

Adversarially Balanced Representation for Continuous Treatment Effect Estimation

Amirreza Kazemi, Martin Ester

School of Computing Science, Simon Fraser University
{aka208, ester}@sfu.ca

Abstract

Individual treatment effect (ITE) estimation requires adjusting for the covariate shift between populations with different treatments, and deep representation learning has shown great promise in learning a balanced representation of covariates. However the existing methods mostly consider the scenario of binary treatments. In this paper, we consider the more practical and challenging scenario in which the treatment is a continuous variable (e.g. dosage of a medication), and we address the two main challenges of this setup. We propose the adversarial counterfactual regression network (ACFR) that adversarially minimizes the representation imbalance in terms of KL divergence, and also maintains the impact of the treatment value on the outcome prediction by leveraging an attention mechanism. Theoretically we demonstrate that ACFR objective function is grounded in an upper bound on counterfactual outcome prediction error. Our experimental evaluation on semi-synthetic datasets demonstrates the empirical superiority of ACFR over a range of state-of-the-art methods.

Introduction

Estimating the individual treatment effect (ITE) from observational datasets has important applications in domains such as personalized medicine (Kent, Steyerberg, and van Klaveren 2018; Prospero et al. 2020), economics (Varian 2016), and recommendation systems (Wang et al. 2020). The observational dataset contains units with their covariates X , the assigned treatment T , and outcome after intervention Y (also known as the factual outcome). Since the treatment assignment policy is unknown, there typically exists an inherent treatment-selection bias stemming from confounding covariates that influence both the treatment assignment and the outcome. Consequently, a causal ITE estimator requires adjusting for the covariate shift among different treatment populations (Bareinboim, Tian, and Pearl 2014; Bareinboim and Pearl 2012).

In recent years, the representation learning approach (Johansson, Shalit, and Sontag 2016; Shalit, Johansson, and Sontag 2017) has demonstrated remarkable success in adjusting for covariate shift. Briefly, the idea is to learn a balanced representation of covariates (rather than balancing the covariates themselves) using an encoder and then predicting

the outcomes from the representation using an outcome prediction network. However, the majority of the representation learning methods have been proposed for binary treatments (Hassanpour and Greiner 2019; Zhang, Liu, and Li 2021; Yao et al. 2018), and incorporating continuous treatments into their architectures is challenging (Chu et al. 2023). We elaborate on the associated challenges with continuous treatments in the following.

i) Balancing representation for continuous treatments

To achieve a balanced representation Z , most existing methods minimize the shift in terms of Integral Probability Metric (IPM) distance¹ (Sriperumbudur et al. 2009) between the distributions of $P(Z, T)$ and $P(Z)P(T)$ as it allows bounding the counterfactual prediction error (Shalit, Johansson, and Sontag 2017). Given that the marginal distribution $P(Z)$ is unknown, (Bellot, Dhir, and Prando 2022) proposed minimizing the IPM distance between $P(Z, T = t)$ and $P(Z, T = -t)$ where t denotes a treatment value in the data and $-t$ denotes all treatments except t . Similarly, (Wang et al. 2022a) suggested discretizing the treatment range into intervals and minimized the maximum IPM distance between the distributions of the two intervals. Despite minimizing an upper bound, these methods involve non-parametric approximations of several IPM distances in practice, which may be inaccurate for high-dimensional representation and small training data (Liang 2019). Furthermore, IPMs are by definition worst-case distances and obtaining a treatment-invariant representation through IPM might be overly restrictive, potentially excluding important confounding factors for outcome prediction (Zhang, Bellot, and van der Schaar 2020).

There are also non-IPM representation learning methods. For instance, (Du et al. 2021; Berrevoets et al. 2020) adopt an adversarial discriminator in order to balance the distributions in the latent representation, (Yao et al. 2018) proposed preserving the local similarity in the representation space, (Hassanpour and Greiner 2019; Wu et al. 2022) learn a disentangled representation to distinguish different latent factors, (Zhang, Bellot, and van der Schaar 2020) enforce invertibility of the encoder function to prevent the loss of covariate information, and (Alaa and van der Schaar 2017)

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹ $IPM_G(p, q) = \sup_{g \in G} \int_S g(s)(p(s) - q(s)) ds$, where specifying G leads to different distributional distance.

introduced a regularization scheme in order to generalize to counterfactual outcomes. However, these methods lack theoretical justification (Du et al. 2021; Hassanpour and Greiner 2019; Yao et al. 2018) or their guarantees are limited to binary treatments (Zhang, Bellot, and van der Schaar 2020; Alaa and van der Schaar 2017).

ii) Treatment impact in outcome prediction network

In order to predict outcomes of different treatments, the outcome prediction network on top of the representation needs to incorporate the treatment value. Considering the treatment variable as an input of the outcome prediction network along with the representation causes overfitting the much higher-dimensional representation and largely limits the impact of the treatment value. Also, a distinct prediction head for each treatment value (as in the case of binary treatments) is not practical. Instead, (Schwab et al. 2020) proposed dose response network (DRNet) which divides the treatment range into intervals and consider a distinct network for the prediction of each interval. (Nie et al. 2021) proposed a varying coefficient network (VCNet) that involves the treatment value in the network parameters through spline functions of treatment. Nonetheless, both networks can not capture the dependency between representation and treatment effectively since the choice of spline functions in VCNet (or the intervals in DRNet) are made irrespective to representation values (Zhang et al. 2022).

In this paper, we propose a representation learning method to accurately predict potential outcomes of a continuous treatment. We address the above challenges through the following contributions:

1. We prove that under certain assumptions the counterfactual error is bounded by the factual error and the KL divergence between $P(Z)P(T)$ and $P(Z, T)$. Unlike the IPM distance, the KL divergence can be estimated parametrically, leading to a more reliable bound.
2. We propose Adversarial Balanced Counterfactual Regression (ACFR) network. ACFR minimizes the KL divergence using an adversarial game extracting a balanced representation for continuous treatments. ACFR also minimizes the factual prediction error by a cross-attention network that captures the complex dependency between treatment and the representation.
3. We conduct an experimental comparison of ACFR against state-of-the-art methods on semi-synthetic datasets, News and TCGA, and analyze the robustness to varying-levels of treatment-selection bias for the methods.

Problem Setup

We assume a dataset of the form $D = \{x_i, t_i, y_i\}_{i=1}^N$, where $x_i \in \mathcal{X} \subseteq \mathbb{R}^d$ denotes the covariates of the i th unit, $t_i \in [0, 1]$ is the continuous treatment that unit i received, and $y_i \in \mathcal{Y} \subseteq \mathbb{R}$ denotes the outcome of interest for unit i after receiving treatment t_i . N is the total number of units, and d is the dimension size of covariates. We are interested in learning a machine learning model to predict the causal quantity $\mu(x, t) = E_{\mathcal{Y}}[Y(t)|X = x]$, which is the potential expected outcome under treatment t for the individual with

covariates x . Note that, unlike binary ITE, the goal is to predict all potential outcomes, not just the difference between them. Similar to previous works, we rely on the following standard assumptions to make treatment effects identifiable from an observational dataset.

[**Assumption 1 - Unconfoundedness**] ($\{Y_t\}_{t \in T} \perp T|X$). In words, given covariates, treatment and potential outcomes are independent.

[**Assumption 2 - Overlap**] ($P(T = t|X = x) > 0, \forall t \in [0, 1], \forall x \in X$). In words, every unit receives treatment level t with a probability greater than zero.

With these assumptions, $\mu(x, t)$ can be rewritten as follows, and we can estimate it:

$$\mu(x, t) = E_{\mathcal{Y}}[Y(t)|X = x] = E_{\mathcal{Y}}[Y|X = x, T = t]$$

Theoretical Analysis

We analyze the properties of encoder function $\phi : \mathcal{X} \rightarrow \mathcal{Z}$, where \mathcal{Z} is the representation space, outcome prediction function $h : \mathcal{Z} \times [0, 1] \rightarrow \mathcal{Y}$ and loss function $L : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^+$.

[**Definition 1**] Define $\ell_{L, h, \phi}(x, t) = L(h(\phi(x), t), y) : \mathcal{X} \times [0, 1] \rightarrow \mathbb{R}^+$ to be the unit-loss for a unit with covariate x that is intervened with treatment t . Unit-loss $\ell_{L, h, \phi}(x, t)$ measures loss L between the predicted outcome $\hat{y} = h(\phi(x), t)$ and the ground-truth outcome $y = \mu(x, t)$.

Using the definition of unit-loss, we can define the expected prediction error of some treatment t by marginalizing over the covariate distribution. As a result of treatment-selection bias, covariate distribution of samples having received treatment t (factual) and samples not having received treatment t (counterfactual) are different. We define factual error $\varepsilon_f^\ell(t)$ by marginalizing over $p(x|t)$ and counterfactual error $\varepsilon_{cf}^\ell(t)$ by marginalizing over $p(x)$ as follows.

$$\begin{aligned} \varepsilon_f^\ell(t) &= \int_{\mathcal{X}} \ell_{L, h, \phi}(x, t) p(x|t) dx \\ \varepsilon_{cf}^\ell(t) &= \int_{\mathcal{T}'=[0,1]-\{t\}} \int_{\mathcal{X}} \ell_{L, h, \phi}(x, t) p(x|t') dx dt' \\ &= \int_{\mathcal{X}} \ell_{L, h, \phi}(x, t) p(x) dx \end{aligned}$$

We also define the expected error of all treatments by marginalizing over their range $[0, 1]$ as follows: $\varepsilon_f^\ell = \int_{[0,1]} \int_{\mathcal{X}} \ell_{L, h, \phi}(x, t) p(x, t) dx dt$ and $\varepsilon_{cf}^\ell = \int_{[0,1]} \int_{\mathcal{X}} \ell_{L, h, \phi}(x, t) p(x) p(t) dx dt$. Note that the expected factual error ε_f^ℓ integrates over joint distribution $p(x, t)$, and expected counterfactual error ε_{cf}^ℓ integrates over $p(x) p(t)$. We aim to reduce the distributional distance in representation space \mathcal{Z} to ensure that minimizing ε_f^ℓ results in minimizing ε_{cf}^ℓ . We need the following assumption on encoder ϕ to ensure balancing properties generalize from the representation space to the covariate space.

[**Assumption 3**] The encoder function ϕ is a twice-differentiable one-to-one mapping and the representation space \mathcal{Z} is the image of \mathcal{X} under ϕ with the induced distribution $p_\phi(z)$.

We also need the following assumption to ensure the unit-loss is not arbitrary large for any (x, t) pair. Constraint on the unit-loss function is required for the IPM distance specification as well.

[Assumption 4] Let G be a class of functions with infinity norm less than 1, $G = \{g : \mathcal{Z} \times [0, 1] \rightarrow \mathbb{R}^+ \mid \|g\|_\infty \leq 1\}$. Then, there exist a constant $C > 0$ such that $\frac{\ell_{L,h,\phi}(x,t)}{C} \in G$. This means for any (x, t) we have $\frac{\ell_{L,h,\phi}(x,t)}{C} \leq 1$.

Note that Assumptions 3 and 4 are common in representation learning literature (Shalit, Johansson, and Sontag 2017; Bellot, Dhir, and Prando 2022; Hassanpour and Greiner 2019). Now we present our main theoretical results which demonstrates a bound on the expected counterfactual error ε_{cf}^ℓ consisting of the expected factual error ε_f^ℓ and the KL divergence between distributions $p_\phi(z, t)$ and $p_\phi(z)p(t)$.

[Proposition 1 - Counterfactual Generalization Bound] Given the one-to-one encoder function $\phi : \mathcal{X} \rightarrow \mathcal{Z}$, the outcome prediction function $h : \mathcal{Z} \times [0, 1] \rightarrow \mathcal{Y}$, and the unit-loss function $\ell_{L,h,\phi}(x, t)$ that satisfies Assumption 4,

$$\varepsilon_{cf}^\ell \leq \varepsilon_f^\ell + C \sqrt{2 D_{KL} \left(p_\phi(z, t) \parallel p_\phi(z)p(t) \right)}$$

Note that the KL divergence is non-negative and becomes zero if and only if two distributions are the same. Therefore, $D_{KL} \left(p_\phi(z, t) \parallel p_\phi(z)p(t) \right) = 0$ implies $p_\phi(z, t) = p_\phi(z)p(t)$ and $\varepsilon_f^\ell = \varepsilon_{cf}^\ell$. We can also interpret the above bound from information theory perspective. Briefly, minimizing RHS of Proposition 1 results in representation z which has maximum mutual information with outcome y given treatment t (equivalent to minimizing ε_f^ℓ) and minimum mutual information with treatment t (equivalent to minimizing KL divergence term).

For some applications, one might be interested in the treatment effect between two different treatments rather than predicting all counterfactual outcomes. For instance, in binary treatment setting it is standard to report the model performance in terms of precision of estimating heterogeneous effect (PEHE) (Hill 2011) which measures the squared difference between ground-truth treatment effect $\tau(x) = \mu(x, 1) - \mu(x, 0)$ and predicted treatment effect $\hat{\tau}(x) = h(\phi(x), 1) - h(\phi(x), 0)$. We define the continuous counterpart $\varepsilon_{pehe}(t_1, t_2)$ between two treatment levels t_1 and t_2 and present an upper bound on it in Proposition 2.

[Definition 2] Define $\varepsilon_{pehe}(t_1, t_2) = \int_{\mathcal{X}} [(\mu(x, t_1) - \mu(x, t_2)) - (h(\phi(x), t_1) - h(\phi(x), t_2))]^2 p(x) dx$ to be the expected precision of estimating heterogeneous effect between treatment levels t_1 and t_2 .

[Proposition 2 - Precision of Estimating Heterogeneous Effect Bound] Given the one-to-one encoder function ϕ and outcome prediction function h as in Proposition 1, and a unit-loss function $\ell_{L,h,\phi}(x, t)$ that satisfies Assump-

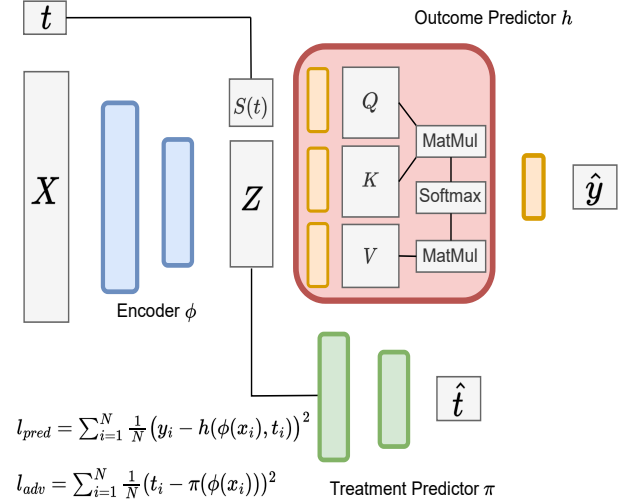


Figure 1: The architecture of Adversarial Counterfactual Regression Network consisting of three sub-networks encoder ϕ , outcome predictor h , and treatment predictor π . Networks ϕ and h are trained to minimize the outcome prediction loss l_{pred} , and networks ϕ and π are trained to maximize / minimize adversarial loss l_{adv} . The encoder and treatment predictor are implemented using linear layers, and the outcome predictor network consists of a cross-attention module followed by a linear layer.

tion 4 and its associated L is squared error $\|\cdot\|^2$,

$$\begin{aligned} \varepsilon_{pehe}(t_1, t_2) &\leq \varepsilon_f^\ell(t_1) + \varepsilon_f^\ell(t_2) \\ &\quad + C \sqrt{2 D_{KL} \left(p_\phi(z) \parallel p_\phi(z|t_1) \right)} \\ &\quad + C \sqrt{2 D_{KL} \left(p_\phi(z) \parallel p_\phi(z|t_2) \right)} \end{aligned}$$

Method

Based on Proposition 1, we can control the counterfactual error by learning an encoder function ϕ and an outcome prediction function h that jointly minimize distribution shift and factual outcome error. We propose our method that implements functions ϕ and h using neural networks and is trained with an objective function inspired by Proposition 1. Figure 1 illustrates the architecture of the ACFR network, consisting of an encoder network ϕ , an outcome prediction network h , and a treatment prediction network π . The key aspects of ACFR, distribution shift minimization and minimization of outcome prediction error, are presented as follows.²

²Note that in the previous section we assumed ϕ is a one-to-one mapping, however, we obtained better results empirically for all methods using a neural network as encoder.




Figure 2: Tsn plot of latent representation Z learned using different distributional distances. After training each method on News dataset, we mapped validation samples into latent representation and plotted them using 2d tsne. We categorized the samples into 4 intervals with respect to their assigned treatment value and each interval corresponds to a color. We consider two important classes of IPM metrics, HSIC and Wasserstein. The treatment value is less distinguishable in the KL divergence representation followed by IPM-ADMIT (minimization with the algorithm proposed in (Wang et al. 2022b)), and IPM (minimization with the procedure proposed in (Bellot, Dhir, and Prando 2022))

Distribution Shift Minimization

As discussed earlier, in order to minimize distribution shift we aim to minimize the KL divergence term with respect to encoder ϕ . The KL divergence can be rewritten as follows:

$$D_{KL} \left(p_{\phi}(z, t) \parallel p_{\phi}(z) p(t) \right) = I(T, Z; \phi) = H(T) - H(T|Z; \phi)$$

where $I(T, Z; \phi)$ is the mutual information between T and Z , and $H(T|Z; \phi)$ is the conditional entropy. Marginal entropy $H(T)$ does not depend on ϕ , thus minimizing KL divergence is equivalent to maximizing the conditional entropy $H(T|Z; \phi) = \mathbb{E}[p_{\phi}(t|z)]$. However, as $p_{\phi}(t|z)$ is intractable we introduce variational distribution $q_{\pi}(t|z)$ defined over the same space to approximate it. For any variational distribution $q_{\pi}(t|z)$ the following holds (Farnia and Tse 2016).

$$\max_{\phi} H(T|Z; \phi) = \max_{\phi} \min_{\pi} \mathbb{E}_{p_{\phi}(t, z)} [-\log q_{\pi}(t|z)]$$

We assume the distribution $q_{\pi}(t|z)$ is a normal distribution with a fixed variance. We can estimate the mean of $q_{\pi}(t|z)$ by a neural network called treatment-prediction network π . By approximating $p_{\phi}(z, t)$ with empirical data, we derive the following mean squared adversarial loss term from the above negative log likelihood.

$$l_{adv} = \max_{\phi} \min_{\pi} \sum_{i=1}^N \frac{1}{N} (t_i - \pi(\phi(x_i)))^2$$

Specifically, the treatment-predictor network π (The green network in Figure 1) is trained to minimize l_{adv} by predicting treatment t from representation $z = \phi(x)$. The encoder network ϕ (The blue network in Figure 1) is trained to maximize l_{adv} by extracting z in such a way that the assigned treatment t is not distinguishable. Therefore KL divergence can be estimated and minimized using two networks and an adversarial loss. Through alternating optimization with respect to ϕ and π , and assuming that the treatment

predictor π reaches the optimum in each iteration, the resulting representation z has a desired property: $\mathbb{E}[t|z] = \mathbb{E}[t]$ (Wang, He, and Katabi 2020). This implies that knowing latent representation z does not provide additional information for predicting the expected treatment. Figure 2 also demonstrates that the representation learned through KL divergence minimization is less predictive of the treatment value compared to representations obtained from two classes of IPM, HSIC (Gretton et al. 2007) and Wasserstein (Villani 2008), thereby showing more effective reduction of the shift.

Factual Outcome Error Minimization

In this section, we discuss the minimization of the factual prediction error ϵ_f^{ℓ} . Recall that:

$$\begin{aligned} \epsilon_f^{\ell} &= \int_{[0,1]} \int_{\mathcal{X}} \ell_{\phi, h}(x, t) p(x, t) dx dt \\ &= \int_{[0,1]} \int_{\mathcal{X}} L(h(\phi(x), t), y) p(x, t) dx dt \end{aligned}$$

Here, outcome y is a continuous variable, and we consider L to be the squared loss. By approximating $p(x, t)$ with empirical distribution, we derive the following outcome prediction loss that needs to be minimized with respect to ϕ and h :

$$l_{pred} = \min_{\phi, h} \sum_{i=1}^N \frac{1}{N} (y_i - h(\phi(x_i), t_i))^2$$

The encoder network ϕ is as defined in the previous section. The outcome prediction network h , however, needs to be particularly designed to maintain the treatment impact on the outcome. We aim to obtain an informative embedding for treatment value, and similar to (Zhang et al. 2022) predict the outcome from the embedding and representation using an attention-based network. (Zhang et al. 2022) proposed to learn the embedding by a neural network. While neural networks are universal function approximators, it has been shown that they can not extract an expressive embedding from a scalar value due to optimization

difficulties (Gorishniy, Rubachev, and Babenko 2023). We construct the treatment embedding applying a set of predefined spline functions to the treatment t shown as $S(t) = [s_1(t), s_2(t), \dots, s_m(t)]$ in Figure 1. Spline functions have been shown to be able to approximate a function in a piecewise manner (Eilers and Marx 1996).

The treatment embedding and the representation are then passed to the cross attention layer (red module in Figure 1) to learn the dependency between treatment and representation. A cross-attention layer has three matrices query Q , key K , and value V , where Q is learned from treatment embedding using h_q parameter and K and V are learned from the representation using h_k and h_v parameters respectively. The output of the cross-attention layer is $\sigma(\frac{Q^T K}{\sqrt{d_k}})V$ where d_k is the dimension of the Q and K matrices and σ denotes the softmax function. We then predict the outcome \hat{y} by a linear layer after the cross attention module.

Unlike ad-hoc architectures VCNet (Nie et al. 2021) and DRNet (Schwab et al. 2020), our outcome prediction network is flexible in terms of the number of spline functions. We can incorporate as many splines as needed without increasing the number of model parameters. This is particularly important in individual effect estimation, because each individual responds differently to a given treatment, and hence different spline functions might be necessary to approximate the treatment-response function for different patients. The proposed architecture can incorporate a large number of spline functions, and the attention layer learns how relevant each spline is for estimating each patient treatment-response function. It is also worth mentioning that by setting h_q to the identity, h_k to the unity and parameterizing h_v with a neural network (which are sub-optimal choices) we recover VCNet and DRNet with a cross-attention layer.

Adversarial Counterfactual Regression

In order to minimize the distribution shift, we derived the adversarial loss l_{adv} and introduced the networks encoder ϕ and treatment-prediction network π to optimize the loss. Similarly, in order to predict factual outcome we derived outcome prediction loss l_{pred} and introduced its associated attention based prediction network h to minimize it. Now, we aim to train the entire network to optimize the following objective function as Proposition 1 suggests:

$$\mathcal{L}_{ACFR} = \max_{\pi} \min_{\phi, h} l_{pred} - \gamma l_{adv}$$

where γ is a tunable parameter. Algorithm 1 presents pseudo-code for the training of the ACFR network. At each iteration a batch of samples is given as input to the network (line 3). At the first stage ACFR predicts the treatment using encoder and treatment-prediction networks, computes l_{adv} and only updates parameters of π for M iterations (line 5-10). At the second stage, ACFR predicts the outcome from the encoded representation of the batch using outcome prediction network h , computes l_{pred} and updates h and ϕ with respect to l_{pred} and $l_{pred} - \gamma l_{adv}$ losses (line 11-14). Finally, the parameters of the encoder and the prediction networks are returned for the inference phase (line 15).

Algorithm 1: Adversarial CounterFactual Regression

```

1: Input: Factual samples  $(x_i, t_i, y_i)_{i=1}^N$ , encoder network
   with initial parameter  $\phi_0$ , treatment-predictor network
   with initial parameters  $\pi_0$ , hypothesis network with initial
   parameters  $h_0$ , batch size  $b$ , iteration number  $T$ , inner
   loop size  $M$ , trade-off parameter  $\gamma$ , and the step
   sizes  $\eta_1$  and  $\eta_2$ .
2: for  $t \leftarrow 0$  to  $T - 1$  do
3:   Sample a mini-batch:  $B = \{i_1, i_2, \dots, i_b\}$ :
4:   Encode into latent representation:  $z_B = \phi_t(x_B)$ 
5:   Initialize  $\omega_0 = \pi_t$ 
6:   for  $m \leftarrow 0$  to  $M - 1$  do
7:     Compute  $l_{adv}$  and update  $\omega_m$ 
8:      $\hat{t}_B = \omega_m(z_B)$   $l_{adv} = \frac{1}{b} \sum_{i \in B} (t_i - \hat{t}_i)^2$ 
9:      $\omega_{m+1} = \omega_m - \eta_2 \gamma \nabla_{\omega} l_{adv}(\omega_m)$ 
10:  end for
11:   $\pi_{t+1} = \omega_m$ 
12:  Compute  $l_{pred}$  and update  $\phi_t$  and  $h_t$ :
13:   $\hat{y} = h_t(z_B)$   $l_{pred} = \frac{1}{b} \sum_{i \in B} (y_i - \hat{y}_i)^2$ 
14:   $\phi_{t+1} = \phi_t - \eta_1 (\nabla_{\phi} l_{pred}(\phi_t) - \gamma \nabla_{\phi} l_{adv}(\phi_t))$ 
15:   $h_{t+1} = h_t - \eta_1 \nabla_h l_{pred}(h_t)$ 
16: end for
17: Return  $\phi_T$  and  $h_T$ 

```

Experiments

Treatment effect estimation methods have to be evaluated for predicting potential outcomes including counterfactuals which are unavailable in real-world observational datasets. Therefore, synthetic or semi-synthetic datasets are commonly used since their treatment assignment mechanism and outcome function are known and hence counterfactual outcomes can be generated. Note that this does not change the fact that only factual outcomes are accessible during training. In this section, we present our experimental results. The code for synthetic data generation and implementation of the methods can be found at here: <https://github.com/amirrezakazemi/acfr>

Setup

Semi-synthetic data generation We used TCGA (Network et al. 2013) and News (Johansson, Shalit, and Sonntag 2016) semi-synthetic datasets. TCGA dataset consists of gene expression measurements of the 4000 most variable genes for 9659 cancer patients. The News dataset which was introduced as a benchmark in (Johansson, Shalit, and Sonntag 2016) consists of 3477 word counts for 5000 randomly sampled news items from the NY times corpus. For each dataset, we first normalized each covariate and then scaled every sample to have a norm 1. We then split the datasets with 68/12/20 ratio into training, validation, and test sets. We followed treatment and outcome generating process of (Bica, Jordon, and van der Schaar 2020), summarized in Table 1. The parameter α in treatment function determines the treatment-selection bias level (α is set 2 in all experiments unless otherwise stated), and v_1, v_2 and v_3 are vectors whose entries are sampled from the normal distribution $\mathcal{N}(0, 1)$,

Method	News		TCGA	
	MISE	PE	MISE	PE
GPS	3.21 ± 0.34	0.39 ± 0.03	6.50 ± 1.21	2.30 ± 0.27
MLP	2.91 ± 0.33	0.31 ± 0.02	4.81 ± 0.54	1.15 ± 0.23
DRNet-HSIC	1.59 ± 0.20	0.21 ± 0.01	2.03 ± 0.27	1.24 ± 0.23
DRNet-Wass	1.64 ± 0.21	0.21 ± 0.01	2.01 ± 0.20	1.29 ± 0.21
VCNet-HSIC	1.28 ± 0.10	0.16 ± 0.01	1.99 ± 0.11	0.94 ± 0.14
VCNet-Wass	1.43 ± 0.11	0.17 ± 0.01	1.76 ± 0.12	0.92 ± 0.14
ADMIT-HSIC	1.25 ± 0.12	0.18 ± 0.01	1.81 ± 0.23	0.86 ± 0.15
ADMIT-Wass	1.35 ± 0.20	0.18 ± 0.01	1.67 ± 0.23	0.81 ± 0.14
SCIGAN	1.21 ± 0.15	0.20 ± 0.01	1.85 ± 0.14	0.97 ± 0.14
ACFR w/o attention	1.58 ± 0.15	0.19 ± 0.01	1.86 ± 0.21	1.01 ± 0.15
ACFR	1.12 ± 0.12	0.18 ± 0.01	1.60 ± 0.20	0.76 ± 0.12

Table 1: Results on News and TCGA datasets for the out-of-sample setting.

Dataset	#Samples	#Covariates	Outcome function	Treatment assignment
TCGA	9659	4000	$y = 10(v_1^T x + 12v_2^T xt - 12v_3^T xt^2)$	$t = \text{Beta}(\alpha, \beta)$
News	5000	3477	$y = 10(v_1^T x + \sin(\frac{v_2^T x}{v_3^T x} \pi t))$	$\beta = \frac{2(\alpha-1)v_2^T x}{v_3^T x} + 2 - \alpha$

Table 2: Datasets and data generating functions.

and then are normalized. Using the functions in Table 1, we assigned the treatment and factual outcome for all samples in the training and validation sets. All methods are then trained on the training set, and the validation set has been used for hyperparameter selection. Same as (Bica, Jordon, and van der Schaar 2020), potential outcomes for a unit are generated using the outcome function given the unit’s covariates and 65 grids in the range $[0, 1]$ as an approximation of the treatment range.

Baselines DRNet (Schwab et al. 2020) and VCNet (Nie et al. 2021) are state-of-the-art neural networks for continuous treatment estimation. Following (Bellot, Dhir, and Prando 2022), we use the following versions of these two methods as the main baselines. HSIC (Gretton et al. 2007) is a version using the Hilbert-Schmidt independence criterion to minimize the distribution shift while Wass (Villani 2008) is a version that uses the Wasserstein distance for that purpose. We consider SCIGAN (Bica, Jordon, and van der Schaar 2020) as state-of-the-art generative method for continuous treatments, and also compare against ADMIT (Wang et al. 2022b) network with their proposed algorithm to estimate IPM distances. Finally we include Generalized Propensity Score (GPS) and a MLP network as baselines. The MLP network consists of two fully connected layers without any attempt to reduce distribution shift.

Metrics Having $\mu(x, t)$ as the ground-truth outcome of the unit with covariate x under treatment t and $f(x, t)$ as the predicted outcome, we report the performance of methods in terms of the two following metrics defined in (Schwab et al. 2020). The Mean Integrated Squared Error (MISE) is the squared error of the predicted outcome averaged over all treatment values and all units. The Policy Error (PE) mea-

sures the average squared error of estimated optimal treatment, where t_i^* and \hat{t}_i^* denote ground-truth and predicted best treatments respectively.

$$\text{MISE} = \frac{1}{N} \sum_{i=1}^N \int_0^1 [\mu(x_i, t) - f(x_i, t)]^2 dt$$

$$\text{PE} = \frac{1}{N} \sum_{i=1}^N [\mu(x_i, t_i^*) - \mu(x_i, \hat{t}_i^*)]^2$$

Results

We performed two sets of experiments for potential outcome prediction, called out-of-sample prediction and within-sample prediction. The out-of-sample experiment shows the ability of models in predicting the potential outcomes for units in the held-out test set, and the within-sample experiment shows the ability for units in the training set.

Prediction error For all methods, we reported the mean and the standard deviation of MISE and PE in the format of mean±std over 20 realizations of each dataset. Table 1 shows that on TCGA dataset, ACFR outperformed the contenders in both metrics, and on News dataset, ACFR achieved the best and second best result in terms of MISE and PE metrics respectively. We can also see the substantial gain of cross-attention layer in the performance of ACFR by comparing it with ACFR w/o attention, demonstrating the effectiveness of proposed outcome prediction network. Comparably, DRNet and VCNet methods have more parameters while ACFR and ADMIT methods are more time-consuming because of their corresponding inner loop to minimize the distribution shift.

Treatment-selection bias robustness We also investigate the robustness of 4 methods (ACFR, VCNet-HSIC, DRNet-HSIC, and ADMIT-HSIC) against varying level of treatment-selection bias. As mentioned earlier, the α parameter of Beta distribution in the treatment generating function controls the amount of bias. As α increases the treatment-selection bias and covariate shift of the observational dataset increase and consequently, we expect the error of methods to increase as well. As shown in Figure 3, ACFR performs consistently and has a notable gap with the baselines at the strong bias level for the out-of-sample setting.

Related Work

Continuous Treatment Effect Estimation These methods can be categorized into those estimating the average effect and those estimating the individual effect. In the first category, (Hirano and Imbens 2004) proposed the generalized propensity score (GPS) that generalizes the notion of propensity score to continuous treatments. (Wu et al. 2021; Fong, Hazlett, and Imai 2018) proposed approaches to matching and covariate balancing, respectively, according to the weights learned using GPS. (Nie et al. 2021) proposed the Varying Coefficient network (VCNet) which extracts a representation sufficient for GPS prediction and predicts the outcome using a network where the treatment value influences the outcome indirectly through parameters instead of being given directly as input. (Bahadori, Tchetgen, and Heckerman 2022) proposes an entropy balancing method to learn more stable weights compared to GPS based weights. Our approach is fundamentally different since sufficiency of (generalized) propensity score theorem holds only for the average effect and we aim to estimate the individual effect.

In the individual effect category, methods are mostly based on learning a balanced representation. DRNet (Schwab et al. 2020) discretized the treatment range into intervals, minimized the pair-wise shift between the populations fall within these intervals in the representation, and proposed a hierarchical multi head network to predict outcomes for the intervals. (Bellot, Dhir, and Prando 2022) demonstrated that minimizing IPM between $P(Z, T)$ and $P(Z)P(T)$ coupled with factual outcome error leads to an upper bound of the counterfactual error. Nonetheless, in practice they minimized the IPM sample-wise since $P(Z)$ is unknown. Similarly, (Wang et al. 2022b) demonstrated that discretizing the treatment range and minimizing the maximum pair-wise IPM bounds the counterfactual error. Our method is different as we learn the balanced representation by minimizing the KL divergence parameterically.

Adversarial Balanced Representation Learning a balanced (invariant) representation using an adversarial discriminator has been studied in the transfer learning literature to align source domain(s) to target domain(s) (Ganin et al. 2016; Tzeng et al. 2017; Wang, He, and Katabi 2020). Similarly, in causal inference (Du et al. 2021; Berrevoets et al. 2020) aimed to balance the distributions of two treatment groups adversarially. (Bica et al. 2020) extended the approach to the multiple time-varying treatment setting. How-

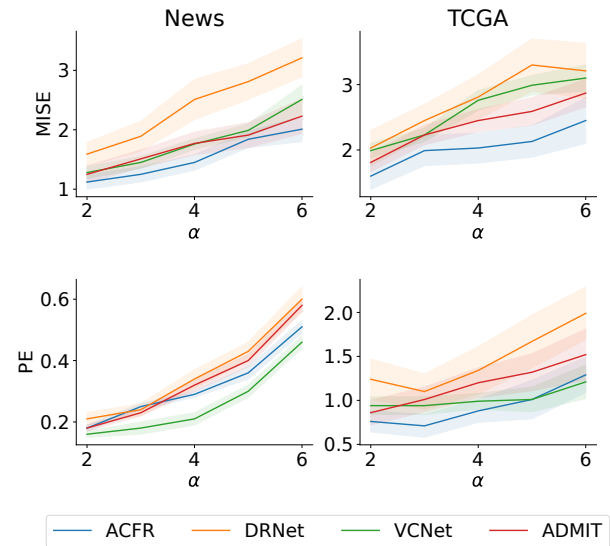


Figure 3: Robustness of ACFR against varying level of selection bias determined by α parameter of treatment assignment distribution. ACFR demonstrates a robust performance in terms of MISE and PE compared to baselines.

ever, the existing methods consider only scenarios with a finite number of treatment options and do not provide theoretical guarantees of their generalization capability.

Conclusion

This paper has investigated the problem of continuous treatment effect estimation and introduced ACFR (Adversarial Counter-Factual Regression) method for predicting potential outcomes. We proved a new bound of the counterfactual error using the KL divergence instead of an IPM distance, which has the benefit that the KL divergence can be estimated parametrically and results in a more reliable bound. Based on the error bound, ACFR uses an adversarial neural network architecture to minimize the KL divergence of the representations and a cross-attention network to minimize the factual prediction error. It is worth mentioning that ACFR is not restricted to continuous treatments, and in future work we plan to extend and evaluate ACFR framework for structured and time series treatments. Nonetheless, we note that ACFR, similar to many treatment effect estimation methods, relies on the strong ignorability assumption, which may not necessarily hold in real-world applications.

Acknowledgements

We would like to thank Oliver Schulte and Sharan Vaswani for feedback on the paper. This research was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant.

References

Alaa, A. M.; and van der Schaar, M. 2017. Bayesian Inference of Individualized Treatment Effects using Multi-task

- Gaussian Processes. arXiv:1704.02801.
- Bahadori, M. T.; Tchetgen, E. T.; and Heckerman, D. E. 2022. End-to-End Balancing for Causal Continuous Treatment-Effect Estimation. arXiv:2107.13068.
- Bareinboim, E.; and Pearl, J. 2012. Controlling selection bias in causal inference. In *Artificial Intelligence and Statistics*, 100–108. PMLR.
- Bareinboim, E.; Tian, J.; and Pearl, J. 2014. Recovering from Selection Bias in Causal and Statistical Inference. *Proceedings of the AAAI Conference on Artificial Intelligence*, 28(1).
- Bellot, A.; Dhir, A.; and Prando, G. 2022. Generalization bounds and algorithms for estimating conditional average treatment effect of dosage. arXiv:2205.14692.
- Berrevoets, J.; Jordon, J.; Bica, I.; gimson, a.; and van der Schaar, M. 2020. OrganITE: Optimal transplant donor organ offering using an individual treatment effect. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 20037–20050. Curran Associates, Inc.
- Bica, I.; Alaa, A. M.; Jordon, J.; and van der Schaar, M. 2020. Estimating Counterfactual Treatment Outcomes over Time Through Adversarially Balanced Representations. arXiv:2002.04083.
- Bica, I.; Jordon, J.; and van der Schaar, M. 2020. Estimating the Effects of Continuous-valued Interventions using Generative Adversarial Networks. arXiv:2002.12326.
- Chu, Z.; Huang, J.; Li, R.; Chu, W.; and Li, S. 2023. Causal Effect Estimation: Recent Advances, Challenges, and Opportunities. arXiv:2302.00848.
- Du, X.; Sun, L.; Duivesteyn, W.; Nikolaev, A.; and Pechenizkiy, M. 2021. Adversarial balancing-based representation learning for causal effect inference with observational data. *Data Mining and Knowledge Discovery*, 35(4): 1713–1738.
- Eilers, P. H.; and Marx, B. D. 1996. Flexible smoothing with B-splines and penalties. *Statistical science*, 11(2): 89–121.
- Farnia, F.; and Tse, D. 2016. A Minimax Approach to Supervised Learning.
- Fong, C.; Hazlett, C.; and 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.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-Adversarial Training of Neural Networks. arXiv:1505.07818.
- Gorishniy, Y.; Rubachev, I.; and Babenko, A. 2023. On Embeddings for Numerical Features in Tabular Deep Learning. arXiv:2203.05556.
- Gretton, A.; Fukumizu, K.; Teo, C. H.; Song, L.; Schölkopf, B.; and Smola, A. 2007. A Kernel Statistical Test of Independence. In *NIPS*.
- Hassanpour, N.; and Greiner, R. 2019. Learning disentangled representations for counterfactual regression. In *International Conference on Learning Representations*.
- Hill, J. L. 2011. Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics*, 20: 217 – 240.
- Hirano, K.; and Imbens, G. W. 2004. The propensity score with continuous treatments. *Applied Bayesian modeling and causal inference from incomplete-data perspectives*, 226164: 73–84.
- Johansson, F.; Shalit, U.; and Sontag, D. 2016. Learning Representations for Counterfactual Inference. In Balcan, M. F.; and Weinberger, K. Q., eds., *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, 3020–3029. New York, New York, USA: PMLR.
- Kent, D. M.; Steyerberg, E.; and van Klaveren, D. 2018. Personalized evidence based medicine: predictive approaches to heterogeneous treatment effects. *Bmj*, 363.
- Liang, T. 2019. Estimating Certain Integral Probability Metric (IPM) is as Hard as Estimating under the IPM. arXiv:1911.00730.
- Network, C.; Weinstein, J.; Collisson, E.; Mills, G.; Shaw, K.; Ozenberger, B.; Ellrott, K.; Shmulevich, I.; Sander, C.; Stuart, J.; and Vandin, F. 2013. The Cancer Genome Atlas Pan-Cancer analysis project. *Nature Genetics*, 45(10): 1113–1120.
- Nie, L.; Ye, M.; Liu, Q.; and Nicolae, D. 2021. Vcnet and functional targeted regularization for learning causal effects of continuous treatments. *arXiv preprint arXiv:2103.07861*.
- Prosperi, M.; Guo, Y.; Sperrin, M.; Koopman, J. S.; Min, J. S.; He, X.; Rich, S.; Wang, M.; Buchan, I. E.; and Bian, J. 2020. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nature Machine Intelligence*, 2(7): 369–375.
- Schwab, P.; Linhardt, L.; Bauer, S.; Buhmann, J. M.; and Karlen, W. 2020. Learning Counterfactual Representations for Estimating Individual Dose-Response Curves. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04): 5612–5619.
- Shalit, U.; Johansson, F. D.; and Sontag, D. 2017. Estimating individual treatment effect: generalization bounds and algorithms. In *International conference on machine learning*, 3076–3085. PMLR.
- Sriperumbudur, B. K.; Fukumizu, K.; Gretton, A.; Schölkopf, B.; and Lanckriet, G. R. G. 2009. On integral probability metrics, ϕ -divergences and binary classification. arXiv:0901.2698.
- Tzeng, E.; Hoffman, J.; Saenko, K.; and Darrell, T. 2017. Adversarial Discriminative Domain Adaptation. arXiv:1702.05464.
- Varian, H. R. 2016. Causal inference in economics and marketing. *Proceedings of the National Academy of Sciences*, 113(27): 7310–7315.
- Villani, C. 2008. Optimal Transport: Old and New. In *Optimal Transport: Old and New*.
- Wang, H.; He, H.; and Katabi, D. 2020. Continuously indexed domain adaptation. *arXiv preprint arXiv:2007.01807*.

- Wang, X.; Lyu, S.; Wu, X.; Wu, T.; and Chen, H. 2022a. Generalization Bounds for Estimating Causal Effects of Continuous Treatments. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 8605–8617. Curran Associates, Inc.
- Wang, X.; Lyu, S.; Wu, X.; Wu, T.; and Chen, H. 2022b. Generalization Bounds for Estimating Causal Effects of Continuous Treatments. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Wang, Y.; Liang, D.; Charlin, L.; and Blei, D. M. 2020. Causal Inference for Recommender Systems. In *Proceedings of the 14th ACM Conference on Recommender Systems, RecSys '20*, 426431. New York, NY, USA: Association for Computing Machinery. ISBN 9781450375832.
- Wu, A.; Yuan, J.; Kuang, K.; Li, B.; Wu, R.; Zhu, Q.; Zhuang, Y.; and Wu, F. 2022. Learning decomposed representations for treatment effect estimation. *IEEE Transactions on Knowledge and Data Engineering*, 35(5): 4989–5001.
- Wu, X.; Mealli, F.; Kioumourtzoglou, M.-A.; Dominici, F.; and Braun, D. 2021. Matching on Generalized Propensity Scores with Continuous Exposures. arXiv:1812.06575.
- Yao, L.; Li, S.; Li, Y.; Huai, M.; Gao, J.; and Zhang, A. 2018. Representation Learning for Treatment Effect Estimation from Observational Data. In *Neural Information Processing Systems*.
- Zhang, W.; Liu, L.; and Li, J. 2021. Treatment effect estimation with disentangled latent factors. arXiv:2001.10652.
- Zhang, Y.; Bellot, A.; and van der Schaar, M. 2020. Learning Overlapping Representations for the Estimation of Individualized Treatment Effects. arXiv:2001.04754.
- Zhang, Y.-F.; Zhang, H.; Lipton, Z. C.; Li, L. E.; and Xing, E. P. 2022. Exploring Transformer Backbones for Heterogeneous Treatment Effect Estimation. arXiv:2202.01336.