

DR-VIDAL - Doubly Robust Variational Information-theoretic Deep Adversarial Learning for Counterfactual Prediction and Treatment Effect Estimation on Real World Data

Shantanu Ghosh, MS¹, Zheng Feng, PhD², Jiang Bian,² Kevin Butler,² Mattia Proserpi,²
¹ University of Pittsburgh, Pittsburgh, PA; ² University of Florida, Gainesville, FL

Abstract

Determining causal effects of interventions onto outcomes from real-world, observational (non-randomized) data, e.g., treatment repurposing using electronic health records, is challenging due to underlying bias. Causal deep learning has improved over traditional techniques for estimating individualized treatment effects (ITE). We present the Doubly Robust Variational Information-theoretic Deep Adversarial Learning (DR-VIDAL), a novel generative framework that combines two joint models of treatment and outcome, ensuring an unbiased ITE estimation even when one of the two is misspecified. DR-VIDAL integrates: (i) a variational autoencoder (VAE) to factorize confounders into latent variables according to causal assumptions; (ii) an information-theoretic generative adversarial network (Info-GAN) to generate counterfactuals; (iii) a doubly robust block incorporating treatment propensities for outcome predictions. On synthetic and real-world datasets (Infant Health and Development Program, Twin Birth Registry, and National Supported Work Program), DR-VIDAL achieves better performance than other non-generative and generative methods. In conclusion, DR-VIDAL uniquely fuses causal assumptions, VAE, Info-GAN, and doubly robustness into a comprehensive, performant framework. Code is available at: <https://github.com/Shantanu48114860/DR-VIDAL-AMIA-22> under MIT license.

Introduction

Understanding causal relationships and evaluating effects of interventions to achieve desired outcomes is key to progress in many fields, especially in medicine and public health. A typical scenario is to determine whether a treatment (e.g., a lipid-lowering medication) is effective to reduce the risk of or cure an illness (e.g., cardiovascular disease). Randomized controlled trials (RCTs) are considered the best practice for evaluating causal effects¹. However, RCTs are not always feasible, due to ethical or operational constraints. For instance, if one wanted to evaluate whether college education is the cause of good salary, it would not be ethical to randomly pick teenagers and randomize their admission to college. So, in many cases, the only usable data sources are observational data, i.e., real-world data collected retrospectively and not randomized. Unfortunately, observational data are often plagued with various biases—since the data generation processes are largely unknown—such as confounding (i.e., spurious causal effects on outcomes by features that are correlated with a true unmeasured cause) and colliders (i.e., mistakenly including effects of an outcome as predictors), making it difficult to infer causal claims². Another problem is that, in both RCTs and observational datasets, only factual outcomes are available, since clearly an individual cannot be treated and non-treated at the same time. Counterfactuals are alternative predictions that respond to the question “*what outcome would have been observed if a person had been given a different treatment?*” If models are biased, counterfactual predictions can be wrong, and interventions can be ineffective or harmful³. In both RCT-based and real-world based studies, two types of treatment effects are usually considered: (i) the average treatment effect (ATE), which is population-based and represents the difference in average treatment outcomes between the treatment and controls; and (ii) the individualized treatment effect (ITE), which represents the difference in treatment outcomes for a single observational unit with the same background covariates⁴. When there is suspected heterogeneity, stratified ATEs, or conditional ATEs, can be calculated. Traditional statistical approaches for estimating treatment effects, taking into account possible bias from pre-treatment characteristics, include propensity score matching (PSM) and inverse probability weighting (IPW)⁵. The propensity score is a scalar estimate representing the conditional probability of receiving the treatment, given a set of measured pre-treatment covariates. By matching (or weighting) treated and control subjects according to their propensity score, a balance in pre-treatment covariates is induced, mimicking a randomization of the treatment assignment. However, the PSM approach only accounts for measured covariates, and latent bias may remain after matching⁶. PSM has been historically implemented with logistic-linear regression, coupled with different feature selection methods in the presence of high-dimensional datasets⁷. A problem with PSM is that it often decreases the sample size due

to matching, while IPW can be affected by skewed, heavy-tailed weight distributions. Machine learning approaches have been introduced more recently, e.g., Bayesian additive regression trees⁸ and counterfactual random forests⁹. Big data also led to the flourishing of causal deep learning¹⁰. Notable examples include the Treatment-Agnostic Representation Network (TARNet)¹¹, Dragonnet¹², Deep Counterfactual Network with Propensity-Dropout (DCN-PD)¹³, Generative Adversarial Nets for inference of Individualized Treatment Effects (GANITE)¹⁴, Causal Effect Variational Autoencoder (CEVAE)¹⁵, and Treatment Effect by Disentangled Variational AutoEncoder (TEDVAE)¹⁶.

Contribution This work introduces a novel deep learning approach for ITE estimation and counterfactual prediction on real-world observational data, named the *Doubly Robust Variational Information-theoretic Deep Adversarial Learning* (DR-VIDAL). Motivated from Makhzani *et al.*¹⁷, we use a lower-dimensional neural representation of the input covariates to generate counterfactuals to improve convergence. We assume a causal graph on top of the covariates where the covariates X are generated from 4 independent latent variables Z_t, Z_{ycf}, Z_{yf} and Z_x indicating latents for treatment, counterfactual, factual outcomes and observed covariates respectively, shown in Figure 1. In generating the representations, we use a variational autoencoder (VAE) to infer the latent variables from the covariates in unsupervised manner and feed the learned lower-dimensional representation from the VAE to a generative adversarial network (GAN). Also, to counter the loss of the predictive information while generating the counterfactuals, we aim to maximize the mutual information between the learned representations and the output of the generator. We add this as a regularizer to the generator loss to obtain more robust counterfactuals. Finally, we incorporate a doubly robust network head to estimate the ITE, improving in loss convergence. As DR-VIDAL generates the counterfactual outcomes, we minimise the supervised loss for both the factual and the counterfactual outcomes to estimate ITE more accurately.




Figure 1: Directed acyclic graph modeling the causal relationships among treatment t , outcome y and pre-treatment covariates X , under the latent space Z .

The main features of DR-VIDAL are, in summary:

- Incorporation of an underlying causal structure where the observed pre-treatment covariate set X is decomposed into four independent latent variables $Z_t, Z_X, Z_{yf}, Z_{ycf}$, inducing confounding on both the treatment and the outcome (Figure 1).
- Latent variables are inferred using a VAE¹⁸.
- A GAN¹⁹ with variational information maximization²⁰ generates (synthetic) complete tuples of covariates, treatment, factual and counterfactual outcomes.
- Individual treatment effects are estimated on complete datasets with a downstream, four-headed deep learning block which is doubly robust^{21,22}.

To our knowledge, this is the first time in which VAE, GAN, information theory and doubly robustness are amalgamated into a counterfactual prediction method. By performing test runs on synthetic and real-world datasets (Infant Health and Development Program, Twin Birth Registry, and National Supported Work Program), we show that DR-VIDAL can outperform a number of state-of-art tools for estimating ITE. DR-VIDAL is implemented in Pytorch and the code is available at: <https://github.com/Shantanu48114860/DR-VIDAL-AMIA-22> under MIT license. In the repository, we also provide an online technical supplement (OTS) with full details on the architectural design, derivation of equations, and additional experimental results.

Problem Formulation

We use the *potential outcomes* framework^{23,24}. Let us consider a treatment t (binary for ease of reading, but the theory can be extended to multiple treatments) that can be prescribed to a population sample of size N . The individuals are characterized by a set of pre-treatment background covariates \mathbf{X} , and a health outcome Y is measured after treatment. We define each subject i with the tuple $\{\mathbf{X}, T, Y\}_{i=1}^N$, where Y_i^0 and Y_i^1 are the potential outcomes when applying treatments $T_i = 0$ and $T_i = 1$, respectively. The ITE $\tau(\mathbf{x})$ for subject i with pre-treatment covariates $\mathbf{X}_i = \mathbf{x}$, is

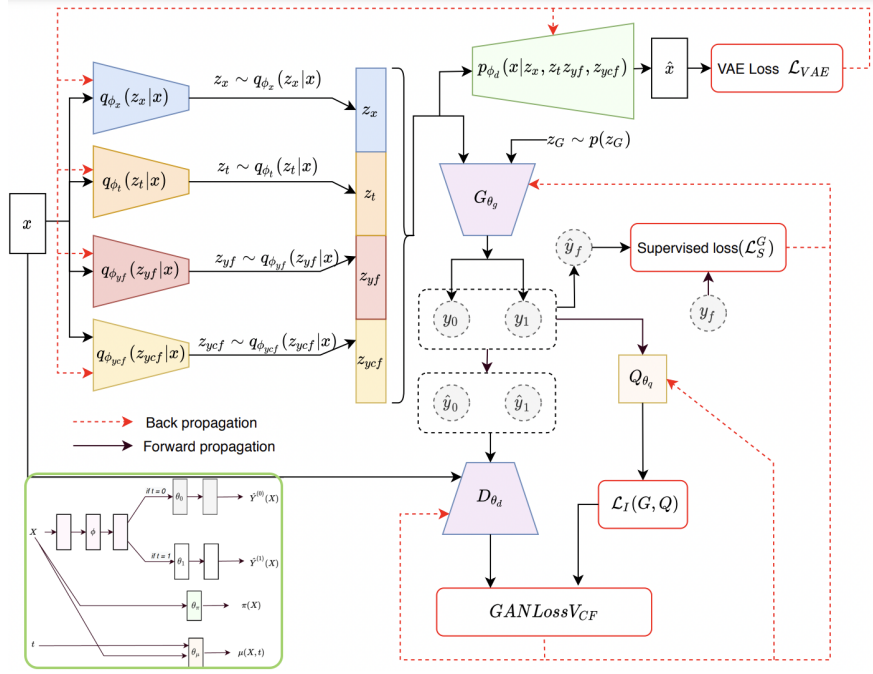


Figure 2: Architecture of DR-VIDAL incorporating the variational autoencoder inferring the latent space (VAE), the generative adversarial network for calculating the counterfactual outcomes (GAN), and the doubly robust module (green box) for estimating ITE.

defined as the difference in the average potential outcomes under both treatment interventions (i.e., treated vs. not treated), conditional on \mathbf{x} , i.e.,

$$\tau(\mathbf{x}) = \mathbb{E}[Y_i^1 - Y_i^0 \mid \mathbf{X}_i = \mathbf{x}] \quad (1)$$

The ITE cannot be calculated directly give the inaccessibility of both potential outcomes, as only factual outcomes can be observed, while the others (counterfactuals) can be considered as missing values. However, when the potential outcomes are made independent of the treatment assignment, conditionally on the pre-treatment covariates, i.e., $\{Y^1, Y^0\} \perp T \mid \mathbf{X}$, the ITE can then be estimated as $\tau(\mathbf{x}) = \mathbb{E}[Y^1 \mid T = 1, \mathbf{X} = \mathbf{x}] - \mathbb{E}[Y^0 \mid T = 0, \mathbf{X} = \mathbf{x}] = \mathbb{E}[Y \mid T = 1, \mathbf{X} = \mathbf{x}] - \mathbb{E}[Y \mid T = 0, \mathbf{X} = \mathbf{x}]$. Such an assumption is called the strongly ignorable treatment assignment (SITA) assumption^{25,26}. By further averaging over the distribution of \mathbf{X} , the ATE τ_{01} can be calculated as

$$\tau_{01} = \mathbb{E}[\tau(\mathbf{X})] = \mathbb{E}[Y \mid T = 1] - \mathbb{E}[Y \mid T = 0] \quad (2)$$

ITE and ATE can be calculated with stratification matching of \mathbf{x} in treatment and control groups, but the calculation becomes unfeasible as the covariate space increases in dimensions. The propensity score $\pi(x)$ represents the probability of receiving the treatment $T = 1$ conditioned on the pre-treatment covariates $X = x$, denoted as $\pi(\mathbf{x}) = P(T = 1 \mid \mathbf{X} = \mathbf{x})$ ²⁴. The propensity score can be calculated using a regression function, e.g., logistic. ITE/ATE can then be calculated by matching (PSM) or weighting (IPW) instances through $\pi(\mathbf{x})$, in a doubly robust way²⁷, or through myriad approaches^{28,9,29,30,27,31,32,33}. In the next section, we describe approaches based on deep learning.

Related Work

Alaa and Van der Schaar³⁴ characterized the conditions and the limits of treatment effect estimation using deep learning. The sample size plays an important role, e.g., estimations on small sample sizes are affected by selection bias, while on large sample sizes, they are affected by algorithmic design. Our work builds up on the ITE estimation approaches of CEVAE¹⁵, DCN-PD¹³, Dragonnet¹², GANITE¹⁴, TARNet¹¹, and TEDVAE¹⁶. DCN-PD is a doubly robust, multitask network for counterfactual prediction, where propensity scores are used to determine a dropout probability of samples to regularize training, carried out in alternating phase, using treated and control batches. CEVAE

uses VAE to identify latent variables from an observed pre-treatment vector and to generate counterfactuals. TARNet aims to provide an upper bound effect estimation by balancing the distributions of treated and controls –with a weight indemnifying group imbalance– within a high dimensional covariate space, but it does not exploit counterfactuals, and only minimises the factual loss function. Dragonnet is a modified TARNet with targeted regularization based on propensity scores. GANITE generates proxies of counterfactual outcomes from covariates and random noise using a GAN, and feeds them to an ITE generator. For both GANITE and TARNet, in presence of high-dimensional data, the loss could be hard to converge. TEDVAE¹⁶ uses a variational autoencoder to infer hidden latent variables from proxies using a causal graph similar to CEVAE. In the next sections, we discuss in detail the novelty of DR-VIDAL and the differences in the architectural design and training mechanisms with respect to the aforementioned approaches.

Proposed Methodology

DR-VIDAL architecture can be summarized in three components: (1) a VAE inferring the latent space, (2) a GAN generating the counterfactual outcomes, and (3) a doubly robust module estimating ITE. The architectural layout is schematized in Figure 2, while the algorithmic details are given in the OTS.

Latent variable inference with VAE. We assume that the observed covariates $\mathbf{X} = \mathbf{x}$ with treatment assignment $T = t$ factual and counterfactual outcomes $Y_f = y_f$ and $Y_{cf} = y_{cf}$ respectively, are generated from an independent latent space \mathbf{z} , composed by $\mathbf{z}_x \sim p(\mathbf{z}_x)$, $z_t \sim p(z_t)$, $\mathbf{z}_{yf} \sim p(\mathbf{z}_{yf})$, and $\mathbf{z}_{ycf} \sim p(\mathbf{z}_{ycf})$, which denote the latent variables for the covariates \mathbf{x} , treatment indicator t , and factual outcomes y_f and y_{cf} , respectively. This decomposition follows the causal structure shown in Figure 1. The goal is to infer the posterior distribution $p(\mathbf{z}_x, \mathbf{z}_t, \mathbf{z}_{yf}, \mathbf{z}_{ycf} | \mathbf{x})$, which is harder to optimize. We use the theory of variational inference³⁵ to learn the variational posteriors $q_{\phi_x}(\mathbf{z}_x | \mathbf{x})$, $q_{\phi_t}(z_t | \mathbf{x})$, $q_{\phi_{yf}}(\mathbf{z}_{yf} | \mathbf{x})$, $q_{\phi_{ycf}}(\mathbf{z}_{ycf} | \mathbf{x})$, using 4 different neural network encoders with parameters $\phi_x, \phi_t, \phi_{yf}$, and ϕ_{ycf} , respectively. Using the latent factors sampled from the learned variational posteriors, we reconstruct \mathbf{x} by estimating the likelihood $p_{\phi_d}(\mathbf{x} | \mathbf{z}_x, \mathbf{z}_t, \mathbf{z}_{yf}, \mathbf{z}_{ycf})$ via a single decoder parameterized by ϕ_d . The latent factors, assumed to be Gaussian, are defined as follows:

$$p(\mathbf{z}_x) = \prod_{i=1}^{D_{z_x}} \mathcal{N}(z_{x_i} | 0, 1); \quad p(\mathbf{z}_t) = \prod_{i=1}^{D_{z_t}} \mathcal{N}(z_{t_i} | 0, 1) \quad (3)$$

$$p(\mathbf{z}_{yf}) = \prod_{i=1}^{D_{z_{yf}}} \mathcal{N}(z_{yf_i} | 0, 1); \quad p(\mathbf{z}_{ycf}) = \prod_{i=1}^{D_{z_{ycf}}} \mathcal{N}(z_{ycf_i} | 0, 1) \quad (4)$$

where $D_{z_x}, D_{z_t}, D_{z_{yf}}, D_{z_{ycf}}$ are the dimensions of the latent factors $\mathbf{z}_x, \mathbf{z}_t, \mathbf{z}_{yf}, \mathbf{z}_{ycf}$, respectively. The variational posteriors of the inference of models are defined as:

$$q_{\phi_x}(\mathbf{z}_x | \mathbf{x}) = \prod_{i=1}^{D_{z_x}} \mathcal{N}(\mu = \hat{\mu}_x, \sigma^2 = \hat{\sigma}_x^2) \quad (5)$$

$$q_{\phi_t}(z_t | \mathbf{x}) = \prod_{i=1}^{D_{z_t}} \mathcal{N}(\mu = \hat{\mu}_t, \sigma^2 = \hat{\sigma}_t^2) \quad (6)$$

$$q_{\phi_{yf}}(\mathbf{z}_{yf} | \mathbf{x}) = \prod_{i=1}^{D_{z_{yf}}} \mathcal{N}(\mu = \hat{\mu}_{yf}, \sigma^2 = \hat{\sigma}_{yf}^2) \quad (7)$$

$$q_{\phi_{ycf}}(\mathbf{z}_{ycf} | \mathbf{x}) = \prod_{i=1}^{D_{z_{ycf}}} \mathcal{N}(\mu = \hat{\mu}_{ycf}, \sigma^2 = \hat{\sigma}_{ycf}^2) \quad (8)$$

where $\hat{\mu}_x, \hat{\mu}_t, \hat{\mu}_{yf}, \hat{\mu}_{ycf}$ and $\hat{\sigma}_x^2, \hat{\sigma}_t^2, \hat{\sigma}_{yf}^2, \hat{\sigma}_{ycf}^2$ are the means and variances of the Gaussian distributions parameterized by encoders $E_{\phi_x}, E_{\phi_t}, E_{\phi_{yf}}, E_{\phi_{ycf}}$ with parameters $\phi_x, \phi_t, \phi_{yf}, \phi_{ycf}$ respectively. The overall evidence lower bound

(ELBO) loss of the VAE is expressed as \mathcal{L}_{ELBO} in the following equation,

$$\begin{aligned} \mathcal{L}_{ELBO}(\phi_x, \phi_t, \phi_{yf}, \phi_{ycf}; \mathbf{x}, \mathbf{z}_x, \mathbf{z}_t, \mathbf{z}_{yf}, \mathbf{z}_{ycf}) &= \mathbb{E}_{q_{\phi_x, \phi_t, \phi_{yf}, \phi_{ycf}}} [\log p_{\phi_d}(\mathbf{x} | \mathbf{z}_x, \mathbf{z}_t, \mathbf{z}_{yf}, \mathbf{z}_{ycf})] \\ &- KL(q_{\phi_x}(\mathbf{z}_x | \mathbf{x}) || p_{\phi_d}(\mathbf{z}_x)) - KL(q_{\phi_t}(\mathbf{z}_t | \mathbf{x}) || p_{\phi_d}(\mathbf{z}_t)) \\ &- KL(q_{\phi_{yf}}(\mathbf{z}_{yf} | \mathbf{x}) || p_{\phi_d}(\mathbf{z}_{yf})) - KL(q_{\phi_{ycf}}(\mathbf{z}_{ycf} | \mathbf{x}) || p_{\phi_d}(\mathbf{z}_{ycf})) \end{aligned}$$

where KL denotes the Kullback–Leibler divergence of two probability distributions. We minimize the optimization function of the VAE as \mathcal{L}_{VAE} to obtain the optimal parameter of the encoders $\phi_x, \phi_t, \phi_{yf}, \phi_{ycf}$, and of the decoder ϕ_d as $\mathcal{L}_{VAE}(\phi_x, \phi_t, \phi_{yf}, \phi_{ycf}; \mathbf{x}, \mathbf{z}_x, \mathbf{z}_t, \mathbf{z}_{yf}, \mathbf{z}_{ycf}) = -\mathcal{L}_{ELBO}(\phi_x, \phi_t, \phi_{yf}, \phi_{ycf}; \mathbf{x}, \mathbf{z}_x, \mathbf{z}_t, \mathbf{z}_{yf}, \mathbf{z}_{ycf})$.

Generation of counterfactuals via GAN. After learning the hidden latent codes $\mathbf{z}_x, \mathbf{z}_t, \mathbf{z}_{yf}, \mathbf{z}_{ycf}$ from the VAE, we concatenate the latent codes to form \mathbf{z}_c , passed to the generator of the GAN block G_{θ_g} , along with a random noise $\mathbf{z}_G \sim \mathcal{N}(0, Id)$. G_{θ_g} is parameterized by θ_g , and it outputs the vector \bar{y} of the potential (factual and counterfactual) outcomes. We replace the factual outcome y_f in the generated outcome vector \bar{y} to form \hat{y}_0 and \hat{y}_1 , which are passed to the counterfactual discriminator D_{θ_d} , along with the true covariate vector \mathbf{x} . D_{θ_d} is parameterized by θ_d , and is responsible to predict the treatment variable, similarly to GANITE. The loss of the GAN block is defined as:

$$V_{GAN}(G, D) = \mathbb{E}_{\mathbf{x}, \mathbf{z}_G, \mathbf{z}_c} [t^T \log D(\mathbf{x}, G(\mathbf{z}_G, \mathbf{z}_c)) + (1-t)^T \log(1 - D(\mathbf{x}, G(\mathbf{z}_G, \mathbf{z}_c)))]$$

where $\mathbf{x} \sim p(\mathbf{x})$, $\mathbf{z}_G \sim p(\mathbf{z}_G)$ and \mathbf{z}_c denote the concatenated latent codes $\mathbf{z}_x \sim q_{\phi_x}(\mathbf{z}_x | \mathbf{x})$, $\mathbf{z}_t \sim q_{\phi_t}(\mathbf{z}_t | \mathbf{x})$, $\mathbf{z}_{yf} \sim q_{\phi_{yf}}(\mathbf{z}_{yf} | \mathbf{x})$ and $\mathbf{z}_{ycf} \sim q_{\phi_{ycf}}(\mathbf{z}_{ycf} | \mathbf{x})$. From \bar{y} , we also calculate the predicted factual outcome \hat{y}_f . As also done in GANITE, we make sure to include the supervised loss $\mathcal{L}_S^G(y_f, \hat{y}_f)$, which enforces the predicted factual outcome \hat{y}_f to be as close as to the true factual outcome y_f .

$$\mathcal{L}_S^G(y_f, \hat{y}_f) = \frac{1}{n} \sum_{i=1}^n (y_f(i) - \hat{y}_f(i))^2 \quad (9)$$

The complete loss function of counterfactual GAN is given by $V_{CF}(G, D) = V_{GAN}(G, D) + \gamma \mathcal{L}_S^G(y_f, \hat{y}_f)$.

We also employ an additional regularization $\lambda I(\mathbf{z}_c; G(\mathbf{z}_G, \mathbf{z}_c))$ to maximize the mutual information between the learned concatenated latent code \mathbf{z}_c and the generated output by the generator $G(\mathbf{z}_G, \mathbf{z}_c)$, as in²⁰. We thus propose to solve the following minimax game:

$$\min_G \max_D V_{CF,I}(G, D) = V_{CF}(G, D) + \lambda I(\mathbf{z}_c; G(\mathbf{z}_G, \mathbf{z}_c)) \quad (10)$$

$I(\mathbf{z}_c; G(\mathbf{z}_G, \mathbf{z}_c))$ is harder to solve because of the presence of the posterior $p(\mathbf{z}_c | \mathbf{x})$ ²⁰, so we obtain the lower bound of it using an auxiliary distribution $Q(\mathbf{z}_c | \mathbf{x})$ to approximate $p(\mathbf{z}_c | \mathbf{x})$.

Finally, the optimization function of the counterfactual information-theoretic GAN *InfoGAN*–incorporating the variational regularization of mutual information and hyperparameter λ is given by:

$$\min_{G, Q} \max_D V_{CF,infoGAN}(G, D, Q) = V_{CF}(G, D) - \lambda \mathcal{L}_I(G, Q) \quad (11)$$

The counterfactual InfoGAN is used to generate the missing counterfactual outcome y_{cf} to form the quadruple $\{\mathbf{x}, t, y_f, y_{cf}\}_{i=1}^N$ and sent to the doubly robust block to estimate the ITE.

Information-theoretic GAN optimization. The GAN generator G_{θ_g} works to fool the discriminator D_{θ_d} . To get the optimal Discriminator $D_{\theta_d}^*$, we maximize $V_{CF,infoGAN}$

$$\max_D \mathcal{L}^D(\theta_d) = V_{CF,infoGAN}(G, D, Q) \quad (12)$$

To get the optimal generator $G_{\theta_g}^*$, we maximize $V_{CF,infoGAN}$

$$\min_{G, Q} \mathcal{L}^G(\theta_g) = V_{CF,infoGAN}(G, D, Q) \quad (13)$$

Doubly robust ITE estimation. As introduced above, the propensity score $\pi(\mathbf{x})$ represents the probability of receiving a treatment $T = 1$ (over the alternative $T = 0$) conditioned on the pre-treatment covariates $X = x$. By combining IPW through $\pi(\mathbf{x})$ with outcome regression by both treatment variable and the covariates, Jonsson defined the doubly robust estimation of causal effect²¹ as follows:

$$\hat{\delta}_{DR} = \frac{1}{n} \sum_{i=1}^n \left[\frac{y_i t_i - (t_i - \pi(x_i)) \mu(x_i, t_i)}{\pi(x_i)} - \frac{y_i (1 - t_i) - (t_i - \pi(x_i)) \mu(x_i, t_i)}{1 - \pi(x_i)} \right] \quad (14)$$

where $\mu(x, t) = \hat{\alpha}_0 + \hat{\alpha}_1 x_1 + \hat{\alpha}_2 x_2 + \dots + \hat{\alpha}_n x_n + \hat{\delta} t$, and $(t_i - \pi(x_i)) \mu(x_i, t_i)$ is used for the IPW estimator.

After getting the counterfactual outcome y_{cf} from the counterfactual GAN to form the quadruple $\{\mathbf{x}, t, y_f, y_{cf}\}_{i=1}^N$, we pass this as the input to the doubly robust multitask network to estimate the ITE, using the architecture shown in Figure 2 (green box). To predict the outcomes $y^{(0)}$ and $y^{(1)}$, we use a configuration similar to TARNet, which contains a number of shared layers, denoted by f_ϕ , parameterized by ϕ , and two outcome-specific heads f_{θ_0} and f_{θ_1} , parameterized by θ_0 and θ_1 .

To ensure doubly robustness, we introduce two more heads that predict the propensity score $\pi(\mathbf{x}) = \mathbb{P}(T = 1 | \mathbf{x})$ and the regressor $\mu(\mathbf{x}, t)$. These two are calculated using two neural networks, parameterized by θ_π and θ_μ respectively. The factual and counterfactual outcome $y_i^{(0)}$ and $y_i^{(1)}$ of the i^{th} sample are then calculated as:

$$\hat{y}_f^{(i)} = t_i (f_{\theta_1}(f_\phi(\mathbf{x}_i))) + (1 - t_i) (f_{\theta_0}(f_\phi(\mathbf{x}_i))) \quad (15)$$

$$\hat{y}_{cf}^{(i)} = (1 - t_i) (f_{\theta_1}(f_\phi(\mathbf{x}_i))) + t_i (f_{\theta_0}(f_\phi(\mathbf{x}_i))) \quad (16)$$

Next, the predicted loss will be

$$\mathcal{L}_i^p(\theta_1, \theta_0, \phi) = (\hat{y}_f^{(i)} - y_f^{(i)})^2 + (\hat{y}_{cf}^{(i)} - y_{cf}^{(i)})^2 + \alpha \text{BinaryCrossEntropy}(\pi(x_i), t_i)$$

where α is a hyperparameter. With the help of the propensity score $\pi(\mathbf{x})$ and the regressor $\mu(\mathbf{x}, T)$, the doubly robust outcomes are calculated as

$$\hat{y}_{f_{DR}}^{(i)} = t_i \left[\frac{t_i \hat{y}_i^{(1)} - (t_i - \pi(\mathbf{x}_i)) \mu(\mathbf{x}_i, t_i)}{\pi(\mathbf{x}_i)} \right] + (1 - t_i) \left[\frac{(1 - t_i) \hat{y}_i^{(0)} - (t_i - \pi(\mathbf{x}_i)) \mu(\mathbf{x}_i, t_i)}{1 - \pi(\mathbf{x}_i)} \right] \quad (17)$$

$$\hat{y}_{cf_{DR}}^{(i)} = (1 - t_i) \left[\frac{(1 - t_i) \hat{y}_i^{(1)} - (t_i - \pi(\mathbf{x}_i)) \mu(\mathbf{x}_i, t_i)}{\pi(\mathbf{x}_i)} \right] + t_i \left[\frac{t_i \hat{y}_i^{(0)} - (t_i - \pi(\mathbf{x}_i)) \mu(\mathbf{x}_i, t_i)}{1 - \pi(\mathbf{x}_i)} \right] \quad (18)$$

The doubly robust loss $\mathcal{L}_i^{DR}(\theta_1, \theta_0, \theta_\phi, \theta_\mu, \phi)$ is calculated as:

$$\mathcal{L}_i^{DR}(\theta_1, \theta_0, \theta_\pi, \theta_\mu, \phi) = (\hat{y}_{f_{DR}}^{(i)} - y_f^{(i)})^2 + (\hat{y}_{cf_{DR}}^{(i)} - y_{cf}^{(i)})^2 \quad (19)$$

Finally, the loss function of the ITE is:

$$\mathcal{L}^{ITE}(\theta_1, \theta_0, \theta_\pi, \theta_\mu, \phi) = \frac{1}{n} \sum_{i=1}^n \left(\mathcal{L}_i^p + \beta \mathcal{L}_i^{DR} \right) \quad (20)$$

where β is a hyperparameter and the whole network is trained using end-to-end strategy.

Experimental Setup

Synthetic datasets. We conduct performance tests on two synthetic data experiments. The first uses the same data generation process devised for CEVAE¹⁵. We generate a marginal distribution \mathbf{x} as a mixture of Gaussians from the 5-dimensional latent variable \mathbf{z} , indicating each mixture component. The details of the synthetic dataset using this process is discussed in the OTS. Datasets of sample size $\{1000, 3000, 5000, 10000, 30000\}$ are generated, and divided




Figure 3: Panel (a): performance (ATE) of DR-VIDAL vs. all other models on samples from the generative process of CEVAE. Panel (b) and (c): performance (PEHE) of DR-VIDAL with or without the doubly robust (DR, w/o DR) block vs. GANITE on samples from the generative process of CEVAE-GANITE.

into 80-20 % train-test split. In the second experimental setting, we amalgamate the synthetic data generation process by CEVAE with that of GANITE¹⁴, to model the more complex causal structure illustrated in Figure 1. We sample 7-, 1-, 1-, and 1-dimensional vectors for \mathbf{z}_x , \mathbf{z}_t , \mathbf{z}_{yf} , and \mathbf{z}_{ycf} from Bernoulli distributions, and then collate them into x . From the covariates x , we simulate the treatment assignment t and the potential outcomes y as described in the GANITE paper. We generate multiple synthetic datasets for sample sizes $\{1000, 3000, 5000, 10000, 30000\}$, also divided into 80-20 % splits. Equations for both data generating processes are provided in the OTS.

Real-world datasets. We use three popular real-world benchmark datasets: the Infant Health and Development Program (IHDP) dataset⁸, the Twins dataset³⁶, and the Jobs dataset³⁷. The IHDP and Twins two are semi-synthetic, and simulated counterfactuals to the real factual data are available. These datasets have been also designed and collated to meet specific treatment overlap condition, nonparallel treatment assignment, and nonlinear outcome surfaces^{8,11,15,14}. In detail, IHDP collates data from a multi-site RCT evaluating early intervention in premature, low birth infants, to decrease unfavorable health outcomes. The dataset is composed by 110 treated subjects and 487 controls, with 25 covariates. The Twins dataset is based on records of twin births in the USA from 1989-1991, where the outcome is mortality in the first year, and treatment is heavier weight, comprising 4553 treated, 4567 controls, with 30 covariates. The Jobs study (1978-1978) investigates if a job training program intervention affects earnings after a two-year period, and comprises 237 treated, 2333 controls, with 17 covariates. For all the real-world datasets, we use the same experimental settings described in GANITE, where the datasets are divided into 56/24/20 % train-validation-test splits. We run 1000, 10 and 100 realizations of IHDP, Jobs and Twins datasets, respectively.

Model fit and test details. Consistent with prior studies^{8,11,14}, we report the error on the ATE ϵ_{ATE} , and the expected Precision in Estimation of Heterogeneous Effect (PEHE), ϵ_{PEHE} , for IHDP and Twins datasets, since factual and the counterfactual outcomes are available. For the Jobs dataset, as the counterfactual outcome does not exist, we report the policy risk $R_{pol}(\pi)$, and the error on the average treatment effect on the treated (ATT) ϵ_{ATT} , as indicated in^{11,14}. The training details and the hyperparameters of the individual networks are given in the OTS. We compared DR-VIDAL with TARNet, CEVAE, and GANITE. In addition, for real-world datasets, we compare: least squares regression with treatment as a covariate (OLS/LR1); separate least squares regression for each treatment (OLS/LR2); balancing linear regression (BLR)¹⁰; k-nearest neighbor (k-NN)³³; Bayesian additive regression trees (BART)²⁸; random and causal forest (R Forest, C Forest)⁹; balancing neural network (BNN)¹⁰; counterfactual regression with Wasserstein distance (CFR_{WASS})¹¹.

Results

Synthetic datasets. Figure 3 (a), (b) and (c) shows ATE/PEHE results of DR-VIDAL vs. all other models according to the two synthetic data generation processes. In the generative process of CEVAE, the doubly robust version of DR-VIDAL demonstrates lower ATE error than all other models at all sample sizes. When comparing PEHE, DR-VIDAL (both with and without the doubly robust feature) largely outperforms GANITE. In the second synthetic dataset, generated under the more complex assumptions, DR-VIDAL (both with and without the doubly robust feature) outperforms GANITE in terms of PEHE. It is worth noting the potential of DR-VIDAL to better infer hidden representations in

	IHDP($\sqrt{\epsilon_{PEHE}}$)		Twins($\sqrt{\epsilon_{PEHE}}$)		Jobs(R_{Pol})	
	Out-Sample	In-Sample	Out-Sample	In-Sample	Out-Sample	In-Sample
OLS/LR1	5.8 ± 0.3*	5.8 ± 0.3*	0.318 ± 0.007	0.319 ± 0.005*	0.23 ± 0.02*	0.22 ± 0.00*
OLS/LR2	2.5 ± 0.1*	2.4 ± 0.1*	0.320 ± 0.003*	0.320 ± 0.001*	0.24 ± 0.01*	0.21 ± 0.00*
BLR	5.8 ± 0.3*	5.8 ± 0.3*	0.323 ± 0.018*	0.312 ± 0.002*	0.25 ± 0.02*	0.22 ± 0.01*
k-NN	4.1 ± 0.2*	2.1 ± 0.1*	0.345 ± 0.007*	0.333 ± 0.003*	0.26 ± 0.02*	0.02 ± 0.00*
BART	2.3 ± 0.1*	2.1 ± 0.2*	0.338 ± 0.016*	0.347 ± 0.009*	0.25 ± 0.00*	0.23 ± 0.02*
R Forest	6.6 ± 0.3*	4.2 ± 0.2*	0.321 ± 0.005*	0.306 ± 0.002	0.28 ± 0.02*	0.23 ± 0.01*
C Forest	3.8 ± 0.2*	3.8 ± 0.2*	0.316 ± 0.011	0.366 ± 0.003*	0.20 ± 0.02*	0.19 ± 0.00*
BNN	2.1 ± 0.1*	2.2 ± 0.1*	0.321 ± 0.018*	0.325 ± 0.003*	0.24 ± 0.02*	0.20 ± 0.01*
TARNET (TensorFlow)	0.95 ± 0.02*	0.88 ± 0.02*	0.315 ± 0.003	0.317 ± 0.007	0.21 ± 0.01*	0.17 ± 0.01*
TARNET (Py- torch)	1.10 ± 0.02*	-	-	-	0.29 ± 0.06*	-
CFR _{WASS}	0.76 ± 0.0*	0.71 ± 0.0*	0.313 ± 0.008	0.315 ± 0.007	0.21 ± 0.01*	0.17 ± 0.01*
GANITE	2.4 ± 0.4*	1.9 ± 0.4*	0.297 ± 0.05	0.289 ± 0.005	0.14 ± 0.01*	0.13 ± 0.01*
CEVAE	2.6 ± 0.1*	2.7 ± 0.1*	n.r	n.r	0.26 ± 0.0*	0.15 ± 0.0*
DR-VIDAL	0.69 ± 0.06	0.69 ± 0.05	0.318 ± 0.008	0.317 ± 0.002	0.10 ± 0.01	0.09 ± 0.005

Table 1: Performance of $\sqrt{\epsilon_{PEHE}}$ and R_{Pol} (mean ± st.dev) of various models (prior tools and DR-VIRDAL) on the IHDP, Twins and Jobs datasets. TARNet was originally developed in TensorFlow. We re-implemented TARNet in Pytorch for IHDP and Jobs dataset. (*) is used to indicate methods that DR-VIDAL shows a statistically significant improvement over




Figure 4: Performance comparison of doubly robust vs. non-doubly robust version of DR-VIDAL. The bar plots show how many times one model setup is better than the other in terms of error on the factual outcome (y_f). Panels, from left to right, show results on IHDP, Jobs and Twins datasets (100, 10, 100 iterations), respectively.

comparison to GANITE irrespective of the presence of the doubly robust module.

Real world datasets. In all three IHDP, Jobs and Twins datasets, across all realizations, the information-theoretic, doubly robust configuration of DR-VIDAL yields the best results against all other configurations –with/without information-theoretic optimization and with/without doubly robust loss. The doubly robust loss seems to be responsible for most of the improvement. The absolute gain is small, in the order of 1%, but the relative gain with respect to the non-doubly robust setup is significant, where the doubly robust module always outperforms its non-doubly robust version, from 55-60% in IHDP to over 80% in Twins and Jobs datasets (Figure 4). Table 1 shows the comparison for the $\sqrt{\epsilon_{PEHE}}$ and R_{Pol} values with the state-of-the-art methods on the three datasets. DR-VIDAL outperforms the other methods on all datasets. On the IHDP and Jobs dataset, DR-VIDAL is the best over all by a larger margin. Instead, performance increment in the Twins dataset is mild. Even if DR-VIDAL has a large number of parameters, the deconfounding of hidden factors and the adversarial training make it appropriate for datasets with relatively small sample size like IHDP. It is worth noting that DR-VIDAL converges much faster than CEVAE and GANITE, possibly due to the doubly robustness.

Conclusions

DR-VIDAL is a new deep learning approach to causal effect estimation and counterfactual prediction that combines adversarial representation learning, information-theoretic optimization, and doubly robust regression. On the benchmark datasets, both the doubly robust property and information-theoretic optimization of DR-VIDAL improve performance over a basic adversarial setup.

The work has some limitations. First, the causal graph, even if more elaborate than CEVAE, could be improved. For instance, by connecting the Z to X and only to their respective t , factual and counterfactual outcome nodes would imply two adjustments set. Another option could be to use the TEDVAE structure in conjunction with our doubly-robust setup. Also, the encoded representation in the VAE does not employ any attention mechanism to identify the most important covariates for the propensity scores, especially with high-dimensional datasets. Finally, one thing that would be worth evaluating is how Dragonnet would perform as a downstream module of DR-VIDAL, substituting it to our current four-head doubly-robust block.

In conclusion, DR-VIDAL framework is a comprehensive approach to predicting counterfactuals and estimating ITE, and its flexibility (modifiable causal structure and modularity) allows for further expansion and improvement.

Acknowledgments

This work was in part funded by NIH awards R21CA245858, R01CA246418, R56AG069880, R01AG076234, R01AI145552, R01AI141810, and NSF 2028221.

References

1. Sibbald B, Roland M. Understanding controlled trials: Why are randomised controlled trials important? *BMJ*. 1998;316(7126):201.
2. Hernán MA, Robins JM. Causal inference: what if. Boca Raton: Chapman & Hall/CRC; 2020.
3. Prosperi M, Guo Y, Sperrin M, Koopman JS, Min JS, He X, et al. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nature Machine Intelligence*. 2020;2:369-75.
4. Hoogland J, IntHout J, Belias M, Rovers MM, Riley RD, E Harrell Jr F, et al. A tutorial on individualized treatment effect prediction from randomized trials with a binary endpoint. *Statistics in Medicine*. 2021;40(26):5961-81. Available from: <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.9154>.
5. Austin PC. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral research*. 2011;46(3):399-424.
6. Garrido M, et al. Methods for Constructing and Assessing Propensity Scores. *Health Services Research*. 2014;49(5):1701-20.
7. Tian Y, Schuemie MJ, Suchard MA. Evaluating large-scale propensity score performance through real-world and synthetic data experiments. *International journal of epidemiology*. 2018;47(6):2005-14.
8. Hill JL. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*. 2011;20(1):217-40.
9. Wager S, Athey S. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*. 2018;113(523):1228-42.
10. Johansson F, Shalit U, Sontag D. Learning representations for counterfactual inference. In: *International conference on machine learning*; 2016. p. 3020-9.
11. Shalit U, Johansson FD, Sontag D. Estimating individual treatment effect: generalization bounds and algorithms. In: Precup D, Teh YW, editors. *Proceedings of the 34th International Conference on Machine Learning*. vol. 70 of *Proceedings of Machine Learning Research*. International Convention Centre, Sydney, Australia: PMLR; 2017. p. 3076-85. Available from: <http://proceedings.mlr.press/v70/shalit17a.html>.
12. Shi C, Blei D, Veitch V. Adapting neural networks for the estimation of treatment effects. In: *Advances in neural information processing systems*; 2019. p. 2507-17.
13. Alaa AM, Weisz M, Van Der Schaar M. Deep counterfactual networks with propensity-dropout. *arXiv preprint arXiv:170605966*. 2017.

14. Yoon J, Jordon J, Van Der Schaar M. GANITE: Estimation of individualized treatment effects using generative adversarial nets. In: International Conference on Learning Representations; 2018. .
15. Louizos C, Shalit U, Mooij JM, Sontag D, Zemel R, Welling M. Causal effect inference with deep latent-variable models. In: Advances in neural information processing systems; 2017. p. 6446-56.
16. Zhang W, Liu L, Li J. Treatment effect estimation with disentangled latent factors. arXiv preprint arXiv:200110652. 2020.
17. Makhzani A, Shlens J, Jaitly N, Goodfellow I, Frey B. Adversarial autoencoders. arXiv preprint arXiv:151105644. 2015.
18. Kingma DP, Welling M. Auto-encoding variational bayes. arXiv preprint arXiv:13126114. 2013.
19. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative adversarial nets. *Advances in neural information processing systems*. 2014;27:2672-80.
20. Chen X, Duan Y, Houthoofd R, Schulman J, Sutskever I, Abbeel P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Advances in neural information processing systems*. 2016;29:2172-80.
21. Funk MJ, Westreich D, Wiesen C, Stürmer T, Brookhart MA, Davidian M. Doubly robust estimation of causal effects. *American journal of epidemiology*. 2011;173(7):761-7.
22. Dudík M, Erhan D, Langford J, Li L, et al. Doubly robust policy evaluation and optimization. *Statistical Science*. 2014;29(4):485-511.
23. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*. 1974;66(5):688-701.
24. Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika*. 1983 04;70(1):41-55.
25. Imbens GW. The role of the propensity score in estimating dose-response functions. *Biometrika*. 2000;87(3):706-10.
26. Pearl J, Glymour M, Jewell NP. *Causal Inference in Statistics: A Primer*. Wiley; 2016. Available from: <https://books.google.com/books?id=L3G-CgAAQBAJ>.
27. Porter KE, Gruber S, Van Der Laan MJ, Sekhon JS. The relative performance of targeted maximum likelihood estimators. *The International Journal of Biostatistics*. 2011;7(1).
28. Chipman HA, George EI, McCulloch RE, et al. BART: Bayesian additive regression trees. *The Annals of Applied Statistics*. 2010;4(1):266-98.
29. Athey S, Imbens G. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*. 2016;113(27):7353-60.
30. Lu M, Sadiq S, Feaster DJ, Ishwaran H. Estimating individual treatment effect in observational data using random forest methods. *Journal of Computational and Graphical Statistics*. 2018;27(1):209-19.
31. Dehejia RH, Wahba S. Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and statistics*. 2002;84(1):151-61.
32. Lunceford JK, Davidian M. Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study. *Statistics in medicine*. 2004;23(19):2937-60.
33. Crump RK, Hotz VJ, Imbens GW, Mitnik OA. Nonparametric tests for treatment effect heterogeneity. *The Review of Economics and Statistics*. 2008;90(3):389-405.
34. Alaa A, van der Schaar M. Limits of Estimating Heterogeneous Treatment Effects: Guidelines for Practical Algorithm Design. In: Dy J, Krause A, editors. *Proceedings of the 35th International Conference on Machine Learning*. vol. 80 of *Proceedings of Machine Learning Research*. Stockholm: PMLR; 2018. p. 129-38. Available from: <http://proceedings.mlr.press/v80/alaa18a.html>.
35. Blei DM, Kucukelbir A, McAuliffe JD. Variational inference: A review for statisticians. *Journal of the American statistical Association*. 2017;112(518):859-77.
36. Almond D, Chay KY, Lee DS. The costs of low birth weight. *The Quarterly Journal of Economics*. 2005;120(3):1031-83.
37. LaLonde RJ. Evaluating the econometric evaluations of training programs with experimental data. *The American economic review*. 1986:604-20.