
Adversarial De-confounding in Individualised Treatment Effects Estimation

Vinod Kumar Chauhan¹

Soheila Molaei¹

Marzia Hoque Tania¹

Anshul Thakur¹ and Tingting Zhu¹ and David A. Clifton^{1,2}

¹Institute of Biomedical Engineering, University of Oxford, UK

²Oxford-Suzhou Centre for Advanced Research, Suzhou, China

Abstract

Observational studies have recently received significant attention from the machine learning community due to the increasingly available non-experimental observational data and the limitations of the experimental studies, such as considerable cost, impracticality, small and less representative sample sizes, etc. In observational studies, de-confounding is a fundamental problem of individualised treatment effects (ITE) estimation. This paper proposes disentangled representations with adversarial training to selectively balance the confounders in the binary treatment setting for the ITE estimation. The adversarial training of treatment policy selectively encourages treatment-agnostic balanced representations for the confounders and helps to estimate the ITE in the observational studies via counterfactual inference. Empirical results on synthetic and real-world datasets, with varying degrees of confounding, prove that our proposed approach improves the state-of-the-art methods in achieving lower error in the ITE estimation.

1 INTRODUCTION

Individualised treatment effects (ITE) estimation is a fundamental problem that is useful for making personalised decisions and estimating their effects. For example, in the intensive care unit (ICU), ITE can be used to decide whether or not to give a medication to a patient, which can be a question of life or death of the patient. The ITE estimation learning requires answering counterfactual questions, such as: “*What would have been the outcome if alternative treatment had been given?*”, i.e., it requires predicting

potential outcomes of unexplored actions (Rubin (2005)). Due to its importance, ITE estimation is studied widely across diverse fields, like medicine, marketing, education and policy-making, etc. (please refer to Bica et al. (2021) for an overview).

Randomised controlled trial (RCT) experiments are the gold standard to evaluate the effectiveness of treatments (Pearl (2009)). However, they are expensive, time-consuming, sometimes unethical and impractical, and have small and less representative sample sizes, etc. On the other hand, non-experimental observational studies are becoming popular for evaluating the effectiveness of treatments due to increasingly available observational data, like electronic health records, and overcoming limitations of the RCT.

The ITE estimation from the observational studies have, recently, received a great attention from the machine learning community, e.g., Shalit et al. (2017); Shi et al. (2019); Hassanpour and Greiner (2019b); Curth and van der Schaar (2021a,b); Wu et al. (2022) etc., and it is different from standard machine learning (Rubin (2005); Pearl (2009); Wager and Athey (2018)). This is because the observational data have outcomes available only for the actions taken, i.e., for the selected treatments (factual outcomes), but the outcomes for alternative treatments are not available (counterfactual outcomes) – which is called *the fundamental problem of causal learning* (Holland (1986)). Moreover, the observational data contain confounders, i.e., covariates which affect outcomes as well as treatment assignment policy, and hence have selection-bias (Imbens and Rubin (2015)) (i.e., $p(T = 0|X = x) \neq p(T = 1|X = x)$), unlike the RCT where treatments are assigned randomly and has $p(T = 0|X = x) = p(T = 1|X = x)$, where $p(T = t|X = x)$ is probability of treatment t for a given patient x in a binary treatment setting). For example, in the above ICU setting scenario, suppose choice of medication is dependent on the age of the patient which also affects the recovery rate.

ITE estimation from observational data is a challenging problem and it involves a fundamental problem of *de-*

confounding, i.e., removing selection-bias which otherwise precludes direct comparison between different treatment groups as they have different distributions. Without careful handling of selection-bias, it could lead to biased estimates of causal effects (Zubizarreta (2015)). Recently, several deep learning architectures and algorithms were proposed to address the de-confounding, which borrow ideas from representation learning (Bengio et al. (2013)), domain adaptation (Johansson et al. (2016)) and disentanglement learning (Hassanpour and Greiner (2019b)).

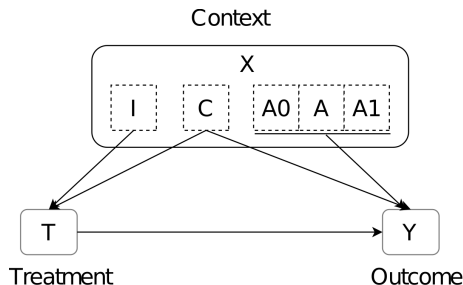


Figure 1: Problem setting: We consider binary treatment setting where the context X can be decomposed into five (or three) subsets as covariates I which affect only the treatment policy π , named instrument variables, covariates $\{A0, A, A1\}$ which affect only the treatment outcomes Y , where $A0$ affects only the control group outcomes μ_0 , $A1$ affects only the treatment group outcomes μ_1 , and A affects both of the outcomes, collectively referred to as adjustment variables and covariates C which affect both treatment policy as well as treatment outcomes.

Representation learning-based methods try to reduce selection-bias by learning a shared representation for outcome and treatment predictions (Johansson et al. (2016)). Similarly, domain adaption-based methods try to learn a common representation space and minimise the discrepancy between control and treatment distributions (Shalit et al. (2017)). Disentanglement learning-based methods decompose the covariates into latent factors corresponding to confounders (i.e., features affecting both treatment and outcomes), instrument variables (i.e., features affecting only treatment decision) and adjustment variables (i.e., features affecting only outcome predictions) for counterfactual interference (Hassanpour and Greiner (2019b)), as shown in Fig. 1 which presents the problem setting.

This paper proposes novel disentangled representations with adversarial training to selectively balance the confounders in the binary treatment setting for ITE estimation. The disentanglement representation learning – using a specific deep learning architecture and constraining orthogonality on each pair of representations – helps to disentangle the latent factors, i.e., representations for the confounders, instrument variables and adjustment variables (refer to Fig.2). The adversarial training jointly optimises the confounder representations to minimise the loss for the

potential outcome predictions and maximise the loss for the treatment assignments. This learns a representation for the confounders, which is treatment-agnostic. Our empirical results on synthetic and real-world datasets, with varying degrees of confounding, show that our proposed approach improves the state-of-the-art methods in achieving lower error in estimating the ITE.

The rest of the paper is organized as: Section 2 presents some related work, Section 3 defines the problem settings, Section 4 describes the proposed methodology and Section 5 presents experiments. Finally, the concluding remarks and future scope are discussed in Section 6.

2 RELATED WORK

The estimation of individualised treatment effects (ITE) from the observational data has recently received growing attention from the machine learning community due to the increasingly available observational data, like electronic health records, and overcoming limitations of randomised controlled trials – the gold standard to estimate treatment effects. ITE estimation from the observational data faces a fundamental problem of de-confounding and the corresponding selection-bias. As discussed below, several techniques have been proposed for de-confounding, such as weighting of outcomes, learning balanced representations, and dis-entanglement learning.

Weighting techniques to address the confounders typically construct some weights to balance the covariates for the control and the treatment groups (Rosenbaum and Rubin (1983)), e.g., inverse probability weighting (IPW). In reality, it is difficult to estimate the correct propensity score values, i.e., the probability of treatment, so these methods are known to be unstable (Hainmueller (2012)). To address this limitation, several alternative weighting schemes are proposed, such as truncated IPW (Crump et al. (2009)), matching weights (Li et al. (2018)) and calculating weights using optimisation program (Li and Fu (2017)) etc. The weighting techniques are known not to scale well to high dimensional problems.

Balanced representation techniques learn a shared representation, i.e., an embedding for the confounders, and force the distributions of the control and treatment arms to be similar. For example, Johansson et al. (2016); Shalit et al. (2017) presented the pioneering work and borrowed the idea of discrepancy from the domain adaptation to introduce a form of regularisation to enforce similarity in the representations for control and treatment groups. Rather than balancing the global distributions for different treatment groups, Yao et al. (2018) extended the idea of representation learning by local similarity preserving technique. The combination of weighting and representation techniques is also explored in Hassanpour and Greiner (2019a),

where authors estimated IPW from the representations to re-weight the regression terms.

Like weighting techniques, representation learning techniques focus on the confounders and balance the representations for the whole context. Recently, VanderWeele (2019); Kuang et al. (2019) highlighted the importance of confounder separation for causal learning and argued that by balancing the complete context, some variables, such as instrument variables, lead to additional bias and variance.

Disentanglement representation techniques address the above concerns of balanced representation learning techniques. They learn to disentangle the context X into latent factors for subsets of X . For example, Hassanpour and Greiner (2019b) proposed disentanglement learning algorithm to learn three latent factors $\{C(X), I(X), A(X)\}$ for confounders C , instrument variables I and adjustment variables A from the context X . However, they didn't guarantee the disentangled representations and used propensity scores for weighting, which runs the risk of instability similar to weighting techniques. The idea of disentanglement is further improved by different researchers and is currently the state-of-the-art technique for ITE estimation. For example, Wu et al. (2022) addressed the above limitations, removed the propensity scores, and introduced orthogonality between each pair of representations, ensuring that each variable in the context X affects only one of the three representations, i.e., $\{I, C, A\}$. Recently, Curth and van der Schaar (2021b) proposed to disentangle the context X into five latent factors $\{C(X), I(X), A0(X), A(X), A1(X)\}$, corresponding to confounders, instrument variables, adjustment variables, where $A0, A, A1$ refer to adjustment variables which affect the control group, control as well as treatment groups and treatment group, respectively, and used orthogonality to enforce that each variable in X affects only one of these five factors.

In this paper, we proposed a novel disentangled representation with adversarial training to selectively balance the confounders. The disentanglement learning helps to separate the confounds and adversarial training (Ganin et al. (2016)) selectively learns balanced representations for the confounders, unlike the balanced representation learning methods which balance the whole context X and may lead to additional bias and variance in counterfactual inference.

3 PROBLEM SETTING

We introduce the problem of ITE estimation for binary treatments from the observational data using the potential outcomes framework (Rubin (2005)). Suppose, $D = \{T_i, X_i, Y_i\}_{i=1}^m$ is a sample of m units, say patients, of the observational data, e.g., electronic health records, taken *i.i.d.* from an unknown distribution \mathbb{P} . $X_i \in \mathbb{R}^d$ is a d -dimensional covariates, i.e., context, e.g., patient history.

$T_i \in \{0, 1\}$ is a binary treatment variable, where $T_i = 1$ refers to a patient i receiving a treatment and $T_i = 0$ refers to a patient not receiving the treatment. $Y_i \in \mathbb{R}$ denotes the patient's outcome.

We use Neyman-Rubin potential outcomes framework (Rubin (2005)) for the individualised treatment effects. Suppose $Y_i(1)$ and $Y_i(0)$ refer to potential outcomes when patient i receives a treatment, i.e., $T_i = 1$ and did not receive the treatment, i.e., $T_i = 0$, respectively. However, due to the *fundamental problem of the causal inference*, we observe only factuals, i.e., the outcome for the selected treatment for a given patient, and not the counterfactuals, i.e., the outcome for the non-selected treatment, as the outcome of the patient is defined as $Y_i = T_i Y_i(1) - (1 - T_i) Y_i(0)$. The ITE is the conditional average treatment effect, also known as CATE¹, as given below.

$$e(x) = \mathbb{E}_{\mathbb{P}} [Y(1) - Y(0) | X = x]. \quad (1)$$

Our work is based on standard assumptions of treatment effects estimation (Imbens and Rubin (2015)), which are given below.

Assumption 1 (Stable Unit Treatment Value:) *The distribution of a patient's outcome depends on the treatment given to it and is independent of the treatment given to other patients.*

Assumption 2 (Ignorability: Rosenbaum and Rubin (1983)) *The treatment assignment policy is independent of the potential outcomes when given the context of a patient,*

$$\{Y(0), Y(1)\} \perp T | X. \quad (2)$$

Assumption 2 is also known as unconfoundedness because it holds if there are no hidden confounders and if these satisfy additional conditions.

Assumption 3 (Overlap: Imbens (2004)) *The treatment assignment policy is stochastic, i.e., each patient has a certain probability of receiving any treatment, i.e.,*

$$0 < \pi(x) < 1, \quad \forall x \in X, \quad (3)$$

where $\pi(x)$ is probability of receiving treatment $T = 1$ of patient x .

The Assumptions 2 and 3 together are called *strong ignorability* assumptions (Imbens and Wooldridge (2009)).

Under the above assumptions, expected potential outcomes and the conditional average treatment effect are identifiable as given below.

$$\mu_t(x) = \mathbb{E}_{\mathbb{P}} [Y | X = x, T = t], \quad (4)$$

¹in machine learning literature, ITE and CATE are interchangeably used to refer to CATE, and is different from $ITE_i = Y_i(1) - Y_i(0)$ in causal inference literature (Johansson et al. (2022))

$$\begin{aligned}
 e(x) &= \mathbb{E}_{\mathbb{P}} [Y(1) - Y(0)|X = x], \\
 &= \mathbb{E}_{\mathbb{P}} [Y|X = x, T = 1] - \mathbb{E}_{\mathbb{P}} [Y|X = x, T = 0], \\
 &= \mu_1(x) - \mu_0(x).
 \end{aligned} \tag{5}$$

It is well-known that the above assumptions are fundamentally untestable from given observational data for ITE estimation (Pearl (2009)), so domain knowledge is necessary to access the validity of these assumptions (Bica et al. (2021)). Moreover, richer datasets can also cover more of the confounding variables.

4 METHODOLOGY

The individualised treatment effects can be estimated from the observational data, $D = \{T_i, X_i, Y_i\}_{i=1}^m$, containing context $X_i \in \mathbb{R}^d$, treatment assignments $T_i \in \{0, 1\}$ and outcomes $Y_i \in \mathbb{R}$ for the patients $i = 1, 2, \dots, m$ by training supervised machine learning model for potential outcome regressions $\mu_t(x) = \mathbb{E}_{\mathbb{P}} [Y|X = x, T = t]$ and using plug-in approach as $e(x) = \mu_1(x) - \mu_0(x)$. However, it needs to address the confounders (i.e., features affecting the treatment policy as well as treatment outcomes) and the resulting selection-bias in the treatment assignment policy (Pearl (2009)). To address the selection-bias caused by the confounders, we separate the confounders using the disentanglement representation learning (Hassanpour and Greiner (2019b); Curth and van der Schaar (2021b); Wu et al. (2022)) and selectively learn balanced confounder representations using the adversarial training (Ganin et al. (2016)) of expected outcome predictions against the propensity score (i.e., probability of treatment) predictions.

4.1 Disentangled Representations

Suppose the context X can be decomposed into five latent factors $\{I(X), C(X), A0(X), A(X), A1(X)\}$ for instrument variables I (variables which affect only the treatment policy), adjustment variables $A, A0$ and $A1$ that affect both the potential outcome regressions μ_0 and μ_1 , only μ_0 and only μ_1 , respectively, and confounder variables C (variables affecting the treatment policy as well as the outcomes). To implement the disentanglement representation learning, it requires a deep learning architecture that has five multi-layer perceptrons (MLP) and takes context X as an input to generate latent representations, i.e., embedding for each of the five factors, $I, A, A0, A1$ and C , as well as orthogonality Curth and van der Schaar (2021b); Wu et al. (2022) between each pair of the representations to ensure each variable in X contributes to only one of the five factors. We present the deep learning architecture to support disentanglement in Fig. 2 and the idea of orthogonality is given below.

The architecture-based representations of

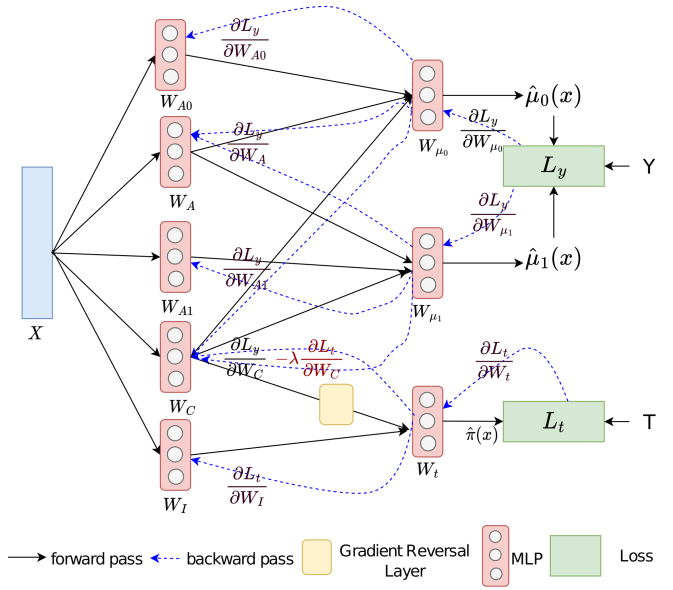


Figure 2: SNet+: proposed disentanglement representation based adversarially balanced architecture. The context X is disentangled into five latent factors $\{I(X), C(X), A0(X), A1(X), A(X)\}$. The treatment classifier is adversarially trained by using a gradient reversal layer between confounder representations $C(X)$ and the treatment classifier to get balanced representations for the confounders.

$\{I, A, A0, A1, C\}$ are insufficient to guarantee disentanglement as data-driven neural networks overfit the training. So, similar to Kuang et al. (2017); Curth and van der Schaar (2021b); Wu et al. (2022), we use orthogonality to enforce variable decompositions among the five representation networks $\{I(X), A(X), A0(X), A1(X), C(X)\}$ corresponding to $\{I, A, A0, A1, C\}$, respectively. Let's take an example of a representation network for confounders $C(X)$ and suppose it takes l layers and weights of the k th layer are denoted by matrix W_k . Then, the contribution of each variable in X on each dimension of representation $C(X)$ can be approximated by computing $W_1 \times W_2 \times \dots \times W_l$ which can be denoted as $W_C \in \mathbb{R}^{d \times n}$, where d is the dimension of X and n is the dimension of $C(X)$. We take the average along the rows of W_C to represent the average contribution of variables in X on representations $C(X)$, denoted as \bar{W}_C . Similarly, we approximate the contributions of each variable in X on representations of $I(X), A(X), A0(X), A1(X)$ as $\bar{W}_I, \bar{W}_A, \bar{W}_{A0}, \bar{W}_{A1}$, respectively. Assuming all representation networks have the same structure, the hard decomposition can be enforced by constraining orthogonality on each pair of them. It can be achieved by adding the following loss term to the

objective function.

$$\begin{aligned}
 L_O = & \bar{W}_I^T \cdot \bar{W}_A + \bar{W}_I^T \cdot \bar{W}_{A0} + \bar{W}_I^T \cdot \bar{W}_{A1} + \bar{W}_I^T \cdot \bar{W}_C \\
 & + \bar{W}_A^T \cdot \bar{W}_{A0} + \bar{W}_A^T \cdot \bar{W}_{A1} + \bar{W}_A^T \cdot \bar{W}_C + \bar{W}_{A0}^T \cdot \bar{W}_{A1} \\
 & + \bar{W}_{A0}^T \cdot \bar{W}_C + \bar{W}_{A1}^T \cdot \bar{W}_C.
 \end{aligned} \quad (6)$$

This loss term can lead to the result $\bar{W}_I^k = \bar{W}_A^k = \bar{W}_{A0}^k = \bar{W}_{A1}^k = \bar{W}_C^k = 0$. So, to avoid this, following regulariser as a soft constraint is used, which make the sum of each of $\bar{W}_I, \bar{W}_A, \bar{W}_{A0}, \bar{W}_{A1}, \bar{W}_C$ to 1, as given below.

$$R_O = \sum_{r \in \{I, A, A0, A1, C\}} \left(\sum_{i=1}^d \bar{W}_r^i - 1 \right)^2. \quad (7)$$

The representation networks for the five latent factors, $\{I(X), A(X), A0(X), A1(X), C(X)\}$, and the above loss term and the regulariser help to learn disentangled representations for separating the confounders from the input context X . Similarly, the context can be decomposed into three latents $\{I(X), A(X), C(X)\}$ corresponding to $\{I, A, C\}$, respectively, and referred to as DRCFR+.

4.2 Adversarial De-confounding

Domain adaptation is the task of building a predictor/classifier model on the source domain while being invariant to the shifts in the target domain, where the source and target domains have similar but different distributions. Johansson et al. (2016); Shalit et al. (2017) presented pioneering work to formulate the ITE estimation problem as domain adaptation and learnt shared representations, i.e., embedding $\phi(X)$, and used integral probability metrics to minimise the distance between control and treatment distributions in the embedding space, thus leading to learning balanced representations to get rid off the selection-bias. However, later VanderWeele (2019); Kuang et al. (2019) pointed out that learning such balanced representations for the context X can lead to additional bias and variance in the predictions.

Here, we adopt another popular idea of adversarial training (Ganin et al. (2016)) from the domain adaptation which is widely adapted in different problem settings, e.g., Bica et al. (2020); Guan and Liu (2021). The main idea of domain adaptation is to obtain invariance between source and target distributions. In our setting, we want invariance of potential outcome regressions μ_t against the treatment policy π , which is also different than the Johansson et al. (2016); Shalit et al. (2017) who didn't consider treatment policy for reducing the selection-bias. Adversarial training helps to learn representations for the confounders such that the representations are discriminative of the potential outcome regressions μ_t but invariant to the treatment policy π . This results in balanced representations for control and

treatment groups, which also avoids the risk of additional bias and variance, as pointed out in VanderWeele (2019); Kuang et al. (2019) from balancing the entire context X .

The balanced representations are achieved by jointly optimising the underlying representations as well as a predictor and a classifier on these representations: (i) potential outcome predictor which is used during the training and test times, and (ii) treatment policy classifier which is used during the training time only. The parameters of the predictor and classifier $\{W_{\mu_0}, W_{\mu_1}, W_t\}$ (refer to figure Fig. 2) are optimised to reduce the training error and parameters of the confounder representation W_C are optimised to be better predictive of potential outcomes, i.e., to *minimise* the potential outcome loss, and to be invariant to the treatment policy, i.e., *maximise* the loss of the treatment classifier and hence push the treatments to be similar to the randomised controlled trials. Thus, the later update works *adversarially* to the treatment classifier, and enforces treatment-invariant balanced representations to the confounders C .

To implement the adversarial training to get the balanced representations for the confounders, a gradient reversal layer is placed between the confounder representations $C(X)$ and the treatment classifier, as shown in Fig. 2. During the forward pass, the layer acts as an identity, but during the backward pass, the layer reverses the direction of the gradient, as the layer's name suggests. A domain adaptation hyperparameter λ is used to control the balancing effect of the gradient reversal layer and is defined as follows.

$$\lambda = \lambda_0 \left(\frac{2}{1 + \exp(-10 \times \text{epoch}/\gamma)} - 1 \right), \quad (8)$$

where epoch is the epoch number, γ controls the variations of λ during the training, and it can take values from one to the total number of epochs. λ can vary from 0 to a large number, which has to be tuned as per the dataset. The constant factor λ_0 is one unless we need λ to be more than one. We observe that the estimator is sensitive to the choice of λ , which controls the balancing effect. Since synthetic and semi-synthetic datasets come with different variants which differ in their outcome surfaces, e.g., the IHDP dataset has 100 variations, it is difficult to find a common value for all the variants which affect the performance of the proposed estimators.

The objective function for the proposed ITE estimator is as given below.

$$\begin{aligned}
 \frac{1}{m} \sum_{i=1}^m [L_y(y_i, \hat{\mu}_t^i(x)) + L_t(t, \hat{\pi}(x))] + \lambda_1 \mathcal{R}_2 \\
 + \lambda_2 (L_O + R_O),
 \end{aligned} \quad (9)$$

where L_y is loss for potential outcomes, L_t is propensity score, R_2 is l_2 -regulariser, and λ_1 and λ_2 are constants.

5 EXPERIMENTS

This section, discusses the experimental settings, metrics, baselines and results on synthetic and real datasets.

5.1 Experimental Setting

Following Curth and van der Schaar (2021b), we fix common hyperparameters for all the ITE estimators. We take a mini-batch of 100, Adam as optimiser, the learning rate of 0.0001, and l_2 -regulariser coefficient of 0.0001. We use early stopping with patience of 50 as stopping criteria (for more details of the hyperparameters and implementation details, please refer to supplementary material). All the experiments are implemented in Pytorch, executed on an Ubuntu machine (64GB RAM, 1 NVIDIA GeForce GPU), and averaged over 10 runs. The code is released and verified at <https://codeocean.com/capsule/1529546>. For baselines, we use the following ITE estimators.

FlexTENet (Curth and van der Schaar (2021a)): It has a flexible architecture that allows flexibility to share information at different layers between potential outcomes for factual and counterfactuals. Moreover, the authors also introduced regularisation-based schemes to encourage similarity between the two outcomes.

SNet (Curth and van der Schaar (2021b)): It decomposes the context X into five latent factors and uses orthogonality to encourage each variable of X to contribute to one of the five latent factors.

DRCRF (Hassanpour and Greiner (2019a); Wu et al. (2022)): It decomposes the context X into three latent factors and uses orthogonality to encourage each variable of X to contribute to one of the three latent factors.

TARNet, CFRNet (Johansson et al. (2016); Shalit et al. (2017)): This pioneering work introduced representation learning and domain adaptation ideas for minimising the distance between control and treatment representations. TARNet, has two heads for the potential outcome regressions, while CFRNet also adds a discrepancy regulariser to TARNet for minimising the distribution distance.

DragonNet, DragonNetTR (Shi et al. (2019)): They used third head for propensity score, in addition to potential outcome regressions (referred to as DragonNet), and introduced targeted regularisation (referred to as DragonNetTR).

SLearner, TLearner (Künzel et al. (2019)): These are basic baselines where the S(single)Learner augments the feature space with the treatment policy and has the advantage of using the whole data, unlike others which use data in potential outcome regression heads as per the treatment group. Similarly, T(two)Learner uses two separate networks for each treatment group and so each network utilises data related to that treatment group only.

TEDVAE (Zhang et al. (2021)): This uses disentangled latent factors in the variational autoencoder.

Moreover, the proposed disentangled representations with adversarial training algorithms are referred to as **SNet+** and **DRCRF+** due to some similarities with corresponding SNet and DRCRF algorithms. In addition, we also present extensions of DragonNet and DragonNetTR using adversarial training, referred to as **DragonNet+**, **DragonNetTR+**. However, adversarial training of the rest of the baselines is not possible and that’s why they don’t have plus-variants.

To assess the performance of the ITE estimators, since we have both factual and counterfactual outcomes in the synthetic and semi-synthetic datasets, we use Precision in the Estimation of Heterogeneous Effects (PEHE) as a metric, as given below.

$$\text{PEHE} = \sqrt{\frac{1}{m} \sum_{i=1}^m (\hat{e}_i - e_i)^2}, \quad (10)$$

where $\hat{e}_i = \hat{\mu}_1^i - \hat{\mu}_0^i$ is the predicted ITE and $e_i = \mu_1^i - \mu_0^i$ is the true ITE.

5.2 Synthetic Experiments

We evaluate the performance of different ITE estimators on simulated data with different sizes and with different number of confounding features. Our data generation is inspired from Curth and van der Schaar (2021b). In all cases, we take $d = 25$ normally distributed covariates, subsets of which decide the potential outcomes $\mu_t(x)$ and the treatments $\pi(x)$. X is generated in disjoint subsets X_s of size d_s , according to $X_s \sim \mathcal{N}(0, 1)$. Suppose, context X is made of covariates X_C, X_O, X_T, X_τ which affect both $\mu_t(x)$ and $\pi(x)$, affect only $\mu_t(x)$, only the treatments $\pi(x)$ and model the treatment effect $e(x)$, respectively. The potential outcomes $\mu_t(x)$ and the treatments $\pi(x)$ are modelled as given below.

$$\mu_0(x) = \mathbf{1}^T X_{CO}^2, \quad (11)$$

$$\mu_1(x) = \mathbf{1}^T X_{CO}^2 + \mathbf{1}^T X_\tau^2, \quad (12)$$

$$\pi(x) = \text{expit} \left(\xi \left(\frac{1}{d_{CT}} \mathbf{1}^T X_{CT}^2 - \omega \right) \right), \quad (13)$$

where $\mathbf{1}$ is a vector of ones, $X_{CO} = [X_C, X_O]$, $X_{CT} = [X_C, X_T]$, ξ controls the selection-bias (in our experiments, it is set to three), and $\omega = \text{median}(\frac{1}{d_{CT}} \mathbf{1}^T X_{CT}^2)$.

In Table 1, we present results on synthetic data, which have 25 covariates of which five, five and, ten covariates affect the outcome, propensity and both, and the rest model the treatment effect, respectively. The dataset has 3000 data points, 30% of which are taken as a test set, and 30% of the remaining are used as a validation set for hyperparameter tuning. The table shows that among the baselines, SNet

Table 1: Comparative study of different ITE estimators on a synthetic dataset. Bold values represent the best values between the extension and the corresponding baselines and PEHE-in and -out refer to PEHE on training data and on the hold-out test dataset. The results are averaged over 10 runs, and the standard error is presented in parentheses.

Name	PEHE-in	PEHE-out
TLearner	2.2822 (0.0041)	2.5233 (0.0130)
SLearner	0.4750 (0.0022)	0.4772 (0.0022)
TARNet	0.5950 (0.0041)	0.5847 (0.0038)
CFRNet	0.5950 (0.0041)	0.5846 (0.0041)
FlexTENet	1.7150 (0.0054)	1.7103 (0.0060)
DragonNet	0.5361 (0.0057)	0.5251 (0.0057)
DragonNet+	0.5414 (0.0066)	0.5220 (0.0063)
DragonNetTR	1.2161 (0.0145)	1.2027 (0.0145)
DragonNetTR+	1.0149 (0.0044)	0.9976 (0.0044)
DRCFR	0.4785 (0.0076)	0.4745 (0.0085)
DRCFR+	0.4742 (0.0063)	0.4688 (0.0076)
SNet	0.4264 (0.0070)	0.4116 (0.0063)
SNet+	0.4147 (0.0066)	0.4008 (0.0073)

and TLearner are the best and worst performers, respectively. This is expected as SNet is one of the recent techniques with a mechanism to separate confounders. On the other hand, TLearner uses two different networks without sharing information and lacks any mechanism to control confounding. Moreover, each of the two neural networks in TLearner also utilises less data as each receives data corresponding to that treatment group only. Interestingly, SLearner performs better than expected. This is likely because of sharing a single network for potential outcomes, which cancels the noise effect and can utilise the entire data for the predictions, unlike the other estimators where each potential outcome predictor utilises the data corresponding to that outcome. Although, SLearner is known to perform poorly in high-dimensional settings.

From the table, it is also clear that adversarial training improves the error in the ITE estimation (see SNet+). It is also observed that the estimators using disentanglement learning (i.e., DRCFR, SNet, and their extensions) perform comparatively better than the other baselines as they have a mechanism to separate the confounders. Moreover, shared representation-based estimators also perform better than the other, i.e., TLearner, which does not share any information between potential outcome predictions.

Fig. 3 compares the performance of different ITE estimators against the scale of the data at 2,000, 3,000, 5,000, and 7,000 data points. For clarity, we removed the ITE estimators with large values, like TLearner, to avoid crowding in the figure (for a detailed comparison, please refer to the supplementary material). From the figure, it is clear that as the data size increases, in general, the error goes

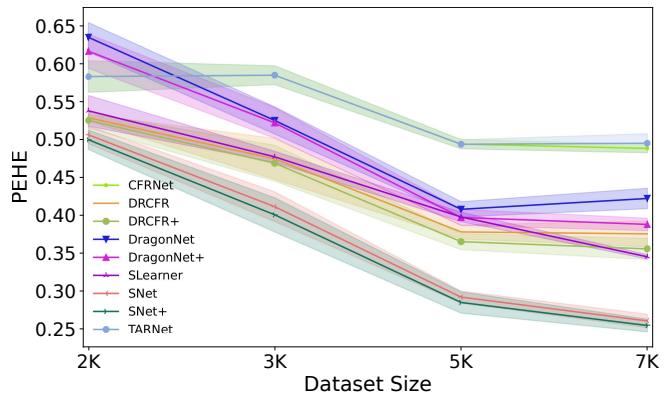


Figure 3: PEHE for different ITE estimators by dataset sizes (shaded area presents one standard deviation).

down. Moreover, the proposed SNet+ is always the best performer, and all the extensions are better than their baselines. It is observed that TARNet and CFRNet perform better than DragonNet and DragonNetTR with small datasets, but as the dataset size increases, they lag behind. Moreover, SNet/SNet+ improve with dataset size compared with the rest of the baselines.

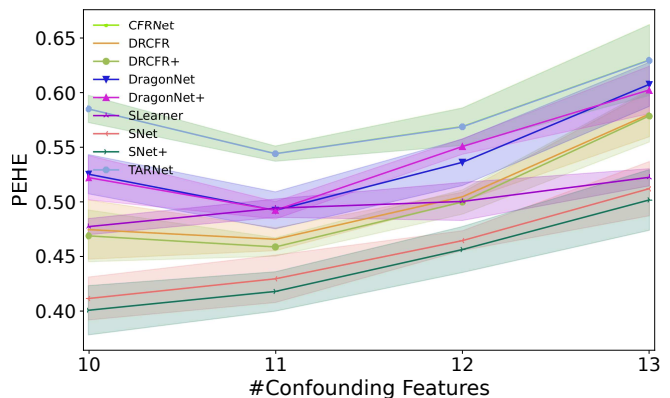


Figure 4: PEHE for different ITE estimators by different number of confounding covariates in the context X (shaded area presents one standard deviation).

Fig. 4 presents the performance of the ITE estimators for an increasing number of confounding features $\{10, 11, 12, 13\}$ with a dataset of 3000 points. All these datasets share the same covariates, i.e., X is the same, although the outcomes and treatment assignment policy are calculated per the data generating process discussed earlier. From the figure, it is clear that, in general, the error increases as the number of confounding covariates increases; however, some estimators decrease at the beginning. Moreover, the proposed SNet+ and other extensions of the baselines maintain their performance and show better results. SNet+ is also the best estimator throughout and TARNet is the worst performer. As observed earlier, SLearner performs well in these settings, although it is not suited for

high-dimensional settings. Once again, the disentanglement learning-based estimators perform better than other baselines.

5.3 Semi-synthetic Benchmark: IHDP

The Infant Health and Development Program (IHDP) Hill (2011) is a semi-simulated benchmark dataset for causal inference for binary treatment problems, where the covariates are real, but outcomes are simulated. It consists of 747 units (139 treated and 608 untreated) and 25 covariates measuring aspects of mother and child. The dataset was prepared from a randomized control trial from the Infant Health and Development Program, which studies the effect of high-quality child care/specialist home visits on future cognitive test scores of the child. A biased set of treated units is removed to make the dataset imbalanced. The dataset provides noiseless true values for both the treatments, i.e., treatment and control, so it can be used to calculate the true ITE. We take IHDP100² with 100 simulation settings from Shalit et al. (2017) for our study. We have used a validation set as 20% of the train set. For more details on hyperparameters, please refer to the supplementary material.

Table 2: Comparative study of different ITE estimators on IHDP dataset. Bold values represent the best values between the extension and the corresponding baselines and PEHE-in, and -out refer to PEHE on training data and on hold-out test dataset. The results are averaged over 10 runs and the standard error presented in parentheses.

Name	PEHE-in	PEHE-out
TLearner	1.1007 (0.0234)	1.3686 (0.0211)
SLearner	2.5331 (0.0120)	2.8332 (0.0101)
TARNet	0.7837 (0.0114)	1.0935 (0.0095)
CFRNet	0.8165 (0.0019)	1.1277 (0.0038)
FlexTENet	1.1238 (0.0130)	1.3954 (0.0123)
TEDVAE	1.5004 (0.0060)	1.6018 (0.0085)
DragonNet	0.8026 (0.0054)	1.1097 (0.0041)
DragonNet+	0.7924 (0.0104)	1.1125 (0.0111)
DragonNetTR	1.9190 (0.0262)	2.1935 (0.0256)
DragonNetTR+	1.7087 (0.0322)	1.9744 (0.0338)
DRCFR	0.8680 (0.0190)	1.1908 (0.0155)
DRCFR+	0.8616 (0.0142)	1.1897 (0.0095)
SNet	1.0586 (0.0066)	1.4030 (0.0060)
SNet+	1.0459 (0.0082)	1.3839 (0.0057)

Table 2 presents experiments with the semi-simulated IHDP dataset. From the table, it is clear that all the extensions of different ITE estimators with adversarial training improve the baseline results. However, disentanglement representation-based estimators are no longer the

best performer. This is in line with earlier observations Curth and van der Schaar (2021b) that the simple estimators perform better than the complex estimators on the IHDP dataset, as it is a very small dataset and the *overlap* assumption holds only partially. TARNet and SLearner are the best and the worst performers, respectively.

The computational cost for the proposed approach is similar to the existing baselines, as the gradient reversal layer acts as an identity in the forward pass and only reverses the gradient during the backward pass by a factor called domain adaptation parameter. So, the proposed approach does not add any additional computational cost.

6 CONCLUSION

This paper proposes disentangled representations with adversarial training to selectively balance the confounders in the binary treatment setting. We also extended the ITE estimators having confounder representations using adversarial training to balance the confounders. Using synthetic and semi-synthetic benchmarks, we showed that the proposed idea improves the error in the ITE estimation. One limitation of the proposed approach is that it introduces an additional hyperparameter, called domain adaptation parameter, which has to be tuned as per the dataset to control the balancing of confounders. However, it can be tuned on the validation dataset along with other machine learning hyperparameters. In this paper, in line with some of the literature, we studied the proposed idea in the plug-in setting for the ITE estimation. Furthermore, it would be interesting to explore this concept in the two-step meta-learners' setting.

Acknowledgments

This work was supported in part by the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre (BRC) and in part by InnoHK Project Programme 3.2: Human Intelligence and AI Integration (HIAI) for the Prediction and Intervention of CVDs: Warning System at Hong Kong Centre for Cerebro-cardiovascular Health Engineering (COCHE). VKC was supported by a Medical Research Council (MRC) Research Grant (MR/W01761X/1). TZ was supported by an RAEng Engineering for Development Research Fellowship. TZ was supported by the RAEng Engineering for Development Research Fellowship. DAC was supported by an NIHR Research Professorship, an RAEng Research Chair, the InnoHK Hong Kong Centre for Cerebro-cardiovascular Health Engineering (COCHE), and the Pandemic Sciences Institute at the University of Oxford. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR, the Department of Health, the InnoHK – ITC, or the University of Oxford.

²<http://www.fredjo.com/>

References

- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- Ioana Bica, Ahmed M Alaa, James Jordon, and Mihaela van der Schaar. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. *arXiv preprint arXiv:2002.04083*, 2020.
- Ioana Bica, Ahmed M Alaa, Craig Lambert, and Mihaela Van Der Schaar. From real-world patient data to individualized treatment effects using machine learning: current and future methods to address underlying challenges. *Clinical Pharmacology & Therapeutics*, 109(1):87–100, 2021.
- Richard K Crump, V Joseph Hotz, Guido W Imbens, and Oscar A Mitnik. Dealing with limited overlap in estimation of average treatment effects. *Biometrika*, 96(1):187–199, 2009.
- Alicia Curth and Mihaela van der Schaar. On inductive biases for heterogeneous treatment effect estimation. *Advances in Neural Information Processing Systems*, 34:15883–15894, 2021a.
- Alicia Curth and Mihaela van der Schaar. Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms. In *International Conference on Artificial Intelligence and Statistics*, pages 1810–1818. PMLR, 2021b.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- Hao Guan and Mingxia Liu. Domain adaptation for medical image analysis: a survey. *IEEE Transactions on Biomedical Engineering*, 69(3):1173–1185, 2021.
- Jens Hainmueller. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political analysis*, 20(1):25–46, 2012.
- Negar Hassanpour and Russell Greiner. Counterfactual regression with importance sampling weights. In *IJCAI*, pages 5880–5887, 2019a.
- Negar Hassanpour and Russell Greiner. Learning disentangled representations for counterfactual regression. In *International Conference on Learning Representations*, 2019b.
- Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Paul W Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.
- Guido W Imbens. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics*, 86(1):4–29, 2004.
- 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 Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International conference on machine learning*, pages 3020–3029. PMLR, 2016.
- Fredrik D Johansson, Uri Shalit, Nathan Kallus, and David Sontag. Generalization bounds and representation learning for estimation of potential outcomes and causal effects. *arXiv preprint arXiv:2001.07426v3*, 2022.
- Kun Kuang, Peng Cui, Bo Li, Meng Jiang, Shiqiang Yang, and Fei Wang. Treatment effect estimation with data-driven variable decomposition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- Kun Kuang, Peng Cui, Bo Li, Meng Jiang, Yashen Wang, Fei Wu, and Shiqiang Yang. Treatment effect estimation via differentiated confounder balancing and regression. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 14(1):1–25, 2019.
- Sören R Künzel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10):4156–4165, 2019.
- Fan Li, Kari Lock Morgan, and Alan M Zaslavsky. Balancing covariates via propensity score weighting. *Journal of the American Statistical Association*, 113(521):390–400, 2018.
- Sheng Li and Yun Fu. Matching on balanced nonlinear representations for treatment effects estimation. *Advances in Neural Information Processing Systems*, 30, 2017.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.

Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pages 3076–3085. PMLR, 2017.

Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems*, 32, 2019.

Tyler J VanderWeele. Principles of confounder selection. *European journal of epidemiology*, 34(3):211–219, 2019.

Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.

Anpeng Wu, Junkun Yuan, Kun Kuang, Bo Li, Runze Wu, Qiang Zhu, Yue Ting Zhuang, and Fei Wu. Learning decomposed representations for treatment effect estimation. *IEEE Transactions on Knowledge and Data Engineering*, 2022.

Liuyi Yao, Sheng Li, Yaliang Li, Mengdi Huai, Jing Gao, and Aidong Zhang. Representation learning for treatment effect estimation from observational data. *Advances in Neural Information Processing Systems*, 31, 2018.

Weijia Zhang, Lin Liu, and Jiuyong Li. Treatment effect estimation with disentangled latent factors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10923–10930, 2021.

José R Zubizarreta. Stable weights that balance covariates for estimation with incomplete outcome data. *Journal of the American Statistical Association*, 110(511):910–922, 2015.

A ADDITIONAL EXPERIMENTS

A.1 Scale and Confounding Comparison for In-distributions

In the main part of the paper, we presented results for scale and confounding feature comparisons for out-distributions only. Here, we presents results for in-distributions, and for the sake of clarity, we have also presented out-distributions results side-by-side to in-distributions. From the Figs. 5 and 6, it is clear that in- and out-distribution results are similar and have same observations, as discussed earlier.

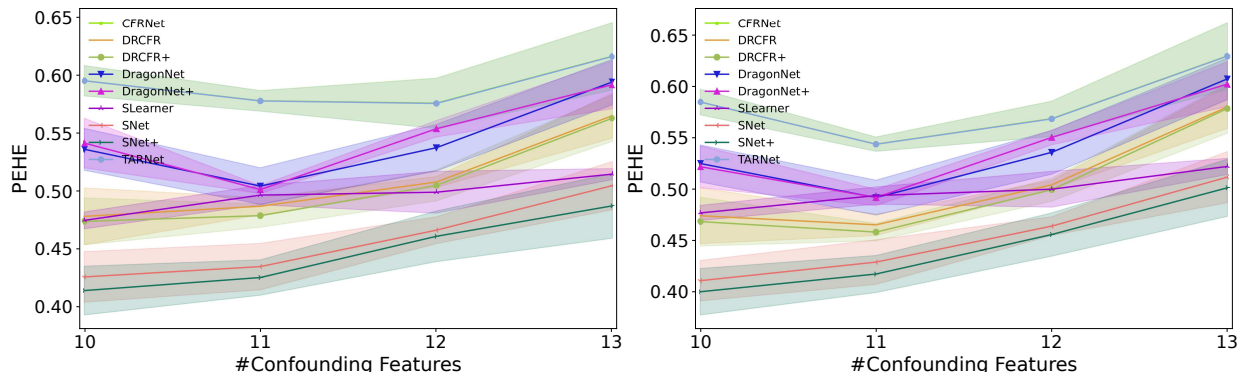


Figure 5: Comparative study for in-distribution (left) and out-distribution (right) PEHE against number of confounding features.

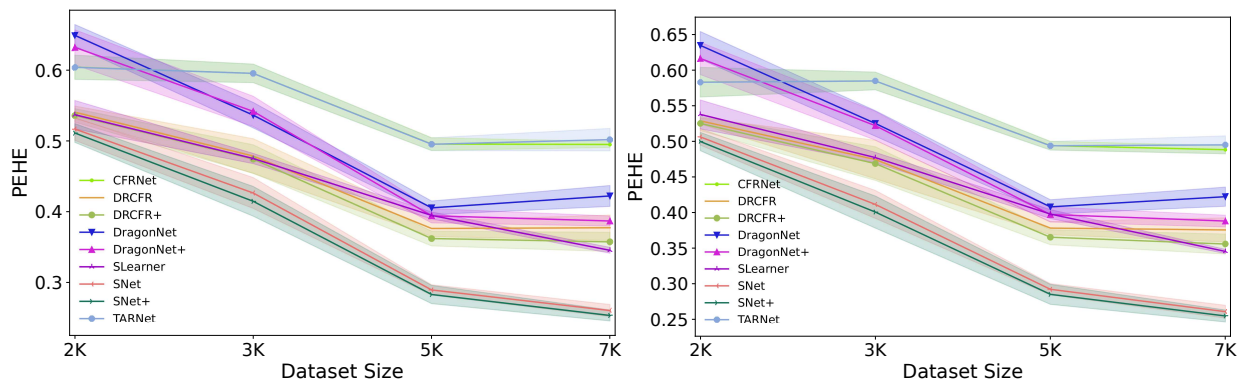


Figure 6: Comparative study for in-distribution (left) and out-distribution (right) PEHE against dataset size.

A.2 Results with coefficients

In the literature, either no coefficients are used to balance the potential outcome loss L_O and the propensity score loss L_T or coefficients are always one for the both. Although, we observed that by putting coefficients around the the two loss terms can help reduce the ITE estimation error. So, in this Subsection, we present results for IHDP dataset using the coefficients. To find the coefficients, we use manual tuning as follows.

$$L = \alpha L_O + (1 - \alpha) L_T + \text{other terms.} \tag{14}$$

We start with a value of 0.5 for α and reduce/increase by 0.1 as long as it keeps improving, in the direction of improvement, and if it reaches 0.1 then to further decrease the value, we divide by 10, but if it reaches 0.9, we add extra digit 9. Table 3 presents results on IHDP dataset with coefficients and the values of α are 0.6, 0.6, 0.99 and 0.99 for SNet/+, DRCFR/+, DragonNet/+ and DragonNetTR/+, respectively. For some of the baselines, there is no difference in the performance because they don't have propensity loss term so no coefficients are used for them. This is especially noted for TARNet

which was earlier the best performer but with coefficients DragonNet+ performs the best. Although, the results still follow the observation from the literature that the simple ITE estimators perform better on IHDP dataset than the complex estimators. Moreover, the proposed extensions are, as observed earlier, always improve the baselines. Following the results with coefficients for IHDP dataset, we used coefficients for reporting results with synthetic data also. Although, to save the overhead to tune the coefficient, we have tuned the coefficients for setting with 3,000 data points only as they are not expected to vary much with change in data points, and values of α are used as 0.2, 0.5, 0.1 and 0.7 for SNet+, DRCFR+, DragonNet/+ and DragonNetTR/+, respectively.

Table 3: Comparative study of different ITE estimators with coefficients on IHDP dataset. Bold values represent the best values between the extension and the corresponding baselines and PEHE-in, and -out refer to PEHE on training data and on hold-out test dataset. The results are averaged over 10 runs and the standard error presented in parentheses.

Name	PEHE-in	PEHE-out
TLearner	1.1007 (0.0234)	1.3686 (0.0212)
SLearner	2.5331 (0.0120)	2.8332 (0.0101)
TARNet	0.7837 (0.0114)	1.0935 (0.0095)
CFRNet	0.8165 (0.0019)	1.1277 (0.0038)
FlexTENet	1.1238 (0.0130)	1.3954 (0.0123)
DragonNet	0.7314 (0.0076)	1.0563 (0.0070)
DragonNet+	0.7269 (0.0092)	1.0478 (0.0111)
DragonNetTR	1.8887 (0.0152)	2.1403 (0.0256)
DragonNetTR+	1.6696 (0.0174)	1.9452 (0.0123)
DRCFR	0.8654 (0.0171)	1.190 (0.0136)
DRCFR+	0.8418 (0.0158)	1.179 (0.0114)
SNet	0.99 (0.0133)	1.30 (0.0107)
SNet+	0.99 (0.0044)	1.29 (0.0107)

A.3 Comparison of Treatment Distributions

Here, we take a random example from synthetic data, to compare the treatment distributions of SNet and SNet+, and a simple baseline of DragonNet. From the Fig. 7, we can infer two points: first that DragonNet performs too much balancing which affects the predictive performance. This is due to the existence of a trade-off in balancing and predictive performance, and balancing the whole context X as pointed out in VanderWeele (2019); Kuang et al. (2019). Secondly, SNet+ improves the treatment probabilities and the density is slightly narrower than the SNet.

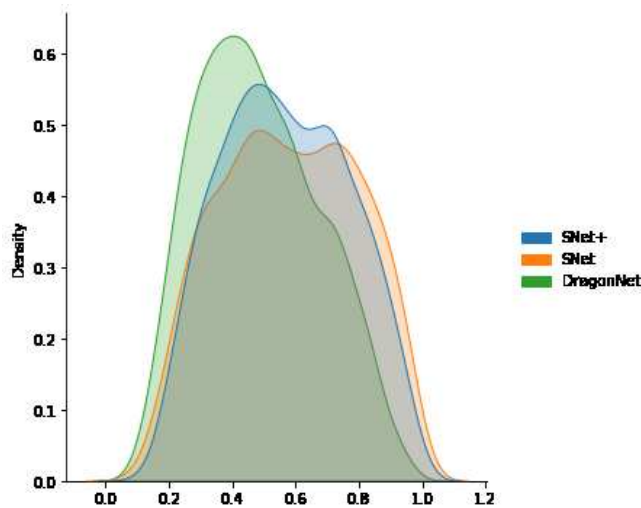


Figure 7: Distribution of treatment classification on synthetic data for DragonNet, SNet and SNet+, with PEHE-out as 0.6066, 0.5728 and 0.5677, respectively.

B IMPLEMENTATION DETAILS

In our implementation, for fair comparison, we use components similar to Curth and van der Schaar (2021b). We set common hyperparameters for all the ITE estimators. We take a mini-batch of 100, Adam as optimiser, the learning rate of 0.0001, and l_2 -regulariser coefficient (λ_1) of 0.0001. Each experiment is run for a maximum of 1000 epochs. We use exponential linear units (ELU) as the activation function for the dense layers and early stopping based on 30% validation (20% for IHDP dataset due to its small size) dataset with patience of 50 as stopping criteria. For Slearner, We have used five layers with 200, 200, 200, 100 and 100 neurons, in addition to the final output layer. For TLearner, we have used five layers with 100, 100, 100, 50 and 50 neurons, in each of the two networks, in addition to the final output layers. For TARNet, CFRNet (maximum mean discrepancy), DragonNet/DragonNet+ and DragonNetTR/DragonNetTR+, we have used five layers with 200, 200 and 200 neurons in the common representation layers and 100, 100 in the potential outcome heads, in addition to the final output layer. For DRCFR/DRCFR+, we have used three representation layers with 150, 50 and 50 neurons for latent factor representations of confounder, outcome and propensity, and two layers of 100 neurons in the output heads, in addition to the final output layer. For FlexTENet, we have used two layers with 100 neurons and two layers of 50 neurons in the outcome heads, in addition to the output layer. Moreover, as suggested in the paper, we have set the $\lambda_1 = 0.0001$, $\lambda_2 = 0.01$ and $\lambda_o = 0.1$. For SNet/SNet+, we have used 100, 100, 50, 50 and 50 neurons in the latent factor representation layers for confounder, treatment, outcome, treatment outcome and control outcome, respectively, and two layers in the outcome heads, in addition to the final output layers. We have set $\lambda_2 = 0$ for IHDP and $\lambda_2 = 0.01$ for DRCFR/DRCFR+ and SNet/SNet+. Our implementations are based on Curth and van der Schaar (2021b,a) which are available online³ in JAX as well as Pytorch. For TEDVAE, we have used the official implementation and the settings from the paper (Zhang et al. (2021)).

Table 4: Domain adaptation parameter settings – λ_0, γ (where w5-c10-o5-3K means 5, 10 and 5 covariates determine treatment, confounders and potential outcomes with dataset size of 3000 data points. IHDP-coeff refers to IHDP dataset with coefficients for the loss terms.)

Dataset	SNet+	DRCFR+	DragonNet+	DragonNetTR+
IHDP	18, 1	1.2, 1	0.2, 600	0.001, 1
IHDP-coeff	1, 1	1, 1	0.001, 1	0.001, 1
w5-c10-o5-2K	1.4, 1	1.8, 10	1.5, 1	0.2, 600
w5-c10-o5-3K	1.7, 1	1.8, 1	4, 1	1, 600
w5-c10-o5-5K	1.7, 1	8, 1	4, 1	1, 600
w5-c10-o5-7K	1.7, 1	15, 1	4, 1	1.5, 600
w5-c11-o5-3K	1.7, 1	4, 1	3, 1	1.5, 1
w5-c12-o5-3K	5, 400	1, 300	3, 1	2, 300
w5-c13-o5-3K	0.6, 600	5, 1	3, 1	0.8, 600

The domain adaptation parameter λ for the adversarial trained ITE estimators is also tuned manually. This requires setting the constant factor λ_0 and the parameter controlling the rate of growth γ . We vary γ between $\{1, 10, 100, 200, 300, 400, 600, 700\}$ and λ_0 is varied between 0.1 to 25, where large values force stronger balancing of the confounder variables. The final selected values for these parameters are presented in the Table 4. We observed that, the training is sensitive to λ , which has to be tuned to each dataset, because large values of λ can lead to divergence of the learning algorithm. Since the datasets have multiple variants, e.g., IHDP has 100 so it was difficult to find a common setting for all the 100 variants because some parameters were working for variant but not for others. So, tuning of λ for each variant on the validation dataset can give even better results at the expense of extra overhead to tune λ . Moreover, we observed that DragonNetTR+ was more sensitive to λ and at sometimes the learning algorithm diverged so we have removed those variants from the evaluations – including for the DragonNetTR. However, we can use gradient clipping to solve this issue of unstable learning.

³<https://github.com/AliciaCurth/CATENets>