

ADAPTIVE CAUSAL BALANCING FOR COLLABORATIVE FILTERING

Anonymous authors

Paper under double-blind review

ABSTRACT

Collaborative filtering builds personalized models from the collected user feedback. However, the collected data is observational rather than experimental, leading to various biases in the data, which can significantly affect the learned model. To address this issue, many studies have focused on propensity-based methods to combat the selection bias by reweighting the sample loss, and demonstrate that balancing is important for debiasing both theoretically and empirically. However, two fundamental questions remain: which function class should be balanced, and how can this balance be effectively achieved? In this paper, we first perform theoretical analysis to show the effect of balancing finite-dimensional function classes on the bias of IPS and DR methods, and based on this, we propose a universal kernel-based balancing method to balance functions on the reproducing kernel Hilbert space. In addition, we propose a novel adaptive **kernel balancing** method during the alternating update between unbiased evaluation and training of the prediction model. Specifically, the prediction loss of the model is projected in the kernel-based covariate function space, and the projection coefficients are used to determine which functions should be prioritized for balancing to reduce the estimation bias. We conduct extensive experiments on three real-world datasets to demonstrate the effectiveness of the proposed approach.

1 INTRODUCTION

In the era of information explosion, recommender systems (RSs) have become a core component of many online platforms, such as social media, e-commerce, and music streaming. The goal of RS is to recommend users content or products that they may be interested in. However, due to the presence of the selection bias, the collected data cannot represent the target population, which poses a great challenge in the RS training phase (Marlin and Zemel, 2009; Schnabel et al., 2016; Lin et al., 2023).

Many methods have been proposed in order to mitigate the selection bias (Schnabel et al., 2016; Dai et al., 2022; Wang et al., 2019; Li et al., 2023a; Saito, 2019). Among them, error imputation based (EIB) methods first impute pseudo-labels to the missing events and then combine the pseudo-labels and observed labels for model training (Marlin et al., 2007; Chang et al., 2010; Steck, 2010). However, it is impractical to obtain accurate pseudo-labelings due to user self-selection and item popularity in practice, thus will introduce bias to the EIB methods. Another frequently-used debiasing methods are based on the inverse propensity score (IPS), which reweights the observed events to achieve unbiased estimation (Oosterhuis, 2022; Saito et al., 2020; Li et al., 2023e; Luo et al., 2021). However, due to data sparsity, propensity scores can hardly be estimated accurately and will always have extreme values, which results in high variance for IPS methods. Doubly robust (DR) methods overcome the shortcomings of EIB and IPS by combining pseudo-labeling and propensity together to achieve smaller biases and variances (Saito, 2020; Guo et al., 2021; Wang et al., 2019; 2022; Li et al., 2023b). However, the DR methods are also biased when the learned pseudo-labeling model and learned propensity model are inaccurate (Wang et al., 2019; Saito, 2020).

The propensity model plays an crucial role in propensity-based debiased methods. Previous causal inference literature have theoretically and empirically demonstrated the importance of balancing property of propensity (Imai and Ratkovic, 2014; Li et al., 2018; 2023c). Nonetheless, the balancing property of propensity is rarely discussed and exploited in the field of debiased recommendation. To fill this gap, a recent work (Li et al., 2023d) proposes a propensity balancing measurement to

regularize the IPS and DR estimators, and achieves the state-of-the-art performance. However, on one hand, it is not realistic to balance all possible functions for a specific model using only finite samples. On the other hand, balancing only an arbitrary function is not sufficient for IPS and DR methods to achieve unbiased learning (see Corollary 1 for the formal theoretical results). Thus, it is necessary to discuss which functions should be more favored to be balanced for the IPS and DR estimators, resulting in smaller estimation biases of the ideal loss and enhanced unbiased learning.

In this paper, we theoretically analyze the advantages of balancing kernel functions, and propose universal kernel-based balancing methods, namely **KBIPS** and **KBDR**, to balance functions on the reproducing kernel Hilbert space (RKHS), which adaptively selects the functions that most need to be balanced to reduce the estimation biases of the previous IPS and DR estimators. Moreover, we propose a novel entropy-based optimization problem to effectively balance the selected functions. Finally, we perform theoretical analysis showing that the learned kernel balanced propensities can reduce the generalization bound to improve the performance of debiased learning.

The main contributions of this paper are shown below:

- We theoretically analyze which functions need to be balanced for IPS and DR methods, and proposes a kernel balancing method to adaptively select balancing functions.
- We further propose an entropy-based optimization approach to effectively balance the selected balancing functions, and derive the generalization error bound.
- We conduct extensive experiment to verify the effectiveness of the proposed estimators and balancing approach on three real-world datasets, including a large-scale industrial dataset.

2 RELATED WORK

Debiased Recommendation. The collected data from RS is observational rather than experimental, leading to various biases in the data, which seriously affect the quality of the learned recommendation model. Many previous methods are proposed to mitigate the selection bias problem (Marlin and Zemel, 2009; Schnabel et al., 2016; Saito, 2019; Chen et al., 2021; Wang et al., 2019). The error imputation based (EIB) methods attempt to impute the missing events, and then train a recommendation model on both observed and imputed data (Chang et al., 2010; Steck, 2010; Hernández-Lobato et al., 2014). Another common debiasing methods are propensity-based, including inverse propensity scoring (IPS) methods (Imbens and Rubin, 2015; Schnabel et al., 2016; Saito et al., 2020; Oosterhuis, 2022; Luo et al., 2021), and doubly robust (DR) methods (Morgan and Winship, 2015; Wang et al., 2019; Saito, 2020). Specifically, IPS adjusts the distribution to all events by reweighting the observed events, while DR combines the EIB and IPS methods, which takes advantage of both, i.e., has lower variance and bias. Based on the above advantages, many competing DR-based methods are proposed, such as MRDR (Guo et al., 2021), DR-BIAS (Dai et al., 2022), ESCM²-DR (Wang et al., 2022), TDR (Li et al., 2023b) and SDR (Li et al., 2023e). Given the widespread of the propensity model, Li et al. (2023d) proposes a propensity balancing measurement to regularize the IPS and DR estimators. In this paper, we extend Li et al. (2023d) by proposing novel kernel balancing-based IPS and DR estimators that adaptively find the balancing functions that contribute the most to reducing the estimation bias.

Covariate Balancing in Causal Inference. Balancing refers to aligning the distribution of covariates in the treatment and control groups, which is a theme central to the estimation of causal effects based on observational studies (Stuart, 2010; Imbens and Rubin, 2015). This is because balancing ensures that units receiving different treatments are comparable directly, and the association becomes causation under unconfoundedness assumption (Imai and Ratkovic, 2014; Hernán and Robins, 2020). In randomized controlled experiments, balance is naturally maintained due to the complete random assignment of treatments. However, in observational studies, treatment groups typically exhibit systematic differences in covariates, which can result in a lack of balance. To obtain accurate estimates of causal effects in observational studies, a wide variety of methods have emerged for balancing the finite order moments of covariates, including matching (Rosenbaum and Rubin, 1983; Stuart, 2010; Wu et al., 2020), stratification (Hernán and Robins, 2020), entropy balancing (Hainmueller, 2012; Zhao and Percival, 2017), covariate balancing propensity (Imai and Ratkovic, 2014), and weighted euclidean balancing (Chen and Zhou, 2023). In recent years, several approaches were developed balancing infinite order moments of covariates (Sant’Anna et al., 2022) or the covariates distributions (Wong and Chan, 2018). However, it is unrealistic to balance infinite order moments with only finite samples, so this paper proposes a novel balancing method that adaptively finds the balancing functions in RKHS that are most important for achieving unbiased learning.

3 PRELIMINARIES

3.1 DEBIASED RECOMMENDATION

Suppose $\mathcal{U} = \{u_1, u_2, \dots\}$, $\mathcal{I} = \{i_1, i_2, \dots\}$ are the user set and item set, and $\mathcal{D} = \{u_1, u_2, \dots, u_m\} \times \{i_1, i_2, \dots, i_n\}$ are sampled from the whole user set and item set, respectively. Let $x_{u,i}$ and $r_{u,i} \in \{0, 1\}$ be the feature and rating of user-item pair (u, i) , where $r_{u,i} = 1$ or 0 represents whether user u likes (or purchases) item i or not. **Let $o_{u,i}$ be a Bernoulli variable indicating whether the true rating $r_{u,i}$ is observed $o_{u,i} = 1$ or missing $o_{u,i} = 0$, and $\mathcal{O} = \{(u, i) \mid (u, i) \in \mathcal{D}, o_{u,i} = 1\}$ be the set of the user-item pairs with observed ratings.** Thus we can only observe $r_{u,i}$ when $o_{u,i} = 1$.

Let $\hat{r}_{u,i} = f(x_{u,i}; \theta)$ be the prediction model that aims to predict all $r_{u,i}$ accurately. If all the ratings are observed, we can train the prediction model directly by minimizing the following ideal loss:

$$\mathcal{L}_{\text{Ideal}}(\theta) = \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} e_{u,i},$$

where $e_{u,i} = \mathcal{L}(\hat{r}_{u,i}, r_{u,i})$ is the prediction error with $\mathcal{L}(\cdot, \cdot)$ be an arbitrary loss function, e.g., mean square loss or cross entropy loss. However, the ratings are missing for all $(u, i) \in \mathcal{D} \setminus \mathcal{O}$, thus $e_{u,i}$ are also not accessible. A naive estimator is to estimate ideal loss directly from the observed samples, which is shown as follows:

$$\mathcal{L}_{\text{Naive}}(\theta) = \frac{1}{|\mathcal{O}|} \sum_{(u,i) \in \mathcal{O}} e_{u,i},$$

However, there is always a discrepancy between observed events \mathcal{O} and all events \mathcal{D} , due to the existence of confounders that affect both treatment and outcome. Therefore, the naive estimator always is not an unbiased estimation of the ideal loss (Wang et al., 2019). The inverse propensity score (IPS) (Schnabel et al., 2016) and doubly robust (DR) learning (Wang et al., 2019; 2021) are two main propensity-based strategies to eliminate such discrepancy. The IPS estimator is given below:

$$\mathcal{L}_{\text{IPS}}(\theta) = \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \frac{o_{u,i} e_{u,i}}{\hat{p}_{u,i}},$$

where $\hat{p}_{u,i} = \pi(x_{u,i}; \phi_p)$ is the estimation of propensity score $p_{u,i} := \mathbb{P}(o_{u,i} = 1 \mid x_{u,i})$. The IPS method is unbiased when all the estimate propensities are correct. The DR estimator further combines the imputed error, which is constructed below:

$$\mathcal{L}_{\text{DR}}(\theta) = \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \left[\hat{e}_{u,i} + \frac{o_{u,i} \cdot (e_{u,i} - \hat{e}_{u,i})}{\hat{p}_{u,i}} \right],$$

where $\hat{e}_{u,i} = m(x_{u,i}; \phi_e)$ is the imputed error. The DR estimator is unbiased when all the estimated propensities or the imputed errors are accurate.

3.2 CAUSAL BALANCING

The propensity score can be used to adjust confounding and recover the distribution. Despite their popularity and theoretical appeal, a main practical difficulty is that the propensity score must be estimated. From a causal perspective, the true propensity score will have balancing property (Imai and Ratkovic, 2014; Imbens and Rubin, 2015; Rosenbaum, 2020; Sant’Anna et al., 2022). Specifically, for all $\phi : \mathcal{X} \rightarrow \mathbb{R}^m$, the true propensity satisfies

$$\mathbb{E} \left[\frac{o_{u,i} \phi(x_{u,i})}{p_{u,i}} \right] = \mathbb{E} \left[\frac{(1 - o_{u,i}) \phi(x_{u,i})}{1 - p_{u,i}} \right] = \mathbb{E}[\phi(x_{u,i})]. \quad (1)$$

Inspiring by this, Li et al. (2023d) adopt the difference of the left term and middle term as a regularization term during the propensity training phase for IPS and DR methods. Meanwhile, they select several $\phi(x)$ manually and verifies the effectiveness of balancing. However, there are infinite $\phi(x)$ need to be balanced, and we often cannot balance too much $\phi(x)$ because of the computational cost. Meanwhile, the balancing property requires exact matching on the estimated propensity score, which is typically impossible (Imai and Ratkovic, 2014). Therefore, it is essential to investigate which function class should be balanced and how to efficiently balance that function class.

4 THE PROPOSED ADAPTIVELY KERNEL BALANCING METHOD

4.1 THE RELATIONSHIP BETWEEN BALANCING PROPERTY AND CROSS ENTROPY LOSS

Many propensity-based methods adopt cross entropy loss in the training phase of propensity model $\pi(x_{u,i}; \phi_p)$. The origin propensity model loss is shown below:

$$\mathcal{L}_p(\phi_p) = \sum_{(u,i) \in \mathcal{D}} -o_{u,i} \log \{\pi(x_{u,i}; \phi_p)\} - (1 - o_{u,i}) \log \{1 - \pi(x_{u,i}; \phi_p)\}.$$

Taking the first derivative of the loss function according to ϕ_p , the following equation is obtained:

$$\frac{\partial \mathcal{L}_p(\phi_p)}{\partial \phi_p} = \sum_{(u,i) \in \mathcal{D}} -\frac{o_{u,i} \pi'(x_{u,i}; \phi_p)}{\pi(x_{u,i}; \phi_p)} + \frac{(1 - o_{u,i}) \pi'(x_{u,i}; \phi_p)}{1 - \pi(x_{u,i}; \phi_p)}.$$

If we minimize $\mathcal{L}_p(\phi_p)$, the first derivative should be zero. Notably, it ensures the balancing property of the propensity with $\phi(x) = \pi'(x_{u,i}; \phi_p)$ in Equation 1. Thus, though not formally discussed, many propensity-based methods can be seen as a special case of propensity balancing.

4.2 BALANCING FINITE FUNCTIONS

Balancing only one $\phi(x)$ is not sufficient for ensuring the propensity balancing property. Before the discussion of how to choose $\phi(x)$, we first propose two general estimators named **Kernel Balancing IPS (KBIPS)** and **Kernel Balancing DR (KBDR)**:

$$\begin{aligned} \mathcal{L}_{\text{KBIPS}}(\theta) &= \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} \hat{w}_{u,i} e_{u,i}, \\ \mathcal{L}_{\text{KBDR}}(\theta) &= \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} [\hat{e}_{u,i} + o_{u,i} \hat{w}_{u,i} (e_{u,i} - \hat{e}_{u,i})]. \end{aligned} \quad (2)$$

When $\hat{w}_{u,i} = \frac{1}{\hat{p}_{u,i}}$, the **KBIPS** and **KBDR** will degenerate to vanilla IPS and DR. The following theorem shows the bias of these two estimators.

Theorem 1. *The biases of the **KBIPS** and **KBDR** are shown as follows:*

$$\begin{aligned} \text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta)) &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} - 1) e_{u,i} \right\}^2, \\ \text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta)) &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} - 1) (e_{u,i} - \hat{e}_{u,i}) \right\}^2. \end{aligned}$$

We then discuss how to choose finite appropriate $\phi(x)$ to learn an appropriate propensity model that **achieves** lower estimation bias. Specifically, let $h^{(j)}(x_{u,i})$, $j = 1, \dots, J$ be the selected balancing functions, **motivated by the widely-adopted entropy balancing method (Hainmueller, 2012)**, we propose to learn the balancing weights $\hat{w}_{u,i} = g(x_{u,i}; \phi_w)$ by solving the optimization problem:

$$\begin{aligned} \max_{\phi_w} & - \sum_{(u,i) \in \mathcal{O}} g(x_{u,i}; \phi_w) \log g(x_{u,i}; \phi_w) \\ \text{s.t.} & g(x_{u,i}; \phi_w) \geq 0 \quad \text{for all } (u,i) \in \mathcal{O}, \\ & \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} g(x_{u,i}; \phi_w) = 1, \\ & \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} g(x_{u,i}; \phi_w) h^{(j)}(x_{u,i}) = \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} h^{(j)}(x_{u,i}), \quad j \in \{1, \dots, J\}. \end{aligned}$$

The above optimization problem consists of three features. First, the objective function is the empirical entropy of the balancing weights, by the principle of maximum entropy (Guiasu and Shenitzer, 1985),

it reaches the maximum value when the balancing weights are uniform, thus effectively avoiding high variance due to extremely small propensities. Second, the balancing constraints are imposed to equalize the selected covariate functions between the observed and missing samples. Third, two normalization constraints are imposed, which implies that the weights sum to the normalization constant of one, and the nonnegativity of the balancing weights, making the empirical entropy as the objective function be well-defined. Remarkably, it is a convex optimization with respect to $\hat{w}_{u,i} = g(x_{u,i}; \phi_w)$, thus we can adopt Lagrange multiplier to solve this problem. The following corollary shows the bias depends on the distance between $e_{u,i}$ and $\mathcal{H}_J = \text{span}\{h^{(1)}(x_{u,i}), \dots, h^{(J)}(x_{u,i})\}$.

Corollary 1 (Main Result). *If $e_{u,i}(x) \in \mathcal{H}_J = \text{span}\{h^{(1)}(x), \dots, h^{(J)}(x)\}$, then proposed KBIPS is unbiased; otherwise, KBIPS is biased. Similarly, if $e_{u,i}(x) - \hat{e}_{u,i}(x) \in \mathcal{H}_J = \text{span}\{h^{(1)}(x), \dots, h^{(J)}(x)\}$, then proposed KBDR is unbiased; otherwise, KBDR is biased.*

Next, taking KBIPS as an example (KBDR follows from a similar argument), two natural questions are how to select those $h(x_{u,i})$ and what should we do when $e_{u,i}(x) \notin \mathcal{H}_J = \text{span}\{h^{(1)}(x_{u,i}), \dots, h^{(J)}(x_{u,i})\}$, but $e_{u,i}(\cdot)$ in a larger hypothesis space instead? Fortunately, kernel balancing gives a feasible solution.

4.3 KERNEL BALANCING

First, we formally provide the definition of the kernel function as below.

Definition 1. *Let \mathcal{X} be a non-empty set. A function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a kernel if there exists a Hilbert space \mathcal{H} and a feature map $\phi : \mathcal{X} \rightarrow \mathcal{H}$ such that $\forall x, x' \in \mathcal{X}$, $K(x, x') := \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}$.*

Typical kernel functions consist of Gaussian kernel and exponential kernel with the explicit forms

$$K^{\text{Gau}}(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\sigma^2}\right) \quad \text{and} \quad K^{\text{Exp}}(x, x') = \exp\left(-\frac{\|x - x'\|}{2\sigma^2}\right).$$

By noticing that $e_{u,i}(x)$ is a continue function of $x_{u,i}$ under mean square loss or cross entropy loss, it is important to have a guarantee that the space \mathcal{H}_J is rich enough, i.e., there is a function in \mathcal{H}_J can approach an arbitrary continue function. We next introduce the definition of universal kernel below:

Definition 2. *For \mathcal{X} compact Hausdorff, a kernel is universal if for any continuous function $e : \mathcal{X} \rightarrow \mathbb{R}$ and $\epsilon > 0$, there exists $f \in \mathcal{H}$ in the corresponding RKHS such that $\sup_{x \in \mathcal{X}} |f(x) - e(x)| \leq \epsilon$.*

The following lemma shows that there are some widely-used kernel functions have such property.

Lemma 1. *Both the Gaussian and exponential kernel are universal (Sriperumbudur et al., 2011).*

Until now, we show that there is a function in RKHS $\mathcal{H} = \text{span}\{K(\cdot, x) \mid x \in \mathcal{X}\}$ that can approach any continuous function when we adopt Gaussian or exponential kernel, where $K(\cdot, x)$ denotes a kernel function. However, \mathcal{H} might be a infinity dimension space with $|\mathcal{X}| = \infty$, which leads to an infinity constrains for the optimization problem. The following representer theorem shows guarantee that kernel methods retain optimality under penalized empirical risk minimization, and provide a form of the best-possible choice of kernel balancing under finite samples.

Lemma 2 (Representer theorem). *If $\Omega = h(\|f\|)$ for some increasing function $h : \mathbb{R}_+ \rightarrow \overline{\mathbb{R}}$, then some empirical risk minimizer must admit the form $f(\cdot) = \sum_{i=1} \alpha_i \kappa(\cdot, x_i)$ for some $\alpha \in \mathbb{R}^n$. If h is strictly increasing, all minimizers admit this form.*

The universal kernel property and representer theorem assure that when we minimize the bias of KBIPS, the following equation holds:

$$\begin{aligned} \min_{w \geq 0} (\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta))) &= \min_{w \geq 0} \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} w_{u,i} - 1) e_{u,i}(x) \right\}^2 \quad (\text{by the bias definition}) \\ &\approx \min_{w \geq 0} \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} w_{u,i} - 1) \left(\sum_{(s,t) \in \mathcal{U} \times \mathcal{I}} \alpha_{s,t} K(x, x_{s,t}) \right) \right\}^2 \quad (\text{by the universal property}) \\ &= \min_{w \geq 0} \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} w_{u,i} - 1) \left(\sum_{(s,t) \in \mathcal{D}} \alpha_{s,t} K(x, x_{s,t}) \right) \right\}^2 \quad (\text{by the representer theorem}), \end{aligned}$$

where $(\alpha_{1,1}, \dots, \alpha_{m,n})$ minimizes the following mean squared loss:

$$(\alpha_{1,1}, \dots, \alpha_{m,n}) = \arg \min_{\alpha \in \mathbb{R}^{|\mathcal{D}|}} \sum_{(u,i) \in \mathcal{D}} \left(e_{u,i} - \sum_{(s,t) \in \mathcal{D}} \alpha_{s,t} K(x_{u,i}, x_{s,t}) \right)^2. \quad (3)$$

Finally, we can construct an optimization problem in Section 4.2 by selecting the kernel function class $\{K(x_{u,i}, \cdot) : (u,i) \in \mathcal{D}\}$ to balance to learn a balancing propensity, which provides a solution for which function class should be balanced.

4.4 THREE PROPOSED **KERNEL BALANCING** METHODS

In practice, it is hard to balance all $|\mathcal{D}|$ functions in a single batch. Thus, we propose three **kernel balancing** methods to effectively balance the kernel functions.

Kernel Balancing. We first propose a simple **kernel balancing** method, **which randomly chooses J** functions from $\text{span}\{K(\cdot, x_{u,i}) : (u,i) \in \mathcal{D}\}$ to balance. However, this simple method regards all kernel functions equally important and is independent of the coefficients $\alpha_{s,t}$ of the kernel function $K(\cdot, x_{s,t})$, which has no guarantee for learning a high quality balancing propensity.

Worst-Case Kernel Balancing. Different from finding the optimal kernel function that needs to be balanced, the worst-case **kernel balancing** method focus on controlling the worst case. Specifically, we propose to play a minimax game as shown below:

$$\min_{w \geq 0} \left[\sup_{e \in \mathcal{H}} \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} w_{u,i} - 1) e(x_{u,i}) \right\}^2 \right] = \min_{w \geq 0} \left[\sup_{e \in \mathcal{H}} \frac{\left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} w_{u,i} - 1) e(x_{u,i}) \right\}^2}{\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} e(x_{u,i})^2} \right],$$

where $\tilde{\mathcal{H}} = \{e(\cdot) \in \mathcal{H} : \|e(\cdot)\|_N^2 = |\mathcal{D}|^{-1} \sum_{(u,i) \in \mathcal{D}} e(x_{u,i})^2 = 1\}$.

By the universal property and representer theorem in Section 4.3, the right hand side of the above equation is the same as the following:

$$\min_{w \geq 0} \left[\sup_{\alpha_{s,t}} \frac{\left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} w_{u,i} - 1) \sum_{(s,t) \in \mathcal{D}} \alpha_{s,t} K(x_{u,i}, x_{s,t}) \right\}^2}{\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} e(x_{u,i})^2} \right],$$

One should be noticed that this method is too conservative and unstable. If we only solve the equation above, $\sum_{(s,t) \in \mathcal{D}} \alpha_{s,t} K(x_{u,i}, x_{s,t})$ will not be able to approach the $e_{u,i}(x)$ because the $\alpha_{s,t}$ cannot minimize the mean square loss in Section 4.3.

Adaptive Kernel Balancing. To overcome the shortcomings of the above two methods, we propose a novel kernel balancing method that can adaptively select which kernel function to balance. Specifically, given current prediction model $f(x_{u,i}; \theta)$, first fit $e_{u,i}(x)$ using $\text{span}\{K(\cdot, x_{u,i}) : (u,i) \in \mathcal{D}\}$, then balance the J functions with maximal $|\alpha_{s,t}|$, where J is a hyper-parameter. This proposed method will balance the kernel functions that contributes the most to the $e_{u,i}$, thus improves the efficiency. The optimization problem is constructed below:

$$\begin{aligned} \min_{\phi_w} \quad & \sum_{(u,i) \in \mathcal{O}} g(x_{u,i}; \phi_w) \log g(x_{u,i}; \phi_w) + \gamma \sum_{j=1}^J \xi_j \\ \text{s.t.} \quad & \xi_j \geq 0, \quad \text{for all } j \in \{1, \dots, J\} \quad \text{and} \quad g(x_{u,i}; \phi_w) \geq 0 \quad \text{for all } (u,i) \in \mathcal{O}, \\ & \sum_{(u,i) \in \mathcal{D}} o_{u,i} g(x_{u,i}; \phi_w) = 1, \\ & \sum_{(u,i) \in \mathcal{D}} o_{u,i} g(x_{u,i}; \phi_w) h^{(j)}(x_{u,i}) - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} h^{(j)}(x_{u,i}) \leq C + \xi_j \quad j \in \{1, \dots, J\}, \\ & \sum_{(u,i) \in \mathcal{D}} o_{u,i} g(x_{u,i}; \phi_w) h^{(j)}(x_{u,i}) - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} h^{(j)}(x_{u,i}) \geq -C + \xi_j \quad j \in \{1, \dots, J\}, \end{aligned}$$

Algorithm 1: The Proposed Adaptive **KBDR** (**AKBDR**) Learning Algorithm**Input:** observed ratings \mathbf{Y}^o , and the hyper-parameter J .

```

1 while stopping criteria is not satisfied do
2   for number of steps for training the imputation model do
3     Sample a batch of user-item pairs  $\{(u_i, i_i)\}_{i=1}^I$  from  $\mathcal{O}$ ;
4     Update  $\phi_e$  by descending along the gradient  $\nabla_{\phi_e} \mathcal{L}_e(\phi_e, \phi_w, \theta)$ ;
5   end
6   for number of steps for training the balancing weight model do
7     Sample a batch of user-item pairs  $\{(u_k, i_k)\}_{k=1}^K$  from  $\mathcal{D}$ ;
8     Solve the Equation 3 for the  $\alpha_{u_k, i_k}$  and select  $J$  pairs  $(u_k, i_k)$  with maximum  $\alpha_{u_k, i_k}$ ;
9     Update  $\phi_w$  by descending along the gradient  $\nabla_{\phi_w} \ell(\phi_w)$ ;
10  end
11  for number of steps for training the prediction model do
12    Sample a batch of user-item pairs  $\{(u_l, i_l)\}_{l=1}^L$  from  $\mathcal{O}$ ;
13    Update  $\theta$  by descending along the gradient  $\nabla_{\theta} \mathcal{L}_{\text{KBDR}}(\phi_e, \phi_w, \theta)$ ;
14  end
15 end

```

which is equivalent to

$$\min_{\phi_w} \ell(\phi_w) = \sum_{(u,i) \in \mathcal{O}} g(x_{u,i}; \phi_w) \log g(x_{u,i}; \phi_w) + \gamma \sum_{j=1}^J \left([-C - \hat{\tau}^{(j)}]_+ + [\hat{\tau}^{(j)} - C]_+ \right), \quad (4)$$

where

$$\hat{\tau}^{(j)} = \sum_{(u,i) \in \mathcal{D}} o_{u,i} g(x_{u,i}; \phi_w) h^{(j)}(x_{u,i}) - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} h^{(j)}(x_{u,i}), \quad j \in \{1, \dots, J\}.$$

Since achieving strict balancing on all balancing functions, i.e., the weighted average balancing functions on the observed samples are exact same with the average of that over all samples, is usually infeasible as J increases, we introduce a slack variable ξ_j for each balancing function and a pre-specified threshold C , which penalizes the loss when the deviation $|\hat{\tau}^j| > C$.

4.5 THE LEARNING ALGORITHM AND THE GENERALIZATION BOUNDS

Taking adaptive **KBDR** (**AKBDR**) as an example, because the balancing weights and $e_{u,i}(x)$ are relying on each other in the Equation 3, thus we adopt a widely used joint learning framework to train the prediction model $\hat{r}_{u,i} = f(x_{u,i}; \theta)$, the balancing weight model $\hat{w}_{u,i} = g(x_{u,i}; \phi_w)$ and imputation model $\hat{e}_{u,i} = m(x_{u,i}; \phi_e)$ alternatively. Specifically, we train the prediction model by minimizing the $\mathcal{L}_{\text{KBDR}}$ loss shown in Equation 2, train the balancing weight model by minimizing the $\ell(\phi_w)$ in Equation 4 and train the imputation model by minimizing the loss function \mathcal{L}_e below:

$$\mathcal{L}_e(\phi_e, \phi_w, \theta) = |\mathcal{D}|^{-1} \sum_{(u,i) \in \mathcal{D}} o_{u,i} w_{u,i} (\hat{e}_{u,i} - e_{u,i})^2, \quad (5)$$

and the whole procedure of the proposed joint learning process is summarized in Alg. 1.

Next, we analyze the generalization bound of the **KBIPS** and **KBDR** methods, which is shown below:

Theorem 2 (Generalization Bounds in RKHS). *Let K be a bounded kernel, $\sup_x \sqrt{K(x, x)} = B < \infty$, and $B_K(M) = \{f \in \mathcal{F} \mid \|f\|_{\mathcal{F}} \leq M\}$ is the corresponding kernel-based hypotheses space. Suppose $\hat{w}_{u,i} \leq C$, $\delta(r, \cdot)$ is L -Lipschitz continuous for all r , and that $E_0 := \sup_r \delta(r, 0) < \infty$. Then with probability at least $1 - \eta$, we have*

$$\begin{aligned} \mathcal{L}_{\text{Ideal}}(\theta) &\leq \mathcal{L}_{\text{KBIPS}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta))| + \frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5C(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}}, \\ \mathcal{L}_{\text{Ideal}}(\theta) &\leq \mathcal{L}_{\text{KBDR}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta))| + (1 + 2C) \left(\frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} \right). \end{aligned}$$

Table 1: Performance on AUC, NDCG@K and F1@K on **Coat**, **Music** and **Product**. The best two result are bolded and the best baseline result is underlined for IPS-based and DR-based methods.

Method	COAT			MUSIC			PRODUCT		
	AUC	NDCG@5	F1@5	AUC	NDCG@5	F1@5	AUC	NDCG@20	F1@20
MF	0.703 \pm 0.006	0.605 \pm 0.012	0.467 \pm 0.007	0.673 \pm 0.001	0.635 \pm 0.002	0.306 \pm 0.002	0.753 \pm 0.001	0.449 \pm 0.002	0.124 \pm 0.002
+ CVIB	0.725 \pm 0.004	0.627 \pm 0.008	0.477 \pm 0.007	0.684 \pm 0.002	0.648 \pm 0.002	0.318 \pm 0.002	0.765 \pm 0.002	0.468 \pm 0.005	0.134 \pm 0.003
+ DIB	0.685 \pm 0.001	0.642 \pm 0.002	0.475 \pm 0.008	0.721 \pm 0.006	0.626 \pm 0.009	0.315 \pm 0.002	0.756 \pm 0.001	0.476 \pm 0.002	0.132 \pm 0.002
+ IPS	0.717 \pm 0.007	0.617 \pm 0.009	0.473 \pm 0.008	0.678 \pm 0.001	0.638 \pm 0.002	0.318 \pm 0.002	0.755 \pm 0.004	0.452 \pm 0.010	0.131 \pm 0.004
+ SNIPS	0.714 \pm 0.012	0.614 \pm 0.012	0.474 \pm 0.009	0.683 \pm 0.002	0.639 \pm 0.002	0.316 \pm 0.002	0.754 \pm 0.003	0.453 \pm 0.004	0.126 \pm 0.003
+ ASIPS	0.719 \pm 0.009	0.618 \pm 0.012	0.476 \pm 0.009	0.679 \pm 0.003	0.640 \pm 0.003	0.319 \pm 0.003	0.757 \pm 0.005	0.474 \pm 0.007	0.130 \pm 0.005
+ IPS-V2	0.726 \pm 0.005	0.627 \pm 0.009	0.479 \pm 0.008	0.685 \pm 0.002	0.646 \pm 0.003	0.320 \pm 0.002	0.764 \pm 0.001	0.476 \pm 0.003	0.135 \pm 0.003
+ KBIPS-Exp	0.714 \pm 0.003	0.618 \pm 0.010	0.474 \pm 0.007	0.676 \pm 0.002	0.642 \pm 0.003	0.318 \pm 0.002	0.763 \pm 0.001	0.463 \pm 0.007	0.134 \pm 0.002
+ KBIPS-Gau	0.715 \pm 0.005	0.619 \pm 0.010	0.475 \pm 0.008	0.678 \pm 0.001	0.640 \pm 0.004	0.315 \pm 0.003	0.760 \pm 0.003	0.470 \pm 0.008	0.133 \pm 0.003
+ WKBIPS-Exp	0.723 \pm 0.004	0.624 \pm 0.009	0.480 \pm 0.007	0.687 \pm 0.002	0.654 \pm 0.002	0.322 \pm 0.002	0.765 \pm 0.003	0.475 \pm 0.007	0.138 \pm 0.003
+ WKBIPS-Gau	0.722 \pm 0.004	0.625 \pm 0.008	0.479 \pm 0.007	0.686 \pm 0.002	0.650 \pm 0.002	0.321 \pm 0.002	0.763 \pm 0.003	0.476 \pm 0.007	0.137 \pm 0.003
+ AKBIPS-Exp	0.732 \pm 0.004	0.636 \pm 0.006	0.483 \pm 0.006	0.689 \pm 0.001	0.658 \pm 0.002	0.324 \pm 0.002	0.766 \pm 0.003	0.478 \pm 0.009	0.138 \pm 0.003
+ AKBIPS-Gau	0.730 \pm 0.003	0.633 \pm 0.008	0.484 \pm 0.007	0.688 \pm 0.003	0.655 \pm 0.003	0.324 \pm 0.002	0.767 \pm 0.003	0.480 \pm 0.009	0.139 \pm 0.003
+ DR	0.718 \pm 0.008	0.623 \pm 0.009	0.474 \pm 0.007	0.684 \pm 0.002	0.658 \pm 0.003	0.326 \pm 0.002	0.755 \pm 0.008	0.462 \pm 0.010	0.135 \pm 0.005
+ DR-JL	0.723 \pm 0.005	0.629 \pm 0.007	0.479 \pm 0.005	0.685 \pm 0.002	0.653 \pm 0.002	0.324 \pm 0.002	0.766 \pm 0.002	0.467 \pm 0.005	0.136 \pm 0.003
+ MRDR-JL	0.727 \pm 0.005	0.627 \pm 0.008	0.480 \pm 0.008	0.684 \pm 0.002	0.652 \pm 0.003	0.325 \pm 0.002	0.768 \pm 0.005	0.473 \pm 0.007	0.139 \pm 0.004
+ DR-BIAS	0.726 \pm 0.004	0.629 \pm 0.009	0.482 \pm 0.007	0.685 \pm 0.002	0.653 \pm 0.002	0.325 \pm 0.003	0.768 \pm 0.003	0.477 \pm 0.006	0.137 \pm 0.004
+ DR-MSE	0.727 \pm 0.007	0.631 \pm 0.008	0.484 \pm 0.007	0.687 \pm 0.002	0.657 \pm 0.003	0.327 \pm 0.003	0.770 \pm 0.003	0.480 \pm 0.006	0.140 \pm 0.003
+ MR	0.724 \pm 0.004	0.636 \pm 0.006	0.481 \pm 0.006	0.691 \pm 0.002	0.647 \pm 0.002	0.316 \pm 0.003	0.776 \pm 0.005	0.483 \pm 0.006	0.142 \pm 0.003
+ TDR	0.714 \pm 0.006	0.634 \pm 0.011	0.483 \pm 0.008	0.688 \pm 0.003	0.662 \pm 0.002	0.329 \pm 0.002	0.772 \pm 0.003	0.486 \pm 0.005	0.140 \pm 0.003
+ TDR-JL	0.731 \pm 0.005	0.639 \pm 0.007	0.484 \pm 0.007	0.689 \pm 0.002	0.656 \pm 0.004	0.327 \pm 0.003	0.772 \pm 0.003	0.489 \pm 0.005	0.142 \pm 0.003
+ SDR	0.735 \pm 0.005	0.640 \pm 0.007	0.484 \pm 0.006	0.688 \pm 0.002	0.661 \pm 0.003	0.329 \pm 0.002	0.773 \pm 0.001	0.491 \pm 0.003	0.143 \pm 0.003
+ DR-V2	0.734 \pm 0.007	0.639 \pm 0.009	0.487 \pm 0.006	0.690 \pm 0.002	0.660 \pm 0.005	0.328 \pm 0.002	0.773 \pm 0.003	0.488 \pm 0.006	0.142 \pm 0.004
+ KBDR-Exp	0.730 \pm 0.003	0.631 \pm 0.005	0.482 \pm 0.006	0.682 \pm 0.002	0.648 \pm 0.003	0.323 \pm 0.002	0.765 \pm 0.004	0.460 \pm 0.006	0.138 \pm 0.003
+ KBDR-Gau	0.726 \pm 0.005	0.630 \pm 0.008	0.480 \pm 0.008	0.683 \pm 0.002	0.652 \pm 0.003	0.325 \pm 0.002	0.766 \pm 0.003	0.469 \pm 0.007	0.134 \pm 0.004
+ WKBDR-Exp	0.735 \pm 0.005	0.637 \pm 0.009	0.483 \pm 0.006	0.685 \pm 0.003	0.654 \pm 0.003	0.325 \pm 0.002	0.773 \pm 0.003	0.489 \pm 0.008	0.142 \pm 0.003
+ WKBDR-Gau	0.732 \pm 0.003	0.638 \pm 0.007	0.483 \pm 0.005	0.687 \pm 0.001	0.655 \pm 0.002	0.327 \pm 0.002	0.773 \pm 0.002	0.490 \pm 0.005	0.142 \pm 0.004
+ AKBDR-Exp	0.745 \pm 0.004	0.645 \pm 0.008	0.493 \pm 0.007	0.692 \pm 0.002	0.661 \pm 0.002	0.328 \pm 0.002	0.782 \pm 0.003	0.498 \pm 0.008	0.147 \pm 0.003
+ AKBDR-Gau	0.746 \pm 0.004	0.646 \pm 0.008	0.492 \pm 0.007	0.694 \pm 0.002	0.664 \pm 0.002	0.332 \pm 0.002	0.782 \pm 0.005	0.503 \pm 0.006	0.148 \pm 0.004

Note: * means statistically significant results (p-value ≤ 0.05) using the paired-t-test compared with the best baseline method.

Remarkably, the above generalization bound in RKHS are able to be greatly reduced by adopting the proposed adaptive KBDR learning approach in Alg. 1. On one hand, the prediction model minimizes the loss $\mathcal{L}_{KBDR}(\theta)$ during the model training phase. On the other hand, as shown in Theorem 1 and Corollary 1, the proposed adaptive kernel balancing method can automatically choose the balancing functions that most need to be balanced to reduce the bias of the KBDR estimator¹.



5 EXPERIMENTS

Dataset and Preprocessing. Following the previous studies (Saito, 2020; Wang et al., 2019; 2021; Chen et al., 2021), we conduct real-world experiments on three widely used benchmark datasets: **Coat**, **Music**, and a large-scale industrial dataset: **Product**. **Coat** consists of 6,960 biased ratings and 4,640 unbiased ratings evaluated by 290 users to 300 items. **Music** consists of 311,704 biased ratings and 54,000 unbiased ratings evaluated by 15,400 users to 1,000 items. **Coat** and **Music** are both five-scale datasets, and we binarize the ratings less than three as 0, otherwise as 1. **Product** consists of 4,676,570 records of video watching ratios from 1,411 users to 3,327 items, and is almost fully exposed. We binarize the video watching ratios less than two as 0, otherwise as 1.

Baselines. In our experiments, we compare the proposed methods implemented with both Gaussian and exponential kernels to the following baseline methods: **MF** (Koren et al., 2009), **DIB** (Liu et al., 2021), **CVIB** (Wang et al., 2020), **IPS** (Schnabel et al., 2016), **SNIPS** (Swaminathan and Joachims, 2015), **ASIPS** (Saito, 2020) and **IPS-V2** (Li et al., 2023d). Meanwhile, we also consider the following DR-based baselines: **DR** (Saito, 2020), **DR-JL** (Wang et al., 2019), **MRDR** (Guo et al., 2021), **DR-BIAS** (Dai et al., 2022), **DR-MSE** (Dai et al., 2022), **MR** (Li et al., 2023a), **TDR** (Li et al., 2023b), **TDR-JL** (Li et al., 2023b), **StableDR** (Li et al., 2023e) and **DR-V2** (Li et al., 2023d).

Experimental Protocols and Details. We adopt widely used evaluation metrics: AUC, NDCG@K, and F1@K to assess the debiasing performance. We set $K = 5$ for **Coat** and **Music**, while $K = 20$ for **Product**. We tune learning rate in $\{0.01, 0.03, 0.05, 0.1\}$, weight decay in $\{1e - 5, 5e - 5, 1e - 4, 5e - 4, 1e - 3, 5e - 3\}$, margin threshold C in $\{0, 1e - 6, 1e - 5, 1e - 4\}$, kernel hyper-parameter σ^2 in $\{0.5, 1, 5\}$ for both Gaussian and exponential kernels, and regularization hyper-parameter γ in $\{0.01, 0.1, 1, 10, 100\}$. We set the batch size to 128 on **Coat** and 2,048 on **Music** and **Product**.

¹Similar arguments also hold for the proposed KBIPS estimator and the corresponding learning approach.

Figure 1: Effects of the value of J on AUC and NDCG@20 on PRODUCT dataset.Figure 2: Effects of hyper-parameter γ on AUC and NDCG@K on MUSIC and PRODUCT datasets.

Performance Comparison. We compare the proposed methods with previous methods shown in Table 1 and have the following findings. First, all the causality inspired methods outperform MF, where CVIB, SNIPS and SDR are the most competitive baselines, while **AKBDR** achieves optimal performance on the three datasets. Second, methods with balancing property such as DR-V2 and **AKBDR** outperform those without, which empirically demonstrates the effectiveness of balancing. Third, among kernel balancing methods, **kernel balancing (KB)** methods perform worst due to the insufficient weights learned by randomly balancing, while adaptive **kernel balancing (AKB)** methods perform best due to the proper priority of balanced functions provided by fitting prediction error.

In-Depth Analysis. We further explore the impact of the value of J on prediction performance on PRODUCT dataset. The proposed kernel balancing methods are compared with moment balancing (MB) methods which balance the former J th order moments and the results are shown in Figure 1. We find the AUC and NDCG@20 metrics for all methods increase monotonically with increasing value of J , because more functions or moments being balanced leads to bias reduction. In addition, kernel balancing methods stably outperform moment balancing methods with varying value of J even if the balanced functions are selected randomly, validating the effectiveness of kernel balancing.

Sensitivity Analysis. To explore the effect of balancing regularization hyper-parameter γ on debiasing performance, we perform sensitivity analysis of **AKB** methods using varying γ in $\{1, 2, 5, 10, 20\}$ on Music and Product datasets, as shown in Figure 2. It can be observed that **AKB** methods stably outperform the baseline methods without balancing property under different regularization strength. Specifically, even when the balancing constraint strength is relatively small, e.g., 1, the **AKB** method can still get clear performance gains, and the optimal performance is achieved around $\gamma = 10$.

6 CONCLUSION

In the information-driven landscape, recommender systems (RSs) are pivotal for various online platforms. However, selection bias in the collected data poses a great challenge for recommendation model training. To mitigate this issue, many methods were developed. However, we theoretically reveal that previous approaches are restricted to balance finite-dimensional pre-specified functions of features. To fill the gap, we first develop two new estimators, **KBIPS** and **KBDR**, which extend the popular IPS and DR estimators in debiased recommendations. Then we propose a universal kernel-based balancing method that adaptively achieve balance for continue functions in a RKHS. Based on it, we further propose an adaptive kernel balancing method. Theoretical analysis demonstrates that the proposed balancing method reduces both estimation bias and the generalization bound. Extensive experiments on real-world datasets validate the effectiveness of our methods.

REFERENCES

- Yin-Wen Chang, Cho-Jui Hsieh, Kai-Wei Chang, Michael Ringgaard, and Chih-Jen Lin. Training and testing low-degree polynomial data mappings via linear svm. *Journal of Machine Learning Research*, 11(4), 2010.
- Jiawei Chen, Hande Dong, Yang Qiu, Xiangnan He, Xin Xin, Liang Chen, Guli Lin, and Keping Yang. Autodebias: Learning to debias for recommendation. In *SIGIR*, 2021.
- Juan Chen and Yingchun Zhou. Weighted euclidean balancing for a matrix exposure in estimating causal effect. *arxiv.org/abs/2303.06812*, 2023.
- Quanyu Dai, Haoxuan Li, Peng Wu, Zhenhua Dong, Xiao-Hua Zhou, Rui Zhang, Xiuqiang He, Rui Zhang, and Jie Sun. A generalized doubly robust learning framework for debiasing post-click conversion rate prediction. In *KDD*, 2022.
- Silviu Guiasu and Abe Shenitzer. The principle of maximum entropy. *The mathematical intelligencer*, 7:42–48, 1985.
- Siyuan Guo, Lixin Zou, Yiding Liu, Wenwen Ye, Suqi Cheng, Shuaiqiang Wang, Hechang Chen, Dawei Yin, and Yi Chang. Enhanced doubly robust learning for debiasing post-click conversion rate estimation. In *SIGIR*, 2021.
- Jens Hainmueller. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20:25–46, 2012.
- Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *WWW*, 2017.
- Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Boca Raton: Chapman and Hall/CRC, 2020.
- José Miguel Hernández-Lobato, Neil Houlsby, and Zoubin Ghahramani. Probabilistic matrix factorization with non-random missing data. In *ICML*, 2014.
- Kosuke Imai and Marc Ratkovic. Covariate balancing propensity score. *Journal of the Royal Statistical Society (Series B)*, 76(1):243–263, 2014.
- Guido W. Imbens and Donald B. Rubin. *Causal Inference For Statistics Social and Biomedical Science*. Cambridge University Press, 2015.
- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- Fan Li, Kari Lock Morgan, and Alan M Zaslavsky. Balancing covariates via propensity score weighting. *Journal of the American Statistical Association*, 2018.
- Haoxuan Li, Quanyu Dai, Yuru Li, Yan Lyu, Zhenhua Dong, Xiao-Hua Zhou, and Peng Wu. Multiple robust learning for recommendation. In *AAAI*, 2023a.
- Haoxuan Li, Yan Lyu, Chunyuan Zheng, and Peng Wu. TDR-CL: Targeted doubly robust collaborative learning for debiased recommendations. In *ICLR*, 2023b.
- Haoxuan Li, Yanghao Xiao, Chunyuan Zheng, and Peng Wu. Balancing unobserved confounding with a few unbiased ratings in debiased recommendations. In *WWW*, 2023c.
- Haoxuan Li, Yanghao Xiao, Chunyuan Zheng, Peng Wu, and Peng Cui. Propensity matters: Measuring and enhancing balancing for recommendation. 2023d.
- Haoxuan Li, Chunyuan Zheng, and Peng Wu. StableDR: Stabilized doubly robust learning for recommendation on data missing not at random. In *ICLR*, 2023e.
- Zinan Lin, Dugang Liu, Weike Pan, Qiang Yang, and Zhong Ming. Transfer learning for collaborative recommendation with biased and unbiased data. *Artificial Intelligence*, 2023.

- Dugang Liu, Pengxiang Cheng, Hong Zhu, Zhenhua Dong, Xiuqiang He, Weike Pan, and Zhong Ming. Mitigating confounding bias in recommendation via information bottleneck. In *RecSys*, 2021.
- Jinwei Luo, Dugang Liu, Weike Pan, and Zhong Ming. Unbiased recommendation model based on improved propensity score estimation. *Journal of Computer Applications*, 2021.
- Benjamin Marlin, Richard S Zemel, Sam Roweis, and Malcolm Slaney. Collaborative filtering and the missing at random assumption. *UAI*, 2007.
- Benjamin M Marlin and Richard S Zemel. Collaborative prediction and ranking with non-random missing data. In *RecSys*, 2009.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- Stephen L. Morgan and Christopher Winship. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge University Press, second edition, 2015.
- Harrie Oosterhuis. Reaching the end of unbiasedness: Uncovering implicit limitations of click-based learning to rank. In *SIGIR*, 2022.
- Paul R. Rosenbaum. *Design of Observational Studies*. Springer Nature Switzerland AG, second edition, 2020.
- Paul R. Rosenbaum and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983.
- Yuta Saito. Unbiased pairwise learning from implicit feedback. In *NeurIPS Workshop*, 2019.
- Yuta Saito. Doubly robust estimator for ranking metrics with post-click conversions. In *RecSys*, pages 92–100, 2020.
- Yuta Saito. Asymmetric tri-training for debiasing missing-not-at-random explicit feedback. In *SIGIR*, 2020.
- Yuta Saito. Doubly robust estimator for ranking metrics with post-click conversions. In *RecSys*, 2020.
- Yuta Saito, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. Unbiased recommender learning from missing-not-at-random implicit feedback. In *WSDM*, 2020.
- Pedro H. C. Sant’Anna, Xiaojun Song, and Qi Xu. Covariate distribution balance via propensity scores. *Journal of Applied Econometrics*, 37(6):1093–1120, 2022.
- Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *ICML*, 2016.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Bharath K Sriperumbudur, Kenji Fukumizu, and Gert RG Lanckriet. Universality, characteristic kernels and rkhs embedding of measures. *Journal of Machine Learning Research*, 12(7), 2011.
- Harald Steck. Training and testing of recommender systems on data missing not at random. In *KDD*, 2010.
- Elizabeth A. Stuart. Matching methods for causal inference: A review and a look forward. *Statistical Science*, 25(1):1–21, 2010.
- Adith Swaminathan and Thorsten Joachims. The self-normalized estimator for counterfactual learning. In *NeurIPS*, 2015.
- Hao Wang, Tai-Wei Chang, Tianqiao Liu, Jianmin Huang, Zhichao Chen, Chao Yu, Ruopeng Li, and Wei Chu. Escm2: Entire space counterfactual multi-task model for post-click conversion rate estimation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 363–372, 2022.

- Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. Doubly robust joint learning for recommendation on data missing not at random. In *ICML*, 2019.
- Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. Combating selection biases in recommender systems with a few unbiased ratings. In *WSDM*, 2021.
- Zifeng Wang, Xi Chen, Rui Wen, Shao-Lun Huang, Ercan E. Kuruoglu, and Yefeng Zheng. Information theoretic counterfactual learning from missing-not-at-random feedback. *NeurIPS*, 2020.
- Raymond K W Wong and Kwun Chuen Gary Chan. Kernel-based covariate functional balancing for observational studies. *Biometrika*, 105(1):199–213, 2018.
- Peng Wu, Qirui Hu, Xingwei Tong, and Min Wu. Learning causal effect using machine learning with application to china’s typhoon. *Acta Mathematicae Applicatae Sinica*, 36(3):702–713, 2020.
- Qingyuan Zhao and Daniel Percival. Entropy balancing is doubly robust. *Journal of Causal Inference*, 5(1):20160010, 2017.

A PROOFS

Theorem 1. *The biases of the KBIPS and KBDR are shown as follows:*

$$\begin{aligned}\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta)) &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} - 1) e_{u,i} \right\}^2, \\ \text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta)) &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} - 1) (e_{u,i} - \hat{e}_{u,i}) \right\}^2.\end{aligned}$$

Proof of Theorem 1. By definition, the bias of the KBIPS estimator is the deviation to the ideal loss

$$\begin{aligned}\text{Bias}^2(\mathcal{L}_{\text{KBIPS}}(\theta)) &= (\mathcal{L}_{\text{KBIPS}}(\theta) - \mathcal{L}_{\text{Ideal}}(\theta))^2 \\ &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} \hat{w}_{u,i} e_{u,i} - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} e_{u,i} \right\}^2 \\ &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} - 1) e_{u,i} \right\}^2.\end{aligned}$$

Similarly, the bias of the KBDR estimator to the ideal loss is

$$\begin{aligned}\text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta)) &= (\mathcal{L}_{\text{KBDR}}(\theta) - \mathcal{L}_{\text{Ideal}}(\theta))^2 \\ &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} [\hat{e}_{u,i} + o_{u,i} \hat{w}_{u,i} (e_{u,i} - \hat{e}_{u,i})] - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} e_{u,i} \right\}^2 \\ &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} - 1) (e_{u,i} - \hat{e}_{u,i}) \right\}^2.\end{aligned}$$

□

Corollary 1. *If $e_{u,i}(x) \in \mathcal{H}_J = \text{span}\{h^{(1)}(x), \dots, h^{(J)}(x)\}$, then proposed KBIPS is unbiased; otherwise, KBIPS is biased. Similarly, if $e_{u,i}(x) - \hat{e}_{u,i}(x) \in \mathcal{H}_J = \text{span}\{h^{(1)}(x), \dots, h^{(J)}(x)\}$, then proposed KBDR is unbiased; otherwise, KBDR is biased.*

Proof of Corollary 1. If $e_{u,i}(x) \in \mathcal{H}_J = \text{span}\{h^{(1)}(x), \dots, h^{(J)}(x)\}$, there exist $\{\alpha_j\}_{j=1}^J$ satisfying $e_{u,i}(x) = \sum_{j=1}^J \alpha_j h^{(j)}(x_{u,i})$. By Theorem 1, the bias of the KBIPS estimator is

$$\begin{aligned}\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta)) &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} - 1) e_{u,i} \right\}^2 \\ &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} - 1) \left(\sum_{j=1}^J \alpha_j h^{(j)}(x_{u,i}) \right) \right\}^2 \\ &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{j=1}^J \alpha_j \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} h^{(j)}(x_{u,i}) - h^{(j)}(x_{u,i})) \right\}^2.\end{aligned}$$

Because the $\hat{w}_{u,i}$ is learned by solving the following optimization problem

$$\begin{aligned} \min_{\phi_w} \quad & \sum_{(u,i) \in \mathcal{O}} w_{u,i} \log w_{u,i} \\ \text{s.t.} \quad & w_{u,i} \geq 0 \quad \text{for all } (u,i) \in \mathcal{O}, \\ & \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} w_{u,i} = 1, \\ & \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} w_{u,i} h^{(j)}(x_{u,i}) = \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} h^{(j)}(x_{u,i}), \quad j \in \{1, \dots, J\}, \end{aligned}$$

thus the learned $\hat{w}_{u,i}$ satisfying the constrains in the optimization problem, which means that

$$\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} h^{(j)}(x_{u,i}) - h^{(j)}(x_{u,i})) = 0, \quad j \in \{1, \dots, J\}.$$

Therefore, the bias of the KBIPS estimator is

$$\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta)) = \left\{ \frac{1}{|\mathcal{D}|} \sum_{j=1}^J \alpha_j \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} h^{(j)}(x_{u,i}) - h^{(j)}(x_{u,i})) \right\}^2.$$

If $e_{u,i}(x) \notin \mathcal{H}_J = \text{span}\{h^{(1)}(x_{u,i}), \dots, h^{(J)}(x_{u,i})\}$, then for all $\{\alpha_j\}_{j=1}^J$, $e_{u,i}(x) = \sum_{j=1}^J \alpha_j h^{(j)}(x) + \epsilon(x_{u,i})$, where $\epsilon(x_{u,i})$ is the non-zero residual term. Therefore, we have

$$\begin{aligned} \text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta)) &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{j=1}^J \alpha_j \sum_{(u,i) \in \mathcal{D}} (o_{u,i} \hat{w}_{u,i} h^{(j)}(x_{u,i}) - h^{(j)}(x_{u,i})) \right\}^2 \\ &= \left\{ \frac{1}{|\mathcal{D}|} \sum_{j=1}^J \alpha_j \sum_{(u,i) \in \mathcal{D}} \epsilon(x_{u,i}) \right\}^2 \neq 0. \end{aligned}$$

Similar argument also holds for the proposed KBDR estimator. \square

Lemma 2 (Representer theorem). *If $\Omega = h(\|f\|)$ for some increasing function $h : \mathbb{R}_+ \rightarrow \overline{\mathbb{R}}$, then some empirical risk minimizer must admit the form $f(\cdot) = \sum_{i=1} \alpha_i \kappa(\cdot, x_i)$ for some $\alpha \in \mathbb{R}^n$. If h is strictly increasing, all minimizers admit this form.*

Proof. The proof can be found in Theorem 6.11 of Mohri et al. (2018). \square

Definition 2 (Empirical Rademacher Complexity (Shalev-Shwartz and Ben-David, 2014)). *Let \mathcal{F} be a family of prediction models mapping from $x \in \mathcal{X}$ to $[a, b]$, and $S = \{x_{u,i} \mid (u,i) \in \mathcal{D}\}$ a fixed sample of size $|\mathcal{D}|$ with elements in \mathcal{X} . Then, the empirical Rademacher complexity of \mathcal{F} with respect to the sample S is defined as:*

$$\mathcal{R}(\mathcal{F}) = \mathbb{E}_{\sigma \sim \{-1, +1\}^{|\mathcal{D}|}} \sup_{f \in \mathcal{F}} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} f(x_{u,i}) \right],$$

where $\sigma = \{\sigma_{u,i} : (u,i) \in \mathcal{D}\}$, and $\sigma_{u,i}$ are independent uniform random variables taking values in $\{-1, +1\}$. The random variables $\sigma_{u,i}$ are called Rademacher variables.

Lemma 3 (Rademacher Comparison Lemma (Shalev-Shwartz and Ben-David, 2014)). *Let \mathcal{F} be a family of real-valued functions on $z \in \mathcal{Z}$ to $[a, b]$, and $S = \{x_{u,i} \mid (u,i) \in \mathcal{D}\}$ a fixed sample of size $|\mathcal{D}|$ with elements in \mathcal{X} . Then*

$$\mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \left[\sup_{f \in \mathcal{F}} \left(\mathbb{E}_{z \sim \mathbb{P}} [f(z)] - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} f(z_{u,i}) \right) \right] \leq 2 \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \mathbb{E}_{\sigma \sim \{-1, +1\}^{|\mathcal{D}|}} \sup_{f \in \mathcal{F}} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} f(z_{u,i}) \right],$$

where $\sigma = \{\sigma_{u,i} : (u,i) \in \mathcal{D}\}$, and $\sigma_{u,i}$ are independent uniform random variables taking values in $\{-1, +1\}$. The random variables $\sigma_{u,i}$ are called Rademacher variables.

Proof of Lemma 3. The proof can be found in Lemma 26.2 of Shalev-Shwartz and Ben-David (2014). \square

Lemma 4 (McDiarmid’s Inequality (Shalev-Shwartz and Ben-David, 2014)). *Let V be some set and let $f : V^m \rightarrow \mathbb{R}$ be a function of m variables such that for some $c > 0$, for all $i \in [m]$ and for all $x_1, \dots, x_m, x'_i \in V$ we have*

$$|f(x_1, \dots, x_m) - f(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_m)| \leq c$$

Let X_1, \dots, X_m be m independent random variables taking values in V . Then, with probability of at least $1 - \delta$ we have

$$|f(X_1, \dots, X_m) - \mathbb{E}[f(X_1, \dots, X_m)]| \leq c \sqrt{\log\left(\frac{2}{\delta}\right) m/2}$$

Proof of Lemma 4. The proof can be found in Lemma 26.4 of (Shalev-Shwartz and Ben-David, 2014). \square

Lemma 5 (Rademacher Calculus (Shalev-Shwartz and Ben-David, 2014)). *For any $A \subset \mathbb{R}^m$, scalar $c \in \mathbb{R}$, and vector $\mathbf{a}_0 \in \mathbb{R}^m$, we have*

$$R(\{c\mathbf{a} + \mathbf{a}_0 : \mathbf{a} \in A\}) \leq |c|R(A).$$

Proof of Lemma 5. The proof can be found in Lemma 26.6 of (Shalev-Shwartz and Ben-David, 2014). \square

Lemma 6 (Talagrand’s Lemma (Mohri et al., 2018)). *Let Φ_1, \dots, Φ_m be L -Lipschitz functions from \mathbb{R} to \mathbb{R} and $\sigma_1, \dots, \sigma_m$ be Rademacher random variables. Then, for any hypothesis set \mathcal{F} of real-valued functions, the following inequality holds:*

$$\frac{1}{m} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \sum_{i=1}^m \sigma_i (\Phi_i \circ f)(x_i) \right] \leq \frac{L}{m} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \sum_{i=1}^m \sigma_i f(x_i) \right] = L\mathcal{R}(\mathcal{F}).$$

In particular, if $\Phi_i = \Phi$ for all $i \in [m]$, then the following holds:

$$\mathcal{R}(\Phi \circ \mathcal{F}) \leq L\mathcal{R}(\mathcal{F}).$$

Proof of Lemma 6. The proof can be found in Section 5.4 of (Mohri et al., 2018). \square

Lemma 7. *Suppose K is a bounded kernel with $\sup_x \sqrt{K(x, x)} = B < \infty$ and let \mathcal{F} be its RKHS. Let $M > 0$ be fixed. Then for any $S = \{x_{u,i} : (u, i) \in \mathcal{D}\}$,*

$$\mathcal{R}(B_K(M)) \leq \frac{MB}{\sqrt{|\mathcal{D}|}},$$

where $B_K(M) = \{f \in \mathcal{F} \mid \|f\|_{\mathcal{F}} \leq M\}$.

Proof of Lemma 7. Fix $S = \{x_{u,i} : (u,i) \in \mathcal{D}\}$. Then

$$\begin{aligned}
\mathcal{R}(B_K(M)) &= \mathbb{E}_\sigma \left[\sup_{f \in B_K(M)} \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} f(x_{u,i}) \right] \\
&= \frac{1}{|\mathcal{D}|} \mathbb{E}_\sigma \left[\sup_{f \in B_K(M)} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} \langle f, K(\cdot, x_{u,i}) \rangle \right] \\
&= \frac{1}{|\mathcal{D}|} \mathbb{E}_\sigma \left[\sup_{f \in B_K(M)} \left\langle f, \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} K(\cdot, x_{u,i}) \right\rangle \right] \\
&= \frac{1}{|\mathcal{D}|} \mathbb{E} \left[\left\langle M \frac{\sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} K(\cdot, x_{u,i})}{\left\| \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} K(\cdot, x_{u,i}) \right\|}, \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} K(\cdot, x_{u,i}) \right\rangle \right] \\
&= \frac{M}{|\mathcal{D}|} \mathbb{E}_\sigma \left[\left\| \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} K(\cdot, x_{u,i}) \right\| \right] \\
&= \frac{M}{|\mathcal{D}|} \mathbb{E}_\sigma \left[\sqrt{\left\| \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} K(\cdot, x_{u,i}) \right\|^2} \right] \\
&\leq \frac{M}{|\mathcal{D}|} \sqrt{\mathbb{E}_\sigma \left\| \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} K(\cdot, x_{u,i}) \right\|^2} \\
&= \frac{M}{|\mathcal{D}|} \sqrt{\sum_{(u,i) \in \mathcal{D}} \|K(\cdot, x_{u,i})\|^2} \\
&= \frac{M}{|\mathcal{D}|} \sqrt{\sum_{(u,i) \in \mathcal{D}} K(x_{u,i}, x_{u,i})} \\
&\leq \frac{M}{|\mathcal{D}|} \sqrt{|\mathcal{D}| B^2} \\
&= \frac{MB}{\sqrt{|\mathcal{D}|}}.
\end{aligned}$$

□

Theorem 2 (Generalization Bounds in RKHS). *Let K be a bounded kernel, $\sup_x \sqrt{K(x, x)} = B < \infty$, and $B_K(M) = \{f \in \mathcal{F} \mid \|f\|_{\mathcal{F}} \leq M\}$ is the corresponding kernel-based hypotheses space. Suppose $\hat{w}_{u,i} \leq C$, $\delta(r, \cdot)$ is L -Lipschitz continuous for all r , and that $E_0 := \sup_r \delta(r, 0) < \infty$. Then with probability at least $1 - \eta$, we have*

$$\mathcal{L}_{\text{Ideal}}(\theta) \leq \mathcal{L}_{\text{KBIPS}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta))| + \frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5C(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}},$$

and

$$\mathcal{L}_{\text{Ideal}}(\theta) \leq \mathcal{L}_{\text{KBDR}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta))| + (1 + 2C) \left(\frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} \right).$$

Proof of Theorem 2. We first prove the generalization bound of the **kernel balancing** IPS estimator, note that the ideal loss can be decomposed as follows.

$$\begin{aligned}\mathcal{L}_{\text{Ideal}}(\theta) &= \mathcal{L}_{\text{KBIPS}}(\theta) + (\mathcal{L}_{\text{Ideal}}(\theta) - \mathbb{E}(\mathcal{L}_{\text{KBIPS}}(\theta))) + (\mathbb{E}[\mathcal{L}_{\text{KBIPS}}(\theta)] - \mathcal{L}_{\text{KBIPS}}(\theta)) \\ &= \mathcal{L}_{\text{KBIPS}}(\theta) + \text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta)) + (\mathbb{E}(\mathcal{L}_{\text{KBIPS}}(\theta)) - \mathcal{L}_{\text{KBIPS}}(\theta)) \\ &\leq \mathcal{L}_{\text{KBIPS}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta))| \\ &\quad + \sup_{f_\theta \in B_K(M)} \left(\mathbb{E} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} \hat{w}_{u,i} e_{u,i} \right] - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} \hat{w}_{u,i} e_{u,i} \right).\end{aligned}$$

For simplicity, we denote the last term in the above formula as

$$\mathcal{B}(\mathcal{F}) = \sup_{f_\theta \in B_K(M)} \left(\mathbb{E} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} \hat{w}_{u,i} e_{u,i} \right] - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} o_{u,i} \hat{w}_{u,i} e_{u,i} \right),$$

we then aim to bound $\mathcal{B}(\mathcal{F})$ in the following.

Note that

$$\mathcal{B}(\mathcal{F}) = \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] + \left\{ \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] \right\},$$

where the first term is $\mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})]$, and by Lemma 2 we have

$$\begin{aligned}\mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] &\leq 2 \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \mathbb{E}_{\sigma \sim \{-1, +1\}^{|\mathcal{D}|}} \sup_{f_\theta \in B_K(M)} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} o_{u,i} \hat{w}_{u,i} e_{u,i} \right] \\ &= 2 \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \{ \mathcal{R}(\mathcal{L}_{\text{KBIPS}}) \},\end{aligned}$$

where

$$\mathcal{R}(\mathcal{L}_{\text{KBIPS}}(\theta)) := \mathbb{E}_{\sigma \sim \{-1, +1\}^{|\mathcal{D}|}} \sup_{f_\theta \in B_K(M)} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} o_{u,i} \hat{w}_{u,i} e_{u,i} \right].$$

By applying McDiarmid's inequality in Lemma 4 and the assumptions that $\hat{w}_{u,i} \leq C$, and $e_{u,i} \leq E_0 + LMB$, let

$$c = \frac{2C(E_0 + LMB)}{|\mathcal{D}|},$$

then with probability at least $1 - \frac{\eta}{2}$,

$$\left| \mathcal{R}(\mathcal{L}_{\text{KBIPS}}(\theta)) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \{ \mathcal{R}(\mathcal{L}_{\text{KBIPS}}(\theta)) \} \right| \leq \frac{2C(E_0 + LMB)}{|\mathcal{D}|} \sqrt{\frac{\log(4/\eta)|\mathcal{D}|}{2}} = 2C(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}}.$$

By the assumption that $\hat{w}_{u,i} \leq C$ and $\delta(r, \cdot)$ is L -Lipschitz continuous for all r , we have

$$\mathcal{R}(\mathcal{L}_{\text{KBIPS}}(\theta)) \leq L\mathcal{R}(\mathcal{F}) \leq \frac{LMB}{\sqrt{|\mathcal{D}|}},$$

where the first inequality is from Lemma 5 and Lemma 6 with L as Lipschitz constant, the second inequality is from Lemma 5, and $\mathcal{R}(\mathcal{F})$ is the empirical Rademacher complexity

$$\mathcal{R}(\mathcal{F}) = \mathbb{E}_{\sigma \sim \{-1, +1\}^{|\mathcal{D}|}} \sup_{f_\theta \in B_K(M)} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} f(x_{u,i}) \right],$$

where $\sigma = \{\sigma_{u,i} : (u,i) \in \mathcal{D}\}$, and $\sigma_{u,i}$ are independent uniform random variables taking values in $\{-1, +1\}$. The random variables $\sigma_{u,i}$ are called Rademacher variables.

For the rest term $\mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}}[\mathcal{B}(\mathcal{F})]$, by applying McDiarmid's inequality in Lemma 4 and the assumptions that $\hat{w}_{u,i} \leq C$ and $e_{u,i} \leq E_0 + LMB$, let

$$c = \frac{E_0 + LMB}{|\mathcal{D}|},$$

then with probability at least $1 - \frac{\eta}{2}$,

$$\left| \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}}[\mathcal{B}(\mathcal{F})] \right| \leq \frac{E_0 + LMB}{|\mathcal{D}|} \sqrt{\frac{\log(4/\eta)|\mathcal{D}|}{2}} = (E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}}.$$

We now bound $\mathcal{B}(\mathcal{F})$ combining the above results. Formally, with probability at least $1 - \eta$, we have

$$\begin{aligned} \mathcal{B}(\mathcal{F}) &= \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}}[\mathcal{B}(\mathcal{F})] + \left\{ \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}}[\mathcal{B}(\mathcal{F})] \right\} \\ &\leq 2 \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \{ \mathcal{R}(\mathcal{L}_{\text{KBIPS}}(\theta)) \} + \left\{ \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}}[\mathcal{B}(\mathcal{F})] \right\} \\ &\leq 2\mathcal{R}(\mathcal{L}_{\text{KBIPS}}(\theta)) + 4C(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} + \left\{ \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}}[\mathcal{B}(\mathcal{F})] \right\} \\ &\leq 2\mathcal{R}(\mathcal{L}_{\text{KBIPS}}(\theta)) + 5C(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} \\ &\leq \frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5C(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}}. \end{aligned}$$

We now bound the ideal loss combining the above results. Formally, with probability at least $1 - \eta$, we have

$$\begin{aligned} \mathcal{L}_{\text{Ideal}}(\theta) &\leq \mathcal{L}_{\text{KBIPS}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta))| + \mathcal{B}(\mathcal{F}) \\ &\leq \mathcal{L}_{\text{KBIPS}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta))| + \frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5C(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}}. \end{aligned}$$

In Theorem 1, we have already prove that

$$|\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta))| = \frac{1}{|\mathcal{D}|} \left| \sum_{(u,i) \in \mathcal{D}} (o_{u,i} w_{u,i} - 1) e_{u,i} \right|,$$

therefore with probability at least $1 - \eta$, we have

$$\mathcal{L}_{\text{Ideal}}(\theta) \leq \mathcal{L}_{\text{KBIPS}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBIPS}}(\theta))| + \frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5C(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}}.$$

We then prove the generalization bound of the **kernel balancing** DR estimator, similarly, the ideal loss can be decomposed as follows.

$$\begin{aligned} \mathcal{L}_{\text{Ideal}}(\theta) &= \mathcal{L}_{\text{KBDR}}(\theta) + (\mathcal{L}_{\text{Ideal}}(\theta) - \mathbb{E}(\mathcal{L}_{\text{KBDR}}(\theta))) + (\mathbb{E}[\mathcal{L}_{\text{KBDR}}(\theta)] - \mathcal{L}_{\text{KBDR}}(\theta)) \\ &= \mathcal{L}_{\text{KBDR}}(\theta) + \text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta)) + (\mathbb{E}(\mathcal{L}_{\text{KBDR}}(\theta)) - \mathcal{L}_{\text{KBDR}}(\theta)) \\ &\leq \mathcal{L}_{\text{KBDR}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta))| \\ &\quad + \sup_{f_\theta \in B_K(M)} \left(\mathbb{E} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} [\hat{e}_{u,i} + o_{u,i} \hat{w}_{u,i} (e_{u,i} - \hat{e}_{u,i})] \right] - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} [\hat{e}_{u,i} + o_{u,i} \hat{w}_{u,i} (e_{u,i} - \hat{e}_{u,i})] \right). \end{aligned}$$

For simplicity, we denote the last term in the above formula as

$$\mathcal{B}(\mathcal{F}) = \sup_{f_\theta \in B_K(M)} \left(\mathbb{E} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} [\hat{e}_{u,i} + o_{u,i} \hat{w}_{u,i} (e_{u,i} - \hat{e}_{u,i})] \right] - \frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} [\hat{e}_{u,i} + o_{u,i} \hat{w}_{u,i} (e_{u,i} - \hat{e}_{u,i})] \right),$$

we then aim to bound $\mathcal{B}(\mathcal{F})$ in the following.

Note that

$$\mathcal{B}(\mathcal{F}) = \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] + \left\{ \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] \right\},$$

where the first term is $\mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})]$, and by Lemma 2 we have

$$\begin{aligned} \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] &\leq 2 \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \mathbb{E}_{\sigma \sim \{-1, +1\}^{|\mathcal{D}|}} \sup_{f_\theta \in B_K(M)} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} \left[\hat{e}_{u,i} + o_{u,i} \hat{w}_{u,i} (e_{u,i} - \hat{e}_{u,i}) \right] \right] \\ &= 2 \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \{ \mathcal{R}(\mathcal{L}_{\text{KBDR}}) \}, \end{aligned}$$

where

$$\mathcal{R}(\mathcal{L}_{\text{KBDR}}(\theta)) := \mathbb{E}_{\sigma \sim \{-1, +1\}^{|\mathcal{D}|}} \sup_{f_\theta \in B_K(M)} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} \left[\hat{e}_{u,i} + o_{u,i} \hat{w}_{u,i} (e_{u,i} - \hat{e}_{u,i}) \right] \right].$$

By applying McDiarmid's inequality in Lemma 4 and the assumptions that $\hat{w}_{u,i} \leq C$, $\hat{e}_{u,i} \leq E_0 + LMB$, and $e_{u,i} \leq E_0 + LMB$, let $c = \frac{2(E_0 + LMB)(1+2C)}{|\mathcal{D}|}$, then with probability at least $1 - \frac{\eta}{2}$,

$$\begin{aligned} \left| \mathcal{R}(\mathcal{L}_{\text{KBDR}}(\theta)) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \{ \mathcal{R}(\mathcal{L}_{\text{KBDR}}(\theta)) \} \right| &\leq \frac{2(E_0 + LMB)(1+2C)}{|\mathcal{D}|} \sqrt{\frac{\log(4/\eta)|\mathcal{D}|}{2}} \\ &= 2(E_0 + LMB)(1+2C) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}}. \end{aligned}$$

By the assumption that $\hat{w}_{u,i} \leq C$ and $\delta(r, \cdot)$ is L -Lipschitz continuous for all r , we have

$$\mathcal{R}(\mathcal{L}_{\text{KBDR}}(\theta)) \leq L(1+2C)\mathcal{R}(\mathcal{F}) \leq (1+2C) \frac{LMB}{\sqrt{|\mathcal{D}|}},$$

where the first inequality is from Lemma 5 and Lemma 6 with $L(1+2C)$ as Lipschitz constant, the second inequality is from Lemma 5, and $\mathcal{R}(\mathcal{F})$ is the empirical Rademacher complexity

$$\mathcal{R}(\mathcal{F}) = \mathbb{E}_{\sigma \sim \{-1, +1\}^{|\mathcal{D}|}} \sup_{f_\theta \in B_K(M)} \left[\frac{1}{|\mathcal{D}|} \sum_{(u,i) \in \mathcal{D}} \sigma_{u,i} f(x_{u,i}) \right],$$

where $\sigma = \{\sigma_{u,i} : (u,i) \in \mathcal{D}\}$, and $\sigma_{u,i}$ are independent uniform random variables taking values in $\{-1, +1\}$. The random variables $\sigma_{u,i}$ are called Rademacher variables.

For the rest term $\mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})]$, by applying McDiarmid's inequality in Lemma 4 and the assumptions that $\hat{w}_{u,i} \leq C$, $\hat{e}_{u,i} \leq E_0 + LMB$, and $e_{u,i} \leq E_0 + LMB$, let

$$c = \frac{(E_0 + LMB)(1+2C)}{|\mathcal{D}|},$$

then with probability at least $1 - \frac{\eta}{2}$,

$$\left| \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] \right| \leq \frac{(E_0 + LMB)(1+2C)}{|\mathcal{D}|} \sqrt{\frac{\log(4/\eta)|\mathcal{D}|}{2}} = (E_0 + LMB)(1+2C) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}}.$$

We now bound $\mathcal{B}(\mathcal{F})$ combining the above results. Formally, with probability at least $1 - \eta$, we have

$$\begin{aligned}
\mathcal{B}(\mathcal{F}) &= \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] + \left\{ \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] \right\} \\
&\leq 2 \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} \{ \mathcal{R}(\mathcal{L}_{\text{KBDR}}(\theta)) \} + \left\{ \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] \right\} \\
&\leq 2\mathcal{R}(\mathcal{L}_{\text{KBDR}}(\theta)) + 4(E_0 + LMB)(1 + 2C) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} + \left\{ \mathcal{B}(\mathcal{F}) - \mathbb{E}_{S \sim \mathbb{P}^{|\mathcal{D}|}} [\mathcal{B}(\mathcal{F})] \right\} \\
&\leq 2\mathcal{R}(\mathcal{L}_{\text{KBDR}}(\theta)) + 5(E_0 + LMB)(1 + 2C) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} \\
&\leq 2(1 + 2C) \frac{LMB}{\sqrt{|\mathcal{D}|}} + 5(E_0 + LMB)(1 + 2C) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} \\
&= (1 + 2C) \left(\frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} \right).
\end{aligned}$$

We now bound the ideal loss combining the above results. Formally, with probability at least $1 - \eta$, we have

$$\begin{aligned}
\mathcal{L}_{\text{Ideal}}(\theta) &\leq \mathcal{L}_{\text{KBDR}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta))| + \mathcal{B}(\mathcal{F}) \\
&\leq \mathcal{L}_{\text{KBDR}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta))| + (1 + 2C) \left(\frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} \right).
\end{aligned}$$

In Theorem 1, we have already prove that

$$|\text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta))| = \frac{1}{|\mathcal{D}|} \left| \sum_{(u,i) \in \mathcal{D}} (o_{u,i} w_{u,i} - 1)(e_{u,i} - \hat{e}_{u,i}) \right|,$$

therefore with probability at least $1 - \eta$, we have

$$\mathcal{L}_{\text{Ideal}}(\theta) \leq \mathcal{L}_{\text{KBDR}}(\theta) + |\text{Bias}(\mathcal{L}_{\text{KBDR}}(\theta))| + (1 + 2C) \left(\frac{2LMB}{\sqrt{|\mathcal{D}|}} + 5(E_0 + LMB) \sqrt{\frac{\log(4/\eta)}{2|\mathcal{D}|}} \right),$$

which yields the stated results. \square

B MORE EXPERIMENTAL RESULTS

Table 2 shows the prediction performance of varying methods by taking **Neural Collaborative Filtering (NCF)** (He et al., 2017) as the base model on the **Product** dataset. Among all the baseline methods, IPS-V2 and DR-V2 demonstrate the most competitive performance, which could balance manually selected functions to learn the balancing weights for unbiased learning. The proposed kernel balancing methods achieve overall performance improvements, which is attribute to the fact that the proposed methods are able to adaptively balance functions that is more worthy of being balanced. This further validates the effectiveness of the proposed methods and demonstrates that the proposed methods are also effective when adopting deep learning-based backbone as the base model.

Table 3 shows the prediction performance and the running time of the proposed KBDR and AKBDR methods on the textscMusic dataset. First, AKBDR outperforms KBDR with varying J , which shows the **effectiveness** of the proposed adaptively balancing method. For **efficiency**, AKBDR requires about **doubled** running time to converge than KBDR, since it needs to sort all the coefficients of kernel functions, which is time-consuming when the number of user-item pairs is large. When $J > 5$, the performance of the AKBDR does not improve as the growth of J , but the running time increases.

Table 2: Performance on AUC, NDCG@K and F1@K on PRODUCT. The best result is bolded and the best baseline is underlined for IPS-based methods and DR-based methods.

Method	AUC	NDCG@20	F1@20
NCF	0.823 \pm 0.001	0.575 \pm 0.002	0.166 \pm 0.002
+ CVIB	0.820 \pm 0.002	0.544 \pm 0.004	0.162 \pm 0.002
+ IPS	0.822 \pm 0.003	0.579 \pm 0.006	0.169 \pm 0.003
+ SNIPS	0.833 \pm 0.002	0.586 \pm 0.002	0.178 \pm 0.002
+ ASIPS	0.832 \pm 0.002	0.583 \pm 0.003	0.178 \pm 0.002
+ IPS-V2	<u>0.835\pm0.001</u>	<u>0.588\pm0.002</u>	<u>0.181\pm0.001</u>
+ KBIPS-Gaussian	0.832 \pm 0.001	0.584 \pm 0.003	0.178 \pm 0.003
+ WKBIPS-Gaussian	0.833 \pm 0.002	0.588 \pm 0.003	0.181 \pm 0.003
+ AKBIPS-Gaussian	0.836* \pm 0.001	0.592* \pm 0.002	0.183* \pm 0.002
+ DR	0.758 \pm 0.004	0.526 \pm 0.003	0.146 \pm 0.003
+ DR-JL	0.832 \pm 0.001	0.581 \pm 0.003	0.178 \pm 0.002
+ MRDR-JL	0.833 \pm 0.001	0.585 \pm 0.002	0.179 \pm 0.001
+ DR-BIAS	0.834 \pm 0.002	0.585 \pm 0.003	0.178 \pm 0.002
+ DR-MSE	0.834 \pm 0.002	0.587 \pm 0.002	0.180 \pm 0.001
+ MR	0.837 \pm 0.001	<u>0.588\pm0.002</u>	0.181 \pm 0.002
+ TDR-JL	0.834 \pm 0.002	0.582 \pm 0.002	0.179 \pm 0.003
+ SDR	0.835 \pm 0.002	0.587 \pm 0.003	0.179 \pm 0.002
+ DR-V2	<u>0.837\pm0.002</u>	0.586 \pm 0.004	<u>0.182\pm0.002</u>
+ KBDR-Gaussian	0.834 \pm 0.002	0.586 \pm 0.003	0.179 \pm 0.003
+ WKBDR-Gaussian	<u>0.837\pm0.002</u>	0.588 \pm 0.003	0.180 \pm 0.002
+ AKBDR-Gaussian	0.840* \pm 0.002	0.590* \pm 0.003	0.183 \pm 0.002

Note: * means statistically significant results (p-value ≤ 0.05) using the paired-t-test compared with the best baseline.

Table 3: Performance and efficiency analysis of the proposed CBDR and KBDR methods on MUSIC.

Number of balancing functions	Metrics	KBDR-Gaussian	AKBDR-Gaussian
$J = 3$	AUC	0.681 \pm 0.003	0.688 \pm 0.003
	NDCG@5	0.650 \pm 0.003	0.657 \pm 0.003
	F1@5	0.323 \pm 0.002	0.326 \pm 0.002
	Time (s)	389.28 \pm 17.51	664.99 \pm 12.37
$J = 5$	AUC	0.683 \pm 0.002	0.694 \pm 0.002
	NDCG@5	0.652 \pm 0.003	0.664 \pm 0.002
	F1@5	0.325 \pm 0.002	0.332 \pm 0.002
	Time (s)	394.51 \pm 16.61	678.13 \pm 15.43
$J = 10$	AUC	0.682 \pm 0.002	0.694 \pm 0.002
	NDCG@5	0.650 \pm 0.003	0.663 \pm 0.002
	F1@5	0.324 \pm 0.003	0.330 \pm 0.003
	Time (s)	399.91 \pm 8.58	719.63 \pm 23.79
$J = 20$	AUC	0.683 \pm 0.002	0.695 \pm 0.002
	NDCG@5	0.651 \pm 0.002	0.664 \pm 0.003
	F1@5	0.325 \pm 0.002	0.332 \pm 0.003
	Time (s)	389.11 \pm 22.39	727.29 \pm 26.76
$J = 50$	AUC	0.684 \pm 0.002	0.695 \pm 0.002
	NDCG@5	0.652 \pm 0.003	0.664 \pm 0.002
	F1@5	0.324 \pm 0.003	0.333 \pm 0.002
	Time (s)	407.89 \pm 11.67	722.43 \pm 25.55