



Moderately-Balanced Representation Learning for Treatment Effects with Orthogonality Information

Yiyan Huang¹, Cheuk Hang Leung¹, Shumin Ma², Qi Wu¹(✉),
Dongdong Wang³, and Zhixiang Huang³

¹ School of Data Science, City University of Hong Kong, Hong Kong, China
yiyhuang3-c@my.cityu.edu.hk, {chleung87,qiwu55}@cityu.edu.hk

² Guangdong Provincial Key Laboratory of Interdisciplinary Research
and Application for Data Science, BNU-HKBU United
International College, Zhuhai, China
shuminma@uic.edu.cn

³ JD Digits, Beijing, China
{wangdongdong9,huangzhixiang}@jd.com

Abstract. Estimating the average treatment effect (ATE) from observational data is challenging due to selection bias. Existing works mainly tackle this challenge in two ways. Some researchers propose constructing a score function that satisfies the orthogonal condition, which guarantees that the established ATE estimator is “orthogonal” to be more robust. The others explore representation learning models to achieve a balanced representation between the treated and the controlled groups. However, existing studies fail to 1) discriminate treated units from controlled ones in the representation space to avoid the over-balanced issue; 2) fully utilize the “orthogonality information”. In this paper, we propose a moderately-balanced representation learning (MBRL) framework based on recent covariates balanced representation learning methods and orthogonal machine learning theory. This framework protects the representation from being over-balanced via multi-task learning. Simultaneously, MBRL incorporates the noise orthogonality information in the training and validation stages to achieve a better ATE estimation. The comprehensive experiments on benchmark and simulated datasets show the superiority and robustness of our method on treatment effect estimations compared with existing state-of-the-art methods.

Keywords: Treatment effects · Causal inference · Representation learning

1 Introduction

Causal inference has drawn a lot of attention across various research areas including statistics [2, 25], economics and finance [3, 7, 15] commercial social network

Y. Huang and C. H. Leung—Co-first authors are in alphabetical order.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022
S. Khanna et al. (Eds.): PRICAI 2022, LNCS 13630, pp. 3–16, 2022.
https://doi.org/10.1007/978-3-031-20865-2_1

applications [5, 10] and health care [8, 12]. One of the main tasks of causal inference is to estimate the *average treatment effect* (ATE). For example, a biotech company must know to what extent a newly developed vaccine can reduce the probability of infection for the whole population. The classical method to acquire the ATE is to conduct randomized controlled trials (RCTs), where the treatment is randomly assigned to the population but not selectively. Then the effect of the vaccine (treatment) on the infection (outcome) is measured by the difference between the average infection rate of the vaccinated group (treated group) and that of the unvaccinated group (controlled group). RCTs are regarded as the golden standard for treatment effect estimation, but conducting RCTs is costly and time-consuming [9, 21]. Thus, estimating the treatment effects in the observational study instead of RCTs becomes more and more tempting.

When it comes to estimating the ATE from the observational data, we need to handle the selection bias. The selection bias exists due to the non-random treatment assignment. The treatment assignment may be further influenced by the covariates that also directly affect the outcome. In the vaccine example, limited vaccines tend to be distributed to vulnerable individuals who are susceptible to infection. Such a non-random treatment assignment mechanism naturally results in a covariate shift phenomenon. That is, the covariates of the treated population can substantially differ from that of the controlled population.

Two classical methods are developed for adjusting the shifted covariates: inverse propensity weighting (IPW) and regression adjustment (see more details in [26]). IPW weights the instances based on the propensity scores to mimic the principle of RCTs to estimate ATE. Nevertheless, the IPW estimators are sensitive to the misspecification of the propensity score. Regression adjustment methods directly estimate the outcome model instead of propensity scores, whereas they would inevitably lead to biased ATE estimations due to overfitting and regularization bias [3]. Researchers improve classical methods from the perspectives of statistics and methodology.

The *orthogonal score function* proposed in [3] is a statistical correction by incorporating both the outcome model and the propensity score estimations. Since such a score function satisfies the *orthogonal condition*, the ATE estimator derived from the score function is consistent as long as one of the two underlying relations is correctly specified. This is also known as the *doubly robust* property. Recently, balanced representation learning techniques have attracted researchers' attention. The intuitive idea is to construct a pair of "twins" in the representation space by minimizing the imbalance between the distributions of the treated and controlled groups [23]. However, such methods mainly focus on the balance but overlook the discrimination between treated and controlled units. If the distributions of the treated and controlled groups in the representation space are too similar to be distinguished, it would be difficult to infer the ATE accurately. Such a trade-off plays a crucial role in identifying the treatment effects [23]. The importance of the undiscriminating problem is also emphasized by [10].

In this paper, with the tool of orthogonal machine learning, we propose a moderately-balanced representation learning (MBRL) framework to estimate

the treatment effects. MBRL trains in a multi-task framework and stops on a perturbation error metric to obtain a moderately-balanced representation. The merits of MBRL include i) preserving predictive information for inferring individual outcomes; ii) designing a multi-task learning framework to achieve a moderately-balanced rather than over-balanced representation; iii) fully utilizing the orthogonality information during the training and validation stages to achieve superior treatment effect estimations.

2 Preliminaries

Potential Outcome Framework. Let \mathbf{Z} be s -dimensional covariates such that $\mathbf{Z} \in \mathcal{Z} \subset \mathbb{R}^s$, where \mathcal{Z} is the sample space of covariates. $D \in \{0, 1\}$ denotes the treatment variable. $Y(0), Y(1)$ represent the potential outcomes for the treatment $D = 0$ and $D = 1$ respectively such that $Y(0), Y(1) \in \mathcal{Y} \subset \mathbb{R}$ with \mathcal{Y} being the sample space of outcome. We denote $w = (\mathbf{z}, d, y)$ as the realizations of the random variables $W = (\mathbf{Z}, D, Y)$. If the observed treatment is d , then the factual outcome Y^F equals $Y(d)$. We suppose the observational dataset contains N individuals and the m^{th} individual is observed as (\mathbf{z}_m, d_m, y_m) . The target quantity ATE τ is defined as $\tau := \mathbb{E}[Y(1) - Y(0)]$.

Identifying the treatment effects under the potential outcome framework [22] requires some fundamental assumptions: Strong Ignorability, Overlap, Consistency and Stable Unit Treatment Value Assumption (SUTVA). These assumptions guarantee that treatment effects can be inferred if we specify the relation $\mathbb{E}[Y | D, \mathbf{Z}]$, which is equivalent to estimating $g_0(D, \mathbf{Z})$ in the following interactive model when the treatment variable takes a binary value [3]:

$$\begin{aligned} Y &= g_0(D, \mathbf{Z}) + \xi, & \mathbb{E}[\xi | D, \mathbf{Z}] &= 0, \\ D &= m_0(\mathbf{Z}) + \nu, & \mathbb{E}[\nu | \mathbf{Z}] &= 0. \end{aligned} \tag{1}$$

Here, g_0 and m_0 are the true *nuisance functions*. ξ and ν are the noise terms. $m_0(\mathbf{Z}) = \mathbb{E}[D | \mathbf{Z}]$ is the *propensity score*. Let i be an element of $\{0, 1\}$. The true causal parameter θ_0^i is defined as $\theta_0^i := \mathbb{E}[Y(i)] = \mathbb{E}[g_0(i, \mathbf{Z})]$ for $i \in \{0, 1\}$, and the true ATE τ is computed by $\tau = \theta_0^1 - \theta_0^0$. We denote the estimated (θ_0^i, g_0, m_0) as $(\hat{\theta}^i, \hat{g}, \hat{m})$, and then the estimated ATE is computed by $\hat{\tau} = \hat{\theta}^1 - \hat{\theta}^0$.

Orthogonal Estimators. We aim to estimate the true causal parameters θ_0^1 and θ_0^0 given N i.i.d. samples $\{W_m = (\mathbf{Z}_m, D_m, Y_m)\}_{m=1}^N$. The standard procedure to acquire the estimated causal parameters $\hat{\theta}^1$ and $\hat{\theta}^0$ is: 1) getting the estimated nuisance functions $\hat{\rho}$, e.g., $\hat{\rho} = (\hat{g}, \hat{m})$; 2) constructing a score function $\psi(W, \theta^i, \rho)$ such that we can derive the estimated causal parameter $\hat{\theta}^i$ by solving $\mathbb{E}[\psi(W, \theta^i, \hat{\rho})] = 0$, where θ^i is a causal parameter that lies in the causal parameter space. According to [3], the estimator $\hat{\theta}^i$ solved from $\mathbb{E}[\psi(W, \theta^i, \hat{\rho})] = 0$ is robust to the estimated nuisance functions $\hat{\rho}$ if the corresponding score function $\psi(W, \theta^i, \rho)$ satisfies the orthogonal condition that is stated in Definition 1.

Definition 1 (Orthogonal Condition). Let $W = (\mathbf{Z}, D, Y)$, $\rho_0 = (h_{0,1}, \dots, h_{0,\gamma})$ be the true nuisance functions and θ_0 be the true causal parameter with θ being a causal parameter that lies in the causal parameter space. A score function $\psi(W, \theta, \rho)$ is said to satisfy the orthogonal condition with respect to $\rho = (h_1, \dots, h_\gamma)$ if

$$\mathbb{E}[\partial_{h_i} \psi(W, \theta, \rho) |_{\rho=\rho_0, \theta=\theta_0} | \mathbf{Z}] = 0 \quad \forall 1 \leq i \leq \gamma.$$

Under the interactive model setup (1), the nuisance functions are (g, m) , and the true ones are (g_0, m_0) . In this case, the orthogonal condition guarantees that the estimator is consistent if either one of the two nuisance functions, but unnecessarily both, is accurately estimated. This is well known as the doubly robust property. In this paper, we introduce two orthogonal estimators $\hat{\theta}_1$ [3] and $\hat{\theta}_2$ [14] in Proposition 1, and we can estimate ATE by plugging the learned nuisance functions into the orthogonal estimators.

Proposition 1 (Orthogonal Estimators). Let the nuisance functions be $\rho = (g, m)$ and the causal parameter be θ^i for $i \in \{0, 1\}$, the score functions $\psi_1(W, \theta^i, \rho)$ and $\psi_2(W, \theta^i, \rho)$ that satisfy the orthogonal condition (Definition 1) are:

$$\psi_1(W, \theta^i, \rho) = \theta^i - g(i, \mathbf{Z}) - (Y - g(i, \mathbf{Z})) \frac{iD + (1-i)(1-D)}{im(\mathbf{Z}) + (1-i)(1-m(\mathbf{Z}))}; \quad (2)$$

$$\psi_2(W, \theta^i, \rho) = \theta^i - g(i, \mathbf{Z}) - (Y(i) - g(i, \mathbf{Z})) \frac{((D - m(\mathbf{Z})) - \mathbb{E}[\nu | \mathbf{Z}])^2}{\mathbb{E}[\nu^2 | \mathbf{Z}]}. \quad (3)$$

The corresponding orthogonal estimators are:

$$\hat{\theta}_1^i \text{ solves } \frac{1}{N} \sum_{m=1}^N \psi_1(W_m, \theta^i, \hat{\rho}) = 0; \quad \hat{\theta}_2^i \text{ solves } \frac{1}{N} \sum_{m=1}^N \psi_2(W_m, \theta^i, \hat{\rho}) = 0.$$

3 Method

In this section, we first introduce the orthogonality information in Sect. 3.1. Then we present the network structure, objective function and model selection criterion of the proposed MBRL method based on the orthogonality information in Sect. 3.2.

3.1 Orthogonality Information

Recall that the ATE estimators $\hat{\theta}_1^i$ and $\hat{\theta}_2^i$ are doubly robust since they are orthogonal estimators. Still, they could be non-orthogonal once the model setup (1) relaxes the restrictions on the noise terms ξ and ν since the score functions ψ_1 and ψ_2 might violate the orthogonal condition. Hence, we propose the *Noise Conditions*, which would enforce the learned nuisance functions adapted to orthogonal estimators.

Proposition 2 (Noise Conditions). Under the interactive model setup (1), the conditions on the noise terms ξ and ν , i.e., $\mathbb{E}[\xi | D, \mathbf{Z}] = 0$ and $\mathbb{E}[\nu | \mathbf{Z}] = 0$, are sufficient conditions for ψ_1 and ψ_2 being orthogonal score functions ($\hat{\theta}_1^i$ and $\hat{\theta}_2^i$ being orthogonal estimators).

Given the noise conditions, we can exploit an essential property, the *noise orthogonality* property.

Property 1 (Noise Orthogonality). Under the interactive model setup (1) and the noise conditions, we have $\mathbb{E}[(Y - g_0(D, \mathbf{Z}))(D - m_0(\mathbf{Z}))] = 0$.

The noise conditions are sufficient conditions for the estimators $\hat{\theta}_1^i$ and $\hat{\theta}_2^i$ being orthogonal, so noise conditions play an important role when we approximate the true nuisance functions (g_0, m_0) with estimated ones (\hat{g}, \hat{m}) . Besides, under the noise conditions, the noise orthogonality can be utilized for our model selection. The decompositions similar to Noise Orthogonality also appeared in [3, 11].

3.2 The Proposed Framework

We propose a moderately-balanced representation learning (MBRL) framework to obtain (\hat{g}, \hat{m}) to estimate ATE, and the MBRL architecture is illustrated in Fig. 1. The MBRL network maps the original covariates space to the representation space (i.e., $\Phi : \mathcal{Z} \rightarrow \mathcal{R}$) such that 1) the representation preserves predictive information for outcomes; 2) the map makes the distributional discrepancy between the treated group and the controlled group small enough; 3) the domain (treated or controlled) of each individual is well discriminated; 4) the orthogonality information is involved.

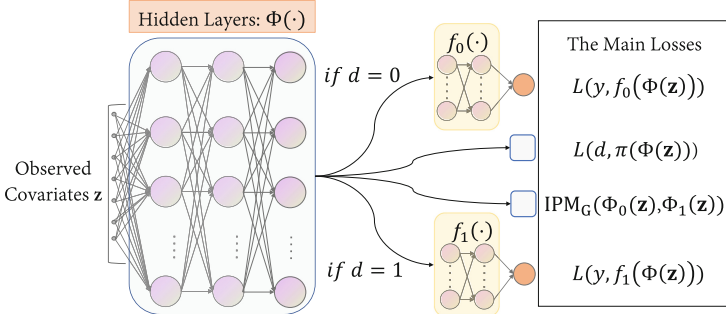


Fig. 1. The MBRL network architecture.

Learning Representation of Covariates. The distributions of the treated group and the controlled group are inherently disparate due to selection bias. Previous works handle this problem using a balanced representation learning method

[16, 23], which forces the distributions of treatment and control groups to be similar enough in the representation space. Specifically, a representation is learned by minimizing the integral probability metrics (IPM), which measures the imbalance between the distributions of the treated population and the controlled population (see the details in [23]):

$$\mathcal{L}_{imb} = \text{IPM}_G(\{\Phi(\mathbf{z}_m)\}_{m:d_m=1}, \{\Phi(\mathbf{z}_m)\}_{m:d_m=0}). \quad (4)$$

The Prediction of Outcome and Treatment. MBRL predicts the outcome by the function $f : \{0, 1\} \times \mathcal{R} \rightarrow \mathcal{Y}$, which is partitioned into two functions f_0 and f_1 :

$$f(d_m, \Phi(\mathbf{z}_m)) = d_m f_1(\Phi(\mathbf{z}_m)) + (1 - d_m) f_0(\Phi(\mathbf{z}_m)). \quad (5)$$

f_1 and f_0 are the output functions that map the representation to the potential outcomes for $D = 1$ and $D = 0$, respectively. $f(d_m, \Phi(\mathbf{z}_m))$ is the predicted factual outcome and we aim to minimize the factual outcome loss \mathcal{L}_{f_o} such that

$$\mathcal{L}_{f_o} = \frac{1}{N} \sum_{m=1}^N [y_m - f(d_m, \Phi(\mathbf{z}_m))]^2. \quad (6)$$

Here, $\hat{g}(d_m, \mathbf{z}_m) = f(d_m, \Phi(\mathbf{z}_m))$ is the estimated factual outcome of the m^{th} unit. Aside from making a low-error prediction over factual outcomes with a small divergence between treated and controlled groups, the distinguishability of the treated units from the controlled ones is also non-negligible. Therefore, we propose to maximize the distinguishability loss \mathcal{L}_{dis} (measured by log-likelihood) such that

$$\mathcal{L}_{dis} = \frac{1}{N} \sum_{m=1}^N [d_m \log \pi(\Phi(\mathbf{z}_m)) + (1 - d_m) \log(1 - \pi(\Phi(\mathbf{z}_m)))]]. \quad (7)$$

Here, $\hat{m}(d_m) = \pi(\Phi(\mathbf{z}_m))$ is the estimated probability of the m^{th} unit being assigned the treatment $D = 1$ (aka the estimated propensity score).

The Noise Regularizations. Recall Proposition 2 that $\mathbb{E}[\xi | D, \mathbf{Z}] = 0$ and $\mathbb{E}[\nu | \mathbf{Z}] = 0$ are sufficient conditions for score functions ψ_1 and ψ_2 being orthogonal. Empirically, we want to involve the following constraints:

$$\begin{aligned} \frac{1}{N} \sum_{m=1}^N [y_m - f(d_m, \Phi(\mathbf{z}_m))] &= 0, \\ \frac{1}{N} \sum_{m=1}^N [d_m - \pi(\Phi(\mathbf{z}_m))] &= 0. \end{aligned} \quad (8)$$

This motivates us to formalize Ω_y and Ω_d such that

$$\begin{aligned} \Omega_y &= \epsilon_y \left| \frac{1}{N} \sum_{m=1}^N [y_m - f(d_m, \Phi(\mathbf{z}_m))] \right|, \\ \Omega_d &= \epsilon_d \left| \frac{1}{N} \sum_{m=1}^N [d_m - \pi(\Phi(\mathbf{z}_m))] \right|. \end{aligned} \quad (9)$$

The partial derivative of Ω_y w.r.t. ϵ_y (or Ω_d w.r.t. ϵ_d) equaling 0 forces the learned nuisance functions to satisfy Eqn. (8). Therefore, minimizing the noise regularizations Ω_y and Ω_d adapts the entire learning process to satisfy the orthogonal score function. This idea corresponds to the targeted regularizations (see more discussions in [11,24]).

Multi-task Learning and Perturbation Error. MBRL learns the nuisance functions through multi-task learning with following three tasks in each iteration:

$$\begin{aligned} \text{Task 1:} & \quad \max_{\pi, \epsilon_d} \mathcal{L}_{dis} - \lambda_1 \Omega_d \\ \text{Task 2:} & \quad \min_{\Phi} \mathcal{L}_{imb} \\ \text{Task 3:} & \quad \min_{\Phi, f, \epsilon_y} \mathcal{L}_{fo} + \lambda_2 \Omega_y \end{aligned} \tag{10}$$

Instead of putting \mathcal{L}_{imb} into Task 3 as a regularization, we let \mathcal{L}_{imb} be one of the multiple tasks. To be specific, Task 1 updates π to produce the propensity scores, and Task 2 achieves a balance between $\{\Phi(\mathbf{z}_m)\}_{m:d_m=1}$ and $\{\Phi(\mathbf{z}_m)\}_{m:d_m=0}$. Additionally, MBRL incorporates a novel model selection criterion, the *Perturbation Error*, according to the noise orthogonality property. It takes advantage of the noise orthogonality information by perturbing the main evaluation metric. For example, if the final model is selected by the metric root-mean-square error ($RMSE = \sqrt{\frac{1}{N} \sum_{m=1}^N (y_m - \hat{y}_m)^2}$), then the perturbation error ϵ_p is defined as

$$\epsilon_p = RMSE + \beta \left| \frac{1}{N} \sum_{m=1}^N (y_m - \hat{y}_m)(d_m - \hat{d}_m) \right|.$$

Here, β is the perturbation coefficient which is a constant; \hat{y}_m and \hat{d}_m are the predicted values of $f(d_m, \Phi(\mathbf{z}_m))$ and $\pi(\Phi(\mathbf{z}_m))$, respectively. The final model is selected on the validation set based on the minimum ϵ_p . If either outcome or propensity score is well specified (i.e., representations are moderately-balanced instead of over-balanced), the second term in ϵ_p would be small.

4 Experiments

In this section, we conduct comprehensive experiments on benchmark datasets to evaluate the performance produced by MBRL and other prevalent causal inference methods. We further test the effectiveness of MBRL on simulated datasets with different levels of selection bias. All the experiments are run on Dell 7920 with 1×16 -core Intel Xeon Gold 6250 3.90 GHz CPU and 3x NVIDIA Quadro RTX 6000 GPU.

4.1 Dataset Description

Since the ground truth of treatment effects are inaccessible for real-world data, it is difficult to evaluate the performance of causal inference methods for ATE estimation. Previous causal inference literatures assess their methods on two prevalent semi-synthetic datasets: IHDP and Twins.

IHDP. The IHDP dataset is a well-known benchmark dataset for causal inference introduced by [13]. It includes 747 samples with 25-dimensional covariates associated with the information of infants and their mothers, such as birth weight and mother’s age. These covariates are collected from a real-world randomized experiment. Our aim is to study the treatment effect of the specialist visits (binary treatment) on the cognitive scores (continuous-valued outcome). The outcome is generated using the NPCI package [6], and the selection bias is created by removing a subset of the treated population. We use the same 1000 IHDP datasets as the ones used in [23], where each dataset is split by the ratio of 63%/27%/10% as training/validation/test sets.

Twins. The Twins dataset [19] collects twin births in the USA between 1989 and 1991 [1]. After the data processing, each unit has 30 covariates relevant to parents, pregnancy and birth [28]. The treatment $D = 1$ indicates the heavier twin while $D = 0$ indicates the lighter twin, and the outcome Y is a binary variable defined as the 1-year mortality. Similar to [28], we only select twins who have the same gender and both weigh less than 2 kg, which finally gives 11440 pairs of twins whose mortality rate is 17.7% for the lighter twin, and 16.1% for the heavier twin. To create the selection bias, we selectively choose one of the two twins as the factual observation based on the covariates of m^{th} individual: $D_m | \mathbf{Z}_m \sim \text{Bernoulli}(\text{Sigmoid}(\mathbf{w}^T \mathbf{Z}_m + n))$, where $\mathbf{w} \sim \mathcal{U}((-0.01, 0.01)^{30 \times 1})$ and $n \sim \mathcal{N}(0, 0.01)$. We repeat the data generating process for 100 times, and the generated 100 Twins datasets are all split by the ratio of 56%/24%/20% as training/validation/test sets.

4.2 Performance Measurement and Experimental Settings

Performance Measurement. Generally, the comparisons are based on the absolute error in ATE: $\epsilon_{ATE} = |\tau - \hat{\tau}|$. Additionally, we also test the performance of MBRL on individual treatment effect (ITE) estimations. For IHDP datasets, we adopt Precision in Estimation of Heterogeneous Effect (PEHE):

$$\epsilon_{PEHE} = \frac{1}{N} \sum_{m=1}^N ([y_m(1) - y_m(0)] - [\hat{y}_m(1) - \hat{y}_m(0)])^2.$$

For Twins datasets, we follow [19] to adopt Area Under ROC Curve (AUC).

Baseline Models. We compare our MBRL method with the following baseline models: linear regression with the treatment as feature (**OLS/LR**₁), separate linear regression for each treatment group (**OLS/LR**₂), k-nearest neighbor (**k-NN**), bayesian additive regression trees (**BART**) [4], causal forest (**CF**) [25], balancing linear regression (**BLR**) [16], balancing neural network (**BNN**) [16], treatment-agnostic representation network (**TARNet**) [23], counterfactual regression with Wasserstein distance (**CFR-WASS**) [23], causal effect variational autoencoders (**CEVAE**) [19], local similarity preserved individual treatment effect (**SITE**) [27], generative adversarial networks for inference of treatment effect (**GANITE**) [28] and (**Dragonnet**) [24].

Experimental Details. In our experiments, $\text{IPM}_{\mathcal{G}}$ is chosen as the Wasserstein distance. Let the empirical distribution of representation be $P(\Phi(\mathbf{Z})) = P(\Phi(\mathbf{Z}) \mid D = 1)$ for the treated group and $Q(\Phi(\mathbf{Z})) = Q(\Phi(\mathbf{Z}) \mid D = 0)$ for the controlled group. Assuming that \mathcal{G} is defined as the functional space of a family of 1-Lipschitz functions, we obtain the 1-Wasserstein distance for $\text{IPM}_{\mathcal{G}}$ [23]:

$$\text{Wass}(P, Q) = \inf_{k \in \mathcal{K}} \int_{\mathbf{h} \in \{\Phi(\mathbf{Z}_m)\}_{m: D_m=1}} \|k(\mathbf{h}) - \mathbf{h}\| P(\mathbf{h}) d\mathbf{h}.$$

Here, $\mathcal{K} = \{k \mid k : Q(k(\Phi(\mathbf{Z}))) = P(\Phi(\mathbf{Z}))\}$ defines the set of push-forward functions that transform the representation distribution of the treated group $P(\Phi(\mathbf{Z}))$ to that of the controlled group $Q(\Phi(\mathbf{Z}))$.

Table 1. Performance comparisons and ablation study with mean \pm standard error on 1000 **IHDP** datasets. ϵ_{ATE} : Lower is better. $\sqrt{\epsilon_{PEHE}}$: Lower is better.

Method	In-sample		Out-of-sample	
	$\sqrt{\epsilon_{PEHE}}$	ϵ_{ATE}	$\sqrt{\epsilon_{PEHE}}$	ϵ_{ATE}
OLS/LR ₁	5.8 \pm .3	.73 \pm .04	5.8 \pm .3	.94 \pm .06
OLS/LR ₂	2.4 \pm .1	.14 \pm .01	2.5 \pm .1	.31 \pm .02
k-NN	2.1 \pm .1	.14 \pm .01	4.1 \pm .2	.79 \pm .05
BART	2.1 \pm .1	.23 \pm .01	2.3 \pm .1	.34 \pm .02
CF	3.8 \pm .2	.18 \pm .01	3.8 \pm .2	.40 \pm .03
CEVAE	2.7 \pm .1	.34 \pm .01	2.6 \pm .1	.46 \pm .02
SITE	.69 \pm .0	.22 \pm .01	.75 \pm .0	.24 \pm .01
GANITE	1.9 \pm .4	.43 \pm .05	2.4 \pm .4	.49 \pm .05
BLR	5.8 \pm .3	.72 \pm .04	5.8 \pm .3	.93 \pm .05
BNN	2.2 \pm .1	.37 \pm .03	2.1 \pm .1	.42 \pm .03
TARNet	.88 \pm .0	.26 \pm .01	.95 \pm .0	.28 \pm .01
CFR-WASS	.71 \pm .0	.25 \pm .01	.76 \pm .0	.27 \pm .01
Dragonnet	1.3 \pm .4	.14 \pm .01	1.3 \pm .5	.20 \pm .05
MBRL	.52 \pm .0	.12 \pm .01	.57 \pm .0	.13 \pm .01
MBRL + $\hat{\theta}_1^i$.52 \pm .0	.10 \pm .00	.57 \pm .0	.17 \pm .01
MBRL + $\hat{\theta}_2^i$.52 \pm .0	.11 \pm .00	.57 \pm .0	.20 \pm .01

In addition, we adopt ELU activation function and set 4 fully connected layers with 200 units for both the representation encoder network $\Phi(\cdot)$ and the discriminator $\pi(\cdot)$, and 3 fully connected layers with 100 units for the outcome prediction networks $f_0(\cdot)$ and $f_1(\cdot)$. The optimizer is chosen as Adam [17], and the learning rate for the optimizer is set to be $1e^{-3}$. We set (batch size, epoch) to be (100, 1000)/(1000, 250) for IHDP/Twins experiments, and the hyper parameters (λ_1, λ_2) to be (0.01, 0.01)/(0.1, 0.1) for IHDP/Twins experiments. The final

model early stops on the metric ϵ_p , and we choose β in ϵ_p as 0.1 and 100 for IHDP experiments and Twins experiments, respectively.

For the baseline models, we follow the same settings of hyperparameters as in their published paper and code. For our MBRL network, the optimal hyperparameters are chosen in the same way as [23]. The searching ranges are reported in Table 4.

4.3 Results Analysis

Table 1 and Table 2 report part of the performances of baseline methods and MBRL on IHDP and Twins datasets. We present the average values and standard errors of ϵ_{ATE} , ϵ_{PEHE} and AUC (mean \pm std). The lower ϵ_{ATE} and ϵ_{PEHE} or the higher AUC, the better. Bold indicates the best method for each dataset.

As stated in Table 1 and Table 2, we have the following observations. 1) MBRL achieves significant improvements in both ITE and ATE estimations across all datasets compared to the baseline models. 2) The advanced representation learning methods that focus on estimating ITE (such as SITE, TARNet and CFR-WASS) show their inapplicability to ATE estimations. By contrast, MBRL not only significantly outperforms these representation learning methods in ITE estimations but also remains among the best ATE results. 3) The state-of-the-art ATE estimation method, Dragonnet, achieves superior ATE estimations across all the baseline models but yields a substantial error in ITE estimations. Although Dragonnet shares a similar basic network architecture to MBRL, MBRL can obtain a substantially lower ϵ_{ATE} than Dragonnet owing to

Table 2. Performance comparisons with mean \pm standard error on 100 **Twins** datasets. ϵ_{ATE} : Lower is better. AUC: Higher is better.

Method	In-sample		Out-of-sample	
	AUC	ϵ_{ATE}	AUC	ϵ_{ATE}
OLS/LR ₁	.660 \pm .005	.004 \pm .003	.500 \pm .028	.007 \pm .006
OLS/LR ₂	.660 \pm .004	.004 \pm .003	.500 \pm .016	.007 \pm .006
k-NN	.609 \pm .010	.003 \pm .002	.492 \pm .012	.005 \pm .004
BART	.506 \pm .014	.121 \pm .024	.500 \pm .011	.127 \pm .024
CEVAE	.845 \pm .003	.022 \pm .002	.841 \pm .004	.032 \pm .003
SITE	.862 \pm .002	.016 \pm .001	.853 \pm .006	.020 \pm .002
BLR	.611 \pm .009	.006 \pm .004	.510 \pm .018	.033 \pm .009
BNN	.690 \pm .008	.006 \pm .003	.676 \pm .008	.020 \pm .007
TARNet	.849 \pm .002	.011 \pm .002	.840 \pm .006	.015 \pm .002
CFR-WASS	.850 \pm .002	.011 \pm .002	.842 \pm .005	.028 \pm .003
MBRL	.879 \pm .000	.003 \pm .000	.874 \pm .001	.007 \pm .001
MBRL + $\hat{\theta}_1^i$.879 \pm .000	.003 \pm .000	.874 \pm .001	.008 \pm .000
MBRL + $\hat{\theta}_2^i$.879 \pm .000	.003 \pm .000	.874 \pm .001	.006 \pm .001

the multi-task learning framework and the utilization for orthogonality information. These observations indicate that the proposed MBRL method is extremely effective for estimating treatment effects.

We further conduct an ablation study on IHDP datasets to test if orthogonality information is practical in real applications. The relevant results are reported in Table 3. We let MBRL* denote MBRL without perturbation error ϵ_p , and MBRL** denote MBRL without any orthogonality information (ϵ_p , Ω_d and Ω_y). We find that incorporating orthogonality information will enhance the power of estimating treatment effects, whether with or without orthogonal estimators. This enhancement is pronounced especially when orthogonal estimators are plugged in for in-sample data.

4.4 Simulation Study

In this part, we mainly investigate two questions. Q1. Does MBRL perform more stably to the level of selection bias than the state-of-the-art model Dragonnet? Q2. Can the noise orthogonality information, the perturbation error ϵ_p , improve ATE estimations regardless of different models/estimators/selection bias levels?

We generate 2500 treated samples whose covariates $\mathbf{Z}^1 \sim \mathcal{N}(\boldsymbol{\mu}^1, 0.5 \times \Sigma \Sigma^T)$, and 5000 controlled whose covariates $\mathbf{Z}^0 \sim \mathcal{N}(\boldsymbol{\mu}^0, 0.5 \times \Sigma \Sigma^T)$, where $\boldsymbol{\mu}^1$ and $\boldsymbol{\mu}^0$ are both 10-dimensional vector and $\Sigma \sim \mathcal{U}((-1, 1)^{10 \times 10})$. The level of selection bias, measured by KL divergence of $\boldsymbol{\mu}^1$ with respect to $\boldsymbol{\mu}^0$, would vary by fixing $\boldsymbol{\mu}^0$ and adjusting $\boldsymbol{\mu}^1$. The potential outcomes of m^{th} individual are generated as $Y(1) | \mathbf{Z}_m \sim (\mathbf{w}_1^T \mathbf{Z}_m + n_1)$, $Y(0) | \mathbf{Z}_m \sim (\mathbf{w}_0^T \mathbf{Z}_m + n_0)$, where $\mathbf{w}_1 \sim \mathcal{U}((-1, 1)^{10 \times 1})$, $\mathbf{w}_0 \sim \mathcal{U}((-1, 1)^{10 \times 1})$, $n_1 \sim \mathcal{N}(0, 0.1)$, $n_0 \sim \mathcal{N}(0, 0.1)$. By adjusting $\boldsymbol{\mu}^1$ and fixing $\boldsymbol{\mu}^0$, we obtain five datasets with different levels of KL divergence in $\{0, 62.85, 141.41, 565.63, 769.89\}$. We run experiments on each dataset 100 times and draw box plots with regard to ϵ_{ATE} on the test set in Fig. 2.

In Fig. 2(a), we first find that MBRL shows stronger robustness and achieves significantly better ATE estimations with regard to different selection bias levels

Table 3. Ablation study on IHDP datasets.

Method	In-sample		Out-of-sample	
	$\sqrt{\epsilon_{PEHE}}$	ϵ_{ATE}	$\sqrt{\epsilon_{PEHE}}$	ϵ_{ATE}
MBRL**	.523 ± .006	.129 ± .005	.568 ± .009	.141 ± .006
MBRL*	.522 ± .006	.128 ± .005	.567 ± .009	.139 ± .006
MBRL	.522 ± .007	.121 ± .005	.565 ± .008	.133 ± .005
MBRL** + θ_1^i	.523 ± .006	.101 ± .004	.568 ± .009	.171 ± .007
MBRL* + θ_1^i	.523 ± .006	.102 ± .004	.567 ± .009	.170 ± .007
MBRL + θ_1^i	.522 ± .007	.102 ± .004	.565 ± .008	.166 ± .007
MBRL** + θ_2^i	.523 ± .006	.122 ± .005	.568 ± .009	.210 ± .008
MBRL* + θ_2^i	.523 ± .006	.121 ± .005	.567 ± .009	.208 ± .008
MBRL + θ_2^i	.522 ± .007	.114 ± .005	.565 ± .008	.204 ± .008

Table 4. The searching ranges of hyperparameters.

Hyperparameters	IHDP	Twins
λ_1, λ_2	0.01, 0.1, 1	0.01, 0.1, 1
Depth of Φ	2, 3, 4	2, 3, 4
Dim of Φ	100, 200	100, 200
Depth of π	2, 3, 4	2, 3, 4
Dim of π	100, 200	100, 200
Depth of f_0, f_1	2, 3, 4	2, 3, 4
Dim of f_0, f_1	100, 200	100, 200
Batch size	100, 300	500, 1000
Epoch	500, 1000	250, 500

compared with Dragonnet. In addition, it is noticeable that choosing the perturbation error ϵ_p as the model selection metric would yield smaller ϵ_{ATE} for any model (Dragonnet or MBRL). Particularly, ϵ_p corrects more errors for MBRL than Dragonnet, which indicates that ϵ_p works better if a model utilizes the orthogonality information in the training stage. In Fig. 2(b), we have two main observations: i) the criterion ϵ_p improves ATE estimations for all estimators across different selection bias levels; ii) the improvement brought by ϵ_p becomes more substantial when selection bias increases.

5 Related Work

Representation Learning. Our work has a strong connection with the balanced representation learning methods proposed in [16, 23], where they mainly focus on minimizing the imbalance between the different treatment groups in the representation space but overlook maximizing the discrimination of each unit’s treatment domain. IGNITE framework is proposed in [10] to infer individual treatment effects from networked data, where they achieve a balanced representation that captures patterns of hidden confounders predictive of treatment assignments. This inspires us to study treatment effects by training a moderately-balanced representation via multi-task learning. Other works relevant to representation learning include [18, 19, 24, 27, 28] and references therein.

Orthogonal Score Function. [3] develop the theory of double/debiased machine learning (DML) from [20]. They define the notion of orthogonal condition, which allows their DML estimator to be doubly robust. Based on the theory of [3], another orthogonal estimator is proposed by [14], aiming to overcome the high

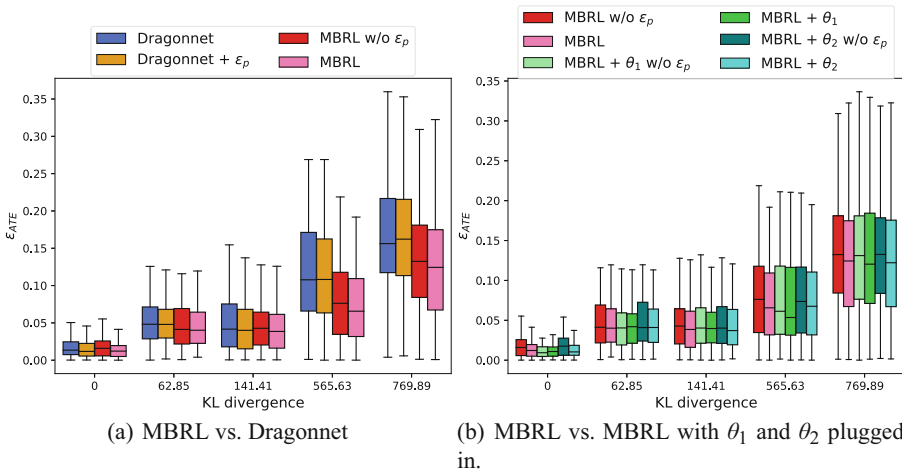


Fig. 2. Comparisons between models with and without ϵ_p w.r.t. varying levels of selection bias.

variance issue suffered by DML due to the misspecified propensity score. Despite the success of orthogonal estimators, the establishment of them requires the noise conditions to guarantee the corresponding score functions satisfying the orthogonal condition. None of the existing literature emphasizes the critical role of noise conditions or utilizes the orthogonality information for the model selection.

6 Conclusion

This paper proposes an effective representation learning method, MBRL, to study the treatment effects. Specifically, MBRL avoids the over-balanced issue by leveraging treatment domains of the representations via multi-task learning. MBRL further takes advantage of the orthogonality information and involves it in the training and validation stages. The extensive experiments show that 1) MBRL has strong predictability for the potential outcomes, distinguishability for the treatment assignment, applicability to orthogonal estimators, and robustness to the selection bias; 2) MRBL achieves substantial improvements on treatment effect estimations compared with existing state-of-the-art methods.

Acknowledgement. Qi Wu acknowledges the support from the Hong Kong Research Grants Council [General Research Fund 14206117, 11219420, and 11200219], CityU SRG-Fd fund 7005300, and the support from the CityU-JD Digits Laboratory in Financial Technology and Engineering, HK Institute of Data Science. The work described in this paper was partially supported by the InnoHK initiative, The Government of the HKSAR, and the Laboratory for AI-Powered Financial Technologies.

References

1. Almond, D., Chay, K.Y., Lee, D.S.: The costs of low birth weight. *Q. J. Econ.* **120**(3), 1031–1083 (2005)
2. Athey, S., Wager, S.: Estimating treatment effects with causal forests: an application. *Observat. Stud.* **5**(2), 37–51 (2019)
3. Chernozhukov, V., et al.: Double/debiased machine learning for treatment and structural parameters. *Econ. J.* **21**(1), C1–C68 (2018)
4. Chipman, H.A., George, E.I., McCulloch, R.E.: Bart: Bayesian additive regression trees. *Ann. Appl. Stat.* **4**(1), 266–298 (2010)
5. Chu, Z., Rathbun, S.L., Li, S.: Graph infomax adversarial learning for treatment effect estimation with networked observational data. In: *KDD*, pp. 176–184 (2021). <https://doi.org/10.1145/3447548.3467302>
6. Dorie, V.: Nonparametric methods for causal inference. <https://github.com/vdorie/npci> (2021)
7. Farrell, M.H.: Robust inference on average treatment effects with possibly more covariates than observations. *J. Econ.* **189**(1), 1–23 (2015)
8. Glass, T.A., Goodman, S.N., Hernán, M.A., Samet, J.M.: Causal inference in public health. *Ann. Rev. Public Health* **34**, 61–75 (2013)
9. Guo, R., Cheng, L., Li, J., Hahn, P.R., Liu, H.: A survey of learning causality with data: problems and methods. *ACM Comput. Surv. (CSUR)* **53**(4), 1–37 (2020)

10. Guo, R., Li, J., Li, Y., Candan, K.S., Raglin, A., Liu, H.: Ignite: a minimax game toward learning individual treatment effects from networked observational data. In: IJCAI, pp. 4534–4540 (2020)
11. Hatt, T., Feuerriegel, S.: Estimating Average Treatment Effects via Orthogonal Regularization, pp. 680–689. Association for Computing Machinery, New York, NY, USA (2021). <https://doi.org/10.1145/3459637.3482339>
12. Hill, J., Su, Y.S.: Assessing lack of common support in causal inference using Bayesian nonparametrics: implications for evaluating the effect of breastfeeding on children’s cognitive outcomes. *The Annals of Applied Statistics*, pp. 1386–1420 (2013)
13. Hill, J.L.: Bayesian nonparametric modeling for causal inference. *J. Comput. Graph. Stat.* **20**(1), 217–240 (2011)
14. Huang, Y., et al.: Robust causal learning for the estimation of average treatment effects. In: 2022 International Joint Conference on Neural Networks (IJCNN 2022). IEEE (2022)
15. Huang, Y., et al.: The causal learning of retail delinquency. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, pp. 204–212 (2021)
16. Johansson, F., Shalit, U., Sontag, D.: Learning representations for counterfactual inference. In: International Conference on Machine Learning, pp. 3020–3029. PMLR (2016)
17. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
18. Li, S., Fu, Y.: Matching on balanced nonlinear representations for treatment effects estimation. In: NIPS (2017)
19. Louizos, C., Shalit, U., Mooij, J., Sontag, D., Zemel, R., Welling, M.: Causal effect inference with deep latent-variable models. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, pp. 6449–6459 (2017)
20. Neyman, J.: $C(\alpha)$ tests and their use. *Sankhyā Indian J. Stat. Ser. A*, 1–21 (1979)
21. Pearl, J.: Causal inference in statistics: an overview. *Stat. Surv.* **3**, 96–146 (2009)
22. Rubin, D.B.: Causal inference using potential outcomes: design, modeling, decisions. *J. Am. Stat. Assoc.* **100**(469), 322–331 (2005)
23. Shalit, U., Johansson, F.D., Sontag, D.: Estimating individual treatment effect: generalization bounds and algorithms. In: International Conference on Machine Learning, pp. 3076–3085. PMLR (2017)
24. Shi, C., Blei, D., Veitch, V.: Adapting neural networks for the estimation of treatment effects. In: Advances in Neural Information Processing Systems, pp. 2503–2513 (2019)
25. Wager, S., Athey, S.: Estimation and inference of heterogeneous treatment effects using random forests. *J. Am. Stat. Assoc.* **113**(523), 1228–1242 (2018)
26. Yao, L., Chu, Z., Li, S., Li, Y., Gao, J., Zhang, A.: A survey on causal inference. *ACM Trans. Knowl. Discov. Data* **15**(5), 1–46 (2021). <https://doi.org/10.1145/3444944>
27. Yao, L., Li, S., Li, Y., Huai, M., Gao, J., Zhang, A.: Representation learning for treatment effect estimation from observational data. In: Advances in Neural Information Processing Systems 31 (2018)
28. Yoon, J., Jordon, J., Van Der Schaar, M.: GANITE: estimation of individualized treatment effects using generative adversarial nets. In: International Conference on Learning Representations (2018)