

Full Length Article

Learning double balancing representation for heterogeneous dose–response curve estimation

Minqin Zhu^a, Anpeng Wu^a, Haoxuan Li^b, Ruoxuan Xiong^c, Bo Li^d, Fei Wu^a, Kun Kuang^{a,*}^a College of Computer Science and Technology, Zhejiang University, China^b Center for Data Science, Peking University, China^c Department of Quantitative Theory and Methods, Emory University, USA^d School of Economics and Management, Tsinghua University, China

ARTICLE INFO

Keywords:

Representation learning

Causal inference

Dose–response curve estimation

Contrastive learning

ABSTRACT

Estimating the individuals' potential response to varying treatment doses is crucial for decision-making in areas such as precision medicine and management science. Most recent studies predict counterfactual outcomes by learning a covariate representation that is independent of the treatment variable. However, such independence constraints neglect much of the covariate information that is useful for counterfactual prediction, especially when the treatment variables are continuous. To tackle the above issue, in this paper, we first theoretically demonstrate the importance of the *balancing* and *prognostic* representations for unbiased estimation of the heterogeneous dose–response curves, that is, the learned representations are constrained to satisfy the conditional independence between the covariates and both of the treatment variables and the potential responses. Based on this, we propose an end-to-end Contrastive balancing Representation learning Network (CRNet) and a three-stage Weighted Double Balancing Network (WDBN) using a partial distance measure, for estimating the heterogeneous dose–response curves without losing the continuity of treatments. Extensive experiments are conducted on synthetic and real-world datasets demonstrating that our proposal significantly outperforms previous methods. Code is available at: <https://github.com/euzmin/Contrastive-Balancing-Representation-Network-CRNet>.

1. Introduction

Causal inference is crucial for individual decision-making, particularly in answering counterfactual questions such as “What would the individual's potential response have been had the person received a different dose of treatment?” (Blum & Blum, 2023; Pan, 2023; Raita, Camargo, Liang, & Hasegawa, 2021; Xu et al., 2023; Yuan & Zhu, 2023). For example, precision medicine is developed by studying the response of drug doses (i.e., continuous treatment) to the potential health state (i.e., potential outcome) of patients (with various medical history information, i.e., covariates) (Shi, Miao, & Tchetgen, 2020). With accessible observational data, an essential challenge for estimating causal effects is the confounding bias. This bias arises from confounders, which are common causes of treatment and outcome, and can lead to spurious correlations between the treatment and outcome, thereby hindering the *unbiasedness* of the estimation. Mealli, Pacini, and Rubin (2011), Pearl (2009). Another challenge is the *heterogeneity* of dose–response curves, implying that individuals with different covariates will

have different responses even with the same dose given (Li et al., 2023, 2023; Schwab, Linhardt, Bauer, Buhmann, & Karlen, 2020; Wager & Athey, 2018; Wu, Kuang, Xiong, Li, & Wu, 2023).

Compared to the binary treatment case, Dose–Response Curve estimation present greater challenges in adjusting for the confounding bias of high-dimensional covariates on continuous treatment (Hirano & Imbens, 2004; Imai & Van Dyk, 2004; Kennedy, Ma, McHugh, & Small, 2017). To tackle this problem, the Generalized Propensity Score (GPS) generalizes the propensity score (Rosenbaum & Rubin, 1983) from the binary treatment case, using a Gaussian distribution to model the treatment conditional density for given covariates (Hirano & Imbens, 2004; Imai & Van Dyk, 2004). Motivated by covariate balancing propensity score (Hainmueller, 2012; Imai & Ratkovic, 2014), the optimal balancing weighting methods focus on learning sample weights to ensure the treatment and covariates are independent in the re-weighted data. Fong, Hazlett, and Imai (2018), Huling, Greifer, and Chen (2024), Vegetabile et al. (2021). Despite these methods aiming to ensure

* Corresponding author.

E-mail addresses: minqinzhu@zju.edu.cn (M. Zhu), anpwu@zju.edu.cn (A. Wu), hxli@stu.pku.edu.cn (H. Li), ruoxuan.xiong@emory.edu (R. Xiong), libo@sem.tsinghua.edu.cn (B. Li), wufei@cs.zju.edu.cn (F. Wu), kunkuang@zju.edu.cn (K. Kuang).<https://doi.org/10.1016/j.neunet.2025.107600>

Received 23 September 2024; Received in revised form 25 March 2025; Accepted 5 May 2025

Available online 22 May 2025

0893-6080/© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

Table 1

A comparison of the constraints employed in various representation learning methods for dose–response curve estimation.

Representation learning method	Formulation
Treatment-Balanced Representation	$T \perp\!\!\!\perp \Phi(X)$
Balancing Representation	$T \perp\!\!\!\perp X \mid \Phi(X)$
Prognostic Representation	$Y(t) \perp\!\!\!\perp X \mid \Phi(X)$
Double Balancing Representation (ours)	$(T, Y) \perp\!\!\!\perp X \mid \Phi(X)$

unbiasedness by constraining treatments and confounders to be conditionally independent on balancing scores, they show limited performance in practice when the covariates are high-dimensional (Nie, Ye, Liu, & Nicolae, 2021; Schwab et al., 2020). Furthermore, these methods neglect the outcome during the modeling process of the balancing weight, potentially omitting confounders necessary for heterogeneous outcome prediction (Hansen, 2008; Stuart, Lee, & Leacy, 2013).

With the progress of deep learning, recent studies apply neural networks to fit dose–response curve of high-dimensional covariates (Bica, Jordon, & van der Schaar, 2020; Nie et al., 2021; Schwab et al., 2020). In deep methods, a critical challenge is how to learn appropriate covariate representations for Heterogeneous Dose–Response Curve (HDRC) estimation (Kallus, 2020). Specifically, DRNet (Schwab et al., 2020) propose to learn treatment-balanced representations (Shalit, Johansson, & Sontag, 2017; Wang et al., 2023; Wu, Kuang, Li & Wu, 2022) which force the learned representations to be independent of continuous treatments. Nonetheless, the method’s unbiasedness hinges on the assumption of invertibility concerning covariate representations, which is stringent for deep methods (Behrmann, Grathwohl, Chen, Duvenaud, & Jacobsen, 2019). In practice, imposing the constraint of independence between treatment assignment and covariate representations runs the risk of neglecting confounder information that is essential for heterogeneous outcome prediction, leading to biased estimates (Assaad et al., 2021). To tackle this problem, VCNet (Nie et al., 2021) employs a propensity score estimator to enforce a *balancing* representation that ensures unbiasedness. Given this representation, the treatment and covariates are conditionally independent. This approach enables the unbiased estimation of the average dose–response curve. However, while the propensity score is the coarsest balancing score (Rosenbaum & Rubin, 1983), it may not be adaptable for the HDRC estimation because of the covariate information loss (Hahn, 1998). In other words, it might not *prognostic* that potential outcomes and confounders are conditionally independent given balancing scores (Hansen, 2008). For the prognostic representation, SCIGAN (Bica et al., 2020) directly models the treatment effect by generative adversarial networks (Goodfellow et al., 2020). However, it is important to note that SCIGAN does not explicitly enforce the balancing representation condition and therefore does not directly ensure unbiasedness.

Overall, obtaining appropriate representations that eliminates confounder bias and retains necessary confounder information for the unbiased heterogeneous dose–response curve is still a challenging problem (Wu et al., 2022). As Table 1 shown, to address this issue, we systematically introduce the double balancing representation, which is a combination of balancing and prognostic representations. This representation is constrained to satisfy the conditional independence between the covariates and both the treatments and observed outcomes, while simultaneously considering the unbiasedness and heterogeneity of the dose–response curve. To model the double balancing representation for unbiased heterogeneous dose–response curves, we propose an end-to-end Contrastive balancing Representation learning Network, referred to as CRNet, and a three-stage Weighted Double Balancing Network, called WDBN.

In the end-to-end CRNet, we propose a novel contrastive regularizer, applying contrastive learning (Chen, Kornblith, Norouzi, & Hinton, 2020; Grill et al., 2020; He, Fan, Wu, Xie, & Girshick, 2020) to monitor the balancing representation condition and maintain treatment continuity. Specifically, we create negative samples by randomly shuffling

the original covariates and treat the original covariates as positive samples (Arbour, Dimmery, & Sondhi, 2021; Cheng et al., 2020). Adaptable to the cross-entropy loss (Chen et al., 2020), we adopt partial distance measure (Székely & Rizzo, 2014) to evaluate the balancing representation condition and design a contrastive regularizer loss to minimize the partial distance measure while discriminating among positive and negative samples. Moreover, we design a mean squared error loss specifically tailored to address the prognostic representation condition, ensuring the preservation of the representation’s predictive power for the outcome. CRNet can be trained directly in an end-to-end manner.

The three-stage WDBN generalizes CRNet by incorporating sample weights for weighted regression. This explicit balancing of the sample distribution aims to improve the model’s performance, particularly when dealing with limited samples. In the WDBN, we first generalize the contrastive regularizer to a double contrastive regularizer. This new regularizer simultaneously enforces both the balancing and the prognostic condition, thereby learning a double balancing representation. This sets the stage for the downstream learning of sample weights. In the second stage, based on the partial distance measure, we propose to directly learn the Double Balancing Sample Weight (DBSW) from the double balancing representation acquired in the first stage. In the third stage, we carry out a weighted regression based on the double balancing representation and the DBSW to finalize the model training. Empirically, we demonstrate that CRNet and WDBN outperforms previous methods on both synthetic and semi-synthetic datasets with different dimensions of continuous treatments.

We summarize our contribution as follows:

- For unbiased heterogeneous dose–response curve estimation, we systematically define a double balancing representation condition which satisfies the conditional independence constraint between the covariates and both of the continuous treatments and the outcomes.
- We propose a end-to-end CRNet and a three-stage WDBN for learning double balancing representations without losing the continuity of treatments. Specifically, we design contrastive losses with the partial distance measure of positive and negative samples and mean square error losses to optimize the CRNet and WDBN. To the best of our knowledge, this is the first paper to apply contrastive learning in the field of heterogeneous dose–response curve estimation.
- Empirically, varying the dimension of continuous treatments and covariates in both simulated and real-world datasets, we demonstrate that the proposed CRNet and WDBN outperform other baseline methods on heterogeneous dose–response curve estimation.

2. Related work

Dose–response curve estimation. For estimating the dose–response curve,¹ traditional methods (Fong et al., 2018; Imai & Van Dyk, 2004; Imbens, 2000; Vegetabile et al., 2021) learn sample weights on selected metrics to achieve the balance of covariates to eliminate the confounding bias. However, these methods neglect the outcome during the modeling of the balancing weight, which might omit confounders that are necessary for outcome prediction (Hansen, 2008; Lee & Lee, 2022; Stuart et al., 2013). Deep methods learn appropriate representations for dose–response curve estimation (Bica et al., 2020; Nie et al., 2021; Schwab et al., 2020). Treatment-balanced representation methods, for instance, DRNet (Schwab et al., 2020) constrains representations independent of continuous treatments. VCNet (Nie et al.,

¹ We only discuss weighting methods because matching and stratification can be considered as particular forms of weighting.

2021) and SCIGAN (Bica et al., 2020) constrain representations by treatment estimators/discriminators. None of them explicitly constrain that the learned representation satisfies both balancing and prognostic representation conditions for dose–response curve estimation. Instead, we propose contrastive regularizers to obtain double balancing representations for unbiased heterogeneous dose–response curve estimation directly.

Contrastive representation learning. Contrastive representation learning (Chen et al., 2020; Gan et al., 2023; Grill et al., 2020; He et al., 2020; Yao et al., 2022; Zhang, Huang, Li, & Wang, 2022; Zhang, Yao, Zhao, Chua, & Wu, 2021) is a self-supervised learning method. It approximates the latent representations by constructing contrastive samples (positive and negative instances) to facilitate instance discrimination (Wu, Xiong, Yu, & Lin, 2018). Through the process of discriminating between contrastive samples, positive instances are closer to the original instance in the projection space, while negative instances are further away from the original instance in the projection space to maximize the lower bound of the mutual information (Huang, Yi, & Zhao, 2021; Wang & Isola, 2020). In this paper, we apply contrastive learning to regularize the representation without breaking the continuity of treatments. To the best of our knowledge, this is the first paper to apply contrastive learning in dose–response curve estimation.

Conditional correlation measure. We conduct an investigation into conditional correlation measures. Specifically, conditional mutual information (Fleuret, 2004; Runge, 2018; Wyner, 1978) has a solid theoretical foundation, but it requires modeling of conditional densities explicitly, which can be challenging to optimize. On the other hand, conditional Hilbert–Schmidt Independence Criterion (Doran, Muandet, Zhang, & Schölkopf, 2014; Zhang, Peters, Janzing, & Schölkopf, 2012) does not require modeling of conditional densities explicitly, but its value does not directly represent the strength of correlation when it is greater than zero. To verify the degree of correlation, an independence test is needed. Partial correlation (Baba, Shibata, & Sibuya, 2004; Lawrance, 1976) is a coefficient that directly measures linear correlation, but it is limited to assessing linear relationships. CODEC (Azadkia & Chatterjee, 2021) can measure non-linear correlations; however, it requires solving the nearest neighbors problem at each iteration and cannot be directly optimized using gradients. Partial distance correlation, a widely used measure (Székely & Rizzo, 2014; Zhen, Meng, Chakraborty, & Singh, 2022), offers the advantage of directly comparing conditional correlations between variables. It stands out for its ability to be expressed as a dissimilarity, making it amenable to direct application in contrastive learning. Moreover, it can capture nonlinear correlations among variables. In our approach, we utilize the power of partial distance correlation by employing contrastive learning to constrain the double balancing representation.

3. Problem setup

For the case of continuous treatments, we observe n units with baseline covariates $X \in \mathcal{X} \subset \mathbb{R}^p$, continuous treatments $T \in \mathcal{T} \subset \mathbb{R}^q$ and outcome $Y \in \mathcal{Y} \subset \mathbb{R}$, where p, q is dimension of covariates and treatments, respectively. We also let $X \in \mathcal{X}^n \subset \mathbb{R}^{n \times p}$, $T \in \mathcal{T}^n \subset \mathbb{R}^{n \times q}$ and $Y \in \mathcal{Y}^n \subset \mathbb{R}^n$ denote all the observed baseline covariates, continuous treatments, and outcomes, respectively. Using Neyman–Rubin potential outcome framework (Rosenbaum & Rubin, 1983; Rubin, 1974), for an observation for unit i with received $T_i = t$, there is a potential outcome $Y_i(t)$.

Throughout this paper, we assume three assumptions that are commonly made in continuous treatment settings (Bica et al., 2020; Imbens, 2000; Nie et al., 2021; Schwab et al., 2020). Specifically, for a unit i , we assume the stable unit treatment value assumption (SUTVA) assumption holds that we can only observe the potential outcome corresponding to the received treatment level t , i.e., $Y_i = Y_i(t)$ and

there should not be alternative forms of the treatment and interference between units, capturing consistency and non-interference. Moreover, we assume the unconfoundedness assumption that $Y(t) \perp\!\!\!\perp T \mid X$ and the positivity assumption that $0 < \mathbb{P}(t|x)$ for $T = t$ and $X = x$. In this paper, $\perp\!\!\!\perp$ denotes (conditional) independence, and \mathbb{P} is the probability density function. We consider estimating the heterogeneous dose–response curve (HDRC):

$$h(t, x) = \mathbb{E}[Y(t) \mid X = x], \quad (1)$$

where \mathbb{E} denotes expectation.

4. Motivation

For estimating the heterogeneous dose–response curve, deep methods require an appropriate criterion to monitor the representation they produce (Kallus, 2020; Schwab et al., 2020). Inspired by the effective balancing score (Hu, Follmann, & Wang, 2014; Huang & Chan, 2017), a linear function of covariates for unbiased causal effect estimation, we turn to define two conditions of representation for unbiased HDRC estimation.

Definition 1 (Balancing Representation Condition). A balancing representation $\Phi(X)$, $X \in \mathcal{X}$, correlated to treatments $T \in \mathcal{T}$ and outcome $Y \in \mathcal{Y}$ satisfies:

$$X \perp\!\!\!\perp T \mid \Phi(X). \quad (2)$$

This condition guarantees that the treatment assignment is ignorable given the balancing representation when the unconfoundedness assumption is satisfied (for a detailed theory and proof, please refer to Theorem 3 of Rosenbaum & Rubin, 1983), which implies that $X \perp\!\!\!\perp T \mid \Phi(X)$ is equivalent to the unbiasedness condition, i.e., $Y(t) \perp\!\!\!\perp T \mid \Phi(X)$. As a result, we can identify the average dose–response curve as follows:

$$\begin{aligned} \mathbb{E}[Y(t)] &= \mathbb{E}_X[\mathbb{E}[Y(t) \mid \Phi(x)]] \\ &= \mathbb{E}_X[\mathbb{E}[Y(t) \mid \Phi(x), T = t]] \\ &= \mathbb{E}_X[\mathbb{E}[Y \mid \Phi(x), T = t]]. \end{aligned} \quad (3)$$

The first equation holds by the iterated expectation, the second equation holds by $Y(t) \perp\!\!\!\perp T \mid \Phi(X)$, and the third equation holds by the consistency assumption in SUTVA.

As we previously mentioned, when a representation satisfies the balancing representation condition, we say that it meets the criterion of unbiasedness. Nevertheless, regressing outcome Y on balancing representation $\Phi(X)$ and treatment T is inadequate for the heterogeneous dose–response curve estimation. This inadequacy can be attributed to the following reasons:

$$\mathbb{E}[Y(t) \mid x] \neq \mathbb{E}[Y(t) \mid \Phi(x)]. \quad (4)$$

For instance, VCNet (Nie et al., 2021), which employs a propensity score constraint on representation, is sufficient for unbiased average dose–response estimation. Nonetheless, it is important to acknowledge that the propensity score is the coarsest balancing score (Rosenbaum & Rubin, 1983). Using representations that are constrained by it might result in the loss of covariate information for outcome prediction (Hahn, 1998). This concern is amplified in situations where VCNet discretizes the continuous treatment variable into discrete intervals and utilizes the cross-entropy loss function for training the propensity score (Li, Xiao, Zheng, Wu & Cui, 2023).

Definition 2 (Prognostic Representation Condition). A prognostic representation $\Phi(X)$, $X \in \mathcal{X}$ correlated to treatments $T \in \mathcal{T}$ and outcome $Y \in \mathcal{Y}$ satisfies:

$$X \perp\!\!\!\perp Y(t) \mid \Phi(X). \quad (5)$$

Under unconfoundedness assumption, this prognostic representation condition is sufficient for unbiased heterogeneous dose-response curve estimation. Consider a unit i with treatment t , we can write the heterogeneous dose-response curve as follows:

$$\mathbb{E}[Y(t) | x] = \mathbb{E}[Y(t) | \Phi(x)] = \mathbb{E}[Y | \Phi(x), T = t]. \quad (6)$$

The first equation holds by Definition 2 and the second equation holds under Definition 2 and the unconfoundedness assumption (Hansen, 2008; Stuart et al., 2013). The above analysis implies that given the prognostic representation, the covariates are ignorable for the outcome prediction (Hansen, 2008).

However, learning the prognostic representation presents a challenge due to the unobservability of potential outcomes $Y(t)$ (Holland, 1986). In practical scenarios, we are constrained to derive representations based on the condition $X \perp\!\!\!\perp Y | \Phi(X)$, which we refer to as the *learnable* prognostic representation condition. The key difference between this and the prognostic representation condition is that the prognostic representation condition $X \perp\!\!\!\perp Y(t) | \Phi(X)$ itself contains all the predictive information of X for the counterfactual outcome $Y(t)$, whereas the learnable prognostic representation only includes predictive information for the observed outcome Y . In other words, the prognostic representation is unbiased but cannot be directly learned from observational data, while the learnable prognostic representation may be biased but can be learned from observational data. Consequently, when representations are constrained by this prognostic representation condition between covariates and observed outcomes, there exists a potential concern regarding ensuring unbiasedness (Hansen, 2008; Huang & Chan, 2017).

To address the challenges posed by both the balancing representation and prognostic representation conditions, where balancing representation condition alone may lead to the loss of essential information for outcome prediction and prognostic representation condition constrained by observed outcomes may introduce bias, we propose the condition of double balancing representation. This condition aims to enhance both unbiasedness through the balancing representation condition and the predictive capacity for heterogeneous outcomes through the learnable prognostic representation condition.

Definition 3 (Double Balancing Representation Condition). A double balancing representation $\Phi(X)$, $X \in \mathcal{X}$ correlated to treatments $T \in \mathcal{T}$ and outcome $Y \in \mathcal{Y}$ satisfies:

$$X \perp\!\!\!\perp T | \Phi(X), \quad X \perp\!\!\!\perp Y | \Phi(X). \quad (7)$$

This condition generalizes the balancing representation and prognostic representation conditions, making it possible to capture representations from observational data that satisfy both unbiasedness and heterogeneity. Based on this, the HDRC admits two equivalent identification formulas:

$$\mathbb{E}[Y(t)|x] = \mathbb{E}[Y|x, T = t] = \mathbb{E}[Y|\Phi(x), T = t]. \quad (8)$$

It becomes evident that the second equation is valid when the condition of double balancing representation is met. In the context of HDRC estimation, the two conditions comprising the double balancing representation are of paramount importance, as they mutually reinforce the effectiveness of this representation. Hence, it is imperative to verify the fulfillment of the double balancing representation condition.

5. End-to-end CRNet

The two conditions in double balancing representation both control confounder information of treatment assignment and retain necessary confounder information for outcome prediction. For unbiased heterogeneous dose-response curve estimation, we design the contrastive regularizer loss and mean square loss to constrain these two conditions.

5.1. Contrastive regularizer

In dose-response curve estimation, the treatments can be multiple and continuous and the covariates are high-dimensional. To ensure the unbiased treatment assignment, it is necessary to quantify the conditional dependence of T and X given $\Phi(X)$ (Rosenbaum & Rubin, 1983). Without loss of generalization, we adopt partial distance measure (Székely & Rizzo, 2014) to achieve this goal. The partial distance measure is a scalar quantity that captures dependence, which equals the conditional correlation in Gaussian scenarios. In the non-Gaussian case, a partial distance of zero does not confirm conditional independence, however, such a measure that is closer to zero indicates a weaker association (refer to Sec 4.2. in Székely & Rizzo, 2014).

Partial distance measure. For all observed data, assuming three variables $\mathbf{A}, \mathbf{B}, \mathbf{Z}$ and their double-centered pairwise distance $\omega(\mathbf{A})$, $\omega(\mathbf{B})$, $\omega(\mathbf{Z})$, we define $[\omega(\mathbf{A})]_{i,j}$ as follows:

$$[\omega(\mathbf{A})]_{i,j} = \|\mathbf{A}_i - \mathbf{A}_j\| - \frac{1}{n} \sum_{k=1}^n \|\mathbf{A}_k - \mathbf{A}_j\| - \frac{1}{n} \sum_{l=1}^n \|\mathbf{A}_i - \mathbf{A}_l\| + \frac{1}{n^2} \sum_{k=1}^n \sum_{l=1}^n \|\mathbf{A}_k - \mathbf{A}_l\| \quad (9)$$

where $\|\cdot\|$ is the Euclidean norm and $\omega(\cdot) \in \mathbb{R}^{n \times n}$. The form of $\omega(\mathbf{B})$, $\omega(\mathbf{Z})$ are similar. We define the inner product of $\omega(\mathbf{A})$ and $\omega(\mathbf{B})$ as follows:

$$\omega(\mathbf{A}) \otimes \omega(\mathbf{B}) = [n(n-3)]^{-1} \sum_{i \neq j} [\omega(\mathbf{A})]_{i,j} \cdot [\omega(\mathbf{B})]_{i,j}, \quad (10)$$

where the inner product $\omega(\mathbf{A}) \otimes \omega(\mathbf{B})$ forms a Hilbert space (refer to Theorem 1 and the corresponding proof on page 6 of Székely & Rizzo, 2014). The double-centered pairwise distance orthogonal projection of \mathbf{A} on \mathbf{Z} is formulated as follows:

$$\text{proj}_{\mathbf{Z}}(\mathbf{A}) = \omega(\mathbf{A}) - \omega(\mathbf{A}) \otimes \omega(\mathbf{Z})[\omega(\mathbf{Z}) \otimes \omega(\mathbf{Z})]^{-1} \omega(\mathbf{Z}), \quad (11)$$

and the projection of $\omega(\mathbf{B})$ on $\omega(\mathbf{Z})$ is similar. Then, we formulate the partial distance measure $D_{\mathbf{Z}}(\mathbf{A}, \mathbf{B})$ as follows:

$$D_{\mathbf{Z}}(\mathbf{A}, \mathbf{B}) = \frac{|\text{proj}_{\mathbf{Z}}(\mathbf{A}) \otimes \text{proj}_{\mathbf{Z}}(\mathbf{B})|}{\|\text{proj}_{\mathbf{Z}}(\mathbf{A})\| \cdot \|\text{proj}_{\mathbf{Z}}(\mathbf{B})\|}, \quad (12)$$

where $|\cdot|$ is the absolute operation. The norm $\|\text{proj}_{\mathbf{Z}}(\mathbf{A})\|$ is defined as:

$$\|\text{proj}_{\mathbf{Z}}(\mathbf{A})\| = \sqrt{(\text{proj}_{\mathbf{Z}}(\mathbf{A}) \otimes \text{proj}_{\mathbf{Z}}(\mathbf{A}))} \quad (13)$$

and the norm $\|\text{proj}_{\mathbf{Z}}(\mathbf{B})\|$ similarly defined.

Given the partial distance measure, a key challenge relates to the design of the loss function for learning the balancing representation (LeCun, Chopra, Hadsell, Ranzato, & Huang, 2006). The motivation for employing contrastive learning stems from the limitations of the evaluation metrics for conditional independence. When measuring conditional associations for three high-dimensional or multi-dimensional indicators, there is currently no metric, including the partial distance measure, that guarantees conditional independence when the conditional association indicator value is zero within a finite sample size (Azadkia & Chatterjee, 2021; Doran et al., 2014; Runge, 2018; Székely & Rizzo, 2014; Zhang et al., 2012; Zhen et al., 2022).

In this paper, we consider this to be a potential mode collapse issue (Goodfellow et al., 2020; Jing, Vincent, LeCun, & Tian, 2021). Mode collapse is a fundamental problem in representation learning (Chen & He, 2021; Chen et al., 2020; He et al., 2020) and arises when a model fails to adequately capture the diverse patterns within the data, instead collapsing them into a single mode or a limited set of modes. For instance, if the partial distance value is zero, but this does not confirm conditional independence, it becomes susceptible to mode collapse. This issue may arise when one naively minimizes partial distance measures for the original observed samples. Such a situation can introduce bias into the learned representation, which in turn affects

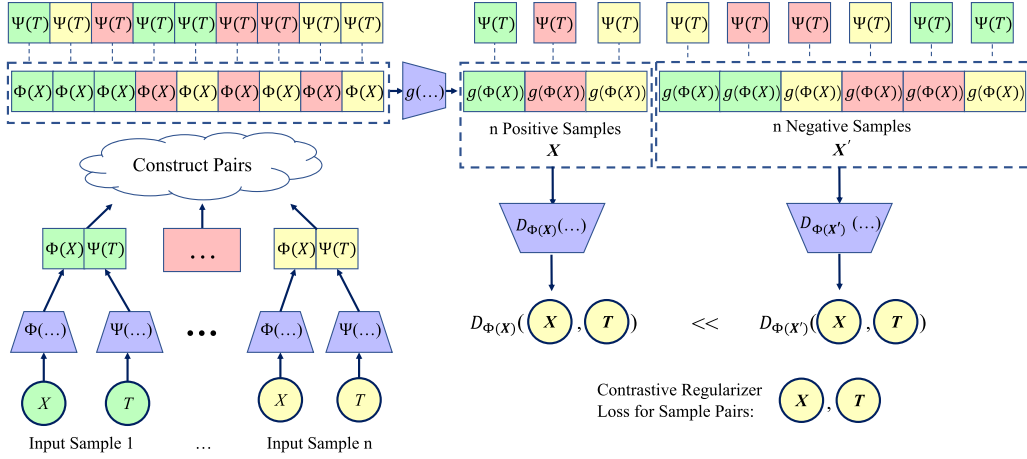


Fig. 1. Contrastive regularizer. The n covariates X undergo a transformation via the encoder Φ resulting in $\Phi(X)$. $\Phi(X)$ transforms to $g(\Phi(X))$ through projection head g (Chen et al., 2020). $g(\Phi(X))$ is directly constrained by $\ell_{\Phi}^{CR}(X, T)$. To simplify notation, we use $\Phi(X)$ in the context to represent $g(\Phi(X))$. $D_{\Phi(X)}$ and $D_{\Phi(X')}$ are partial distance measure of positive/negative samples (Székely & Rizzo, 2014).

the estimation of heterogeneous dose–response curves (Li, Xiao, et al., 2023). To obviate it, we adopt the contrastive learning (Chen et al., 2020) to alleviate the issue of insufficient correlation measurement by introducing positive and negative sample pair constraints to the double balancing representation.

We propose a novel Contrastive Regularizer (CR) to constrain the balancing representation. Since directly employing partial distance correlation carries the risk of misjudging some conditionally associated samples as conditionally independent, we introduce positive and negative sample pairs. In these pairs, the ideal representation of the positive samples satisfies conditional independence, while the ideal representation of the negative samples does not. By enforcing the constraint that the partial distance correlation of the positive sample representation is significantly smaller than that of the negative sample representation, we can enhance the evaluation capability of partial distance correlation for conditional independence given a specific representation. To construct the positive and negative samples for contrastive learning, we randomly shuffle X in the original data m times to get $n * m$ permuted data X' that $X' \perp\!\!\!\perp T$ (Arbour et al., 2021; Cheng et al., 2020). Then we name X positive samples and X' negative samples and define their partial distance measure as follows:

$$D_{\Phi(X)}(X, T) = \frac{|\text{proj}_{\Phi(X)}(X) \cdot \text{proj}_{\Phi(X)}(T)|}{\|\text{proj}_{\Phi(X)}(X)\| \cdot \|\text{proj}_{\Phi(X)}(T)\|}, \quad (14)$$

$$D_{\Phi(X')}(X, T) = \frac{|\text{proj}_{\Phi(X')}(X) \cdot \text{proj}_{\Phi(X')}(T)|}{\|\text{proj}_{\Phi(X')}(X)\| \cdot \|\text{proj}_{\Phi(X')}(T)\|}.$$

As Fig. 1 shown, given the positive representation $\Phi(X)$ and negative representation $\Phi(X')$ from n observed covariates X and n shuffled covariates X' , the correctly specified function Φ should satisfy that $D_{\Phi(X)}(X, T) \ll D_{\Phi(X')}(X, T)$. Then we propose to perform contrastive learning for the positive samples X and negative samples X' . The contrastive regularizer loss is formulated as follows:

$$\ell_{\Phi}^{CR}(X, T) = D_{\Phi(X)}(X, T) - \log \sum_{j=1}^m \exp(D_{\Phi(X'_j)}(X, T)), \quad (15)$$

where m represents the number of shuffles, and X'_j denotes the shuffled covariates from the n negative samples acquired during the j shuffle. The total count of negative samples is given by $n * m$. During the training procedure, for each batch of samples, we perform random shuffling of the original covariates within the batch a total of m times (Cheng et al., 2020).

It is worth noting that the contrastive regularizer serves a dual purpose: it preserves the balancing representation condition and ensures

the continuity of treatments, thereby enhancing the quality of prognostic representation. More specifically, given that all discrimination operates at the instance level (Wu et al., 2018), there is no necessity to discretize treatment variables into bins (Bica et al., 2020; Nie et al., 2021; Schwab et al., 2020). Consequently, the continuity of treatments is inherently maintained. Furthermore, the contrastive regularizer with positive and negative samples effectively captures diverse information from covariates X for the learned representation. This aspect aligns with the requirement of addressing prognostic representation condition, aiming to capture the heterogeneity in causal effects among various study subjects (Hansen, 2008).

5.2. End-to-end CRNet

Different from the balancing representation condition, which focuses on the treatment assignment. The prognostic representation condition focus on the outcome prediction power in representation. In this paper, we firstly design a end-to-end two-head neural network, which encodes the treatments T through Ψ and covariates X through Φ for representations $\Psi(T)$ and $\Phi(X)$. Then, we adopt a mean square error loss (MSE) to directly constrain the condition in the double balancing representation that $Y \perp\!\!\!\perp X | \Phi(X)$. In particular, for a unit i , the MSE loss is formulated as follows:

$$\ell^{\text{MSE}}(X_i, T_i, Y_i) = (Y_i - h(\Phi(X_i), \Psi(T_i)))^2. \quad (16)$$

Although the MSE loss has been commonly employed in previous works for outcome prediction (Bica et al., 2020; Nie et al., 2021; Schwab et al., 2020), it is important to note that most of these approaches do not constrain the prognostic representation condition effectively. To elaborate, DRNet (Schwab et al., 2020) imposes the MSE loss on a treatment-balanced representation that is independent of treatments T . However, as treatments T is correlated with covariates X , this approach may lead to a loss of essential confounder information for outcome prediction. VCNet (Nie et al., 2021) utilizes the MSE loss on representations constrained by a propensity score estimator, which is considered the coarsest balancing score (Rosenbaum & Rubin, 1983). Representations subject to this constraint may also fail to satisfy the prognostic representation condition (Hahn, 1998). Furthermore, these methods tend to neglect the issue of mode collapse, which can compromise their ability to estimate heterogeneous causal effects effectively. Different from them, our MSE loss is imposed on double balancing representations. We achieve this by employing the contrastive regularizer to preserve the confounder information of X . This facilitates the CRNet to learn prognostic representation as much as possible.




Fig. 2. The end-to-end CRNet. For the training procedure, the representations $\Phi(X)$ constrained by contrastive loss $\ell_{\Phi}^{\text{CR}}(X, T)$ are concatenated and input to MLPs h to obtain the estimated outcome \hat{Y} by the final loss in Eq. (17). The final objective is to minimize the loss. For the inference procedure, the estimated HDRC is obtained by $h(\Phi(X), \Psi(T))$.

In summary, we propose an end-to-end neural network framework called CRNet for the estimation of HDRC. As depicted in Fig. 2, the overall architecture of CRNet comprises three distinct blocks: (a) Two-Head Encoder. The first head, denoted as Φ , encodes covariates X into representation $\Phi(X)$. The role of the second head, denoted by Ψ , is to encode the treatments to the representation $\Psi(T)$. (b) Projection Head. The projection head, denoted by g , project the covariate representation $\Phi(X)$ into $g(\Phi(X))$ for the partial distance measure involving covariates X and treatments T . (c) Outcome Estimator. The outcome estimator h takes the concatenated representations of covariates and treatments, $\Phi(X)$ and $\Psi(T)$, as input and transforms them into $h(\Phi(X), \Psi(T))$. This estimated outcome $h(\Phi(X), \Psi(T))$ approximates the observed Y by the regularized regression loss $\ell_{\text{CR}}^{\text{final}}(X, T, Y)$, which is formulated as follows:

$$\ell_{\text{CR}}^{\text{final}}(X, T, Y) = \sum_{i=1}^n \ell^{\text{MSE}}(X_i, T_i, Y_i) + \alpha * \ell_{\Phi}^{\text{CR}}(X, T), \quad (17)$$

where α represents the hyperparameter. The use of both the MSE loss ℓ^{MSE} and the contrastive regularizer loss ℓ_{Φ}^{CR} is necessary for unbiased estimation of HDRC. The loss ℓ^{MSE} is used to minimize the difference between the predicted outcome and the observed outcome, thus ensuring that the model can make reliable predictions. Without this loss, the CRNet model would not be able to accurately predict the outcomes, resulting in loss of prediction capacity. On the other hand, the loss ℓ_{Φ}^{CR} helps to prevent treatment assignment bias and outcome overfitting by comparing the representation of positive/negative samples. Without this loss, the CRNet model might induce treatment assignment bias or mode collapse, leading to inaccurate outcome prediction. In short, the two losses in $\ell_{\text{CR}}^{\text{final}}$ complement each other.

Algorithm 1 The end-to-end training procedure of CRNet

Require: Training set data: \mathcal{O} ; Loss function: $\ell_{\text{CR}}^{\text{final}}$; Learning rate: ϵ ; Number of iterations: I ; CRNet: $\hat{\pi}$; All parameters of CRNet: Θ^{CRNet} .

```

1: for  $i = 1$  to  $I$  do
2:   Train CRNet  $\hat{\pi}$  on  $\mathcal{O}$ ;
3:    $\Theta^{\text{CRNet}} \leftarrow \Theta^{\text{CRNet}} - \epsilon \nabla \ell_{\text{CR}}^{\text{final}}(\Theta^{\text{CRNet}}; \mathcal{O})$ .
4: end for

```

Ensure: CRNet $\hat{\pi}$.

The end-to-end algorithm of CRNet training is presented in Algorithm 1. The advantage of this CRNet lies in its simplicity, making it ideal to serve as the backbone for new modules. The method we will introduce in the next section provides an example of how CRNet can be utilized for expansion.

6. Three-stage WDBN

The end-to-end CRNet aims to capture confounder information as much as possible by learning a double balancing representation. Ideally, if the confounder distribution is learned well enough, it can estimate unbiased heterogeneous dose-response curves (HDRC). However, in settings with high-dimensional covariates and multi-dimensional continuous treatments, learning a sufficiently accurate confounder distribution from a limited sample size is challenging. To further improve the model's performance, following (Fong et al., 2018; Huling et al.,

2024; Imai & Van Dyk, 2004; Vegetabile et al., 2021), we can also learn a sample weight based on the double balancing representation to explicitly balance the distribution of samples.

6.1. Double contrastive regularizer

In CRNet, we use ℓ_{Φ}^{CR} and ℓ^{MSE} to respectively constrain the balancing representation condition $X \perp\!\!\!\perp T | \Phi(X)$ and learnable prognostic representation condition $X \perp\!\!\!\perp Y | \Phi(X)$. This is reasonable in an end-to-end training strategy because the model's predictions satisfy both unbiasedness and retain the predictive information for outcomes. However, when we aim to directly learn the sample weight from the representation $\Phi(X)$, we prefer it to explicitly satisfy the double balancing representation condition, $X \perp\!\!\!\perp T | \Phi(X)$, $X \perp\!\!\!\perp Y | \Phi(X)$, rather than merely constraining $X \perp\!\!\!\perp T | \Phi(X)$ while regressing outcomes. Therefore, we extend the contrastive regularizer to the double contrastive regularizer.

To construct the new positive and negative samples for directly constraining the double balancing representation condition, we randomly shuffle the n units of X in the original data m times to get $n * m$ permuted data X' that $X' \perp\!\!\!\perp S$ where $S = T \oplus Y$, means the concatenation of treatment and outcomes. \oplus is the concatenation operation. Then, their partial distance measure are formulated as follows:

$$D_{\Phi(X)}(X, S) = \frac{|\text{proj}_{\Phi(X)}(X) \cdot \text{proj}_{\Phi(X)}(S)|}{\|\text{proj}_{\Phi(X)}(X)\| \cdot \|\text{proj}_{\Phi(X)}(S)\|}, \quad (18)$$

$$D_{\Phi(X')}(X, S) = \frac{|\text{proj}_{\Phi(X')}(X) \cdot \text{proj}_{\Phi(X')}(S)|}{\|\text{proj}_{\Phi(X')}(X)\| \cdot \|\text{proj}_{\Phi(X')}(S)\|}.$$

To perform contrastive learning for the positive samples X and negative samples X' , we modify the contrastive regularizer loss ℓ_{Φ}^{CR} in Eq. (15) to the double contrastive regularizer loss, which is formulated as follows:

$$\ell_{\Phi}^{\text{DCR}}(X, S) = \sum_{i=1}^n \ell^{\text{MSE}}(X_i, T_i, Y_i) + \alpha (D_{\Phi(X)}(X, S) - \log \sum_{j=1}^m \exp(D_{\Phi(X'_j)}(X, S))). \quad (19)$$

Similar to the contrastive regularizer, the double contrastive regularizer serves a dual purpose. It not only explicitly preserving the double balancing representation condition but also ensuring the continuity of treatments, thereby benefiting the prediction of outcomes, owing to its enhancement of representation quality. The key difference between the double contrastive regularizer and the contrastive regularizer is that the former explicitly enforces the double balancing representation condition, rather than just the balancing representation condition. This is done in order to prepare the sample weight learning task for downstream applications.

6.2. Double balancing sample weight

To explicitly balance the treatment assignment and thereby improve the model's ability to predict HDRC, we propose learning the sample

weights based on the double balancing representation, which are then used for weighted regression on the outcomes. First, we define the Double Balancing Sample Weight (DBSW) as $\mathbf{w} = (w_1, \dots, w_n)$. Then, we define the weighted double balancing representation $\Phi(\tilde{X})$ as:

$$\Phi(\tilde{X}) = \mathbf{w}\Phi(X). \quad (20)$$

And the double balancing sample weight is determined by the following objective function:

$$\mathbf{w} \in \underset{\mathbf{w}=(w_1, \dots, w_n)}{\operatorname{argmin}} -D_{\Phi(\tilde{X})}(X, S) + \beta \frac{1}{n^2} \sum_{i=1}^n w_i^2, \quad (21)$$

such that $\sum_{i=1}^n w_i = n, w_i \geq 0,$

where $\beta > 0$ is the hyperparameter. In this objective function, the first term implies that we desire the weighted samples to satisfy the condition $X \perp\!\!\!\perp S | \Phi(\tilde{X})$. The underlying intuition here is that by constraining $X \perp\!\!\!\perp S | \Phi(X)$ to satisfy the double balancing representation condition, the learned $\Phi(X)$ aims to encompass as much information about all confounders as possible. Consequently, we need the weighted representation $\Phi(\tilde{X})$ to not satisfy this condition, in order to balance the distribution of confounders. This ensures that the predictions based on the representation $\Phi(\tilde{X})$ are unaffected by the distribution of confounders. The second term in the objective function serves as a penalty term, the purpose of which is to constrain the variance of the sample weight distribution. By avoiding extreme values, it aims to enhance the stability of the sample weights. Furthermore, the purpose of defining the two conditions, $\sum_{i=1}^n w_i = n, w_i \geq 0$, is to prevent the sample weights from falling into trivial solutions.

We apply a neural network module b to learn the DBSW from the representation $\Phi(X)$. Specifically, b is composed of a single-layer MLP followed by an ELU activation function and another single-layer MLP. We then obtain the DBSW w_i from the following formula:

$$w_i = \frac{\tilde{w}_i}{\sum_{i=1}^n \tilde{w}_i / n}, \text{ where } \tilde{w}_i = \text{Softplus}(b(\Phi(X_i))). \quad (22)$$

where \tilde{w}_i is the unnormalized sample weight. This equation implies that $b(\Phi(X_i))$ will pass through a Softplus activation function to ensure $w_i \geq 0$, and then be normalized to satisfy the condition $\sum_{i=1}^n w_i = n$. Then, the loss function for learning the DBSW is formulated as:

$$\ell^W(X, S, \mathbf{w}) = -D_{\Phi(\tilde{X})}(X, S) + \beta \frac{1}{n^2} \sum_{i=1}^n w_i^2. \quad (23)$$

After minimizing ℓ^W , we can apply the learned DBSW for the final weighted regression of the outcomes.

Algorithm 2 The three-stage training procedure of WDBN

Require: Training set data: O ; Loss functions: $\ell_{\Phi}^{\text{DCR}}, \ell^W, \ell^{\text{WMSE}}$; Learning rate: ϵ ; Number of iterations in three stages: I_1, I_2, I_3 ; WDBN: $\hat{\tau}$; Parameters of WDBN: $\Theta^{\Phi}, \Theta^{\Psi}, \Theta^g, \Theta^h, \Theta^b$.

```

1: for  $i = 1$  to  $I_1$  do
2:   Train DCR modules  $\Phi, \Psi, g$  on  $O$ ; {First stage}
3:    $\Theta^{\Phi} \leftarrow \Theta^{\Phi} - \epsilon \nabla_{\Theta^{\Phi}} \ell_{\Phi}^{\text{DCR}}(\Theta^{\Phi}; O)$ ;
4:    $\Theta^{\Psi} \leftarrow \Theta^{\Psi} - \epsilon \nabla_{\Theta^{\Psi}} \ell_{\Phi}^{\text{DCR}}(\Theta^{\Psi}; O)$ ;
5:    $\Theta^g \leftarrow \Theta^g - \epsilon \nabla_{\Theta^g} \ell_{\Phi}^{\text{DCR}}(\Theta^g; O)$ ;
6: end for
7: for  $i = 1$  to  $I_2$  do
8:   Train DBSW module  $b$  on  $O$ ; {Second stage}
9:    $\Theta^b \leftarrow \Theta^b - \epsilon \nabla_{\Theta^b} \ell^W(\Theta^b; O)$ ;
10: end for
11: for  $i = 1$  to  $I_3$  do
12:   Train WMSE module  $h$  on  $O$ ; {Third stage}
13:    $\Theta^h \leftarrow \Theta^h - \epsilon \nabla_{\Theta^h} \ell^{\text{WMSE}}(\Theta^h; O)$ ;
14: end for

```

Ensure: WDBN $\hat{\tau}$.

6.3. Three-stage WDBN

Based on the proposed double contrastive regularizer and double balancing sample weight, we can design a Weighted Double Balancing Network (WDBN) to implement the weighted regression for HDRC.

As Fig. 3 shown, the overall architecture of WDBN comprises four distinct blocks: (a) Two-Head Encoder. The first head, denoted as Φ , encodes covariates X into representation $\Phi(X)$. The second head, denoted by Ψ , encodes treatments T into representation $\Psi(T)$; (b) Projection Head. The projection head, denoted by g , project the covariate representation $\Phi(X)$ into $g(\Phi(X))$ for the partial distance measure involving covariates X and the concatenation S of treatments and outcomes; (c) Weight Learner. The weight learner b learns the DBSW from the double balancing representation, which is constrained by ℓ^W . (d) Outcome Estimator. The outcome estimator h takes the concatenated representations of covariates and treatments, $\Phi(X)$ and $\Psi(T)$, as input and transforms them into $h(\Phi(X), \Psi(T))$. This estimated outcome $h(\Phi(X), \Psi(T))$ approximates the observed Y by the Weighted Mean Square Error (WMSE) loss, which is formulated as follows:

$$\ell^{\text{WMSE}}(X, S, \mathbf{w}) = \sum_{i=1}^n w_i (Y_i - h(\Phi(X_i), \Psi(T_i)))^2. \quad (24)$$

The three-stage algorithm of WDBN training is presented in Algorithm 2. The inference procedure is similar to CRNet that the estimated HDRC is obtained by $h(\Phi(X), \Psi(T))$.

In comparison to the end-to-end CRNet, the three-stage WDBN introduces sample weights to explicitly balance the sample distribution through weighted regression, thereby enhancing the prediction capability for HDRC within limited samples. To ensure the stability of this model, we need to train each module in three separate stages, which reduces the training efficiency compared to the end-to-end CRNet. In addition to reducing training efficiency compared to the end-to-end CRNet, WDBN also incurs higher training costs, especially for high-dimensional datasets. Therefore, the choice between CRNet and WDBN can be seen as a trade-off between performance and cost (including both training and inference costs). In summary, CRNet is suitable for clean and abundant datasets, as it has lower training costs, can be deployed quickly, and delivers effective results. In contrast, WDBN is more appropriate for datasets with insufficient data and imbalanced distributions, as it better balances the data.

7. Experiments

Since there are no widely used open-source real-world datasets for continuous treatment effect estimation, we follow the approach of VCNet and SCIGAN (Bica et al., 2020; Nie et al., 2021) by conducting semi-synthetic experiments on the IHDP² and News³ datasets to assess the real-world applicability of our method.

The IHDP dataset is a real-world dataset in a binary treatment setting, used to study the relationship between “whether an expert visits a household” and “an infant’s later cognitive test scores.” We include IHDP to assess the applicability of our method in the education domain. The News dataset, on the other hand, is a real-world dataset in a multi-treatment setting, designed to analyze the impact of different devices on user retention time. We include News to explore the potential of our method in marketing applications.

Overall, we simulate 4 synthetic datasets and 5 semi-synthetic datasets to verify the performance of our proposed CRNet and WDBN. All experiments are conducted on Intel(R) Xeon(R) Gold 6240 CPU @ 2.60 GHz and NVIDIA GeForce RTX 3090.

² <https://www.fredjio.com>

³ <https://paperdatasets.s3.amazonaws.com/news.db>

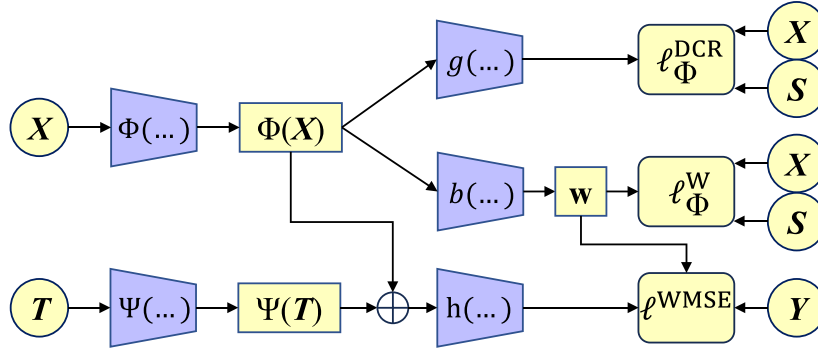


Fig. 3. The three-stage WDBN. For the training procedure, the representations $\Phi(X)$ and $\Psi(T)$ are constrained by double contrastive loss $\ell_{\Phi}^{\text{DCR}}(X, T)$. Subsequently, $\Phi(X)$ and $\Psi(T)$ are concatenated and input to the MLPs b to obtain the DBSW w using the loss ℓ_{Φ}^w , and input to the MLPs h to obtain the estimated outcome \hat{Y} by the WMSE loss ℓ^{WMSE} . For the inference procedure, the estimated HDRC is obtained by $h(\Phi(X), \Psi(T))$.

7.1. Experimental setup

Synthetic data generation. We simulate synthetic data in the following manner. For each unit $i \in \{1, 2, \dots, n\}$, we generate $p = 100$ covariates from an independent and identically distributed (i.i.d.) normal distribution, i.e., $X_i \sim \mathcal{N}(\mathbf{0}_p, E_p)$, where $\mathbf{0}_p$ denotes a p -dimensional vector with all elements equal 0, and E_p represents p -order identity matrix. We generate q treatments using the following rules:

$$T_{i,j} = 0.2 \sum_{j=1}^5 W_j X_{i,j} + \frac{1}{p-10} \sum_{j=11}^p W_j X_{i,j}^2 + \bar{T}_{i,j} + 0.5 \bar{T}_{i,j}^3 X_{i,p-j}. \quad (25)$$

Here, we denote $\bar{T}_{i,j} \sim \mathcal{N}(0, 1)$, $W_j \sim U(0.5, 1)$ for $j \in \{1, 2, \dots, q\}$. The outcome is generated according to the following rules:

$$Y_i = 0.5 \sum_{j=1}^q W_j^T |T_{i,j}| + 0.5 \sum_{j=6}^{10} W_j X_{i,j}^2 + \sum_{j=11}^p W X_i + 0.5 \sum_j T_{i,j} X_{i,q-j-10}. \quad (26)$$

Here, we denote $W_j^T \sim U(0.5, 1)$. Then, we design 4 simulation datasets and name them Data- q where q means the dimension of T and X (e.g., Data-1 means a simulation with 1 treatment, 100 covariates). Then we sample 2100/600/300 units for training/validation/test for each data.

Semi-synthetic data generation. We proceed to perform semi-simulation experiments with the aim of demonstrating the robustness of our method across a range of settings. In our semi-synthetic experiments, we retain the original covariates of these datasets but extend the treatment variables to multi-dimensional continuous treatments to align with our study's setting. To evaluate the results, we generated the true causal effects based on the data and covariates. We sample units from the IHDP data to create the training, validation, and test sets, with 522/150/75 units for each data split. For the News dataset, we perform data splits into training, validation, and test sets with 2100/600/300 units, respectively.

Baselines and evaluation. We compare our model with the following baselines in the above datasets: For statistical methods, we use (1) **Causal Forest** (Wager & Athey, 2018), a random forest algorithm for causal inference. (2) **GPS** (Imbens, 2000), a generalized propensity score for continuous treatments. (3) **CBGPS** (Fong et al., 2018), a generalized covariate balancing propensity score (Imai & Ratkovic, 2014) for continuous treatments. For deep methods, we apply (4) **SCIGAN** (Bica et al., 2020), a hierarchical generative adversarial network (Goodfellow et al., 2020). (5) **DRNet** (Schwab et al., 2020), a multi-head deep model stratified according to treatment. (6) **VCNet** (Nie et al., 2021), a varying coefficient neural network with functional targeted regularization.

Additionally, we incorporate three classic deep methods designed for binary treatment settings to highlight the necessity of specially designed model structures for continuous treatment scenarios: (7) **TARNet** (Shalit et al., 2017), a two-headed neural network model, where each head corresponds to the outcome for the treatment and control groups, respectively. (8) **CFRNet** (Shalit et al., 2017), a neural network model based on TARNet that introduces a constraint ensuring independence between covariate representations and treatment variables. (9) **DragonNet** (Shi, Blei, & Veitch, 2019), a doubly robust three-headed neural network, where each head corresponds to the propensity score, the treatment group outcome, and the control group outcome, respectively. We define the treatment variable for these methods as follows:

$$\bar{T}_i = \mathbb{I} \left(\frac{1}{q} \sum_{j=1}^q T_{ij} > \frac{1}{nq} \sum_{i=1}^n \sum_{j=1}^q T_{ij} \right), \quad (27)$$

where $\mathbb{I}(\cdot)$ represents the treatment variable input for binary treatment methods. This formulation effectively reduces the multi-dimensional treatment values to a one-dimensional proxy by taking the average across all dimensions. If the proxy treatment value for the i th sample is greater than the overall sample mean of the proxy treatment, \bar{T}_i is assigned a value of 1; otherwise, it is assigned 0.

For all experiments, we perform 30 replications to report the mean integrated square error (MISE) and the standard deviations (std) of HDRC estimation:

$$\text{MISE} = s^{-1} \sum_{i=1}^s \int_a^b (h(t, X_i) - \hat{h}(t, X_i))^2 dt, \quad (28)$$

where s is the test sample size and $[a, b]$ is the interval of treatment values.

7.2. Results

Performance comparison. We carry out simulation and semi-simulation experiments, as presented in Tables 2 and 3. In this table, bold text indicates the optimal performance, while underlined text represents suboptimal performance. In these experiments, as the dimensionality of treatments increases, traditional statistical methods GPS, CBGPS and Causal Forest tend to fail, highlighting their limitations in handling complex, high-dimensional data. The binary treatment methods perform poorly across all dataset settings, which indicates that directly applying binary treatment methods to model data from multi-dimensional continuous treatment scenarios results in significant information loss, emphasizing the necessity of designing models specifically for multi-dimensional continuous treatment scenarios. SCIGAN's poor performance in high-dimensional data reflects the instability inherent in generative adversarial networks while also emphasizing the necessity of balancing representation condition.

Table 2Performance comparison (MISE \pm std) and ablation studies on simulation data.

Method	Data-1	Data-2	Data-5	Data-10
GPS (Imai & Van Dyk, 2004)	57.7 \pm 18	57.8 \pm 14	57.4 \pm 15	78.0 \pm 19
CBGPS (Fong et al., 2018)	57.8 \pm 18	57.8 \pm 14	57.3 \pm 15	70.5 \pm 19
Causal Forest (Wager & Athey, 2018)	1.83 \pm 0.6	2.50 \pm 0.7	5.16 \pm 0.9	14.9 \pm 2.4
TARNet (Shalit et al., 2017)	57.7 \pm 18	57.8 \pm 14	57.2 \pm 15	69.8 \pm 19
CFRNet (Shalit et al., 2017)	57.8 \pm 18	57.7 \pm 14	57.1 \pm 15	69.6 \pm 19
DragonNet (Shi et al., 2019)	113 \pm 38	119 \pm 42	117 \pm 30	142 \pm 30
DRNet (Schwab et al., 2020)	2.35 \pm 0.7	3.49 \pm 1.4	6.39 \pm 2.1	18.5 \pm 4.7
SCIGAN (Bica et al., 2020)	15.0 \pm 13	26.1 \pm 13	43.6 \pm 15	59.6 \pm 26
VCNet (Nie et al., 2021)	5.79 \pm 4.8	6.41 \pm 4.7	13.7 \pm 5.7	28.2 \pm 7.1
CRNet	<u>1.69 \pm 0.5</u>	<u>2.07 \pm 0.8</u>	<u>3.05 \pm 0.7</u>	<u>7.55 \pm 2.6</u>
w/o BR	2.04 \pm 0.5	2.56 \pm 1.0	4.76 \pm 1.2	9.69 \pm 5.1
w/o PR	52.9 \pm 16	53.9 \pm 14	51.0 \pm 13	55.2 \pm 16
WDBN	1.20 \pm 0.4	1.57 \pm 0.6	2.57 \pm 0.6	5.81 \pm 1.5
w/o DBSW	1.88 \pm 0.7	2.35 \pm 0.7	3.99 \pm 1.1	9.27 \pm 2.1

Table 3Performance comparison (MISE \pm std) and ablation studies on real-world data IHDP- q and News- q .

Method	IHDP-1	News-2	News-4	News-8	News-16
GPS (Imai & Van Dyk, 2004)	0.98 \pm 0.4	84.3 \pm 4.6	83.8 \pm 4.5	87.8 \pm 3.8	89.0 \pm 4.2
CBGPS (Fong et al., 2018)	1.03 \pm 0.4	84.1 \pm 4.4	83.6 \pm 4.6	85.5 \pm 5.0	86.7 \pm 4.5
Causal Forest (Wager & Athey, 2018)	0.79 \pm 0.3	26.8 \pm 15	26.9 \pm 11	47.9 \pm 21	82.7 \pm 81
TARNet (Shalit et al., 2017)	1.01 \pm 0.4	264 \pm 9.1	266 \pm 7.8	270 \pm 10	277 \pm 22
CFRNet (Shalit et al., 2017)	1.01 \pm 0.4	265 \pm 9.8	264 \pm 6.0	269 \pm 10	277 \pm 19
DragonNet (Shi et al., 2019)	6.18 \pm 9.7	479 \pm 31	470 \pm 37	479 \pm 49	493 \pm 40
DRNet (Schwab et al., 2020)	1.29 \pm 0.4	18.0 \pm 8.9	18.6 \pm 10	33.3 \pm 65	26.1 \pm 10
SCIGAN (Bica et al., 2020)	0.65 \pm 0.3	233 \pm 218	163 \pm 151	254 \pm 365	200 \pm 248
VCNet (Nie et al., 2021)	1.28 \pm 0.7	11.3 \pm 6.0	9.80 \pm 3.3	26.5 \pm 51	25.3 \pm 31
CRNet	<u>0.22 \pm 0.1</u>	<u>3.21 \pm 1.4</u>	<u>5.19 \pm 2.3</u>	<u>8.35 \pm 5.0</u>	<u>9.18 \pm 2.9</u>
w/o BR	0.63 \pm 0.4	6.03 \pm 4.6	5.60 \pm 3.4	9.88 \pm 5.9	15.9 \pm 23
w/o PR	0.92 \pm 0.4	35.3 \pm 17	36.1 \pm 17	36.1 \pm 15	38.3 \pm 14
WDBN	0.20 \pm 0.1	2.89 \pm 0.7	3.00 \pm 1.2	4.19 \pm 0.7	8.72 \pm 1.5
w/o DBSW	0.46 \pm 0.2	4.20 \pm 1.4	5.47 \pm 1.3	9.26 \pm 6.6	13.4 \pm 7.5

The CRNet outperforms both DRNet and VCNet, highlighting that solely relying on treatment-balanced representation or balancing representation can indeed result in a decrease in heterogeneous dose-response curve estimation. This underscores the effectiveness of the double balancing representation condition in enhancing both the constraint on unbiasedness and the predictive capacity for heterogeneous outcomes. The three-stage WDBN achieves a state-of-the-art performance level in all experiments. This demonstrates that weighted regression based on the DBSW can indeed further improve the model's performance in estimating heterogeneous dose-response curves on the basis of double balancing representation. The superior performance of CRNet and WDBN on the IHDP and News datasets confirms its potential application value in fields such as education and marketing.

Ablation studies. To verify the performance of balancing representation, we conduct the w/o balancing representation (BR) ablation study on CRNet, setting the hyperparameter $\alpha = 0$. To assess the performance of the prognostic representation, we conduct the w/o prognostic representation (PR) ablation study on CRNet, which employs a two-stage training strategy: Only loss ℓ_{ϕ}^{CR} is used in the first stage, and only ℓ^{MSE} loss is used in the second stage. To validate the performance of double balancing sample weight (DBSW), we perform a w/o DBSW ablation study on WDBN, wherein the DBSW is omitted and all sample weights are set to a constant value of 1.

The results are shown in Tables 2 and 3. Although w/o balancing representation achieved good performance in most settings, its performance was still significantly degraded compared to CRNet. On the other hand, w/o prognostic representation performed poorly in all settings. This result aligns with our expectations since the model's predictive accuracy deteriorates when the prognostic representation condition is unsatisfied, and the sole reliance on prognostic representation proves biased in practice. In w/o DBSW, we find that even without DBSW, the

Table 4Ablation studies (MISE \pm std) of CRNet varying values m of the number of negative sample augmentations.

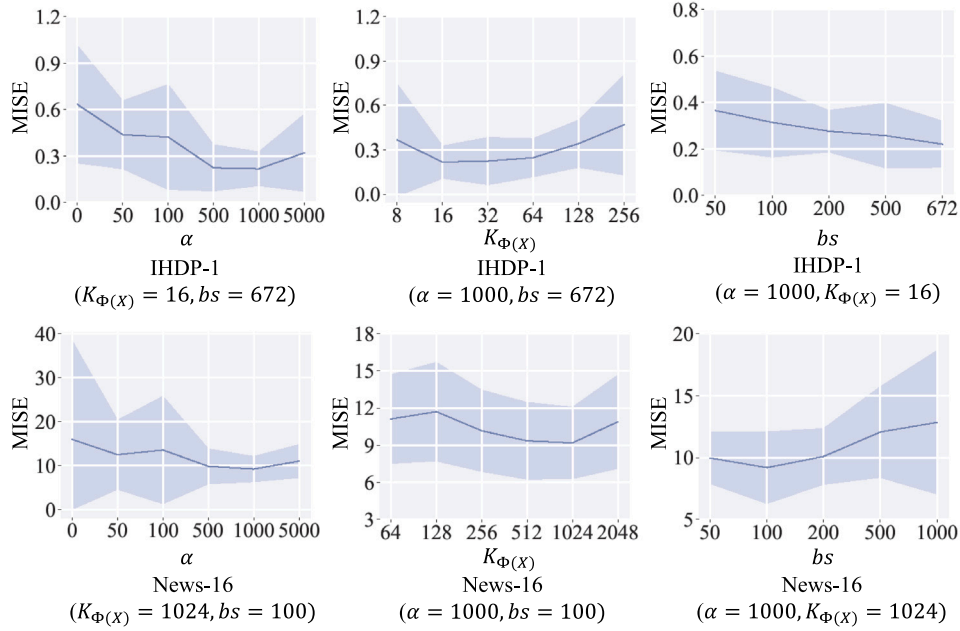
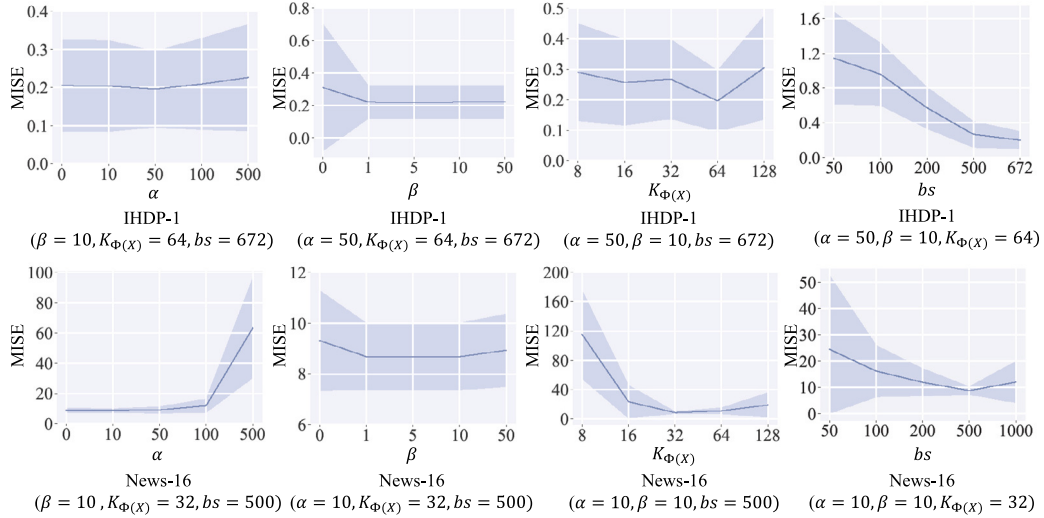
$n * m$	Data-1	Data-10	IHDP-1	News-16
$m = 0$	1.99 \pm 0.6	11.0 \pm 3.1	0.58 \pm 0.3	10.2 \pm 2.9
$m = 1$	1.69 \pm 0.5	7.55 \pm 2.6	0.22 \pm 0.1	9.18 \pm 2.9
$m = 2$	1.69 \pm 0.5	7.06 \pm 2.4	0.24 \pm 0.1	10.1 \pm 4.0
$m = 3$	<u>1.69 \pm 0.5</u>	<u>7.38 \pm 3.2</u>	0.23 \pm 0.1	10.9 \pm 5.0
$m = 5$	1.71 \pm 0.6	8.08 \pm 3.1	0.23 \pm 0.1	10.5 \pm 4.1
$m = 10$	1.70 \pm 0.6	7.36 \pm 2.7	<u>0.23 \pm 0.1</u>	11.5 \pm 6.4

Table 5Ablation studies (MISE \pm std) of WDBN varying values m of the number of negative sample augmentations.

$n * m$	Data-1	Data-10	IHDP-1	News-16
$m = 0$	1.50 \pm 0.5	6.76 \pm 1.5	0.45 \pm 0.2	11.1 \pm 5.1
$m = 1$	1.35 \pm 0.5	5.95 \pm 1.4	0.20 \pm 0.1	8.88 \pm 1.3
$m = 2$	1.28 \pm 0.5	<u>5.87 \pm 1.2</u>	0.22 \pm 0.1	8.72 \pm 1.5
$m = 3$	<u>1.24 \pm 0.4</u>	5.81 \pm 1.5	0.22 \pm 0.1	9.31 \pm 2.9
$m = 5$	1.26 \pm 0.5	5.97 \pm 1.3	0.23 \pm 0.1	8.99 \pm 2.0
$m = 10$	1.20 \pm 0.4	6.33 \pm 1.5	<u>0.21 \pm 0.1</u>	9.06 \pm 1.6

model's performance remains comparable to previous methods (except for CRNet). When DBSW is added, the model's performance is further enhanced. This demonstrates that DBSW can indeed further improve the accuracy of HDRC estimation.

Hyperparameters tuning. For the performance of CRNet, we conduct experiments to evaluate the impact of hyperparameters α in Eq. (17), the dimension of double balancing representation $K_{\phi(X)}$ and the batch size bs . As Fig. 4 shown, we found that a large α improves estimation performance. Nevertheless, when α is too large, it will be an obstacle to

Fig. 4. The sensitivity experiments (MISE \pm std) of CRNet.Fig. 5. The sensitivity experiments (MISE \pm std) of WDBN.

fitting the outcome. Moreover, we found that increasing the dimension $K_{\Phi(x)}$ does not lead to a substantial improvement in estimation performance, which implies that CRNet is not sensitive to the dimension. The choice of batch size is highly dependent on the dataset. For the low-dimensional IHDP dataset, model performance improves progressively as the batch size increases, with the optimal batch size being the entire training set. However, for the high-dimensional News dataset, the situation is different—while an excessively small batch size still results in performance degradation, indiscriminately increasing the batch size can also harm model performance. Therefore, selecting a more moderate batch size is necessary for the News dataset.

We further conduct experiments, as shown in Table 4, by increasing the number of shuffle times m from 0 to 10. The results show a degradation in performance when $m = 0$, highlighting that naive minimization of the partial distance measure can induce mode collapse. Conversely,

increasing m to 1 significantly improves performance, indicating the effectiveness of our designed negative sample constraint. This issue is particularly pronounced in high-dimensional datasets such as Data-10 and News-16.

For the performance of WDBN, we conduct experiments to evaluate the impact of hyperparameters α , β in Eqs. (19) and (23), the dimension of double balancing representation $K_{\Phi(x)}$ and the augmentation of negative samples m . As Fig. 5 and Table 5 shown, we observe that large α , β and $K_{\Phi(x)}$ improves estimation performance. However, when they are too large, they pose an obstacle to fitting the outcome. The impact of batch size on WDBN is similar to its effect on CRNet. However, WDBN exhibits a more pronounced performance improvement as batch size increases. This is likely due to WDBN's more complex training process compared to CRNet, requiring a larger amount of data to achieve optimal performance. Moreover, the trend of m changes is

also quite similar to that of CRNet (see Table 4): there is a significant improvement when the value of m increases from 0 to 1, while further increasing the value of m yields diminishing returns.

8. Discussion

This study has some potential limitations and areas for future expansion. We assume unconfoundedness assumption and although we make efforts to control for potential confounding factors, there remains a possibility that unmeasured or unknown confounders may influence the results. Fortunately, since our method does not require a specially designed neural network structure, it can be easily extended to any gradient-based framework. Therefore, it can also be applied to scenarios where the above three assumptions are violated. For example, our representation learning method can be embedded into the DWR framework (Zhao et al., 2022) to relax the SUTVA assumption, incorporated into IV methods (Hartford, Lewis, Leyton-Brown, & Taddy, 2017) or negative control approaches (Xu, Kanagawa, & Gretton, 2021) to relax the unconfoundedness assumption, or integrated with beta-intact-VAE-based methods (Wu & Fukumizu, 2022) to relax the overlap assumption. Additionally, different conditional dependence measures may lead to different conclusions about the performance of our method. To focus on the challenges of heterogeneous causal effect estimation, we do not explore the relaxation of these assumptions or additional conditional independence criteria in depth. However, we look forward to future work extending the double balancing representation to broader settings.

While no well-established evaluation metrics exist for continuous treatment effects, in the field of binary treatment estimation using RCT data, the Area Under the Uplift Curve (AUUC) and QINI coefficient are commonly used as proxies for true causal effects (Zhang, Li, & Liu, 2020). Future research could explore the possibility of integrating observational data with randomized controlled trial (RCT) data for real-world dataset evaluation. One potential approach is to use observational data for model training while evaluating the model on RCT data using AUUC and QINI. However, two key challenges remain: How to collect RCT data in scenarios with multi-dimensional continuous treatment variables and how to extend AUUC and QINI metrics from binary or multi-valued settings to the multi-dimensional continuous domain remain open challenges. We look forward to future work addressing these issues.

Given that research on continuous causal effect estimation is still in its early stages, we have only preliminarily validated the potential applications of our method in education and marketing through semi-synthetic experiments. However, our approach is not limited to these fields. The advantage of our method lies in its ability to provide unbiased estimation of heterogeneous causal effects in complex scenarios involving high-dimensional covariates and multi-dimensional continuous treatments. This capability has broad applications in fields such as healthcare, law, and management. For example, in healthcare, our method could be used to analyze the effects of different drug dosages on patients. We look forward to future studies testing the effectiveness of our method in these application domains.

9. Conclusion

For estimating heterogeneous dose-response curves, we propose an end-to-end neural network called CRNet and a three-stage neural network named WDBN. Without disrupting the continuity of treatments, these networks distinguish between positive samples X and negative samples X' using a partial distance measure applied to a double balancing representation. By leveraging these networks to constrain the double balancing representation, we enforce unbiasedness in our causal estimates and effectively capture the heterogeneous dose-response curves from individuals.

CRedit authorship contribution statement

Minqin Zhu: Writing – review & editing, Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Anpeng Wu:** Writing – review & editing, Writing – original draft, Methodology. **Haoxuan Li:** Writing – review & editing, Methodology. **Ruoxuan Xiong:** Writing – review & editing, Methodology. **Bo Li:** Writing – review & editing, Conceptualization. **Fei Wu:** Funding acquisition, Data curation, Conceptualization. **Kun Kuang:** Writing – review & editing, Methodology, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (62376243, 62441605, 62037001), and the Starry Night Science Fund of Zhejiang University Shanghai Institute for Advanced Study, China (SN-ZJU-SIAS-0010). All opinions in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

Data availability

Data will be made available on request.

References

- Arbour, D., Dimmery, D., & Sondhi, A. (2021). Permutation weighting. In *International conference on machine learning* (pp. 331–341). PMLR.
- Assaad, S., Zeng, S., Tao, C., Datta, S., Mehta, N., Henao, R., et al. (2021). Counterfactual representation learning with balancing weights. In *International conference on artificial intelligence and statistics* (pp. 1972–1980). PMLR.
- Azadkia, M., & Chatterjee, S. (2021). A simple measure of conditional dependence. *The Annals of Statistics*, 49(6), 3070–3102.
- Baba, K., Shibata, R., & Sibuya, M. (2004). Partial correlation and conditional correlation as measures of conditional independence. *Australian & New Zealand Journal of Statistics*, 46(4), 657–664.
- Behrmann, J., Grathwohl, W., Chen, R. T., Duvenaud, D., & Jacobsen, J.-H. (2019). Invertible residual networks. In *International conference on machine learning* (pp. 573–582). PMLR.
- Bica, I., Jordon, J., & van der Schaar, M. (2020). Estimating the effects of continuous-valued interventions using generative adversarial networks. *Advances in Neural Information Processing Systems*, 33, 16434–16445.
- Blum, L., & Blum, M. (2023). A theoretical computer science perspective on consciousness and artificial general intelligence. *Engineering*, 25, 12–16.
- Chen, X., & He, K. (2021). Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 15750–15758).
- Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International conference on machine learning* (pp. 1597–1607). PMLR.
- Cheng, P., Hao, W., Dai, S., Liu, J., Gan, Z., & Carin, L. (2020). Club: A contrastive log-ratio upper bound of mutual information. In *International conference on machine learning* (pp. 1779–1788). PMLR.
- Doran, G., Muandet, K., Zhang, K., & Schölkopf, B. (2014). A permutation-based kernel conditional independence test. In *UAI* (pp. 132–141).
- Fleuret, F. (2004). Fast binary feature selection with conditional mutual information. *Journal of Machine Learning Research*, 5(9).
- Fong, C., Hazlett, C., & Imai, K. (2018). Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements. *The Annals of Applied Statistics*, 12(1), 156–177.
- Gan, L., Li, B., Kuang, K., Zhang, Y., Wang, L., Luu, A., et al. (2023). Exploiting contrastive learning and numerical evidence for confusing legal judgment prediction. In *Findings of the association for computational linguistics: EMNLP 2023* (pp. 12174–12185).

- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144.
- Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., et al. (2020). Bootstrap your own latent-a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33, 21271–21284.
- Hahn, J. (1998). On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, 315–331.
- Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20(1), 25–46.
- Hansen, B. B. (2008). The prognostic analogue of the propensity score. *Biometrika*, 95(2), 481–488.
- Hartford, J. S., Lewis, G., Leyton-Brown, K., & Taddy, M. (2017). Deep IV: A flexible approach for counterfactual prediction. In *International conference on machine learning*. URL: <https://api.semanticscholar.org/CorpusID:2951417>.
- He, K., Fan, H., Wu, Y., Xie, S., & Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9729–9738).
- Hirano, K., & Imbens, G. W. (2004). The propensity score with continuous treatments. *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, 226164, 73–84.
- Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, 81(396), 945–960.
- Hu, Z., Follmann, D. A., & Wang, N. (2014). Estimation of mean response via the effective balancing score. *Biometrika*, 101(3), 613–624.
- Huang, M.-Y., & Chan, K. C. G. (2017). Joint sufficient dimension reduction and estimation of conditional and average treatment effects. *Biometrika*, 104(3), 583–596.
- Huang, W., Yi, M., & Zhao, X. (2021). Towards the generalization of contrastive self-supervised learning. arXiv preprint arXiv:2111.00743.
- Huling, J. D., Greifer, N., & Chen, G. (2024). Independence weights for causal inference with continuous treatments. *Journal of the American Statistical Association*, 119(546), 1657–1670.
- Imai, K., & Ratkovic, M. (2014). Covariate balancing propensity score. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 76(1), 243–263.
- Imai, K., & Van Dyk, D. A. (2004). Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, 99(467), 854–866.
- Imbens, G. W. (2000). The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3), 706–710.
- Jing, L., Vincent, P., LeCun, Y., & Tian, Y. (2021). Understanding dimensional collapse in contrastive self-supervised learning. In *International conference on learning representations*.
- Kallus, N. (2020). Deepmatch: Balancing deep covariate representations for causal inference using adversarial training. In *International conference on machine learning* (pp. 5067–5077). PMLR.
- Kennedy, E. H., Ma, Z., McHugh, M. D., & Small, D. S. (2017). Non-parametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 79(4), 1229–1245.
- Lawrance, A. (1976). On conditional and partial correlation. *The American Statistician*, 30(3), 146–149.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., & Huang, F. (2006). A tutorial on energy-based learning. *Predicting Structured Data*, 1.
- Lee, M.-J., & Lee, S. (2022). Review and comparison of treatment effect estimators using propensity and prognostic scores. *The International Journal of Biostatistics*, 18(2), 357–380.
- Li, H., Xiao, Y., Zheng, C., Wu, P., & Cui, P. (2023). Propensity matters: Measuring and enhancing balancing for recommendation. In *International conference on machine learning* (pp. 20182–20194). PMLR.
- Li, H., Zheng, C., Cao, Y., Geng, Z., Liu, Y., & Wu, P. (2023). Trustworthy policy learning under the counterfactual no-harm criterion. In *International conference on machine learning* (pp. 20575–20598). PMLR.
- Li, H., Zheng, C., Wu, P., Kuang, K., Liu, Y., & Cui, P. (2023). Who should be given incentives? counterfactual optimal treatment regimes learning for recommendation. In *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining* (pp. 1235–1247).
- Mealli, F., Pacini, B., & Rubin, D. B. (2011). Statistical inference for causal effects. *Modern Analysis of Customer Surveys: With Applications using R*, 171–192.
- Nie, L., Ye, M., Liu, Q., & Nicolae, D. (2021). Vcnet and functional targeted regularization for learning causal effects of continuous treatments. In *International conference on learning representations*.
- Pan, Y. (2023). Structure analysis of crowd intelligence systems. *Engineering*, 25, 17–20.
- Pearl, J. (2009). *Causality*. Cambridge University Press.
- Raita, Y., Camargo, C. A., Jr., Liang, L., & Hasegawa, K. (2021). Leveraging “big data” in respiratory medicine—data science, causal inference, and precision medicine. *Expert Review of Respiratory Medicine*, 15(6), 717–721.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688.
- Runge, J. (2018). Conditional independence testing based on a nearest-neighbor estimator of conditional mutual information. In *International conference on artificial intelligence and statistics* (pp. 938–947). PMLR.
- Schwab, P., Linhardt, L., Bauer, S., Buhmann, J. M., & Karlen, W. (2020). Learning counterfactual representations for estimating individual dose-response curves. In *Proceedings of the AAAI conference on artificial intelligence: vol. 34*, (04), (pp. 5612–5619).
- Shalit, U., Johansson, F. D., & Sontag, D. (2017). Estimating individual treatment effect: generalization bounds and algorithms. In *International conference on machine learning* (pp. 3076–3085). PMLR.
- Shi, C., Blei, D., & Veitch, V. (2019). Adapting neural networks for the estimation of treatment effects. *Advances in Neural Information Processing Systems*, 32.
- Shi, X., Miao, W., & Tchetgen, E. T. (2020). A selective review of negative control methods in epidemiology. *Current Epidemiology Reports*, 7(4), 190–202.
- Stuart, E. A., Lee, B. K., & Leacy, F. P. (2013). Prognostic score-based balance measures can be a useful diagnostic for propensity score methods in comparative effectiveness research. *Journal of Clinical Epidemiology*, 66(8), S84–S90.
- Székely, G. J., & Rizzo, M. L. (2014). Partial distance correlation with methods for dissimilarities. *The Annals of Statistics*, 42(6), 2382–2412.
- Vegetabile, B. G., Griffin, B. A., Coffman, D. L., Cefalu, M., Robbins, M. W., & McCaffrey, D. F. (2021). Nonparametric estimation of population average dose-response curves using entropy balancing weights for continuous exposures. *Health Services and Outcomes Research Methodology*, 21(1), 69–110.
- Wager, S., & Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523), 1228–1242.
- Wang, H., Chen, Z., Fan, J., Li, H., Liu, T., Liu, W., et al. (2023). Optimal transport for treatment effect estimation. *Advances in Neural Information Processing Systems*.
- Wang, T., & Isola, P. (2020). Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning* (pp. 9929–9939). PMLR.
- Wu, P. A., & Fukumizu, K. (2022). β -Intact-VAE: Identifying and estimating causal effects under limited overlap. In *International conference on learning representations*. URL: <https://api.semanticscholar.org/CorpusID:265099243>.
- Wu, A., Kuang, K., Li, B., & Wu, F. (2022). Instrumental variable regression with confounder balancing. In *International conference on machine learning* (pp. 24056–24075). PMLR.
- Wu, A., Kuang, K., Xiong, R., Li, B., & Wu, F. (2023). Stable estimation of heterogeneous treatment effects. In *International conference on machine learning* (pp. 37496–37510). PMLR.
- Wu, Z., Xiong, Y., Yu, S. X., & Lin, D. (2018). Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 3733–3742).
- Wu, A., Yuan, J., Kuang, K., Li, B., Wu, R., Zhu, Q., et al. (2022). Learning decomposed representations for treatment effect estimation. *IEEE Transactions on Knowledge and Data Engineering*, 35(5), 4989–5001.
- Wyner, A. D. (1978). A definition of conditional mutual information for arbitrary ensembles. *Information and Control*, 38(1), 51–59.
- Xu, J., Hong, N., Xu, Z., Zhao, Z., Wu, C., Kuang, K., et al. (2023). Data-driven learning for data rights, data pricing, and privacy computing. *Engineering*, 25, 66–76.
- Xu, L., Kanagawa, H., & Gretton, A. (2021). Deep proxy causal learning and its application to confounded bandit policy evaluation. ArXiv arXiv:2106.03907, URL: <https://api.semanticscholar.org/CorpusID:235367652>.
- Yao, D., Zhao, Z., Zhang, S., Zhu, J., Zhu, Y., Zhang, R., et al. (2022). Contrastive learning with positive-negative frame mask for music representation. In *WWW '22: the ACM web conference 2022, virtual event, lyon, France, April 25 - 29, 2022* (pp. 2906–2915). ACM.
- Yuan, L., & Zhu, S.-C. (2023). Communicative learning: A unified learning formalism. *Engineering*, 25, 77–100.
- Zhang, M., Huang, S., Li, W., & Wang, D. (2022). Tree structure-aware few-shot image classification via hierarchical aggregation. In *European conference on computer vision* (pp. 453–470). Springer.
- Zhang, W., Li, J., & Liu, L. (2020). A unified survey of treatment effect heterogeneity modelling and uplift modelling. *ACM Computing Surveys*, 54, 1–36, URL: <https://api.semanticscholar.org/CorpusID:220793171>.
- Zhang, K., Peters, J., Janzing, D., & Schölkopf, B. (2012). Kernel-based conditional independence test and application in causal discovery. arXiv preprint arXiv:1202.3775.
- Zhang, S., Yao, D., Zhao, Z., Chua, T., & Wu, F. (2021). CauseRec: Counterfactual user sequence synthesis for sequential recommendation. In *SIGIR '21: the 44th international ACM SIGIR conference on research and development in information retrieval, virtual event, Canada, July 11–15, 2021* (pp. 367–377). ACM.
- Zhao, Z., Bai, Y., Xiong, R., Cao, Q., Ma, C., Jiang, N., et al. (2022). Learning individual treatment effects under heterogeneous interference in networks. *ACM Transactions on Knowledge Discovery from Data*, URL: <https://api.semanticscholar.org/CorpusID:253107774>.
- Zhen, X., Meng, Z., Chakraborty, R., & Singh, V. (2022). On the versatile uses of partial distance correlation in deep learning. In *Computer vision—ECCV 2022: 17th European conference, tel aviv, Israel, October 23–27, 2022, proceedings, part XXVI* (pp. 327–346). Springer.