DR-VAE: Debiased and Representation-enhanced Variational Autoencoder for Collaborative Recommendation

Fan Wang[1], Chaochao Chen[1], Weiming Liu[1], Minye Lei[2], Jintao Chen[1], Yuwen Liu[3], Xiaolin Zheng[1]*, Jianwei Yin[1]

[1]College of Computer Science and Technology, Zhejiang University, China
[2]College of Computer Science and Technology, Jilin University, China
[3]College of Computer Science and Technology, China University of Petroleum (East China), China
{fanwang97, zjuccc, 21831010} @zju.edu.cn, Minye.Lei@outlook.com, chenjintao@zju.edu.cn, yuwenliu97@gmail.com, xlzheng@zju.edu.cn, zjuyjw@cs.zju.edu.cn

Abstract

Recommender Systems (RSs) are widely applied for navigating information, and Collaborative Filtering (CF) is one of prominent recommendation techniques due to the advantages of domain independence and easy interpretation. Among the numerous CF methods, Variational Autoencoders (VAE), benefiting from modeling in a probabilitistic way, stands out in capturing user preferences through representation learning. Despite the superiority, VAE-based CF models still suffer from two challenging problems: (1) *Exposure bias*: models in training state are narrowly exposed to a limited, biased sample of data, leading to a skewed understanding of users' true preferences; (2) *Posterior collapse*: models excessively simplify the learned latent variable distributions, generating na"ive representations that are unable to encapsulate the complex data patterns and thereby resulting improper recommendations. In this paper, we propose a Debiased and Representation-enhanced Variational AutoEncoder (`DR-VAE`) framework for collaborative recommendations. Specifically, for *exposure bias* problem, `DR-VAE` incorporates a Debiasing Estimator, mitigating the impact of exposure bias. For *poster collapse* issue, `DR-VAE` innovatively introduces a Flow-based Representation Enhancement module, ensuring us to encapsulate complex data patterns by fitting complex and intricate posterior distributions directly. We provide experimental validations over four datasets to substantiate the efficacy of our `DR-VAE` framework.

Introduction

As essential instruments to counteract information overload problems, Recommender Systems (RSs) have been extensively applied in the areas of e-commerce, education, health, etc (Qi et al. 2024; Chen et al. 2024f,g; Gong et al. 2022; Chen et al. 2024a). Among the various techniques employed in RSs, Collaborative Filtering (CF) (Li et al. 2023; Liu et al. 2024c; Chen et al. 2024b) stands out as a pivotal method due to its accuracy and interpretability. The core idea of CF (e.g., user-based CF) is to learn user representations from past user-item interactions (Liu et al. 2024a,b) and then predict users' preferences towards items.

*Corresponding author
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

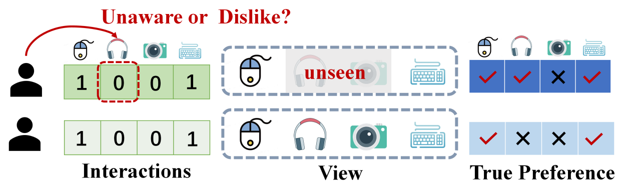

Figure 1: Given the limited view of the user, it remains uncertain whether an unobserved interaction is due to the user simply being unaware of the item, or if they have seen it and decided they don't like it.

Among representation-based CF techniques, Variational Autoencoders (VAE) (Liang et al. 2018) have achieved significant performance improvements and drawn extensive research attention for CF because of their superior modeling capability (Kingma and Welling 2013; Liang et al. 2018; Truong, Salah, and Lauw 2021). Specifically, VAE introduces a probabilistic framework that uses stochastic encoders and decoders to learn the latent probability distributions rather than fixed points in the latent space. This approach allows VAE to capture the uncertainty and variability in user preferences, enabling effective generalization on unseen data and ensuring robust recommendations (Truong, Salah, and Lauw 2021).

However, current VAE-based CF models often face two challenges. **CH1-***exposure bias:* This refers to the fact that users are exposed to only a subset of items, so unobserved interactions do not necessarily indicate negative preferences (Chen et al. 2023), as illustrated in Fig. 1. Moreover, popular items are more likely to be seen by users while long-tail items struggle to attract attention. Consequently, the learned user representations might not accurately reflect true preferences, leading to suboptimal recommendations. Although some researchers have explored debiasing in RSs (Saito et al. 2020; Li et al. 2022; Chen et al. 2021; Wang et al. 2023), they often focus solely on fixed latent embeddings, overlooking the data's inherent probabilistic nature, which can weaken model robustness and generalization. **CH2-***posterior collapse:* This phenomenon occurs when the learned latent representation ignores the input data and relies entirely on the prior distribution, causing the model to miss capturing useful data features. VAE learns latent representations primarily by maximizing the Evidence Lower BOund (ELBO) (Kingma and Welling 2013; Liang et al. 2018). However, the standard ELBO enforces that all posterior distributions conform to a fixed, overly simple prior (e.g., an isotropic Gaussian), leading to posterior collapse (Wang and Ziyin 2022). Such a strict constraint may result in underfitting or uninformative representations, ultimately yielding suboptimal recommendation outcomes (Truong, Salah, and Lauw 2021).

Considering the above challenges, in this paper, we propose a novel **D**ebiased and **R**epresentation-enhanced **V**ariational **A**uto-**E**ncoder, `DR-VAE`, for collaborative recommendations, including a Debiasing Estimator (abbreviated as `D-VAE`) for **CH1** and a Flow-based Representation Enhancement module (abbreviated as `R-VAE`) for **CH2**. `DR-VAE` not only retains the probabilistic nature of VAE, but also enhances VAE models to learn debiased and insightful representations. Specifically, `D-VAE` improves the reconstruction term of the vanilla VAE by devising unbiased estimators tailored to different distributions, thereby capturing users' unbiased and true preferences. `R-VAE` innovatively rewrites the KL alignment in the regularization term of vanilla VAEs using Ordinary Differential Equations (ODE), overcoming posterior collapse and learning insightful and informative representations.

In summary, the main contributions of `DR-VAE` are: (1) For **CH1**, we establish a Debiasing Estimator (`D-VAE`) tailored to VAE, designed to learn representations that reflect users' true preferences. Theoretical analyses verify that the estimator eliminates exposure bias in RSs. (2) For **CH2**, we introduce a Flow-based Representation Enhancement (`R-VAE`) module to alleviate the posterior collapse problem in VAE, enabling more informative and effective representations for high-quality recommendations. (3) Extensive experiments on four datasets demonstrate the superiority of `DR-VAE` over state-of-the-art baselines.

Problem Formulation and Preliminaries

Problem Formulation

We begin by introducing some basic notations that will be used in our problem. Let $\mathbf{X} \in \mathbb{R}^{U \times I}$ denote the user-item interaction matrix including U users and I items, where $x_{u,i}$ is an entry in $\mathbf{X}$ representing the interaction record from user u to item i. In this paper, we mainly focus on implicit feedback, i.e., indirect indications of user preferences (e.g., click-throughs, purchase actions, etc.) rather than explicit ratings or reviews. Thus, $x_{u,i} \in 0, 1$ is a Bernoulli random variable, where 1 indicates that user u has interacted with item i, and 0 indicates that they have not. Additionally, $\mathbf{x}u \in \mathbb{R}^I$ denotes the u-th row of $\mathbf{X}$, i.e., $\mathbf{X} = \mathbf{x}u_{u=1}^U$; $\mathbf{z}_u$ denotes the representation of user u that we aim to learn. In summary, the main problem we aim to solve is: *Given the provided user-item interaction matrix* $\mathbf{X}$*, how to learn unbiased and informative representations* $\mathbf{z}_u$ *that capture users' true preference, in order to generate the reconstructed interaction matrix* $\hat{\mathbf{X}}$ *for better recommendations.*

Preliminaries

Variational Autoencoders for Recommendations As we all know, VAE is a mainstream and classical generative model (Liang et al. 2018; Kingma and Welling 2013). In general, the goal of recommendation task based on generative model lies in maximizing the probability of generating the user-item interaction matrix by introducing latent representations as $p(\mathbf{x}_u) = \sum_{\mathbf{z}_u} p(\mathbf{x}_u|\mathbf{z}_u)p(\mathbf{z}_u)$. To learn the representation $\mathbf{z}_u$, VAE introduces the true posterior distribution $p(\mathbf{z}_u|\mathbf{x}_u)$. Ideally, accurate representations can be obtained through sampling from the true posteriors. However, true posterior distributions are generally intractable (Kingma and Welling 2013). In view of this, VAE adopts distributions $q_\phi(\mathbf{z}_u|\mathbf{x}_u)$ with parameter ϕ to approximate the true posteriors. Then, VAE achieves Evidence Lower Bound (ELBO), a proxy of $p(\mathbf{x}_u)$, to be maximized:

$$\log p(\mathbf{x}_u) \geq \mathbb{E}_{q_\phi(\mathbf{z}_u|\mathbf{x}_u)}[\log\ p_\psi(\mathbf{x}_u|\mathbf{z}_u)] - \mathrm{KL}(\mathrm{q}_\phi(\mathbf{z}_\mathrm{u}|\mathbf{x}_\mathrm{u})\|\mathrm{p}(\mathbf{z}_\mathrm{u})), \tag{1}$$

where the first part in Eq. (1) is the reconstruction term, measuring the probability of reconstructing the user-item interaction matrix. Without loss of generality, $p_\psi(\mathbf{x}_u|\mathbf{z}_u)$ follows a univariate exponential family, such as Gaussian, Bernoulli, or Poisson distributions (Truong, Salah, and Lauw 2021). The second part of Eq. (1) is a regularizer, constraining the nature of the approximate posteriors.

Based on the two terms, we recall the challenges of vanilla VAE in two aspects. **CH1**-*exposure bias*: When using the reconstruction term to measure the discrepancy between the observed and reconstructed data, items that have not been observed by users are often misinterpreted as items that users are not interested in, leading to inaccurate estimations of user preferences; **CH2**-*posterior collapse*: When employing the regularizer to learn approximate posteriors, the approximate posterior $q_\phi(\mathbf{z}_u|\mathbf{x}_u)$ is constrained to align with pre-specified and oversimplified prior distributions $p(\mathbf{z}_u)$ (e.g., standard Gaussian), making it challenging to capture complex and accurate features.

The Proposed Method: `DR-VAE`

In this section, we propose a novel generative method, `DR-VAE`, to learn debiased and informative representations for collaborative recommendations. Fig. 2 illustrates the whole framework of `DR-VAE`. As shown in Fig. 2, given the observed user-item interactions $\mathbf{X}$, `DR-VAE` aims to generate reconstructed data $\hat{\mathbf{X}}$ that closely resembles $\mathbf{X}$. Different from vanilla VAE, `DR-VAE` improve $\mathbf{z}_u$ from two perspectives: (a) It improves the reconstruction term of the vanilla VAE through the Debiasing Estimator (`D-VAE`) module to overcome **CH1**-*exposure bias*, thereby obtaining debiased representations; (b) It enhances the regularization term of the vanilla VAE through the Flow-based Representation Enhancement module (`R-VAE`) to overcome **CH2**-*posterior collapse*, resulting in informative representations.

`D-VAE`: Debiasing Estimator

To begin with, we give the basic assumptions about user-item interactions that are used for unbiased estimator (Saito et al. 2020).

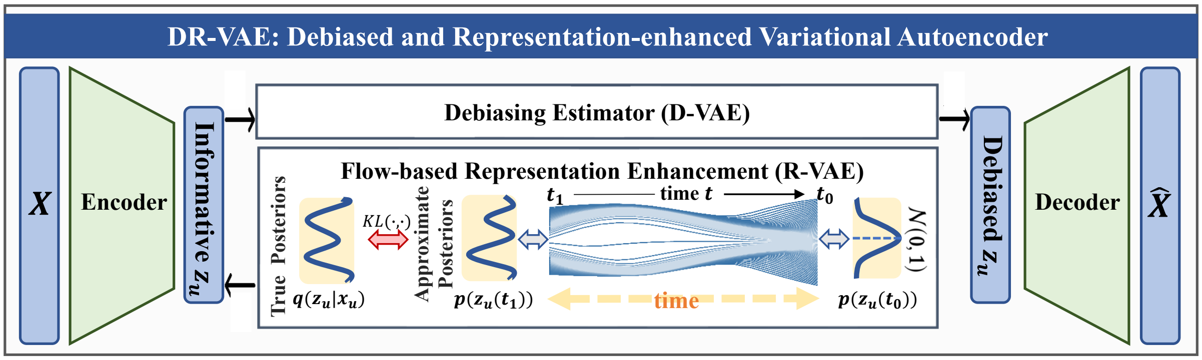

Figure 2: The overview of DR-VAE, including two modules: (a) Debiasing Estimater (D-VAE) module for **CH1**-*exposure bias*; (b) Flow-based Representation Enhancement (R-VAE) module for **CH2**-*posterior collapse*.

Distribution	$\zeta^{(\mathbf{r})}(\hat{\mathbf{x}})\ \mathbf{r}=\{0,1\}$	$\mathcal{L}_{\text{naïve}}$	$\mathcal{L}_{\text{ideal}}$	$\mathcal{L}_{\text{unbiased}}$
General	–	$\sum[x\zeta^{(1)}(\hat{x})+(1-x)\,\zeta^{(0)}(\hat{x})]$	$\sum[\gamma\zeta^{(1)}(\hat{x})+(1-\gamma)\,\zeta^{(0)}(\hat{x})]$	$\sum[\frac{x}{\theta}\zeta^{(1)}(\hat{x})+\left(1-\frac{x}{\theta}\right)\zeta^{(0)}(\hat{x})]$
Poisson	$\hat{x}-r\log(\hat{x})$	$\sum[\hat{x}-x\log(\hat{x})]$	$\sum[\hat{x}-\gamma\log(\hat{x})]$	$\sum[\hat{x}-\frac{x}{\theta}\log(\hat{x})]$
Multinomial	$-r\log(\hat{x})$	$-\sum x\log(\hat{x})$	$-\sum\gamma\log(\hat{x})$	$-\sum\frac{x}{\theta}\log(\hat{x})$
Gaussian	$(r-\hat{x})^2$	$\sum(x-\hat{x})^2$	$\sum\gamma-2\gamma\hat{x}+\hat{x}^2$	$\sum\frac{x}{\theta}-2\frac{x}{\theta}\hat{x}+\hat{x}^2$
Bernoulli	$-(r\log(\hat{x})+(1-r)\log(1-\hat{x}))$	$-\sum x\log\hat{x}+(1-x)\log(1-\hat{x})$	$-\sum\gamma\log\hat{x}+(1-\gamma)\log(1-\hat{x})$	$-\sum\frac{x}{\theta}\log\hat{x}+(1-\frac{x}{\theta})\log(1-\hat{x})$

Table 1: Specification of unbiased estimators tailored to VAE's reconstruction term w.r.t. four typical distributions.

Assumption* 1.** ***Interaction Disentanglement: *The implicit feedback $x_{u,i}=\{0,1\}$ generally depends on two aspects: (a) whether item i has been exposed to user u, denoted as $o_{u,i}=\{0,1\}$; and (b) whether user u is interested in item i, denoted as $r_{u,i}=\{0,1\}$, that is:*

$$x_{u,i}=o_{u,i}\cdot r_{u,i},\quad P(x_{u,i}=1)=P(o_{u,i}=1)P(r_{u,i}=1)=\theta_{u,i}\cdot\gamma_{u,i}. \tag{2}$$

From the assumption, we can clarify two statements: (1) User u interacts with item i only if he (or she) has seen and is interested in it; (2) The probability of interaction is determined by both the exposure probability ($\theta_{u,i}$) and the relevance degree $\gamma_{u,i}$. Therefore, an interaction or not (i.e., $x_{u,i}=1$ or $x_{u,i}=0$) does not always imply relevance, highlighting the prevalence of exposure bias in user-item interaction data.

We now present the naïve pointwise loss function[1] for estimating user preferences.

Definition* 1.** ***Naïve Loss: *Let $\hat{x}_{u,i}$ be the reconstructed value of $x_{u,i}$, and $\zeta^{(r)}(\cdot)$ ($r\in\{0,1\}$) be local loss, the naïve loss is defined as:*

$$\mathcal{L}_{\textit{naïve}}(\hat{X})=\sum_{u,i}\left[x_{u,i}\zeta^{(1)}(\hat{x}_{u,i})+(1-x_{u,i})\,\zeta^{(0)}(\hat{x}_{u,i})\right]. \tag{3}$$

Compared to the naïve loss that focuses on interaction $x_{u,i}$ which includes exposure bias, the ideal pointwise loss function focuses on the true relevance between u and i.

[1]We primarily focus on pointwise loss because it often aligns with the loss function associated with VAE's reconstruction term.

Definition* 2.** ***Ideal Loss: *Let $\hat{x}_{u,i}$ be the reconstructed value of $x_{u,i}$, $\gamma_{u,i}$ the relevance between u and i, and $\zeta^{(r)}(\cdot)$ ($r\in\{0,1\}$) the local loss. The ideal loss is defined as:*

$$\mathcal{L}_{ideal}(\hat{X})=\sum_{u,i}\left[\gamma_{u,i}\zeta^{(1)}(\hat{x}_{u,i})+(1-\gamma_{u,i})\,\zeta^{(0)}(\hat{x}_{u,i})\right]. \tag{4}$$

Note that Eqs. (3) and (4) represent the general form of ideal pointwise loss, which can be specified by defining the local loss $\zeta^{(r)}(\cdot)$. We will specify $\zeta^{(r)}(\cdot)$ later in this section to correspond with VAE's reconstruction term.

Considering that $\gamma_{u,i}$ is intractable, we utilize propensity score (Schnabel et al. 2016; Saito et al. 2020; Seaman and White 2013) to implement our unbiased estimator.

Theorem* 1.** ***Unbiased Estimator: *Given the propensity score $\theta_{u,i}$=$P(o_{u,i}=1)$=$P(x_{u,i}=1|r_{u,i}=1)$, we can achieve unbiased estimator with lemma 1:*

$$\mathcal{L}_{unbiased}(\hat{X})=\sum_{u,i}\left[\frac{x_{u,i}}{\theta_{u,i}}\zeta^{(1)}(\hat{x}_{u,i})+\left(1-\frac{x_{u,i}}{\theta_{u,i}}\right)\zeta^{(0)}(\hat{x}_{u,i})\right]. \tag{5}$$

Lemma* 1.** ***Unbiased Estimator Tailored to VAE*. Considering the observed interactions conditional on the latent variables follows a univariate exponential family, we specify $\xi^{r}(\cdot)$ to obtain unbiased estimator tailored to VAE in terms of four typical distributions, i.e., Poisson, Multinomial, Gaussian and Bernoulli, as illustrated in Tab. 1.*

From Tab. 1, we can observe that by substituting distribution-specific $\xi^{(r)}(\cdot)$ into general loss function $\mathcal{L}_{\text{naïve}}$, the obtained loss functions can align with the reconstruction terms specific to different distributions in vanilla CF-based

VAE. This consistency validates that our R-VAE unifies with vanilla CF-based VAE. Moreover, the distribution-specific unbiased estimator ($\mathcal{L}_{unbiased}$) is unbiased against the corresponding ideal loss. Due to space constraints, we primarily follow (Truong, Salah, and Lauw 2021) and focus on the Poisson distribution in the remainder of this paper. Here, we provide a proof that the unbiased estimator specific to the Poisson distribution can converge to the ideal loss.

Proof.

$$\begin{aligned}\mathbb{E}\left[\mathcal{L}_{unbiased}^{Poiss}\right] &= \mathbb{E}\left[\sum_{u,i}\left[\hat{x}_{u,i} - \frac{x_{u,i}}{\theta_{u,i}}\log(\hat{x}_{u,i})\right]\right] \\ &= \sum_{u,i}\left[\hat{x}_{u,i} - \frac{\mathbb{E}[x_{u,i}]}{\theta_{u,i}}\log(\hat{x}_{u,i})\right] \\ &= \sum_{u,i}\left[\hat{x}_{u,i} - \gamma_{u,i}\log(\hat{x}_{u,i})\right] = \mathcal{L}_{ideal}^{Poiss}.\end{aligned}$$

□

In practice, there are numerous well-established studies on estimating propensity scores (Austin 2011; Zigler et al. 2013; Fan et al. 2022a). Here, we recruit a classical one of them for propensity score estimation in this paper:

$$\theta_{u,i} = \left(\sum_{u=1}^{U} x_{u,i} / \max_j \left\{\sum_{u=1}^{U} x_{u,j}\right\}\right)^{\eta}. \tag{6}$$

where hyperparameter η is employed to prevent $\theta_{u,i}$ from becoming too small, since too small $\theta_{u,i}$ could result in a denominator close to zero during weighting operations.

R-VAE: Flow-based Representation Enhancement

In this section, we focus on refining the regularizer in vanilla VAEs to address *posterior collapse*. *Posterior collapse* occurs when alignment of the regularizer oversimplifies the posterior distribution of learned representations to approximate the simple prior distribution. As a result, the representations may fail to capture useful information from complex data and accurately reflect user preferences.

To overcome the posterior collapse problem, we innovatively introduce Continuous Normalizing Flow (CNF). CNF can effectively fit arbitrarily complex distributions (e.g., complex posteriors in this paper) by initiating from a simple distribution and undergoing a series of reversible and continuous transformations, which are solved using Ordinary Differential Equations (ODE), and can be applied to areas such as domain adaptation (Liao et al. 2024b; Chen et al. 2024e), time series analysis (Liu et al. 2023; Song et al. 2025), and more. Accordingly, the transformations of representations adhering to different distributions during the contious process also constitutes a continuous dynamic. In concrete, to start with, we formulate representation $\mathbf{z}_u$ at different moments during the continuous process as $z_u(t)$. Then, we design ODEs $\frac{\partial \mathbf{z}_u(t)}{\partial t} = f_u(\mathbf{z}_u(t), t, \xi)$ to characterize the continuous dynamics of the hidden representations, where f_u is a neural network parameterized by ξ. Similarly, following the instantaneous change of variable (Chen et al. 2018b; Grathwohl et al. 2018), we can model the continuous dynamics of log-distribution as the following ODE:

$$\frac{\partial \log p_\xi(\mathbf{z}_u(t))}{\partial t} = -\operatorname{Tr}\left(\frac{\partial f_u(\mathbf{z}_u(t), t, \xi)}{\partial \mathbf{z}_u(t)}\right), \tag{7}$$

where $\operatorname{Tr}(\cdot)$ represents trace operation.

Till now, we can solve the initial value problems of the above ODEs. Specifically, we retain the encoder mechanism of VAE and consider the $\mathbf{z}_u$ obtained through reparametrization trick as the target representations, i.e., $z_u(t_1) = \mathbf{z}_u$. Next, we take $z_u(t_1)$ as the initial value and solve ODE:

$$\mathbf{z}_u(t_0) = \mathbf{z}_u(t_1) + \int_{t_1}^{t_0} f_u(\mathbf{z}_u(t), t, \xi)dt = \text{ODESolve}(\mathbf{z}_u(t_1), f_u, t_1, t_0, \xi), \tag{8}$$

where ODE solver serves as a black box to calculate the value. Here, we adopt the inverse process of normalizing flow (from t_1 to t_0) to obtain the representation $z_u(t_0)$ of initial status. Given that CNFs typically assume a simple base distribution, we hypothesize that the distribution of $z_u(t_0)$ follows a standard Gaussian distribution. After substituting $z_u(t_0)$ into the simple distribution, we obtain $p_\xi(z_u(t_0))$. Based on Eq. (7) and the initial value $p_\xi(z_u(t_0))$, we solve the forward process of the normalizing flow using ODE:

$$\begin{aligned}\log p_\xi\left(\mathbf{z}_u\left(t_1\right)\right) &= \log p_\xi(\mathbf{z}_u(t_0)) - \int_{t_0}^{t_1} \operatorname{Tr}\left(\frac{\partial f(\mathbf{z}_u(t), t, \xi)}{\partial \mathbf{z}_u(t)}\right) dt \\ &= \text{ODESolve}(\log p_\xi(\mathbf{z}_u(t_0)), f_u, t_0, t_1, \xi).\end{aligned} \tag{9}$$

We now obtain the log-density of posterior distribution $p_\xi(z_u(t_1))$ through continuous normalizing flow. The posterior $p_\xi(z_u(t_1))$ is expected to align with the approximate posterior $q_\phi(\mathbf{z}_u|\mathbf{x_u})$ derived from the observation data $\mathbf{X}$. Hence, we rewrite loss w.r.t. regularizer as follows:

$$\begin{aligned}\mathcal{L}_{regu} &= KL(q_\phi(\mathbf{z}_u|\mathbf{x}_u) \| p_\xi(z_u(t_1))) \\ &= \mathbb{E}_{q_\phi}\left[\log q_\phi(\mathbf{z}_u|\mathbf{x}_u) - \log p_\xi(\mathbf{z}_u(t_1))\right] \\ &= \mathbb{E}_{q_\phi}\left[\log q_\phi(\mathbf{z}_u|\mathbf{x}_u) - \log p_\xi(\mathbf{z}_u(t_0)) + \int_{t_0}^{t_1} \operatorname{Tr}\left(\frac{\partial f(\mathbf{z}_u(t), t, \xi)}{\partial \mathbf{z}_u(t)}\right) dt\right].\end{aligned} \tag{10}$$

Based on the refined regularizer's loss in Eq. (10), we are able to bypass the stringent constraint in traditional VAEs that necessitates the alignment of the posterior distribution with an overly simplistic prior distribution, thereby alleviating the posterior collapse issue. Furthermore, after drawing approximate posterior $q_\phi(\mathbf{z}_u|\mathbf{x_u})$ using encoder in VAE, CNF refits the complex posterior $p_\xi(z_u(t_1))$ to make the representation learning more robust. Overall, these enhancements ensure acquisition of effective and accurate representations that reflect users' true preferences.

Put Them Together

In summary, the unbiased reconstruction loss and refined regularizer's loss jointly constitute the final optimization goal of our DR-VAE in this paper:

$$\mathcal{L} = \mathcal{L}_{unbiased} + \beta \cdot \mathcal{L}_{regu}, \tag{11}$$

where β is a parameter to maintain a trade-off between the reconstruction term and regularizer term. Overall, DR-VAE is indeed a representation learning framework (Zheng et al. 2023; Liao et al. 2024a), and with minimal adjustments, it can also be applied to downstream tasks beyond recommendations, such as graph-based (Chen et al. 2024d; Gong et al. 2025; Mao et al. 2024) and time-series tasks (Fan et al. 2023, 2022b).

Experiments

In this section, we carry out extensive experiments to answer the four main questions: **RQ1**: Can our DR-VAE effectively learn users' true preferences compared to state-of-the-art debiasing baselines? **RQ2**: Does our R-VAE improve existing VAEs in terms of representational ability? **RQ3**: What are the effecs of the different components of our proposal? **RQ4**: How does the performance of DR-VAE vary w.r.t. different values of the hyper-parameters?

Experimental Settings

Datasets We conduct experiments on two groups of datasets: (1) *Semi-synthetic* datasets from Recbole-Debias[2], including **ML 100K** (Harper and Konstan 2015) and **KuaiRec** (Gao et al. 2022), where 50% of the data is biased normal and 50% unbiased intervened. The datasets are split into training, validation, and test sets as per Recbole-Debias. (2) *Real-world* datasets, including **Amazon Toys** (Ruining and Julian 2016) and **ModCloth** (Misra, Wan, and McAuley 2018), with no intervention. These two datasets are split into 8:1:1 for training, validation, and test. The general statistics of the four datasets are presented in Tab. 3.

Comparison Methods (SOTA) We compare our proposal with three categories of state-of-the-art approaches: (1) *Debiasing recommendation strategies*, including **IPS-MF**, **Rel-MF** (Saito et al. 2020), **MACR** (Wei et al. 2021), **BC LOSS** (Zhang et al. 2022), **S-DRO** (Wen et al. 2022), and **InvCF** (Zhang et al. 2023); (2) *VAE and its enhanced versions* for CF, including **MultVAE** (Liang et al. 2018), **RecVAE** (Shenbin et al. 2020) and **BiVAE** (Truong, Salah, and Lauw 2021); (3) *Classical CF methods*, including **MF** and **BPR**(Rendle et al. 2009).

Comparison Methods (Ours) We consider three variants of DR-VAE: (1) D-VAE: VAE improved only by Debiasing Estimator; (2) R-VAE: VAE improved only by Flow-based Representation Enhancement module; (3) DR-VAE: VAE improved by both Debiasing Estimator and Flow-based Representation Enhancement module, i.e., the complete DR-VAE framework.

Evaluation Metrics We adopt four widely-used evaluation metrics, i.e., Recall, Precision, Normalized Discounted Cumulative Gain (NDCG) and Hit Rate (HR), to assess the recommendation performances. We repeat each experiment five times and report the average results of top 20/50 candidates following (Zhang et al. 2021).

Implementation Details Experiments were conducted on an NVIDIA RTX3090 GPU. We optimized the models using the Adam optimizer with a learning rate of 0.001 and weight decay λ of 0.01. The latent dimension D was set to 300, and the batch size N to 32. For the hyperparameters in DR-VAE, we set $\beta = 0.01$ and $\eta = 0.3$ for all datasets.

[2]https://github.com/JingsenZhang/Recbole-Debias/

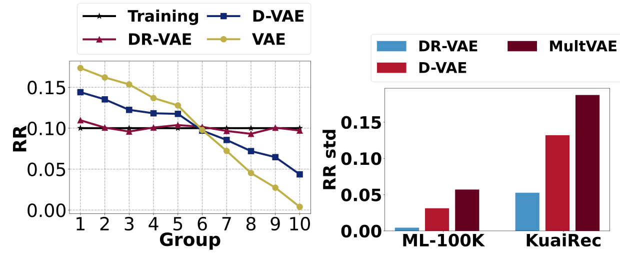

(a) ML-100K (b) std. dev.

Figure 3: Recommendation rate (RR) over item groups.

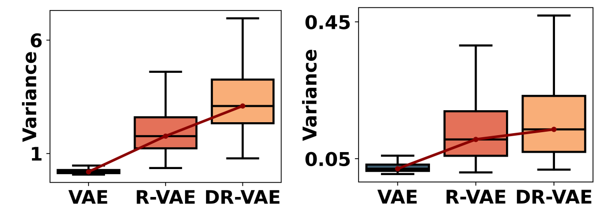

(a) ML-100K (b) KuaiRec

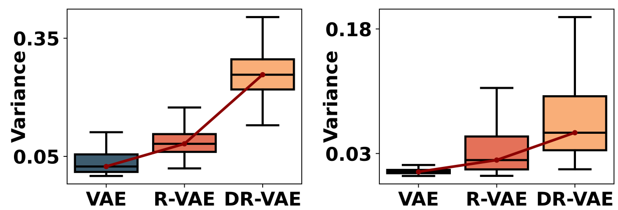

(c) Amazon Toys (d) ModCloth

Figure 4: Variance of latent variable estimation.

Experimental Results

Debiasing Performance Comparison (RQ1) To measure debiasing capability, we need biased normal data for training and unbiased intervened data for testing. Thus, *semi-synthetic* datasets are adopted to compare our proposal with state-of-the-art (SOTA) *debiasing recommendation strategies*. Additionally, MF is used as a baseline to observe the effect of methods that do not consider bias. The concrete results are shown in Tab. 2. From these results, we have two observations: **(1) DR-VAE surpasses all baselines across all metrics**, confirming the effectiveness of our proposal. This can be attributed to the Debiasing Estimator, which counters exposure bias, and the Flow-based Representation Enhancement module, which boosts representation learning in complex scenarios. **(2) All *debiasing recommendation strategies* outperform MF to varying degrees**, indicating the necessity of bias removal for providing better QoE (Quality of Experience (Chen et al. 2024c; Pang et al. 2023)) by capturing users' true preferences.

Datasets	Top@K	Metric	MF	IPS-MF	Rel-MF	MACR	BC Loss	S-DRO	InvCF.	DR-VAE
ML-100K	Top@20	Recall	0.1425	0.1438	0.1549	0.1626	0.1682	0.1746*	0.1720	**0.1796**
		Precision	0.0668	0.0706	0.0697	0.0746	0.0766	0.0813*	0.0805	**0.0825**
		NDCG	0.1062	0.1131	0.1166	0.1228	0.1291	0.1362*	0.1330	**0.1407**
		HR	0.6877	0.6974	0.7165	0.7352	0.7519	0.7575	0.7586*	**0.7812**
	Top@50	Recall	0.3008	0.3056	0.3305	0.3174	0.3176	0.3434*	0.3406	**0.3515**
		Precision	0.0610	0.0620	0.0644	0.0620	0.0631	0.0678*	0.0676	**0.0681**
		NDCG	0.1635	0.1690	0.1789	0.1755	0.1819	0.1953*	0.1919	**0.1995**
		HR	0.8483	0.8999	0.9203	0.9021	0.9087	0.9221*	0.9210	**0.9235**
KuaiRec	Top@20	Recall	0.0904	0.1602	0.1346	0.1706	0.1788	0.1911	0.1928*	**0.2060**
		Precision	0.1880	0.2164	0.2365	0.2522	0.2629	0.2774	0.2856*	**0.3152**
		NDCG	0.2072	0.2618	0.2922	0.2962	0.3326	0.3314	0.3352*	**0.3605**
		HR	0.8483	0.8894	0.9780	0.9844*	0.9681	0.9638	0.9723	**0.9951**
	Top@50	Recall	0.1672	0.2237	0.2022	0.2465	0.2352	0.2598	0.2654*	**0.2754**
		Precision	0.1482	0.1487	0.1609	0.1730	0.1673	0.1819	0.1908*	**0.2023**
		NDCG	0.2069	0.2563	0.2673	0.2868	0.3060	0.3136	0.3196*	**0.3321**
		HR	0.9440	0.9631	0.9936	0.9986*	0.9894	0.9936	0.9943	**0.9993**

Table 2: Overall recommendation performance on two *semi-synthetic* datasets. The best performance data of all approaches are highlighted in bold, while the optimal performance values of all competitive state-of-the-art baselines are marked with *.

Datasets	#Users	#Items	#Interactions	sparsity
ML-100K	943	1682	74,817	95.29%
KuaiRec	7,176	10,612	1,153,787	98.49%
Amazon Toys	2,726	13,620	52,587	99.86%
ModCloth	47,958	1,378	82,790	99.87%

Table 3: Datasets descriptions.

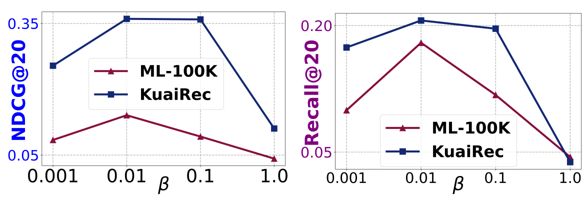

(a) NDCG@20 (b) Recall@20

Figure 5: Parameter analysis w.r.t. β.

Representational Capacity Comparison (RQ2) we focus on how our approach enhances representational capability compared to existing VAEs, without considering debiasing. Therefore, we only use the variant R-VAE and compare it with *VAE and its enhanced versions* as well as *Classical CF methods*. The latter serves as a benchmark for other representation-based recommendation approaches beyond VAEs. The results are shown in Tab. 4, from which we draw the following conclusions: **(1) Our DR-VAE outperforms MultVAE and its variants**, highlighting the inherent posterior collapse issue in vanilla VAE and demonstrating our solution's ability to mitigate it for better representations. **(2) BiVAE achieves remarkable results** due to its Constrained Adaptive Prior (CAP) module, yet it still falls short of DR-VAE, showing our method's superior handling of posterior collapse. **(3) Although the two datasets are sparse, VAE and its enhanced variants generally perform better than *classical CF methods***, owing to VAE's probabilistic framework, which enables robust generalization and generative capacity even for unseen data.

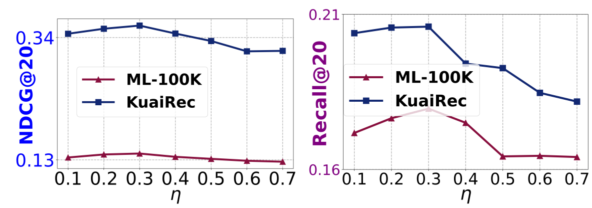

(a) NDCG@20 (b) Recall@20

Figure 6: Parameter analysis w.r.t. η.

Analysis w.r.t. Different Components (RQ3)

Firstly, for D-VAE component, we claim that D-VAE can mitigate exposure bias. To verify this, we analyzed recommendation lists following the methodology outlined by (Zhang et al. 2021). Specifically, items were divided into ten groups through two steps: first, we sorted all items in descending order based on their exposure probabilities; second, we divided the items into ten groups, ensuring the sum of exposure probabilities within each group was equal. Consequently, Group 1 has the highest average exposure probability, which decreases in subsequent groups. We then calculated the recommendation rate (RR) for each group, i.e., the probability that the items in each group are recommended. Fig. 3(a) shows the RR performance on ML-100K, while Fig. 3(b) depicts the standard deviation of RR. The black line in Fig. 3(a) indicates uniform recommendation rates during training, meaning that any large deviation suggests amplified exposure bias. From Fig. 3, we make the following observations: (1) Exposure bias is significantly amplified by the vanilla VAE, which tends to recommend the items that are easily exposed while seldom recommend the items with low exposure probabilities. This validates that the vanilla VAE is unable to handle exposure bias, thereby needing to be further improved in future research. (2) Our models, DR-VAE and D-VAEhave flatter curves and stay closer to the black line, indicating the Debiasing Estimator's effectiveness in mitigating exposure bias. (3) DR-VAE shows the most stable RR in Fig 3(a), and has the smallest variance in Fig 3(b). This stability is because the R-VAE in DR-VAE can effectively

Datasets	Top@K	Metric	MF	BPR	MultVAE	RecVAE	BiVAE	R-VAE	RI
Amazon Toys	Top@20	Recall	0.0440	0.0662	0.0647	0.0783	0.0890*	**0.1045**	17.42%
		NDCG	0.0334	0.0385	0.0311	0.0422*	0.0418	**0.0490**	16.11%
	Top@50	Recall	0.0820	0.1178	0.1315	0.1358	0.1675*	**0.2022**	20.72%
		NDCG	0.0334	0.0523	0.0469	0.0563	0.0605*	**0.0715**	18.18%
ModCloth	Top@20	Recall	0.1167	0.1897	0.2376	0.2766	0.3000*	**0.3273**	9.10%
		NDCG	0.0474	0.0830	0.1080	0.1334	0.1423*	**0.1541**	8.29%
	Top@50	Recall	0.2235	0.3198	0.4535	0.4617	0.4621*	**0.5569**	20.52%
		NDCG	0.0686	0.1089	0.1527	0.1698	0.1764*	**0.2019**	14.46%

Table 4: Overall recommendation performances on two *real-world* datasets. The best performance data of all approaches are highlighted in bold, while the optimal performance values of all competitive state-of-the-art baselines are marked with *.

address the issue of posterior collapse, allowing DR-VAE to learn information-rich representations even with insufficient items due to low exposure probabilities, thus achieving robust recommendation.

Secondly, for R-VAE component, we claim that R-VAE can handle posterior collapse and enhance the model's representational capacity. We calculate the average variance of latent variables for vanilla VAE, R-VAE and DR-VAE, which reflect the extent of variance vanishing. The experimental results on the four datasets are shown in Fig. 4. From Fig. 4, it is evident that the VAE exhibits vanishing variance across all four datasets, indicating a model collapse in the representation space. In contrast, our models, R-VAE and DR-VAE, display considerably higher average variance across these datasets, thereby validating the effectiveness of our methods in addressing the issue of posterior collapse.

Parameter Sensitivity Analyses (RQ4) We investigate how different hyperparameter values affect DR-VAE on ML-100K and KuaiRec. **Firstly**, we vary β in Eq. (11) from $0.001, 0.01, 0.1, 1$, observing in Fig. 5 that DR-VAE converges around $\beta = 0.01$. This is because an appropriate regularizer can strengthen representational power, whereas a suboptimal β can degrade performance, so we set $\beta = 0.01$ in all datasets. **Secondly**, we study the effect of η in Eq. (6), where η ranges from 0.1 to 0.7 with a step size of 0.2. As shown in Fig. 6, DR-VAE converges at $\eta = 0.3$. This is because setting η too low blurs the distinction among propensity scores, while setting it too high causes excessive variance in the Debiasing Estimator. Thus, we set $\eta = 0.3$ for the semi-synthetic datasets.

Related Work

Exposure Bias in Recommender Systems

Exposure bias is a longstanding issue in RSs, attracting significant research attention (Abdollahpouri and Mansoury 2020). In our discussion, we consider three main approaches to mitigating this bias: CF methods, Inverse Probability Weighting (IPS), and causal inference. **Firstly**, CF methods address exposure bias by incorporating additional data (e.g., social and consumption behaviors (Chen et al. 2018a)) or modifying algorithms (Liang et al. 2016). **Secondly**, IPS-based methods counteract the bias by reweighting observed data. (Schnabel et al. 2016) pioneered this approach, introducing a novel debiasing framework. Subsequent work has adapted IPS for correcting biases in Missing-Not-At-Random implicit feedback (Saito et al. 2020), extending it to pairwise recommendation, and integrating it into contrastive representation learning for CF (Lee et al. 2023). **Thirdly**, causal inference methods explore the roots of exposure bias (Yang et al. 2021; Bonner and Vasile 2018; Zheng et al. 2021), often using Structural Causal Models to untangle complex dependencies. Popularity bias, often seen as a form of exposure bias (Zhang and Shen 2023), has also been addressed through causal inference (Zhang et al. 2021; Wei et al. 2021). Despite extensive research, a significant gap remains in mitigating exposure bias while preserving the probabilistic nature of data to improve recommendation quality.

VAE and its Posterior Collapse

VAE has made significant advancements in CF as well as various CF-based smart applications such as recommender systems (Liang et al. 2018; Karamanolakis et al. 2018). The generative capacity of VAE is attributed to its ability to represent data as a distribution, rather than a fixed vector, allowing it to accommodate high uncertainty in latent space (Truong, Salah, and Lauw 2021). However, VAEs are inherently prone to the issue of posterior collapse due to their reliance on simplistic prior distributions and difficulties in learning complex data distributions (Zhao, Song, and Ermon 2019). Several researchers have studied this issue. For instance, (Kim et al. 2018) made strides in enhancing learning but didn't eradicate the posterior collapse, (He et al. 2018) addressed lagging inference networks but encountered a significant increase in time overhead, (Kingma et al. 2016) enhanced latent space flexibility but struggled with complex data distributions. However, these methods still face challenges in robustly adapting to complex data and extensive scalability, hindering their application in large-scale recommendation systems.

Conclusion

VAE improves traditional CF recommendation approaches significantly. However, current VAE-based CF models still suffer from two challenges: *exposure bias* and *posterior collapse*. In view of these challenges, we propose a DR-VAE framework, which (1) tackled *exposure bias* by introducing a Debiasing Estimator and (2) alleviated *posterior collapse* issue by introducing a Flow-based Representation Enhancement module. Through experiments on four datasets, we prove that DR-VAE significantly improves the accuracy and reliability of RSs over state-of-the-art baselines.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (No.62172362).

References

Abdollahpouri, H.; and Mansoury, M. 2020. Multi-sided Exposure Bias in Recommendation. In *KDD*.

Austin, P. C. 2011. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral research*, 46(3): 399–424.

Bonner, S.; and Vasile, F. 2018. Causal embeddings for recommendation. In *RecSys*, 104–112.

Chen, C.; Feng, X.; Li, Y.; Lyu, L.; Zhou, J.; Zheng, X.; and Yin, J. 2024a. Integration of large language models and federated learning. *Patterns*, 5(12): 101098.

Chen, C.; Zhang, Y.; Li, Y.; Wang, J.; Qi, L.; Xu, X.; Zheng, X.; and Yin, J. 2024b. Post-training attribute unlearning in recommender systems. *ACM Transactions on Information Systems*.

Chen, J.; Dong, H.; Qiu, Y.; He, X.; Xin, X.; Chen, L.; Lin, G.; and Yang, K. 2021. AutoDebias: Learning to debias for recommendation. In *SIGIR*, 21–30.

Chen, J.; Dong, H.; Wang, X.; Feng, F.; Wang, M.; and He, X. 2023. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems*, 41(3): 1–39.

Chen, J.; Feng, Y.; Ester, M.; Zhou, S.; Chen, C.; and Wang, C. 2018a. Modeling Users' Exposure with Social Knowledge Influence and Consumption Influence for Recommendation. In *CIKM*, 953–962.

Chen, J.; Pang, S.; Xi, M.; Zhao, T.; Deng, S.; and Yin, J. 2024c. Service Regulation Analysis Framework for Service Design Time: A Case Study of Internet Healthcare Service. *IEEE Transactions on Services Computing*.

Chen, J.; Wang, F.; Pang, S.; Tan, S.; Chen, M.; Zhao, T.; Xi, M.; and Yin, J. 2024d. UniGM: Unifying Multiple Pretrained Graph Models via Adaptive Knowledge Aggregation. In *MM*, 8556–8565.

Chen, R. T.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018b. Neural ordinary differential equations. *NeurIPS*, 31.

Chen, W.; Liu, Y.; Zhang, Z.; Zhuang, F.; and Zhong, J. 2024e. Modeling Adaptive Inter-Task Feature Interactions via Sentiment-Aware Contrastive Learning for Joint Aspect-Sentiment Prediction. In *AAAI*, volume 38, 17781–17789.

Chen, W.; Wu, Y.; Zhang, Z.; Zhuang, F.; He, Z.; Xie, R.; and Xia, F. 2024f. FairGap: Fairness-aware Recommendation via Generating Counterfactual Graph. *ACM Transactions on Information Systems*, 42(4): 1–25.

Chen, W.; Yuan, M.; Zhang, Z.; Xie, R.; Zhuang, F.; Wang, D.; and Liu, R. 2024g. FairDgcl: Fairness-aware Recommendation with Dynamic Graph Contrastive Learning. *arXiv preprint arXiv:2410.17555*.

Fan, J.; Imai, K.; Lee, I.; Liu, H.; Ning, Y.; and Yang, X. 2022a. Optimal covariate balancing conditions in propensity score estimation. *Journal of Business & Economic Statistics*, 41(1): 97–110.

Fan, W.; Wang, P.; Wang, D.; Wang, D.; Zhou, Y.; and Fu, Y. 2023. Dish-TS: A General Paradigm for Alleviating Distribution Shift in Time Series Forecasting. In *AAAI*, volume 37, 7522–7529.

Fan, W.; Zheng, S.; Yi, X.; Cao, W.; Fu, Y.; Bian, J.; and Liu, T.-Y. 2022b. DEPTS: Deep Expansion Learning for Periodic Time Series Forecasting. In *ICLR*.

Gao, C.; Li, S.; Lei, W.; Chen, J.; Li, B.; Jiang, P.; He, X.; Mao, J.; and Chua, T.-S. 2022. KuaiRec: A Fully-Observed Dataset and Insights for Evaluating Recommender Systems. In *CIKM*, 540–550.

Gong, W.; Geng, Y.; Zhang, D.; Zhu, Y.; Xu, X.; Xiang, H.; Beheshti, A.; Zhang, X.; and Qi, L. 2025. DivGCL: A Graph Contrastive Learning Model for Diverse Recommendation. In *AAAI*.

Gong, W.; Zhang, X.; Chen, Y.; He, Q.; Beheshti, A.; Xu, X.; Yan, C.; and Qi, L. 2022. Dawar: Diversity-aware web apis recommendation for mashup creation based on correlation graph. In *SIGIR*, 395–404.

Grathwohl, W.; Chen, R. T.; Bettencourt, J.; Sutskever, I.; and Duvenaud, D. 2018. FFJORD: Free-Form Continuous Dynamics for Scalable Reversible Generative Models. In *ICLR*.

Harper, F. M.; and Konstan, J. A. 2015. The MovieLens Datasets: History and Context. 5(4).

He, J.; Spokoyny, D.; Neubig, G.; and Berg-Kirkpatrick, T. 2018. Lagging Inference Networks and Posterior Collapse in Variational Autoencoders. In *ICLR*.

Karamanolakis, G.; Cherian, K. R.; Narayan, A. R.; Yuan, J.; Tang, D.; and Jebara, T. 2018. Item recommendation with variational autoencoders and heterogeneous priors. In *RecSys Workshop*, 10–14.

Kim, Y.; Wiseman, S.; Miller, A.; Sontag, D.; and Rush, A. 2018. Semi-amortized variational autoencoders. In *ICML*, 2678–2687. PMLR.

Kingma, D. P.; Salimans, T.; Jozefowicz, R.; Chen, X.; Sutskever, I.; and Welling, M. 2016. Improved variational inference with inverse autoregressive flow. *NeurIPS*, 29.

Kingma, D. P.; and Welling, M. 2013. Auto-Encoding Variational Bayes. *CoRR*, abs/1312.6114.

Lee, J.-w.; Park, S.; Yoon, M.; and Lee, J. 2023. In *SIGIR*, 2456–2460.

Li, H.; Lyu, Y.; Zheng, C.; and Wu, P. 2022. TDR-CL: Targeted Doubly Robust Collaborative Learning for Debiased Recommendations. *ICLR*.

Li, Y.; Chen, C.; Zhang, Y.; Liu, W.; Lyu, L.; Zheng, X.; Meng, D.; and Wang, J. 2023. Ultrare: Enhancing receraser for recommendation unlearning via error decomposition. *NeurIPS*, 36.

Liang, D.; Charlin, L.; McInerney, J.; and Blei, D. M. 2016. Modeling user exposure in recommendation. In *WWW*, 951–961.

Liang, D.; Krishnan, R. G.; Hoffman, M. D.; and Jebara, T. 2018. Variational autoencoders for collaborative filtering. In *WWW*, 689–698.

Liao, X.; Liu, W.; Chen*, C.; Zhou, P.; Yu, F.; Zhu, H.; Yao, B.; Wang, T.; Zheng, X.; and Tan, Y. 2024a. Rethinking the Representation in Federated Unsupervised Learning with Non-IID Data. In *CVPR*.

Liao, X.; Liu, W.; Zhou, P.; Yu, F.; Xu, J.; Wang, J.; Wang, W.; Chen, C.; and Zheng, X. 2024b. FOOGD: Federated Collaboration for Both Out-of-distribution Generalization and Detection. In *NeurIPS*.

Liu, W.; Chen, C.; Liao, X.; Hu, M.; Tan, Y.; Wang, F.; Zheng, X.; and Ong, Y.-S. 2024a. Learning Accurate and Bidirectional Transformation via Dynamic Embedding Transportation for Cross-Domain Recommendation. In *AAAI*.

Liu, W.; Zheng, X.; Chen, C.; Su, J.; Liao, X.; Hu, M.; and Tan, Y. 2023. Joint Internal Multi-Interest Exploration and External Domain Alignment for Cross Domain Sequential Recommendation. In *WWW*, 383–394.

Liu, W.; Zheng, X.; Chen, C.; Xu, J.; Liao, X.; Wang, F.; Tan, Y.; and Ong, Y.-S. 2024b. Reducing Item Discrepancy via Differentially Private Robust Embedding Alignment for Privacy-Preserving Cross Domain Recommendation. In *ICML*.

Liu, Y.; Qi, L.; Liu, W.; Xu, X.; Zhang, X.; and Dou, W. 2024c. GraphSAGE-based POI Recommendation via Continuous-Time Modeling. In *WWW*, 585–588.

Mao, X.; Liu, Y.; Qi, L.; Duan, L.; Xu, X.; Zhang, X.; Dou, W.; Beheshti, A.; and Zhou, X. 2024. Cluster-driven Personalized Federated Recommendation with Interest-aware Graph Convolution Network for Multimedia. In *MM*, 5614–5622.

Misra, R.; Wan, M.; and McAuley, J. 2018. Decomposing Fit Semantics for Product Size Recommendation in Metric Spaces. 422–426.

Pang, S.; Zhao, X.; Luo, J.; Zheng, B.; Yin, J.; and Zheng, X. 2023. Incentive-Driven Pricing Game for Multi-Edge Service Providers towards Optimal Profits. In *ICWS*, 350–359. IEEE.

Qi, L.; Liu, Y.; Liu, W.; Pei, S.; Xu, X.; Zhang, X.; Wang, Y.; and Dou, W. 2024. Counterfactual user sequence synthesis augmented with continuous time dynamic preference modeling for sequential POI recommendation. In *IJCAI*, 2306–2314.

Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. 452–461.

Ruining, H.; and Julian, M. 2016. Matrix factorization techniques for recommender systems. In *WWW*.

Saito, Y.; Yaginuma, S.; Nishino, Y.; Sakata, H.; and Nakata, K. 2020. Unbiased recommender learning from missing-not-at-random implicit feedback. In *WSDM*, 501–509.

Schnabel, T.; Swaminathan, A.; Singh, A.; Chandak, N.; and Joachims, T. 2016. Recommendations as treatments: Debiasing learning and evaluation. In *ICML*, 1670–1679.

Seaman, S. R.; and White, I. R. 2013. Review of inverse probability weighting for dealing with missing data. *Statistical methods in medical research*, 22(3): 278–295.

Shenbin, I.; Alekseev, A.; Tutubalina, E.; Malykh, V.; and Nikolenko, S. I. 2020. RecVAE: A New Variational Autoencoder for Top-N Recommendations with Implicit Feedback. 528–536.

Song, T.; Qi, L.; Liu, W.; Wang, F.; Xu, X.; Zhang, X.; Beheshti, A.; Zhou, X.; and Dou, W. 2025. Enhancing Diffusion Model with Auxiliary Information Mining-Exploration and Efficient Sampling Mechanism for Sequential Recommendation. In *AAAI*.

Truong, Q.-T.; Salah, A.; and Lauw, H. W. 2021. Bilateral variational autoencoder for collaborative filtering. In *WSDM*, 292–300.

Wang, W.; Lin, X.; Wang, L.; Feng, F.; Ma, Y.; and Chua, T.-S. 2023. Causal Disentangled Recommendation Against User Preference Shifts. *ToIS*.

Wang, Z.; and Ziyin, L. 2022. Posterior collapse of a linear latent variable model. *NeurIPS*, 35: 37537–37548.

Wei, T.; Feng, F.; Chen, J.; Wu, Z.; Yi, J.; and He, X. 2021. Model-Agnostic Counterfactual Reasoning for Eliminating Popularity Bias in Recommender System. 1791–1800.

Wen, H.; Yi, X.; Yao, T.; Tang, J.; Hong, L.; and Chi, E. H. 2022. Distributionally-Robust Recommendations for Improving Worst-Case User Experience. 3606–3610.

Yang, Y.; Li, M.; Hu, X.; Pan, G.; Huang, W.; Wang, J.; and Wang, Y. 2021. Exploring Exposure Bias in Recommender Systems from Causality Perspective. In *QRS-C*, 425–432.

Zhang, A.; Ma, W.; Wang, X.; and Chua, T.-S. 2022. Incorporating bias-aware margins into contrastive loss for collaborative filtering. *NeurIPS*, 35: 7866–7878.

Zhang, A.; Zheng, J.; Wang, X.; Yuan, Y.; and Chua, T.-S. 2023. Invariant Collaborative Filtering to Popularity Distribution Shift. 1240–1251.

Zhang, F.; and Shen, Q. 2023. A Model-Agnostic Popularity Debias Training Framework for Click-Through Rate Prediction in Recommender System. In *SIGIR*, 1760–1764.

Zhang, Y.; Feng, F.; He, X.; Wei, T.; Song, C.; Ling, G.; and Zhang, Y. 2021. Causal intervention for leveraging popularity bias in recommendation. In *SIGIR*, 11–20.

Zhao, S.; Song, J.; and Ermon, S. 2019. Infovae: Balancing learning and inference in variational autoencoders. In *AAAI*, volume 33, 5885–5892.

Zheng, X.; Hu, M.; Liu, W.; Chen, C.; and Liao, X. 2023. Robust Representation Learning with Reliable Pseudo-labels Generation via Self-Adaptive Optimal Transport for Short Text Clustering. In *ACL*.

Zheng, Y.; Gao, C.; Li, X.; He, X.; Li, Y.; and Jin, D. 2021. Disentangling user interest and conformity for recommendation with causal embedding. In *WWW*, 2980–2991.

Zigler, C. M.; Watts, K.; Yeh, R. W.; Wang, Y.; Coull, B. A.; and Dominici, F. 2013. Model feedback in Bayesian propensity score estimation. *Biometrics*, 69(1): 263–273.