impossible to minimize the Rademacher complexity directly.

To solve Eq. (7), we empirically approximate $\psi_{\hat{\mathbf{R}}, \mathcal{H}}$ by solving

$$\max_{\hat{\mathbf{R}} \in \mathcal{H}} \hat{\mathcal{L}}_{naive}^\ell(\hat{\mathbf{R}}, \hat{\mathbf{R}}'; \mathbf{O}_{MCAR}) - \hat{\mathcal{L}}_{naive}^\ell(\hat{\mathbf{R}}, \hat{\mathbf{R}}'; \mathbf{O}_{MNAR}). \quad (8)$$

This optimization corresponds to accurately estimating PMD from observable data.

Given the aforementioned details, Eq. (7) is now reduced to:

$$\begin{aligned} & \min_{\hat{\mathbf{R}} \in \mathcal{H}} \frac{1}{M} \underbrace{\sum_{(u,i): O_{u,i}=1} \ell(R_{u,i}, \hat{R}_{u,i})}_{\text{empirical loss on MNAR training data}} \\ & \quad + \beta \cdot \underbrace{\{\hat{\mathcal{L}}^\ell(\hat{\mathbf{R}}, \hat{\mathbf{R}}^*; \mathbf{O}_{MCAR}) - \hat{\mathcal{L}}^\ell(\hat{\mathbf{R}}, \hat{\mathbf{R}}^*; \mathbf{O}_{MNAR})\}}_{\text{approximated PMD between MCAR and MNAR}} \\ & \quad + \lambda \cdot \underbrace{\Omega(\hat{\mathbf{R}})}_{\text{regularization}}, \end{aligned} \quad (9)$$

where $\hat{\mathbf{R}}^*$ is the solution of Eq. (8).

Algorithm 1 describes the detailed procedure of our DAMF, which is based on the matrix factorization (MF) model, where a prediction matrix is defined as the product of user-item latent factors $\mathbf{U} \in \mathbb{R}^{m \times d}$ and $\mathbf{V} \in \mathbb{R}^{n \times d}$. Given that the predictions are modelled by MF, a possible regularization is $\Omega(\hat{\mathbf{R}}) := \|\mathbf{U}\|_2^2 + \|\mathbf{V}\|_2^2$, where $\|\cdot\|_2$ is the L2 norm.

4 Experiments

This section empirically evaluates the proposed method on public real-world datasets. Appendix B describes detailed experimental settings such as dataset description, evaluation metrics, and hyper-parameter tuning. Our code is available at <https://github.com/usaito/ijcai2022-adversarial-mf>.

4.1 Experimental Setups

Datasets. We use Yahoo! R3³ and Coat⁴, which separately contain MNAR and MCAR datasets. We randomly selected 10% of the MNAR training set as the validation set, which is used to perform hyper-parameter tuning.

Baseline Methods and Propensity Estimators. We compare our proposed algorithm with the following baselines:

- **MF** [Koren *et al.*, 2009], which optimizes its model parameters by minimizing the naive loss in Eq. (2) and does not depend on the propensity score.

³<http://webscope.sandbox.yahoo.com/>

⁴<https://www.cs.cornell.edu/~schnabts/mnar/>

Datasets	Methods	MSE (\pm StdDev)
Yahoo! R3	MF	1.7343 (\pm 0.0309)
	MF-IPS	1.7320 (\pm 0.0311)
	MF-DR	1.7445 (\pm 0.0189)
	DAMF (ours)	1.2787 (\pm 0.0126)
	CausE	1.7390 (\pm 0.0283)
	MF-IPS (true)	1.1281 (\pm 0.0161)
Coat	MF-DR (true)	1.0435 (\pm 0.0166)
	MF	1.2166 (\pm 0.0007)
	MF-IPS	1.2044 (\pm 0.0012)
	MF-DR	1.2108 (\pm 0.0016)
	DAMF (ours)	1.1371 (\pm 0.0025)
	CausE	1.2801 (\pm 0.0043)
MF-IPS (true)	1.0675 (\pm 0.0043)	
MF-DR (true)	1.0760 (\pm 0.0037)	

Table 1: Rating Prediction Performance

Note: The bold font indicates the best performance in each metric and dataset among methods using only MNAR data.

- **MF-IPS** [Schnabel *et al.*, 2016], which optimizes its model parameters by minimizing the IPS loss in Eq. (3).
- **MF-DR** [Wang *et al.*, 2019], which optimizes its model parameters by minimizing the DR loss in Eq. (4).
- **CausE** [Bonner and Vasile, 2018], which minimizes the sum of the naive loss and the domain regularization term $\Omega_{domain}(\cdot)$, which measures the divergence between the MNAR and MCAR data. To calculate its loss function, we sample 5% of the MCAR test data. MCAR data are in general unavailable in real-world settings, and thus we report the results of CausE as a reference.

For MF-IPS and MF-DR, we use *user propensity*, *item propensity*, *user-item propensity*, and *1BitMC* [Ma and Chen, 2019] as the variants of the propensity estimator and report the results with the best estimator for each dataset. Note that these four propensity estimators (formally defined in the Appendix) are usable in real-world recommender systems because they use only MNAR data.

We also report the results of MF-IPS and MF-DR with the *true propensity score*. We follow previous studies [Schnabel *et al.*, 2016; Wang *et al.*, 2019] and calculate this true propensity using 5% of the MCAR test data. Note that the true propensity is unavailable in real-world situations, as it requires MCAR explicit feedback. Therefore, we report the results of MF-IPS and MF-DR with the true propensity score as a reference.

4.2 Results and Discussion

Here, we report and discuss the experimental results with respect to four research questions (RQ1-RQ4). We run the experiments 10 times with different random seeds and report the averaged results in Tables 1-4.

RQ1. How well does the proposed algorithm perform on the rating prediction task? We evaluate and compare the prediction performance of DAMF and the baselines on Yahoo! R3 and Coat. Table 3.2 provides the MSEs and their standard

deviations (StdDev). For Yahoo! R3, DAMF reveals the best rating prediction performance among the methods using only MNAR data. Specifically, it outperforms MF-IPS by 26.3% and MF-DR by 26.7% in terms of MSE. For Coat, DAMF outperforms MF-IPS by 5.59% and MF-DR by 6.08%, even though the distributional shift of Coat is significantly smaller than that of Yahoo! R3. Note that it is reasonable to assume that there exists a user-item pair with zero observed probability on Coat. This is because the training set was collected via workers’ self-selection, and male workers did not provide the ratings of women’s coats and vice versa. Thus, the results also suggest the stability and adaptability of the proposed method to the data with a user-item pair with $P_{u,i} = 0$. In contrast, the performance of the propensity-based methods on Coat is not theoretically grounded, as the training and test sets may not overlap.

RQ2. How do propensity-based methods perform with different propensity estimators? Here, we evaluate the sensitivity of MF-IPS and MF-DR to the choice of the propensity estimator. First, consistent with previous studies [Schnabel *et al.*, 2016; Wang *et al.*, 2019], Table 3.2 shows that MF-IPS and MF-DR with the true propensity exhibit better rating prediction performance compared to the other methods. However, Table 3.2 shows that they perform poorly with other propensity estimators, including *1BitMC*, and especially for Yahoo! R3. In some cases, they even underperform MF, which is based on the naive loss. Therefore, although propensity-based methods are potentially high-performing with the true propensity, their prediction performances are highly sensitive to the choice of propensity estimator and negatively affected by the propensity estimation bias. These empirical observations imply the *propensity contradiction* of the previous methods.

RQ3. How well does the proposed algorithm perform on the ranking task? Next, we test the ranking performance of DAMF on Yahoo! R3 and Coat. Tables 3.2 and 3.2 describe the ranking metrics computed on the MCAR test data.

For Yahoo! R3, DAMF demonstrates the best ranking performance among all methods, including the methods with the true propensity score. It outperforms MF-IPS by 3.39% and MF-DR by 4.15%, MF-IPS (true) by 6.04%, and MF-DR (true) by 3.61% in NDCG. The results suggest that DAMF is powerful and useful for improving the recommendation quality and user experience with only biased rating data. Moreover, DAMF is the best in Coat, improving MF-IPS by 3.48%, MF-DR by 1.94%, MF-IPS (true) by 1.10%, and MF-DR (true) by 1.32% in NDCG. These results demonstrate that the proposed method works satisfactorily in the ranking task even when the dataset size is small and the degree of bias is not significant.

RQ4 Does adversarial learning really minimize the ideal loss? Finally, we investigate the validity of minimizing the propensity-independent upper bound in Eq. (6) as a solution to minimize the ideal loss function in Eq. (1) using only MNAR data. Figure 1 depicts the sum of the first two terms of the upper bound (blue) and ideal loss (orange) during the training of DAMF on Yahoo! R3 and Coat. First, it is evident from the figures that our adversarial learning procedure effectively minimizes the upper bound of the ideal loss during training. Moreover, the figures suggest that the ideal loss function is

Datasets	Methods	<i>user propensity</i>	<i>item propensity</i>	<i>user-item propensity</i>	<i>1BitMC</i>
Yahoo! R3	MF-IPS	1.8216 (+61.5%)	1.8083 (+60.3%)	1.8714 (+65.9%)	1.7343 (+53.8%)
	MF-DR	1.7445 (+67.2%)	1.9074 (+82.8%)	1.8815 (+80.4%)	1.8953 (+81.6%)
Coat	MF-IPS	1.2171 (+14.0%)	1.2161 (+13.9%)	1.2044 (+12.8%)	1.2087 (+13.2%)
	MF-DR	1.2141 (+12.8%)	1.2108 (+12.5%)	1.2181 (+13.2%)	1.2195 (+13.3%)

Table 2: Rating Prediction Performance of MF-IPS and MF-DR with Different Propensity Estimators

Note: Performances relative to the true propensity score are in parentheses. The result suggests that the performances of MF-IPS and MF-DR using only MNAR data are significantly worse than the performances of those using the true propensity.

Methods	NDCG@5 (\pm StdDev)	Recall@5 (\pm StdDev)
MF	0.7645 (\pm 0.0008)	0.5616 (\pm 0.0006)
MF-IPS	0.7645 (\pm 0.0008)	0.5616 (\pm 0.0006)
MF-DR	0.7589 (\pm 0.0007)	0.5602 (\pm 0.0005)
DAMF (ours)	0.7904 (\pm 0.0011)	0.5737 (\pm 0.0004)
CausE	0.7645 (\pm 0.0007)	0.5617 (\pm 0.0006)
MF-IPS (true)	0.7454 (\pm 0.0021)	0.5565 (\pm 0.0011)
MF-DR (true)	0.7629 (\pm 0.0032)	0.5624 (\pm 0.0015)

Table 3: Ranking Performance on **Yahoo! R3**

Methods	NDCG@5 (\pm StdDev)	Recall@5 (\pm StdDev)
MF	0.6617 (\pm 0.0029)	0.3858 (\pm 0.0008)
MF-IPS	0.6717 (\pm 0.0026)	0.3844 (\pm 0.001)
MF-DR	0.6789 (\pm 0.0018)	0.3887 (\pm 0.0016)
DAMF (ours)	0.6847 (\pm 0.0034)	0.3901 (\pm 0.0011)
CausE	0.6332 (\pm 0.0039)	0.3752 (\pm 0.0013)
MF-IPS (true)	0.6772 (\pm 0.002)	0.3906 (\pm 0.0009)
MF-DR (true)	0.6758 (\pm 0.0014)	0.3900 (\pm 0.0012)

Table 4: Ranking Performance on **Coat**


Note: The bold font indicates the best performance in each metric and dataset among methods using only MNAR data.

also minimized well by our method. These results demonstrate that minimizing the theoretical upper bound in Eq. (6) is a valid approach for improving the recommendation quality on the MCAR test set.


5 Conclusion

We explored the problem of offline recommender learning from *only* MNAR explicit rating feedback. To this end, we derived the *propensity independent generalization error bound* of the loss function and proposed an algorithm to minimize the bound via adversarial learning. Our algorithm enables an accurate rating prediction without estimating the propensity score, thereby resolving the *propensity contradiction* of the existing literature. Extensive experiments demonstrate the advantages of the proposed framework in terms of rating prediction and ranking measures when true propensities are inaccessible.

Our work represents an initial attempt for offline recommender learning via adversarial training, and there are still many problems to study in this direction. For example, our theory leads to a propensity-independent upper bound of the (ideal) rating prediction loss. It would be valuable to construct a corresponding theory for the ranking task, such as showing the propensity-independent upper bound of the additive



(a) Yahoo! R3



(b) Coat

Figure 1: Theoretical Upper Bound and Ideal Loss Function during DAMF Training.

ranking metrics. Moreover, testing our method in an online experiment will support the reliability of our framework. We leave the investigation of these directions for future studies.

A Omitted Proofs

Throughout this section, we use the following *Rademacher complexity* [Bartlett and Mendelson, 2002; Mohri *et al.*, 2018], which captures the complexity of a class of functions by measuring its ability to correlate with random noise [Kuroki *et al.*, 2019].

Definition 2. (*Rademacher complexity*) Let \mathcal{H} be any set of real-valued matrices. Given *i.i.d* samples with observed ratings $\{(u, i, R_{u,i}) \mid O_{u,i} = 1, (u, i) \in \mathcal{D}\}$, the Rademacher

complexity of \mathcal{H} is

$$\mathfrak{R}_{\mathbf{P},M}(\mathcal{H}) := \mathbb{E}_{\mathbf{O} \sim \mathbf{P}} \mathbb{E}_{\boldsymbol{\xi}} \left[\sup_{\widehat{\mathbf{R}} \in \mathcal{H}} \frac{1}{M} \sum_{(u,i): O_{u,i}=1} \xi_{u,i} \widehat{R}_{u,i} \right],$$

where $\boldsymbol{\xi} = (\xi_1, \dots, \xi_M)$ denotes a set of independent uniform random variables taking values in $\{+1, -1\}$.

A.1 Uniform Derivation Bound

Lemma 1. (Rademacher Generalization Bound; A modified version of Theorem 3.3 in [Mohri et al., 2018]) Let $\mathcal{F} = \{f : \mathcal{D} \rightarrow [0, \Delta]\}$ be a class of bounded functions where $\Delta > 0$ is a positive constant and $\{(u, i, R_{u,i}) \mid O_{u,i} = 1, (u, i) \in \mathcal{D}\}$ be any i.i.d. sample drawn from \mathbf{P} of size M . Then, for any $\delta \in (0, 1)$, the following inequality holds with probability of at least $1 - \delta$

$$\begin{aligned} \sup_{f \in \mathcal{F}} \left| \underbrace{\frac{1}{M} \sum_{(u,i): O_{u,i}=1} f(u, i)}_{(a)} - \underbrace{\frac{1}{mn} \sum_{(u,i) \in \mathcal{D}} P_{u,i} \cdot f(u, i)}_{(b)} \right| \\ \leq 2\mathfrak{R}_{\mathbf{P},M}(\mathcal{F}) + \Delta \sqrt{\frac{\log(2/\delta)}{2M}}, \end{aligned} \quad (10)$$

where (a) is an empirical mean of a function (f), and (b) is its expectation over \mathbf{P} .

A.2 Deviation Bound of PMD

The true propensity matrices (\mathbf{P} and \mathbf{P}') are unknown, and it is necessary to estimate the PMD using their realizations (\mathbf{O} and \mathbf{O}'). The following lemma shows the deviation bound of PMD.

Lemma 2. Suppose any pair of propensity matrices (\mathbf{P} and \mathbf{P}') and their realizations (\mathbf{O} and \mathbf{O}') are given. Accordingly, for any $\delta \in (0, 1)$ and \mathcal{H} , the following inequality holds with a probability of at least $1 - \delta$:

$$\begin{aligned} \left| \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{P}, \mathbf{P}') - \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{O}, \mathbf{O}') \right| \\ \leq 4L (\mathfrak{R}_{\mathbf{P},M}(\mathcal{H}) + \mathfrak{R}_{\mathbf{P}',M}(\mathcal{H})) + \Delta \sqrt{\frac{2 \log(4/\delta)}{M}}. \end{aligned}$$

Proof. For any given real-valued prediction matrix $\widehat{\mathbf{R}}$, we

have

$$\begin{aligned} & \left| \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{P}, \mathbf{P}') - \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{O}, \mathbf{O}') \right| \\ &= \left| \sup_{\widehat{\mathbf{R}}' \in \mathcal{H}} \{ \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}') \} \right. \\ & \quad \left. - \sup_{\widehat{\mathbf{R}}' \in \mathcal{H}} \{ \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}') \} \right| \\ &\leq \sup_{\widehat{\mathbf{R}}' \in \mathcal{H}} \left| \{ \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}') \} \right. \\ & \quad \left. - \{ \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}') \} \right| \\ &= \sup_{\widehat{\mathbf{R}}' \in \mathcal{H}} \left| \{ \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}) \} \right. \\ & \quad \left. - \{ \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}') - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}') \} \right| \\ &= \sup_{\widehat{\mathbf{R}}, \widehat{\mathbf{R}}' \in \mathcal{H}} \left| \{ \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}) \} \right. \\ & \quad \left. - \{ \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}') - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}') \} \right| \\ &\leq \sup_{\widehat{\mathbf{R}}, \widehat{\mathbf{R}}' \in \mathcal{H}} \left| \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}) \right| \\ & \quad + \sup_{\widehat{\mathbf{R}}, \widehat{\mathbf{R}}' \in \mathcal{H}} \left| \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}') - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}') \right|. \end{aligned} \quad (11)$$

The deviations in the last line can be bounded by using Lemma 1, and the following inequalities hold with a probability of at least $1 - \delta/2$.

$$\begin{aligned} \sup_{\widehat{\mathbf{R}}' \in \mathcal{H}} \left| \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}) \right| \\ \leq 2\mathfrak{R}_{\mathbf{P},M}(\mathcal{H}') + \Delta \sqrt{\frac{\log(4/\delta)}{2M}}, \end{aligned} \quad (12)$$

$$\begin{aligned} \sup_{\widehat{\mathbf{R}}' \in \mathcal{H}} \left| \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}') - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}') \right| \\ \leq 2\mathfrak{R}_{\mathbf{P}',M}(\mathcal{H}') + \Delta \sqrt{\frac{\log(4/\delta)}{2M}}. \end{aligned} \quad (13)$$

where we regard $\mathcal{H}' := \{(u, i) \rightarrow \ell(\widehat{R}_{u,i}, \widehat{R}'_{u,i}) \mid \widehat{\mathbf{R}}, \widehat{\mathbf{R}}' \in \mathcal{H}\}$ as \mathcal{F} in Lemma 1. Then, we have

$$\begin{aligned} \sup_{\widehat{\mathbf{R}}, \widehat{\mathbf{R}}' \in \mathcal{H}} \left| \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}) \right| \\ \leq 4L \mathfrak{R}_{\mathbf{P},M}(\mathcal{H}) + \Delta \sqrt{\frac{\log(4/\delta)}{2M}}, \end{aligned} \quad (14)$$

$$\begin{aligned} \sup_{\widehat{\mathbf{R}}, \widehat{\mathbf{R}}' \in \mathcal{H}} \left| \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{P}') - \mathcal{L}^\ell(\widehat{\mathbf{R}}, \widehat{\mathbf{R}}'; \mathbf{O}') \right| \\ \leq 4L \mathfrak{R}_{\mathbf{P},M}(\mathcal{H}) + \Delta \sqrt{\frac{\log(4/\delta)}{2M}}. \end{aligned} \quad (15)$$

where $\mathfrak{R}_{\mathbf{P},M}(\mathcal{H}') \leq 2L \mathfrak{R}_{\mathbf{P},M}(\mathcal{H})$ by using the result of Corollary 5 of [Mansour et al., 2009]. Finally, combining Eq. (11), Eq. (14), and Eq. (15) with the union bound completes the proof. \square

A.3 Proof of Additional Lemmas

Here, we state the generalization error bound when using the naive loss.

Lemma 3. (Generalization Error Bound) *An MCAR-observation matrix $\mathbf{O} \sim \mathbf{P}$ and any matrix as predictions $\widehat{\mathbf{R}} \in \mathcal{H}$ are given. Then, for any $\delta \in (0, 1)$, the following inequality holds with a probability of at least $1 - \delta$:*

$$\mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}) \leq \widehat{\mathcal{L}}_{naive}^\ell(\widehat{\mathbf{R}}; \mathbf{O}) + 2L\mathfrak{R}_{\mathbf{P}, M}(\mathcal{H}) + \Delta \sqrt{\frac{\log(2/\delta)}{2M}}. \quad (16)$$

Proof. By using Lemma 1, we have

$$\sup_{\widehat{\mathbf{R}} \in \mathcal{H}} \left| \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}) - \widehat{\mathcal{L}}_{naive}^\ell(\widehat{\mathbf{R}}; \mathbf{O}) \right| \leq 2\mathfrak{R}_{\mathbf{P}, M}(\ell \circ \mathcal{H}) + \Delta \sqrt{\frac{\log(2/\delta)}{2M}}.$$

with a probability of at least $1 - \delta$ for any $\delta \in (0, 1)$. We regard $\ell \circ \mathcal{H} = \{\mathcal{D} \rightarrow \ell(R_{u,i}, \widehat{R}_{u,i}) \mid \widehat{\mathbf{R}} \in \mathcal{H}\}$ as \mathcal{F} in Lemma 1. Then, by using the Talagrand's lemma (Lemma 5.7 of [Mohri *et al.*, 2018]), we have

$$\mathfrak{R}_{\mathbf{P}, M}(\ell \circ \mathcal{H}) \leq L\mathfrak{R}_{\mathbf{P}, M}(\mathcal{H}).$$

as ℓ is L -lipschitz. \square

Lemma 4. *For any given predicted rating matrix $\widehat{\mathbf{R}} \in \mathcal{H}$ and two propensity matrices (\mathbf{P} and \mathbf{P}'), the following inequality holds*

$$\mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}) \leq \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}') + \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{P}, \mathbf{P}').$$

Proof. By the definition of PMD, we have

$$\begin{aligned} \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}) &= \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}') - \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}') + \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}) \\ &\leq \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}') + \sup_{\widehat{\mathbf{R}}' \in \mathcal{H}} \{\mathcal{L}^\ell(\widehat{\mathbf{R}}, \mathbf{P}) - \mathcal{L}^\ell(\widehat{\mathbf{R}}'; \mathbf{P}')\} \\ &= \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}') + \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{P}, \mathbf{P}'). \end{aligned} \quad \square$$

A.4 Proof of Theorem 1

Proof. First, we obtain the following inequality by replacing \mathbf{P} and \mathbf{P}' for \mathbf{P}_{MCAR} and \mathbf{P}_{MNAR} in Lemma 4.

$$\mathcal{L}_{ideal}^\ell(\widehat{\mathbf{R}}) \leq \widehat{\mathcal{L}}_{naive}^\ell(\widehat{\mathbf{R}}; \mathbf{P}_{MNAR}) + \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{P}_{MCAR}, \mathbf{P}_{MNAR}), \quad (17)$$

where $\mathcal{L}_{ideal}^\ell(\widehat{\mathbf{R}}) = \mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}_{MCAR})$ by definition. Then, from Lemma 2 and Lemma 3, the following inequalities hold with a probability of at least $1 - 2\delta/3$ and $1 - \delta/3$, respectively.

$$\mathcal{L}^\ell(\widehat{\mathbf{R}}; \mathbf{P}_{MNAR}) \leq \widehat{\mathcal{L}}_{naive}^\ell(\widehat{\mathbf{R}}; \mathbf{O}_{MNAR}) + 2L\mathfrak{R}_{\mathbf{P}, M}(\mathcal{H}) + \Delta \sqrt{\frac{\log(6/\delta)}{2M}}, \quad (18)$$

$$\begin{aligned} &\left| \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{P}_{MCAR}, \mathbf{P}_{MNAR}) - \psi_{\widehat{\mathbf{R}}, \mathcal{H}}(\mathbf{O}_{MCAR}, \mathbf{O}_{MNAR}) \right| \\ &\leq 4L(\mathfrak{R}_{\mathbf{P}, M}(\mathcal{H}) + \mathfrak{R}_{\mathbf{P}', M}(\mathcal{H})) + 2\Delta \sqrt{\frac{\log(6/\delta)}{2M}}. \end{aligned} \quad (19)$$

Combining Eq. (17), Eq. (18), and Eq. (19) with the union bound completes the proof. \square

B Experiment Details

Datasets. We used the following real-world datasets.

- Yahoo! R3⁵: It contains five-star user-song ratings. The training data contain approximately 300,000 MNAR ratings from 15,400 users for 1,000 songs, and the test data were collected by a subset of users to rate 10 randomly sampled songs.
- Coat⁶: It contains five-star user-coat ratings from 290 Amazon Mechanical Turk workers on an inventory of 300 coats. The training data contain 6,500 MNAR ratings collected via self-selection by Turk workers. Conversely, the test data were collected by asking Turk workers to rate 16 randomly selected coats.

Note here that Yahoo! R3 and Coat are the only publicity available real-world datasets that contain the test sets with MCAR ratings. Figure 2 shows the training and test rating distributions of Yahoo! R3 and Coat. The figures clearly show the distributional shifts between the training and test sets.

Definition of Propensity Estimators. For MF-IPS and MF-DR, we use *user propensity*, *item propensity*, *user-item propensity*, and *1BitMC* [Ma and Chen, 2019] as the variants of the propensity estimator and report the results with the best estimator for each dataset. Note that these four propensity estimators are usable in real-world recommender systems because they use only MNAR data. These four estimators are defined as follows:

$$\text{user propensity} : \widehat{P}_{u,*} := \frac{\sum_{i \in \mathcal{I}} O_{u,i}}{\max_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} O_{u,i}},$$

$$\text{item propensity} : \widehat{P}_{*,i} := \frac{\sum_{u \in \mathcal{U}} O_{u,i}}{\max_{i \in \mathcal{I}} \sum_{u \in \mathcal{U}} O_{u,i}},$$

$$\text{user-item propensity} : \widehat{P}_{u,i} := \widehat{P}_{u,*} \cdot \widehat{P}_{*,i},$$

1BitMC :

$$\begin{aligned} \widehat{\Gamma}_{u,i} &:= \arg \min_{\Gamma \in \mathcal{F}_{\tau, \gamma}} \sum_{(u,i) \in \mathcal{D}} \{O_{u,i} \log(\sigma(\Gamma_{u,i})) \\ &\quad + (1 - O_{u,i}) \log(1 - \sigma(\Gamma_{u,i}))\} \end{aligned}$$

where $\mathcal{F}_{\tau, \gamma} := \{\Gamma \in \mathbb{R}^{m \times n} \mid \|\Gamma\|_* \leq \tau \sqrt{mn}, \|\Gamma\|_{\max} \leq \gamma\}$, $\|\cdot\|_*$ ($\tau, \gamma > 0$) denotes the nuclear norm, and $\|\cdot\|_{\max}$ is the entry-wise max norm. $\sigma(\cdot)$ is the sigmoid function, and $\widehat{P}_{u,i} := \sigma(\widehat{\Gamma}_{u,i})$. We use the implementation provided by Ma *et al.* [Ma and Chen, 2019] for 1BitMC. Note that 1BitMC is a sophisticated version of the logistic regression to estimate the propensity score described in Schnabel *et al.* [Schnabel *et al.*, 2016].

hyper-parameter Tuning. Table 5 shows the hyper-parameter search space used in the experiments. We tune the dimensions of the latent factors d in the range of $\{5, 10, \dots, 40\}$ and the L2-regularization hyper-parameter λ

⁵<http://webscope.sandbox.yahoo.com/>

⁶<https://www.cs.cornell.edu/~schnabts/mnar/>




Figure 2: Rating Distributions of Training and Test Sets of Yahoo! R3 and Coat.

Note: The rating distributions are significantly different between the training and test sets for both datasets. Note that $KL-div$ is the Kullback–Leibler divergence of the rating distributions between training and test sets. Therefore, the distributional shift of Yahoo! R3 dataset is relatively large compared to that of the Coat dataset.

in the range of $[10^{-4}, 1]$ for all methods. For DAMF and CausE, the trade-off hyper-parameter β is also tuned in the range of $[10^{-2}, 1]$. We identify the best values for these hyper-parameters by running an adaptive hyper-parameter tuning procedure implemented in *Optuna* [Akiba *et al.*, 2019] on a validation set.

Datasets	Methods	d	λ	β
Yahoo! R3	MF-IPS	{5, 10, ..., 40}	$[10^{-4}, 1]$	-
	MF-DR			-
	CausE			-
	DAMF (ours)			$[10^{-2}, 1]$
Coat	MF-IPS	{5, 10, ..., 40}	$[10^{-8}, 1]$	-
	MF-DR			-
	CausE			$[10^{-10}, 1]$
	DAMF (ours)			$[10^{-2}, 1]$

Table 5: Hyper-parameter Search Space

Note: d is the dimension of the latent factors, λ is the hyperparameter for the L2-regularization, and β is the trade-off hyperparameter. We also used the Adam optimizer with 0.01 as its initial learning rate for all methods.

Performance Measures. The following defines the performance measures used in the experiments.

- **MSE** evaluates how far the predicted ratings are away from the true rating and is defined as

$$MSE := \frac{1}{|\mathcal{D}_{te}|} \sum_{(u,i) \in \mathcal{D}_{te}} (R_{u,i} - \hat{R}_{u,i})^2,$$

where \mathcal{D}_{te} is the user–item pairs in the test set.

- **Normalized discounted cumulative gain (NDCG)** measures the ranking quality and is defined as:

$$DCG@K := \frac{1}{m} \sum_{(u,i)} \frac{2^{(R_{u,i}-1)} \cdot \mathbb{I}\{rank(u,i) \leq K\}}{\log_2(rank(u,i) + 1)}.$$

where $\mathbb{I}\{\cdot\}$ is the indicator function, $rank(u,i)$ is a ranking of i for u induced by $\hat{\mathbf{R}}$. Then, $NDCG@K := DCG@K / IDC@K$, where $IDC@K$ is the maximum possible $DCG@K$.

- **Recall** evaluates how many relevant items are selected and is defined as

$$Recall@K := \frac{1}{m} \sum_{(u,i)} \frac{R_{u,i} \cdot \mathbb{I}\{rank(u,i) \leq K\}}{\sum_{i \in [n]} R_{u,i}}.$$

C Related Work

This section summarizes the related literature.

Offline Recommender Learning from MNAR Feedback:

To address the selection bias of MNAR explicit feedback, propensity-based methods have been explored [Liang *et al.*, 2016; Schnabel *et al.*, 2016; Wang *et al.*, 2019]. Among these, the most basic method is the IPS estimation, which was originally established in causal inference [Imbens and Rubin, 2015; Rosenbaum and Rubin, 1983; Rubin, 1974; Saito and Yasui, 2020]. This method provides an unbiased estimator of the true metric of interest by weighting each data point using the inverse of its propensity. The rating predictor based on the IPS estimator empirically outperformed both the naive MF [Koren *et al.*, 2009] and probabilistic generative models [Hernández-Lobato *et al.*, 2014]. IPS can reasonably remove the bias of naive methods, however, their performance depends on the quality of the propensity score estimation [Saito *et al.*, 2020c; Saito, 2020a; Saito, 2020c; Saito *et al.*, 2020b]. Specifically, it is challenging to ensure the performance of the propensity estimation in real-world recommendations, as users are independent in selecting which items to rate, and one cannot control the missing mechanism [Marlin and Zemel, 2009]. In addition to the simple IPS estimator, Wang *et al.* [Wang *et al.*, 2019] proposed the doubly robust (DR) variant to decrease the variance of the propensity-weighting approach. The DR estimator utilizes both the error imputation model and the propensity model, and theoretically improves the bias and estimation error bound compared to the IPS counterpart. However, the proposed joint learning algorithm still requires pre-estimated propensity scores [Wang *et al.*, 2019]. Furthermore, the estimation performance of the DR estimator is significantly degraded when error imputation and propensity models are both misspecified [Saito *et al.*, 2019; Saito, 2020b; Saito *et al.*, 2020a; Dudík *et al.*, 2011; Wang *et al.*, 2019; Kiyohara *et al.*, 2022;

Kallus *et al.*, 2021; Saito *et al.*, 2019]. In the empirical evaluation of the propensity-based methods, MCAR test data were used to estimate the propensity score [Schnabel *et al.*, 2016; Wang *et al.*, 2019]. However, MCAR data are unavailable in most practical situations, as gathering a sufficient amount of MCAR data requires more time and cost requirements for the annotation process [Gilotte *et al.*, 2018; Saito *et al.*, 2020a; Saito and Joachims, 2022; Saito and Joachims, 2021; Saito *et al.*, 2021].

Currently, there are two studies that address issues related to conventional propensity-based methods. Bonner and Vasile [Bonner and Vasile, 2018] proposed CausE, a domain adaptation inspired method that introduces a regularizer term on the discrepancy between latent factors obtained from MCAR and MNAR data. However, this method requires some MCAR training data, which are generally unavailable. Moreover, there is no theoretical guarantee for the proposed loss function. Therefore, our method is more desirable than CausE in that (i) Our method does not use any MCAR data in its training process and is feasible in a realistic situation with no MCAR data. (ii) Our method is theoretically refined as it is designed to minimize the propensity-independent upper bound of the ideal loss function.

Next, Ma *et al.* [Ma and Chen, 2019] proposed a propensity estimation method, 1BitMC, which does not require MCAR data. The authors constructed the theoretical guarantee for the consistency of 1BitMC. However, it presupposes the use of **inverse** propensity weighting to debias downstream recommenders; it cannot be used when there exists a user–item pair with zero observed probability. Furthermore, its experiments use only small-size recommendation datasets. In addition, the performance of the recommendation methods with 1BitMC was evaluated using only prediction accuracy measures. We address the limitations of Ma *et al.* [Ma and Chen, 2019] as follows: (i) Our proposed method and theory do not depend on the propensity score and are applicable to settings where there exists a user–item pair with zero observed probability. (ii) Through comprehensive experiments on some moderately sized datasets, we demonstrate that our proposed method performs better than MF-IPS and MF-DR with 1BitMC for both rating prediction and ranking tasks.

Unsupervised Domain Adaptation (UDA): The aim of UDA is to train a predictor that works well on a target domain by using only labeled source data and unlabeled target data during training [Kuroki *et al.*, 2019; Saito *et al.*, 2017]. However, the major challenge of UDA is that the feature distributions and labeling functions can differ between the source and target domains. Thus, a predictor trained using only labeled source data does not generalize well to the target domain. Therefore, it is essential to measure the dissimilarity between the domains to achieve a desired performance in the target domain [Kuroki *et al.*, 2019; Lee *et al.*, 2019]. Several discrepancy measures have been proposed to measure the difference in feature distributions between the source and target domains [Ben-David *et al.*, 2010; Kuroki *et al.*, 2019; Lee *et al.*, 2019; Zhang *et al.*, 2019]. For example, \mathcal{H} -divergence and $\mathcal{H}\Delta\mathcal{H}$ -divergence [Ben-David *et al.*, 2010; Ben-David *et al.*, 2007] have been used to construct many

prediction methods for UDA, such as DANN, ADDA, and MCD [Ganin and Lempitsky, 2015; Ganin *et al.*, 2016; Tzeng *et al.*, 2017; Saito *et al.*, 2018]. These methods are based on the adversarial learning framework and can be theoretically explained as minimizing empirical errors and discrepancy measures between the source and target domains. Note that the optimization of these methods does not depend on the propensity score. Thus, UDA is useful in constructing an effective recommender with biased rating feedback, given the absence of access to the true propensity score.

This study extended the idea of using a discrepancy measure to quantify the difference between two propensity score matrices and derive a propensity-independent generalization error bound, which is non-trivial given only the UDA literature, as it does not handle matrices with different distributions. Moreover, we provided an algorithm to optimize the upper bound of the ideal loss function via adversarial learning and matrix factorization.

References

- [Akiba *et al.*, 2019] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD ’19, pages 2623–2631, 2019.
- [Bartlett and Mendelson, 2002] Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- [Ben-David *et al.*, 2007] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *Advances in neural information processing systems*, pages 137–144, 2007.
- [Ben-David *et al.*, 2010] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- [Bonner and Vasile, 2018] Stephen Bonner and Flavian Vasile. Causal embeddings for recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems*, RecSys ’18, pages 104–112, New York, NY, USA, 2018. ACM.
- [Chen *et al.*, 2020] 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.
- [Dudík *et al.*, 2011] 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*, pages 1097–1104, 2011.
- [Ganin and Lempitsky, 2015] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings*

- of *Machine Learning Research*, pages 1180–1189. PMLR, 2015.
- [Ganin *et al.*, 2016] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- [Gilotte *et al.*, 2018] Alexandre Gilotte, Clément Calauzènes, Thomas Nedelec, Alexandre Abraham, and Simon Dollé. Offline a/b testing for recommender systems. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 198–206. ACM, 2018.
- [Hernández-Lobato *et al.*, 2014] 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, 2014.
- [Imbens and Rubin, 2015] Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- [Joachims *et al.*, 2017] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. Unbiased learning-to-rank with biased feedback. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 781–789. ACM, 2017.
- [Kallus *et al.*, 2021] Nathan Kallus, Yuta Saito, and Masatoshi Uehara. Optimal off-policy evaluation from multiple logging policies. In *International Conference on Machine Learning*, pages 5247–5256. PMLR, 2021.
- [Kiyohara *et al.*, 2022] Haruka Kiyohara, Yuta Saito, Tatsuya Matsuhira, Yusuke Narita, Nobuyuki Shimizu, and Yasuo Yamamoto. Doubly robust off-policy evaluation for ranking policies under the cascade behavior model. In *Proceedings of the 15th International Conference on Web Search and Data Mining*, pages 487–497, 2022.
- [Koren *et al.*, 2009] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [Kuroki *et al.*, 2019] Seiichi Kuroki, Nontawat Charoenphakdee, Han Bao, Junya Honda, Issei Sato, and Masashi Sugiyama. Unsupervised domain adaptation based on source-guided discrepancy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4122–4129, 2019.
- [Lee *et al.*, 2019] Jongyeong Lee, Nontawat Charoenphakdee, Seiichi Kuroki, and Masashi Sugiyama. Domain discrepancy measure using complex models in unsupervised domain adaptation. *arXiv preprint arXiv:1901.10654*, 2019.
- [Liang *et al.*, 2016] Dawen Liang, Laurent Charlin, and David M Blei. Causal inference for recommendation. In *Causation: Foundation to Application, Workshop at UAI*, 2016.
- [Ma and Chen, 2019] 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. In *Advances in Neural Information Processing Systems*, pages 14871–14880, 2019.
- [Mansour *et al.*, 2009] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. In *Proceedings of The 22nd Annual Conference on Learning Theory (COLT 2009)*, Montréal, Canada, 2009.
- [Marlin and Zemel, 2009] 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. ACM, 2009.
- [Mnih and Salakhutdinov, 2008] Andriy Mnih and Ruslan R Salakhutdinov. Probabilistic matrix factorization. In *Advances in neural information processing systems*, pages 1257–1264, 2008.
- [Mohri *et al.*, 2018] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [Rosenbaum and Rubin, 1983] Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- [Rubin, 1974] Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- [Saito and Joachims, 2021] Yuta Saito and Thorsten Joachims. Counterfactual learning and evaluation for recommender systems: Foundations, implementations, and recent advances. In *Fifteenth ACM Conference on Recommender Systems*, pages 828–830, 2021.
- [Saito and Joachims, 2022] Yuta Saito and Thorsten Joachims. Off-policy evaluation for large action spaces via embeddings. *arXiv preprint arXiv:2202.06317*, 2022.
- [Saito and Yasui, 2020] Yuta Saito and Shota Yasui. Counterfactual cross-validation: Stable model selection procedure for causal inference models. In *International Conference on Machine Learning*, pages 8398–8407. PMLR, 2020.
- [Saito *et al.*, 2017] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2988–2997. PMLR, 2017.
- [Saito *et al.*, 2018] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2018.
- [Saito *et al.*, 2019] Yuta Saito, Hayato Sakata, and Kazuhide Nakata. Doubly robust prediction and evaluation methods

- improve uplift modeling for observational data. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, pages 468–476. SIAM, 2019.
- [Saito *et al.*, 2020a] Yuta Saito, Shunsuke Aihara, Megumi Matsutani, and Yusuke Narita. Open bandit dataset and pipeline: Towards realistic and reproducible off-policy evaluation. *arXiv preprint arXiv:2008.07146*, 2020.
- [Saito *et al.*, 2020b] Yuta Saito, Gota Morishihta, and Shota Yasui. Dual learning algorithm for delayed conversions. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1849–1852, 2020.
- [Saito *et al.*, 2020c] Yuta Saito, Yaginuma Suguru, Yuta Nishino, Hayato Sakata, and Nakata Kazuhide. Unbiased recommender learning from missing-not-at-random implicit feedback. In *Proceedings of the Thirteenth ACM International Conference on Web Search and Data Mining*. ACM, 2020.
- [Saito *et al.*, 2021] Yuta Saito, Takuma Udagawa, Haruka Kiyohara, Kazuki Mogi, Yusuke Narita, and Kei Tateno. Evaluating the robustness of off-policy evaluation. In *Fifteenth ACM Conference on Recommender Systems*, pages 114–123, 2021.
- [Saito, 2020a] 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.
- [Saito, 2020b] Yuta Saito. Doubly robust estimator for ranking metrics with post-click conversions. In *Fourteenth ACM Conference on Recommender Systems*, pages 92–100, 2020.
- [Saito, 2020c] Yuta Saito. Unbiased pairwise learning from biased implicit feedback. In *Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval*, pages 5–12, 2020.
- [Schnabel *et al.*, 2016] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 1670–1679. PMLR, 2016.
- [Steck, 2010] 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. ACM, 2010.
- [Tzeng *et al.*, 2017] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017.
- [Wang *et al.*, 2019] 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.
- [Zhang *et al.*, 2019] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *International Conference on Machine Learning*, pages 7404–7413, 2019.