
Reducing Selection Bias in Counterfactual Reasoning for Individual Treatment Effects Estimation

Zichen Zhang[†], Qingfeng Lan[†], Lei Ding[‡], Yue Wang[‡], Negar Hassanpour[†], Russ Greiner[†]

[†]Department of Computing Science

[‡]Department of Mathematical and Statistical Sciences

University of Alberta

{zichen2, qlan3, lding1, yue9, hassanpo, rgreiner}@ualberta.ca

Abstract

Counterfactual reasoning is an important paradigm applicable in many fields, such as healthcare, economics, and education. In this work, we propose a novel method to address the issue of sample *selection bias*. We learn two groups of latent random variables, where one group corresponds to variables that only cause selection bias, and the other group is relevant for outcome prediction. They are learned by an auto-encoder where an additional regularized loss based on Pearson Correlation Coefficient (PCC) encourages the de-correlation between the two groups of random variables. This allows for explicitly alleviating selection bias by only keeping the latent variables that are relevant for estimating individual treatment effects. Experimental results on a synthetic toy dataset and a benchmark dataset show that our algorithm is able to achieve state-of-the-art performance and improve the result of its counterpart that does not explicitly model the selection bias.

1 Introduction

Studying the causal effect of different treatments on individuals to assist in decision making is an essential problem in various fields. Examples include a doctor deciding the most effective medical treatment for a specific patient, a company deciding the most profitable commercial advertisement for a specific product, etc.

In this paper, we focus on understanding individual-level causal effects in healthcare. Access of many observational data in this field allows us to develop methods for predicting individual-level causal effects. However, many challenges remain. The first challenge is that the observational data at the individual level tells us only the outcome of received treatments (the *factuals*), whereas the responses of the alternative treatments (the *counterfactuals*) are never available. For example, if a patient is given surgery, we would not be able to observe the true effect of applying medication instead. This setting is called *counterfactual reasoning*, that is, to predict the individual treatment effect of the counterfactual treatment. The second challenge is that the data often exhibits *selection bias* [Imbens and Rubin, 2015]. For example, patient living in the rural area may not have access to a certain medication. Consequently, there are only a few, if any, patients receiving that medication in the dataset, i.e., the observational data have a *selection bias*. In this case, home address affects only the treatment. In other cases, there are factors that affect both the treatment and the respective outcome, called *confounder*. It also causes the difficulty of predicting the causal effects since it partially leads to the *selection bias*. For example, it is more likely for a doctor to prescribe surgery to younger patients while to give medication to older patients. On the other hand, ages may affect the potential outcome regardless of the treatment given.

We represent the observed features of each patient, such as age and gender, as a random vector \mathbf{X} . For simplicity, we assume that there are only two treatments, denoted as a binary variable $T \in \{0, 1\}$.




Figure 1: The proposed probabilistic graphical model for individual treatment effect estimation

The patients who receive treatment $T = 0$ or $T = 1$ are in the control group and treatment group, respectively. Consequently, there are two possible outcomes Y^0 and Y^1 corresponding to each treatment option: $T = 0$ and $T = 1$. However, we do not have access to both of the outcomes. For each patient, we only observe the outcome corresponding to the received treatment. We denote all the observed outcomes (the factials) as Y^f and all the unobserved outcomes (the counterfactuals) as Y^{cf} . Moreover, the selection bias can be expressed as $p(T|\mathbf{X}) \neq p(T)$. The goal is to estimate the Individual Treatment Effect (ITE), i.e., $Y^1 - Y^0$ for each individual.

The graphical model in our analysis is shown in Figure 1, inspired by Hassanpour and Greiner [2019]. We assume that the covariate \mathbf{X} is generated by three types of latent variables. The first type **A** includes the latent variables that only affect treatment selection procedure but do not determine outcomes. Type **B** are the confounders which influence both treatments and outcomes. The last type **C** only affects outcomes.

Our main contribution is that we propose a novel method that separates the learned feature representations into two parts, corresponding to **A** and **BC** described above. Then we reduce the selection bias by using only the representation of type **BC** to predict the outcomes. We test our algorithm on two datasets: a synthetic toy dataset and a benchmark dataset simulated from real-world data. The results show that our method helps to improve the prediction performance in many settings.

Assumptions Similar to the work by Shalit et al. [2017], we assume that there exists a joint distribution $p(\mathbf{X}, T, Y^0, Y^1)$ with “strong ignorability” assumption: $Y \perp T | \mathbf{X}$ and $0 < p(t = 1|x) < 1, \forall x \in \mathcal{X}$. This is sufficient for the ITE to be identifiable [Imbens and Wooldridge, 2009]. We also assume that sample points $(\mathbf{x}_1, t_1, y_1), \dots, (\mathbf{x}_n, t_n, y_n)$ are generated from $y^i \sim p(y^{t_i} | \mathbf{x}_i)$.

2 Related works

In the work of Johansson et al. [2016], ideas from domain adaptation and representation learning were combined. Besides learning a latent representation for outcome prediction, the discrepancy distance [Mansour et al., 2009] was introduced as a distribution distance metric. By minimizing the discrepancy distance, the distributions of populations with different treatments were balanced to reduce the selection bias. On the network structure, they simply concatenated the treatment t with the representation Φ which easily led to information loss of treatment t .

Shalit et al. [2017] improved this line of work and proved a generalization error-bound for estimating ITE. The improvements they made were two-fold. First, they introduced a branching network structure where one branch of the network learned the prediction of treated outcome ($t = 1$) and the other branch learned the prediction of the outcome under control ($t = 0$). This new network structure solved the issue in the work of Johansson et al. [2016] of losing the influence of t when the dimension of the representation was large. Second, they introduced a measure of distance between two distributions $p(\mathbf{x}|t = 1)$ and $p(\mathbf{x}|t = 0)$, called Integral Probability Metric (IPM). They showed that the expected error of the ITE prediction was upper bounded by the error of learning Y^1 and Y^0 , plus the IPM term. This IPM measure was therefore used in the loss function to encourage that the two distributions of representations being closer.

From the perspective of network structure, our work is closely related to Atan et al. [2018]. In this work, a latent representation is learned by using an auto-encoder. By jointly minimizing the




Figure 2: The architecture of the proposed model RSB-Net. The representation learning network is an auto-encoder that learns the bias variables \mathbf{A} and the variables \mathbf{BC} that are relevant for prediction. The outcome prediction network has a branching structure predicting the outcome y^t based on the treatment t and the representation of \mathbf{BC} .

reconstruction loss and the distribution distance between different representation groups, it balanced between information loss and bias reduction. However, the entire learned representation was then used for outcome prediction. This could inevitably contain features that is not useful for outcome prediction therefore counteracting the effect of bias reduction in the first step. We address this issue by learning the features that only causes selection bias and discard them during the outcome prediction.

In terms of learning two independent categories of representations, we are inspired by the work from Cheung et al. [2014]. This work was in the domain of image classification, where they learned the features of class-independent variations Z apart from the features for classification. They introduced a cross-covariance penalty (XCov) for this purpose. It disentangled factors like the hand-writing style from the digits labels. We improve their work by introducing a penalty term that better reflects the correlation between random variables, as detailed in the next section.

3 Proposed method

3.1 Network Architecture

Following the model of deep-treat in [Atan et al., 2018] and the Counter Factual Regression (CFR) framework in [Shalit et al., 2017], the overall structure of our method is an autoencoder where the representation learning stage is also followed by an outcome prediction stage that branches based on the treatment t of the input sample \mathbf{x} .

In order to reduce the selection bias, we explicitly model the bias (latent variables type \mathbf{A} , note that it is in **bold** denoting a vector of random variables) in the learned representation $\Phi(\mathbf{x})$ and separate it from the rest of the features (latent variables type \mathbf{BC}) that are relevant for the outcome prediction. Since the bias variables \mathbf{A} do not play a part in the outcome prediction, only the variables \mathbf{BC} are then used as the input to the downstream prediction network to predict the outcomes for various treatments $\hat{y}^1(\mathbf{x})$ and $\hat{y}^0(\mathbf{x})$ (we consider binary treatment in this work, i.e, $t = 1$ or $t = 0$).

The architecture of the proposed method, named RSB-Net (stands for Reducing Selection Bias), is illustrated in Fig. 2. In the next subsection, we explain how this network can be trained to explicitly learn the two groups of latent variables \mathbf{A} and \mathbf{BC} .

3.2 Loss Function

On the high level, the proposed neural network is trained end-to-end with a hybrid loss that is a weighted sum of the following: a prediction loss L_{pred} , a distribution loss L_{IPM} , an input reconstruction loss L_{recons} and a loss based on Pearson correlation coefficient L_{pcc} .

$$\mathcal{L} = L_{\text{pred}} + \alpha L_{\text{IPM}} + \beta L_{\text{recon}} + \gamma L_{\text{pcc}} + \lambda R(W) \quad (1)$$

where $R(W)$ is the regularization term of the network weights W that penalizes complexity of the network, and $\alpha, \beta, \gamma, \lambda$ are the weights of the loss terms.

Prediction Loss and Distribution Loss This is the supervised loss proposed in Shalit et al. [2017]. For a batch of data samples \mathbf{x}_i, t_i , we aim to predict the factual outcome $y_i^{t_i}$. The prediction loss is defined on the factual outcome $y_i^{t_i}$ using a weighted squared loss as

$$L_{pred} = \frac{1}{N} w_i \|h^{t_i}(\Phi_{BC}(\mathbf{x}_i)) - y_i^{t_i}\|_2^2 \quad (2)$$

where N is the sample size, $w_i = \frac{t_i}{2u} + \frac{1-t_i}{2(1-u)}$, $u = \frac{1}{N} \sum_i t_i$. Note that u is the probability of choosing treatment $t = 1$ in the entire population, i.e. $u = p(t = 1)$. w_i compensates for the size difference in different treatment arms. A more detailed explanation can be found in Rubin [1974]. The distribution loss, using Integral Probability Metric (IPM) is defined as

$$L_{IPM} = \text{IPM}(\Phi_{BC}(\mathbf{x}_i)_{i:t_i=0}, \Phi_{BC}(\mathbf{x}_i)_{i:t_i=1}) \quad (3)$$

It measures the distribution distance of the latent representation of the treated group and the control group, i.e. between $\Phi(\mathbf{x}|t = 1)$ and $\Phi(\mathbf{x}|t = 0)$.

Reconstruction Loss Inspired by the approach in Cheung et al. [2014] and Atan et al. [2018], we use an auto-encoder with a squared L_2 loss to learn a good set of latent representation for both groups of random variables \mathbf{A} and \mathbf{BC} . The loss is defined as $L_{recons} = \|\hat{\mathbf{x}} - \mathbf{x}\|_2^2$, where $\hat{\mathbf{x}}$ is the reconstruction of the input features of the sample \mathbf{x} .

PCC Loss To explicitly learn the random variables \mathbf{A} and \mathbf{BC} , the latent representation $\Phi(\mathbf{x})$ is first split into two parts (the ratio is a hyperparameter), denoted as $\Phi_{\mathbf{A}}(\mathbf{x}) \in \mathbb{R}^m$ and $\Phi_{\mathbf{BC}}(\mathbf{x}) \in \mathbb{R}^n$, corresponding to \mathbf{A} and \mathbf{BC} respectively. m and n denote the dimension of the vectors for each sample \mathbf{x} . We would like $\Phi_{\mathbf{A}}(\mathbf{x})$ and $\Phi_{\mathbf{BC}}(\mathbf{x})$ to be de-correlated in the learned representation.

To this end, we define a loss based on Pearson correlation coefficient (PCC):

$$L_{pcc} = \frac{1}{2mn} \sum_{i=1}^m \sum_{j=1}^n \left[\frac{\frac{1}{N} \sum_{k=1}^N (\Phi_{\mathbf{A}}(\mathbf{x}_k)_i - \overline{\Phi_{\mathbf{A}_i}})(\Phi_{\mathbf{BC}}(\mathbf{x}_k)_j - \overline{\Phi_{\mathbf{BC}_j}})}{\sigma(\Phi_{\mathbf{A}_i})\sigma(\Phi_{\mathbf{BC}_j})} \right]^2 \quad (4)$$

where $\Phi_{\mathbf{A}}(\mathbf{x}_k)_i$ is the i -th element of vector $\Phi_{\mathbf{A}}(\mathbf{x}_k)$ for sample k , $\overline{\Phi_{\mathbf{A}_i}}$ is the mean value of the i -th element of vector $\Phi_{\mathbf{A}}$ for all samples. We use similar notations for $\Phi_{\mathbf{BC}}$. The idea of this loss is to take the mean of the squared PCC between every pair of random variables formed by one entry in vector $\Phi_{\mathbf{A}}$ and one entry $\Phi_{\mathbf{BC}}$. Since $\text{PCC} \in [-1, 1]$, we have the range of this loss $L_{pcc} \in [0, 0.5]$, reaching the minimum when features $\Phi_{\mathbf{A}}$ and $\Phi_{\mathbf{BC}}$ are linearly independent.

Algorithm 1: RSB-Net

Input: Factual samples $\{(\mathbf{x}_1, t_1, y_1^f), \dots, (\mathbf{x}_N, t_N, y_N^f)\}$, coefficients of the loss terms: $\alpha, \beta, \gamma, \lambda$

Compute $u = \frac{1}{N} \sum_{i=1}^N t_i$ and the sample weight: $w_i = \frac{t_i}{2u} + \frac{1-t_i}{2(1-u)}$ for $i = 1, \dots, N$;

Initialize the weights \mathbf{W} in neural networks;

repeat

 Sample a mini-batch $\{i_1, i_2, \dots, i_m\} \subseteq \{1, 2, \dots, N\}$;

foreach sample (\mathbf{x}, y, t) **do**

 /* Representation Learning */

 Compute $\Phi(\mathbf{x})$ and the reconstructed \mathbf{x} : $\hat{\mathbf{x}} = \Psi(\Phi(\mathbf{x}))$;

 /* Outcome Prediction */

 Split the representation $\Phi(\mathbf{x})$ into bias $\Phi_{\mathbf{A}}(\mathbf{x})$ and representation $\Phi_{\mathbf{BC}}(\mathbf{x})$;

 Compute the predicted outcome corresponding to treatment t , $y^t = h^t(\Phi_{\mathbf{BC}}(\mathbf{x}))$;

end

 Compute the prediction loss $L_{pred} = \frac{1}{m} \sum_{j=1}^m w_{i_j} (\ell_2(\hat{y}_{i_j}^{t_{i_j}}, y_{i_j}^{t_{i_j}}))^2$;

 Compute the distribution loss $L_{IPM} = \ell_{IPM}(\Phi_{\mathbf{BC}}(\mathbf{x}_i)_{i:t_i=0}, \Phi_{\mathbf{BC}}(\mathbf{x}_i)_{i:t_i=1})$;

 Compute the reconstruction loss $L_{recon} = \frac{1}{m} \sum_{j=1}^m (\ell_2(\hat{\mathbf{x}}_{i_j}, \mathbf{x}_{i_j}))^2$;

 Compute the Pearson correlation coefficient loss L_{pcc} defined in Eq. 4;

 Sum up the above loss functions and add regularization R to get the total loss \mathcal{L}

$$\mathcal{L} = L_{pred} + \alpha L_{IPM} + \beta L_{recon} + \gamma L_{pcc} + \lambda R(\mathbf{W})$$

 Optimize all weights \mathbf{W} in the neural networks;

until max iterations;

Output: Neural network weights \mathbf{W}

4 Experiments

In this section, we present the experimental results of our proposed method on a new synthetic toy dataset and a benchmark dataset. We compare our method with the following baseline methods: k-nearest neighbor (kNN), Bayesian Additive Regressoin Trees (BART) [Chipman et al., 2010], Balancing Neural Network (BNN) [Johansson et al., 2016], Deep-Treat [Atan et al., 2018], Treatment-Agnostic Representation Network (TARNET) [Shalit et al., 2017], Counterfactual Regression with Wasserstein metric (CFRW) [Shalit et al., 2017], Counterfactual Regression with Importance Sampling Weights (CFR-ISW) [Hassanpour and Greiner, 2019], , Causal Effect Variational Autoencoder (CEVAE) [Louizos et al., 2017].

4.1 Evaluation Metrics

Our goal is to estimate the Individual Treatment Effects (ITE), which measures the difference between possible outcomes for each patient. The ground truth ITE and estimated ITE are defined as follows:

$$\tau(x) := \mathbb{E}[Y^1 - Y^0|x] \quad \text{and} \quad \hat{\tau}(x) = h^1(\Phi(x)) - h^0(\Phi(x))$$

where $\Phi(\mathbf{x})$ is the representation function of the form $\Phi : \mathcal{X} \rightarrow \mathcal{R}$ that transforms \mathbf{x} from the sample space \mathcal{X} into the learned representation space \mathcal{R} , and h is the hypothesis function $\mathcal{R} \times \{0, 1\} \rightarrow \mathcal{Y}$, defined over the representation space \mathcal{R} and the treatment $t \in \{0, 1\}$, mapping to the output space \mathcal{Y} . Following the work by Shalit et al. [2017], we use two metrics to evaluate estimated ITE. The first one is Precision in Estimation of Heterogeneous Effect (PEHE) defined as:

$$\epsilon_{\text{PEHE}} = \frac{1}{N} \sum_i^N (\hat{\tau}(\mathbf{x}_i) - \tau(\mathbf{x}_i))^2$$

where N is the sample size. This measures the mean squared difference between the estimated ITE and true ITE. Note that ϵ_{PEHE} is originally defined on a continuous distribution [Hill, 2011]. Here we use the discrete version for finite samples.

Another metric is the bias of the Average Treatment Effect (ATE):

$$\epsilon_{\text{ATE}} = |\widehat{\text{ATE}} - \text{ATE}| = \left| \frac{1}{N} \sum_{i=1}^N (\hat{\tau}(\mathbf{x}_i) - \tau(\mathbf{x}_i)) \right|$$

where $\text{ATE} = \mathbb{E}[\tau(\mathbf{x}_i)] = \frac{1}{N} \sum_{i=1}^N \tau(\mathbf{x}_i)$. This measures the population difference between the expectation of the estimated ITE and true ITE.

For all experiments, we report the within-sample and out-of-sample mean and standard errors of $\sqrt{\epsilon_{\text{PEHE}}}$ and ϵ_{ATE} . Within-sample takes into account the entire training data, including the training and validation split. Out-of-sample result measures the performance on the hold-out test dataset.

4.2 Implementation details

We implemented our RSB net using TensorFlow, based on the code¹ provided by Johansson et al. [2016].

Preprocessing of the data have not been commonly used or mentioned in the literature. In the hyper-parameter tuning, we tested different preprocessing methods such as Z-score Standardization and min-max Normalization. Our empirical result across multiple datasets suggests that min-max Normalization either improves or shows no impact on the result, compared to the alternatives like Z-score and using raw data. Hence we suggest to always apply this step.

During training, the weights of the neural network were initialized randomly and optimized using Adam [Kingma and Ba, 2014]. The maximum iteration was 5k for all experiments. Early stopping wa performed based on the validation loss.

For multiple realizations, a random split is performed once per realization, to prevent over-fitting. ‘‘Realization’’ refers to the randomized experiments for each input features \mathbf{X} and treatment T .

¹<https://github.com/clinicalml/cfrnet>

Algorithm 2: Synthetic Data Generation

Input: Dimension for each latent features M_1, M_2, M_3 ; Sample size N ; Covariate Dimension D

```
foreach sample  $(\mathbf{x}, t, y^{CF}, y^F, \mu^1, \mu^0)$  do
  foreach dimension  $h$  of  $x$ ,  $h \in \{1:D\}$ , do
     $\{A_i\}_{i=1}^{M_1} \sim \mathcal{N}(\mu_A, 2)$ , where  $\mu_A \sim \mathcal{N}(1, 1)$ ;
     $\{B_j\}_{j=1}^{M_2} \sim \mathcal{N}(\mu_B, 1)$ , where  $\mu_B \sim \mathcal{N}(2, 1)$ ;
     $\{C_k\}_{k=1}^{M_3} \sim \mathcal{N}(\mu_C, 1)$ , where  $\mu_C \sim \mathcal{N}(3, 1)$ ;
     $\mathbf{x}_h \sim \mathcal{N}(\mu, \sigma^2)$ , where  $\mu = \frac{1}{3}\{\overline{A_i}\} + \frac{1}{3}\{\overline{B_j}\} + \frac{1}{3}\{\overline{C_k}\}$  and
       $\sigma^2 = \frac{1}{3}Var(\{A_i\}) + \frac{1}{3}Var(\{B_j\}) + \frac{1}{3}Var(\{C_k\})$ ;
    end
    foreach each realization under  $\mathbf{x}$  do
       $t(\{A_i\}, \{B_j\}) \sim \text{Bernoulli}(p(t=1))$ , where
         $p(t=1) = 1 - \text{Sigmoid}(0.7 * \{A_i\} + 0.3 * \{B_j\})$ ;
      end
      foreach each realization under  $\mathbf{x}$  do
        if  $t = 0$  then
           $a = 0.3 * \{\overline{B_j}\} + 0.7 * \{\overline{C_k}\}$ ;  $b = 2 * Var(\{B_j\}) + 8 * Var(\{C_k\})$ ;
           $\mu^1 \sim \Gamma(\alpha, \beta)$ , where  $\alpha = \frac{a^2}{b}$ ,  $\beta = \frac{b}{a}$ ;  $\mu^0 \sim \mu^1 - \mathcal{N}(2, 1)$ ;
           $y^{CF}(\{B_j\}, \{C_k\}) = \mu^1 + \mathcal{N}(0, 1)$ ;  $y^F(\{B_j\}, \{C_k\}) = \mu^0 + \mathcal{N}(0, 1)$ ;
        end
        else if  $t = 1$  then
           $a = 0.1 * \{\overline{B_j}\} + 0.9 * \{\overline{C_k}\}$ ;  $b = 3 * Var(\{B_j\}) + 6 * Var(\{C_k\})$ ;
           $\mu^1 \sim \Gamma(\alpha, \beta)$ , where  $\alpha = \frac{a^2}{b}$ ,  $\beta = \frac{b}{a}$ ;  $\mu^0 \sim \mu^1 - \mathcal{N}(2, 1)$ ;
           $y^F(\{B_j\}, \{C_k\}) = \mu^1 + \mathcal{N}(0, 1)$ ;  $y^{CF}(\{B_j\}, \{C_k\}) = \mu^0 + \mathcal{N}(0, 1)$ ;
        end
      end
    end
  end
end
Output: Sample set  $\{(\mathbf{x}_i, t_i, y_i^{CF}, y_i^F, \mu_i^1, \mu_i^0)\}_{i=1}^N$ 
```

Table 1: Performance comparison on the synthetic dataset over 100 realizations. The metrics are mean and standard errors of $\sqrt{\epsilon_{PEHE}}$ and ϵ_{ATE} . Better result is highlighted in blue.

Methods	Within-sample		Out-of-sample	
	$\sqrt{\epsilon_{PEHE}}$	ϵ_{ATE}	$\sqrt{\epsilon_{PEHE}}$	ϵ_{ATE}
CFRW	6.335 \pm 0.032	0.392 \pm 0.047	6.304 \pm 0.088	0.596 \pm 0.074
RSB(Ours)	6.331 \pm 0.032	0.370 \pm 0.042	6.297 \pm 0.088	0.574 \pm 0.070

Unless mentioned otherwise, we run hyperparameter selection based on the nearest neighbor version of PEHE defined in [Shalit et al., 2017], on the validation set. This metric is used since we do not have access to true PEHE in real-world settings. $\epsilon_{PEHE_{nn}} = \frac{1}{N} \sum_{i=1}^N ((1 - 2t_i)(y_{j(i)} - y_i) - (\hat{y}^1 - \hat{y}^0))^2$ where $j(i)$ is the index of the nearest neighbor to sample i in the opposite treatment group.

4.3 Experiment on Synthetic Dataset

We experiment on a synthetic toy dataset to evaluate how well our model handles hidden confounders, by which we mean that the covariates \mathbf{X} serves as a proxy to all confounders. In the generation of the synthetic dataset, for each sample, X_i is a mixture of Gaussian distributions with hidden variables A_i, B_i , and C_i determining the mixture component. Note that i is the sample index, each of A_i, B_i , and C_i are vectors instead of scalars. For example, A_1 can be a 10x1 vector, representing 10 random variables. The length of the vectors A_i, B_i , and C_i per sample can be arbitrary. We used 10,15,10 respectively. These hidden variables are generated using the hierarchical Bayes method.

Since treatment T is binary and only affected by hidden variables \mathbf{A} and \mathbf{B} in our graphical model, a Bernoulli distribution is used to generate T and the probability p is calculated as a *sigmoid* function applied to a weighted sum of \mathbf{A} and \mathbf{B} to map the value to $[0, 1]$. Outcomes Y^1 and Y^0 are generated by Gamma distributions. The shape and location parameters of Gamma distribution α and β are generated by a weighted average of \mathbf{B} and \mathbf{C} . A more detailed explanation can be found

Table 2: Performance comparison on the IHDP dataset over 100 realizations. The metrics are mean and standard errors of $\sqrt{\epsilon_{\text{PEHE}}}$ and ϵ_{ATE} . Best result is highlighted in blue. Entry '-': not reported in the paper.

Methods	Within-sample		Out-of-sample	
	$\sqrt{\epsilon_{\text{PEHE}}}$	ϵ_{ATE}	$\sqrt{\epsilon_{\text{PEHE}}}$	ϵ_{ATE}
BNN	-	-	2.20 ± 0.130	-
Deep-Treat	-	-	1.93 ± 0.070	-
CFRW	-	-	0.88 ± 0.010	0.20 ± 0.003
CFR-ISW	-	-	0.77 ± 0.010	0.19 ± 0.003
RSB(Ours)	0.63 ± 0.025	0.25 ± 0.033	0.67 ± 0.043	0.26 ± 0.035

Table 3: Performance comparison on the IHDP dataset over 1000 realizations. The metrics are mean and standard errors of $\sqrt{\epsilon_{\text{PEHE}}}$ and ϵ_{ATE} . Best result is highlighted in blue.

Methods	Within-sample		Out-of-sample	
	$\sqrt{\epsilon_{\text{PEHE}}}$	ϵ_{ATE}	$\sqrt{\epsilon_{\text{PEHE}}}$	ϵ_{ATE}
k-NN	2.1 ± 0.1	0.14 ± 0.01	4.1 ± 0.2	0.79 ± 0.05
BART	2.1 ± 0.1	0.23 ± 0.01	2.3 ± 0.1	0.34 ± 0.02
BNN	2.2 ± 0.1	0.37 ± 0.03	2.1 ± 0.1	0.42 ± 0.03
TARNET	0.88 ± 0.0	0.26 ± 0.01	0.95 ± 0.0	0.28 ± 0.01
CFRW	0.71 ± 0.0	0.25 ± 0.01	0.76 ± 0.0	0.27 ± 0.01
CEVAE	2.7 ± 0.1	0.25 ± 0.01	2.6 ± 0.1	0.46 ± 0.02
RSB(Ours)	0.66 ± 0.0	0.26 ± 0.01	0.68 ± 0.0	0.27 ± 0.01

in Algorithm 2. The total number of samples is 1000, each with 25 covariates. We use a 63/27/10 train/valid/test split.

We compare our method RSB with its counterpart CFRW [Shalit et al., 2017], to evaluate how well reducing the selection bias helps in the presence of hidden confounders. CFRW is considered as the counterpart since our model RSB follows the same network architecture and can be viewed roughly as CFRW with two additional loss terms: reconstruction loss L_{recon} and Pearson correlation coefficient based loss L_{pcc} . The comparison is shown in Table 1. The result shows that in a dataset that is generated with selection bias, reducing the selection bias using our method helps to improve the counterfactual prediction a bit, although not significant. Note that we ran extensive hyperparameter tuning on both methods. Although we know the dimensions of the hidden variables **A**, **B**, **C** when we generate the data, we did not use that information to select the best hyper-parameter.

4.4 Benchmark dataset - IHDP

To further evaluate our method, we benchmark our method in the real-world setting, using a semi-simulated dataset based on the Infant Health and Development Program (IHDP), introduced by Hill [2011]. The data have features from a real randomized experiment, studying the effect of high-quality childcare and home visits on future cognitive test scores. The IHDP dataset uses a simulated outcome and it also artificially introduces sample selection bias by removing a biased subset of the treated population. The dataset has 747 samples in total (139 treated and 608 control). For each sample, there are multiple realizations of the outcomes corresponding to either of the available treatments². We use the same 63/27/10 train/valid/test split as in the literature. For evaluation, following the setup in [Shalit et al., 2017], we use the noiseless outcomes μ^1 and μ^0 as the ground truth so that $\text{ITE}(\mathbf{x}) = \mu^1(\mathbf{x}) - \mu^0(\mathbf{x})$.

For the comparison with baselines, we test our method under both 100 and 1000 realizations, using the dataset IHDP-100 and IHDP-1000 provided by Johansson et al. [2016]. The results are shown in Table 2 and 3. Under 100 realizations, we compare with four other neural network based methods: BNN and Deep-Treat whose results are replicated from [Atan et al., 2018], CFRW and CFR-ISW for which the results from [Hassanpour and Greiner, 2019] are replicated (results with the hyperparameter selected based on $\epsilon_{\text{PEHE}_{\text{nn}}}$ for a fair comparison). Under 1000 realizations, we compare with all baseline methods described in the beginning of Sec.4, except for CFR-ISW [Hassanpour and Greiner, 2019] and Deep-Treat [Atan et al., 2018] which only reported results under 100 realizations.

Our method achieves state-of-the-art performance in most metrics. Since this dataset is simulated from real-world observational data, it is not clear what categories of hidden features are present. The

²the counterfactual outcomes are only used for evaluation purposes

experimental results show that our method is able to perform well on real-world datasets where the underlying structure of the hidden variables is unknown.

5 Conclusion and future work

In this paper, we proposed a novel and intuitive method to reduce selection bias in the problem of estimating the individual treatment effect. We modeled the input features as generated by three types of latent variables **A**, **B**, **C** (in Figure 1). The variables of type **A** only cause selection bias while not contributing to the outcome Y . Discarding it would help to alleviate selection bias. In order to learn the representation of **A**, an auto-encoder is used to learn the representations of features. We view the learned representations as two random vectors corresponding to categories **A** and **BC**. We then apply a loss based on Pearson correlation coefficient between any pair of random variables between these two vectors to encourage the linear independence of **A** and **BC**. This allows us to explicitly discard the category (**A**) that partially induces selection bias and only use the relevant features (**BC**) for the outcome prediction. We tested our approach on both synthetic and simulated real-world tasks, showing that our method achieved state-of-the-art results. Future work include but are not limited to testing on more experimental results, further separating the latent variables **B** and **C**, and modeling the relationship between treatments and outcomes.

References

- Onur Atan, J Jordan, and Mihaela van der Schaar. Deep-treat: Learning optimal personalized treatments from observational data using neural networks. In *AAAI*. AAAI, 2018.
- Brian Cheung, Jesse A. Livezey, Arjun K. Bansal, and Bruno A. Olshausen. Discovering Hidden Factors of Variation in Deep Networks. *arXiv:1412.6583 [cs]*, December 2014. URL <http://arxiv.org/abs/1412.6583>. 00075 arXiv: 1412.6583.
- Hugh A Chipman, Edward I George, Robert E McCulloch, et al. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.
- Negar Hassanpour and Russell Greiner. Counterfactual regression with importance sampling weights. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5880–5887, 7 2019.
- Jennifer L. Hill. Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, January 2011. ISSN 1061-8600, 1537-2715.
- Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- Guido W Imbens and Jeffrey M Wooldridge. Recent developments in the econometrics of program evaluation. *Journal of economic literature*, 47(1):5–86, 2009.
- Fredrik D Johansson, Uri Shalit, and David Sontag. Learning Representations for Counterfactual Inference. In *International Conference on Machine Learning*, pages 3020–3029, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *Advances in Neural Information Processing Systems*, pages 6446–6456, 2017.
- Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009.
- Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974. ISSN 0022-0663.
- Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning (ICML)*, pages 3076–3085, 2017.