VARIATIONAL AUTO-ENCODER ARCHITECTURES THAT EXCEL AT CAUSAL INFERENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper provides a generative approach for causal inference using data from observational studies. Inspired by the work of Kingma et al. (2014), we propose a sequence of three architectures (namely Series, Parallel, and Hybrid) that each incorporate their M1 and M2 models as building blocks. Each architecture is an improvement over the previous one in terms of estimating causal effect, culminating in the Hybrid model. The Hybrid model is designed to encourage decomposing the underlying factors of any observational dataset; this in turn, helps to accurately estimate all treatment outcomes. Our empirical results demonstrate the superiority of all three proposed architectures compared to both state-of-the-art discriminative as well as other generative approaches in the literature.

1 INTRODUCTION

As one of the main tasks in studying causality (Peters et al., 2017; Guo et al., 2018), the goal of Causal Inference is to figure out **how much** the value of a certain variable would change (*i.e.*, the *effect*) had another certain variable (*i.e.*, the cause) changed its value. A prominent example is the counterfactual question (Rubin, 1974; Pearl, 2009) "Would this patient have *lived longer* [and by **how much**], had she received an alternative treatment?". Such question is often asked in the context of precision medicine, which attempts to identify which medical procedure $t \in \mathcal{T}$ will benefit a certain patient x the most, in terms of the treatment outcome $y \in \mathbb{R}$ (*e.g.*, survival time).

A fundamental problem in causal inference is the unobservablity of the counterfactual outcomes (Holland, 1986). That is, for each subject i, any real-world dataset can only contain the outcome of the administered treatment (aka the *observed* outcome: y_i), but not the outcome(s) of the alternative treatment(s) (aka the *counterfactual* outcome(s)) — *i.e.*, y_i^t for $t \in \mathcal{T} \setminus \{t_i\}$. In other words, the causal effect is never observed (*i.e.*, missing in any training data) and cannot be used to train predictive models, nor can it be used to evaluated a proposed model. This makes estimating causal effects a more difficult problem than that of generalization in the supervised learning paradigm.

In general, we can categorize most machine learning algorithms into two general approaches, which differ in how the input features x and their target values y are modeled (Ng & Jordan, 2002):

Discriminative methods focus solely on modeling the conditional distribution $p(y|x)$ with the goal of direct prediction of y for each instance x. For prediction tasks, discriminative approaches are often more accurate since they use the model parameters more efficiently than generative approaches. Most of the current causal inference methods are discriminative, including the Balancing Neural Network (BNN) (Johansson et al., 2016), CounterFactual Regression Network (CFR-Net) (Shalit et al., 2017), and CFR-Net's extensions — *cf.*, (Yao et al., 2018; Hassanpour & Greiner, 2019; 2020) — as well as Dragon-Net (Shi et al., 2019).

Generative methods, on the other hand, describe the relationship between x and y by their joint probability distribution $p(x, y)$. This, in turn, would allow the generative model to answer arbitrary queries, including coping with missing features x using the marginal distribution $p(x)$ or [similar to discriminative models] predicting the unknown target values y via $p(y|x)$. A promising direction forward for causal inference is developing *generative* models, using either Generative Adverserial Network (GAN) (Goodfellow et al., 2014) or Variational Auto-Encoder (VAE) (Kingma & Welling, 2014; Rezende et al., 2014). This has led to two generative approaches for causal inference: GANs for inference of Individualised Treatment Effects (GANITE) (Yoon et al., 2018) and Causal Effect

VAE (CEVAE) Louizos et al. (2017). However, neither of the two achieve competitive performance in terms of treatment effect estimation compared to the discriminative approaches.

Although discriminative models have excellent predictive performance, they suffer from two drawbacks: (i) overfitting, and (ii) making highly-confident predictions, even for instances that are "far" from the observed training data. Generative models based on Bayesian inference, on the other hand, can handle both of these drawbacks: issue (i) can be minimized by taking an average over the posterior distribution of model parameters; and issue (ii) can be addressed by explicitly providing model uncertainty via the posterior (Gordon & Hernández-Lobato, 2020). Although the exact inference is often intractable, efficient approximations to the parameter posterior distribution is possible through variational methods. Here, we use the Variational Auto-Encoder (VAE) (Kingma & Welling, 2014; Rezende et al., 2014) for the Bayesian inference component of our causal inference method.

Contribution: In this paper, we propose three interrelated Bayesian model architectures (namely Series, Parallel, and Hybrid) that employ the VAE framework to address the task of causal inference for binary treatments. We find that the best performing architecture is the Hybrid model, that is [partially] successful in decomposing the underlying factors of any observational dataset. This is a valuable property, as that means it can accurately estimate all all treatment outcomes. We demonstrate that these models significantly outperform the state-of-the-art in terms of treatment effect estimation performance on two publicly available benchmarks, as well as a fully synthetic dataset that allows for detailed performance analyses.

2 Related works

CFR-Net Shalit et al. (2017) considered the binary treatment task and attempted to learn a representation space Φ that reduces selection bias by making $\Pr(\,\Phi(x)\,|\,t{=}0\,)$ and $\Pr(\,\Phi(x)\,|\,t{=}1\,)$ as close to each other as possible, provided that $\Phi(\,x\,)$ retains enough information that the learned regressors $\{h^t\big(\Phi(\cdot)\big)\,:\;t\in\{0,1\}\}$ can generalize well on the observed outcomes. Their objective function includes $L\big[\,y_i,\,h^{t_i}\big(\Phi(x_i)\big)\big]$, which is the loss of predicting the observed outcome for sample i (described as x_i), weighted by $\omega_i = \frac{t_i}{2u} + \frac{1-t_i}{2(1-u)}$, where $u = \Pr(\,t{=}1\,)$. This is effectively setting $\omega_i = \frac{1}{2\,\Pr(\,t_i\,)}$ where $\Pr(\,t_i\,)$ is the probability of selecting treatment t_i over the entire population.

DR-CFR Hassanpour & Greiner (2020) argued against the standard implicit assumption that *all* of the covariates X are confounders (*i.e.*, contribute to both treatment assignment and outcome determination). Instead, they proposed a graphical model similar to that in Figure 1 and designed a discriminative causal inference approach accordingly — built on top of the CFR-Net. Specifically, their model, named Disentangled Representations for CFR (DR-CFR), includes three representation networks, each trained with constraints to insure that each component corresponds to its respective underlying factor. While the idea behind DR-CFR provides an interesting intuition, it is known that only generative models (and not discriminative ones) can truly identify the underlying data generating mechanism. This paper is a step in this direction.

Dragon-Net Shi et al. (2019)'s main objective was to estimate the Average Treatment Effect (ATE), which they explain requires a two stage procedure: (i) fit models that predict the outcomes for both treatments; and (ii) find a downstream estimator of the effect. Their method is based on a classic result from strong ignorability — *i.e.*, Theorem 3 in (Rosenbaum & Rubin, 1983) — that states:

$$\begin{array}{lcl} (y^1, y^0) \perp\!\!\!\perp t \,|\, x & \& & \Pr(\,t=1\,|\,x\,) \in (0,1) \quad \Longrightarrow \\ (y^1, y^0) \perp\!\!\!\perp t \,|\, b(x) & \& & \Pr(\,t=1\,|\,b(x)\,) \in (0,1) \end{array}$$

where $b(x)$ is a balancing score[1]. They consider propensity score as a balancing score and argue that only the parts of X relevant for predicting T are required for the estimation of the causal effect [2]. This theorem only provides a way to *match* treated and control instances though — *i.e.*, it helps finding potential counterfactuals from the alternative group to calculate ATE. Shi et al. (2019), however, used this theorem to derive minimal representations on which to *regress* to estimate the outcomes.

[1] That is, $X \perp\!\!\!\perp T \,|\, b(X)$ (Rosenbaum & Rubin, 1983)

[2] The authors acknowledge that this would hurt the predictive performance for individual outcomes. As a result, this yields inaccurate estimation of Individual Treatment Effects (ITEs).

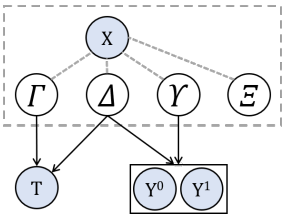

Figure 1: Underlying factors of any observational dataset (Hassanpour & Greiner, 2020)

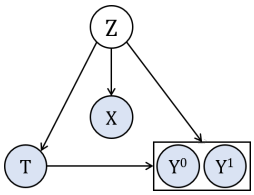

Figure 2: Graphical model of CEVAE (Louizos et al., 2017)

GANITE Yoon et al. (2018) proposed the counterfactual GAN, whose generator G, given $\{x, t, y^t\}$, estimates the counterfactual outcomes ($\hat{y}^{\neg t}$); and whose discriminator D tries to identify which of $\{[x, 0, y^0],\ [x, 1, y^1]\}$ is the factual outcome. It is, however, unclear why this requires that G must produce samples that are indistinguishable from the factual outcomes, especially as D can just learn the *treatment selection mechanism* instead of distinguishing the factual outcomes from counterfactuals. Although this work is among the few generative approaches for causal inference, our empirical results (in Section 4) show that it does not effectively estimate counterfactual outcomes.

CEVAE Louizos et al. (2017) used VAE to extract latent confounders from their observed proxies in X. While this is an interesting step in the right direction, empirical results show that it does not always accurately estimate treatment effect (see Section 4). The authors note that this may be because CEVAE is not able to address the problem of selection bias. Another reason that we think contributes to CEVAE's sub-optimal performance is its assumed graphical model of the underlying data generating mechanism (depicted in Figure 2). This model assumes that there is only one latent variable Z (confounding T and Y) that generates the entire observational data; however, we know from (Kuang et al., 2017) and (Hassanpour & Greiner, 2020) that there must be more (see Figure 1).

3 METHOD

Following (Hassanpour & Greiner, 2020) and without loss of generality, we assume that the random variable X follows a(n unknown) joint probability distribution $\Pr(\, X \mid \Gamma, \Delta, \Upsilon, \Xi \,)$, where Γ, Δ, Υ, and Ξ are non-overlapping independent factors. Moreover, we assume that treatment T follows $\Pr(\, T \mid \Gamma, \Delta \,)$ (*i.e.*, Γ and Δ are the responsible factors for selection bias) and outcome Y^T follows $\Pr_T(\, Y^T \mid \Delta, \Upsilon \,)$; see Figure 1. Observe that the factor Γ (resp., Υ) partially determines only T (resp., Y), but not Y (resp., T); and Δ includes the confounding factors between T and Y.

Our goal is to design generative model architectures that encourage learning disentangled representations of these four underlying latent factors (see Figure 1). In other words, it is an attempt to decompose and separately learn the underlying factors that are responsible for determining T and Y. To achieve this, we propose three architectures (as illustrated in Figures 3(a), 3(b), and 3(c)), each employing a VAE (Kingma & Welling, 2014; Rezende et al., 2014) that each include a decoder (generative model) and an encoder (variational posterior). Specifically, we use the M1 and M2 models from (Kingma et al., 2014) as our building blocks, leading to a Series architecture, a Parallel architecture, and a Hybrid one. An overview of the M1 and M2 VAEs is presented in Appendix A.1; however, as a brief reminder, the M1 model helps learning latent representations from x alone, while the M2 model allows the treatment as well as the outcome information to guide the representation learning process. Each component is parametrized as a deep neural network.

3.1 THE VARIATIONAL AUTO-ENCODER COMPONENT

3.1.1 THE SERIES ARCHITECTURE

The architecture of the Series model is illustrated in Figure 3(a). Louizos et al. (2015) proposed a similar architecture to address fairness in machine learning, but using a binary sensitive variable S (*e.g.*, gender, race, etc.) rather than the treatment T. Here, we employ this architecture for causal inference and explain why it should work. We hypothesize that this structure functions as a distillation tower: the bottom M2 VAE attempts to decompose Γ (guided by T) from Δ and Υ (captured by Z_1); and the top M2 VAE attempts to learn Δ and Υ (guided by Y). Decoder and encoder components of the Series model (parametrized by θ_s and ϕ_s respectively) involve the following distributions:

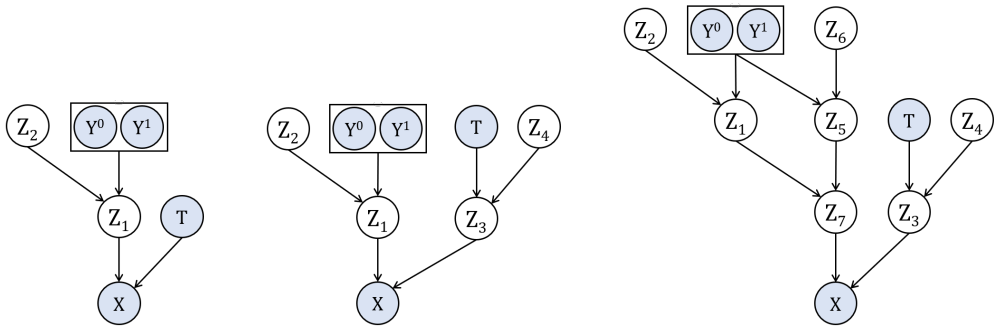

(a) The **Series** model (b) The **Parallel** model (c) The **Hybrid** model

Figure 3: Belief nets of the proposed architectures.

Priors	**Likelihood**	**Posteriors**
$p_{\theta_s}(z_2)$	$p_{\theta_s}(x\|z_1,t)$	$q_{\phi_s}(z_1\|x,t)$
$p_{\theta_s}(z_1\|y,z_2)$		$q_{\phi_s}(y\|z_1)$
		$q_{\phi_s}(z_2\|y,z_1)$

The goal is to maximize the conditional log-likelihood of the observed data (left-hand-side of the following inequality) by maximizing the Evidence Lower BOund (ELBO; right-hand-side) — *i.e.*,

$$\sum_{i=1}^{N} \log p(x_i|t_i,y_i) \;\geq\; \sum_{i=1}^{N} \mathbb{E}_{q_{\phi_s}(z_1|x,t)}\big[\log p_{\theta_s}(x_i|z_{1i},t_i)\big] \tag{1}$$

$$-\,\mathrm{KL}\big(\,q_{\phi_s}(z_1|x,t)\,||\,p_{\theta_s}(z_1|y,z_2)\,\big) - \mathrm{KL}\big(\,q_{\phi_s}(z_2|y,z_1)\,||\,p_{\theta_s}(z_2)\,\big) \tag{2}$$

where KL denotes the Kullback-Leibler divergence, $p_{\theta_s}(z_2)$ is the unit multivariate Gaussian (*i.e.*, $\mathcal{N}(0,\mathbb{I})$), and the other distributions are parameterized as deep neural networks.

3.1.2 The Parallel Architecture

The Series model is composed of two M2 stacked models. However, Kingma et al. (2014) showed that an M1+M2 stacked architecture learns better representations than an M2 model alone for a downstream prediction task. This motivated us to design a double M1+M2 Parallel model; where one arm is for the outcome to guide the representation learning via Z_1 and another for the treatment to guide that via Z_3. This architecture is illustrated in Figure 3(b). We hypothesize that Z_1 would learn Δ and Υ, and Z_3 would learn Γ (and perhaps partially Δ). Decoder and encoder components of the Parallel model (parametrized by θ_p and ϕ_p respectively) involve the following distributions:

Priors		**Likelihood**	**Posteriors**	
$p_{\theta_p}(z_2)$	$p_{\theta_p}(z_4)$	$p_{\theta_p}(x\|z_1,z_3)$	$q_{\phi_p}(z_1\|x,t)$	$q_{\phi_p}(z_3\|x,y)$
$p_{\theta_p}(z_1\|y,z_2)$	$p_{\theta_p}(z_3\|t,z_4)$		$q_{\phi_p}(y\|z_1)$	$q_{\phi_p}(t\|z_3)$
			$q_{\phi_p}(z_2\|y,z_1)$	$q_{\phi_p}(z_4\|t,z_3)$

Here, the conditional log-likelihood can be upper bounded by:

$$\sum_{i=1}^{N} \log p(x_i|t_i,y_i) \;\geq\; \sum_{i=1}^{N} \mathbb{E}_{q_{\phi_p}(z_1,z_3|x,t,y)}\big[\log p_{\theta_p}(x_i|z_{1i},z_{3i})\big] \tag{3}$$

$$-\,\mathrm{KL}\big(\,q_{\phi_p}(z_1|x,t)\,||\,p_{\theta_p}(z_1|y,z_2)\,\big) - \mathrm{KL}\big(\,q_{\phi_p}(z_2|y,z_1)\,||\,p_{\theta_p}(z_2)\,\big) \tag{4}$$

$$-\,\mathrm{KL}\big(\,q_{\phi_p}(z_3|x,y)\,||\,p_{\theta_p}(z_3|t,z_4)\,\big) - \mathrm{KL}\big(\,q_{\phi_p}(z_4|t,z_3)\,||\,p_{\theta_p}(z_4)\,\big) \tag{5}$$

3.1.3 The Hybrid Architecture

The final architecture, Hybrid, attempts to get the best capabilities of the previous two architectures. The backbone of the Hybrid model has a Series architecture, that separates Γ (factors related to the treatment T; captured by the right module with Z_3 as its head), from Δ and Υ (factors related to the outcome Y; captured by the left module with Z_7 as its head). The left module, itself, consists of a Parallel model that attempts to proceed one step further and decompose Δ from Υ. This is done with the help of a discrepancy penalty (see Section 3.3). Figure 3(c) illustrates our designed architecture

for the Hybrid model. Decoder and encoder components of the Hybrid model (parametrized by θ_h and ϕ_h respectively) involve the following distributions:

Priors			**Likelihood**	**Posteriors**		
$p_{\theta_h}(z_2)$	$p_{\theta_h}(z_4)$	$p_{\theta_h}(z_6)$	$p_{\theta_h}(x\|z_3, z_7)$	$q_{\phi_h}(z_7\|x,t)$	$q_{\phi_h}(z_1\|z_7)$	$q_{\phi_h}(z_5\|z_7)$
$p_{\theta_h}(z_1\|y,z_2)$	$p_{\theta_h}(z_3\|t,z_4)$	$p_{\theta_h}(z_5\|y,z_6)$		$q_{\phi_h}(z_3\|x,y)$	$q_{\phi_h}(y\|z_1,z_5)$	$q_{\phi_h}(t\|z_3)$
$p_{\theta_h}(z_7\|z_1,z_5)$				$q_{\phi_h}(z_2\|y,z_1)$	$q_{\phi_h}(z_6\|y,z_5)$	$q_{\phi_h}(z_4\|t,z_3)$

Here, the conditional log-likelihood can be upper bounded by:

$$\sum_{i=1}^{N} \log p(x_i|t_i, y_i) \;\geq\; \sum_{i=1}^{N} \mathbb{E}_{q_{\phi_h}(z_3,z_7|x,t,y)}\big[\log p_{\theta_h}(x_i|z_{3i}, z_{7i})\big] \tag{6}$$

$$-\,\mathrm{KL}\big(\, q_{\phi_h}(z_1|z_7) \,||\, p_{\theta_h}(z_1|y,z_2)\,\big) - \mathrm{KL}\big(\, q_{\phi_h}(z_2|y,z_1) \,||\, p_{\theta_h}(z_2)\,\big) \tag{7}$$

$$-\,\mathrm{KL}\big(\, q_{\phi_h}(z_3|x,y) \,||\, p_{\theta_h}(z_3|t,z_4)\,\big) - \mathrm{KL}\big(\, q_{\phi_h}(z_4|t,z_3) \,||\, p_{\theta_h}(z_4)\,\big) \tag{8}$$

$$-\,\mathrm{KL}\big(\, q_{\phi_h}(z_5|z_7) \,||\, p_{\theta_h}(z_5|y,z_6)\,\big) - \mathrm{KL}\big(\, q_{\phi_h}(z_6|y,z_5) \,||\, p_{\theta_h}(z_6)\,\big) \tag{9}$$

$$-\,\mathrm{KL}\big(\, q_{\phi_h}(z_7|x,t) \,||\, p_{\theta_h}(z_7|z_1,z_5)\,\big) \tag{10}$$

The first term in the ELBO (*i.e.*, right-hand-side of Equations (1), (3), or (6)) is called the Reconstruction Loss (RecL) and the next term(s) (*i.e.*, Equation (2), summation of Equations (4) and (5), or summation of Equations (7), (8), (9), and (10)) is referred to as the KL Distance (KLD). Concisely, the ELBO can be written as: RecL $-$ KLD, which is to be maximized.

3.2 Disentanglement with β-VAE

As mentioned earlier, we want the learned latent variables to be disentangled, to match our assumption of non-overlapping factors Γ, Δ, and Υ. To ensure this, we employ the β-VAE (Higgins et al., 2017) which adds a hyperparameter β as a multiplier of the KLD part of the ELBO. This adjustable hyperparameter facilitates a trade-off that helps balance the latent channel capacity and independence constraints (handled by the KL terms) with the reconstruction accuracy — *i.e.*, including the β hyperarameter grants a better control over the level of disentanglement in the learned representations (Burgess et al., 2018). Therefore, the generative objective to be minimized becomes:

$$\mathcal{L}_{\text{VAE}} \;=\; -\text{RecL} \,+\, \beta \cdot \text{KLD} \tag{11}$$

Although Higgins et al. (2017) suggest the β to be set greater than 1 in most applications, Hoffman et al. (2017) show that having a $\beta < 1$ weight on the KL term can be interpreted as optimizing the ELBO under an alternative prior, which functions as a regularization term to prevent degeneracy.

3.3 Discrepancy

Although all the three proposed graphical models suggest that T and Z_1 are statistically independent (see, for example, the collider structure (at X): $T \rightarrow X \leftarrow Z_1$ in Figure 3(a)), an information leak is quite possible due to the correlation between the outcome y and treatment t in the data. We therefore require an extra regularization term on $q_\phi(z_1|t)$ in order to penalize the discrepancy (denoted by `disc`) between the conditional distributions of z_1 given $t=0$ versus given $t=1$. To achieve this regularization, we calculate the `disc` using an Integral Probability Metric (IPM) (Mansour et al., 2009) [3] that measures the distance between the two above-mentioned distributions:

$$\mathcal{L}_{\texttt{disc}} \;=\; \text{IPM}\big(\, \{z_1\}_{i:t_i=0},\; \{z_1\}_{i:t_i=1} \,\big) \tag{12}$$

3.4 Predictive Loss

Note, however, that neither the VAE nor the `disc` losses contribute to training a predictive model for outcomes. To remedy this, we extend the objective function to include a discriminative term for the regression loss of predicting y: [4]

$$\mathcal{L}_{\text{pred}} \;=\; \frac{1}{N}\sum_{i=1}^{N} \omega_i \cdot \mathcal{L}\big[\, y_i,\, \hat{y}_i \big] \tag{13}$$

[3] *E.g.*, Maximum Mean Discrepancy (MMD) (Gretton et al., 2012).

[4] This is similar to the way (Kingma et al., 2014) included a classification loss in their Equation (9).

where the predicted outcome $\hat{y}_i$ is derived as the mean of the $q_\phi^{t_i}(y_i|z_{1i})$ posterior trained for the respective treatment t_i; $\mathcal{L}\big[\, y_i,\, \hat{y}_i \big]$ is the factual loss (*i.e.*, L2 loss for real-valued outcomes and log loss for binary-valued outcomes); and ω_i represent the weights as derived similar to those in CFR-Net.

3.5 Final Model(s)

Putting everything together, the overall objective function to be minimized is then:

$$\mathcal{J} \quad = \quad \mathcal{L}_{\text{pred}} \;+\; \alpha \cdot \mathcal{L}_{\texttt{disc}} \;+\; \gamma \cdot \mathcal{L}_{\text{VAE}} \;+\; \lambda \cdot \mathfrak{Reg} \tag{14}$$

where $\mathfrak{Reg}$ penalizes the model complexity.

This objective function is motivated by the work of McCallum et al. (2006), which suggested optimizing a convex combination of discriminative and generative losses would indeed improve predictive performance. As an empirical verification, note that for $\gamma = 0$, the Series and Parallel models effectively reduce to CFR-Net. However, our empirical results (*cf.*, Section 4) suggest that the generative term in the objective function helps learning representations that embed more relevant information for estimating outcomes than that of Φ in CFR-Net.

We refer to the family of our proposed methods as VAE-CI (Variational Auto-Encoder for Causal Inference); specifically: **{S, P, H}-VAE-CI**, for **S**eries, **P**arallel, and **H**ybrid respectively.

4 Experiments, Results, and Discussion

4.1 Benchmarks

Infant Health and Development Program (IHDP) The original IHDP randomized controlled trial was designed to evaluate the effect of specialist home visits on future cognitive test scores of premature infants. Hill (2011) induced selection bias by removing a non-random subset of the treated population. The dataset contains 747 instances (608 control and 139 treated) with 25 covariates. We use the same benchmark (with 100 realizations of outcomes) provided by and used in (Johansson et al., 2016) and (Shalit et al., 2017).

Atlantic Causal Inference Conference 2018 (ACIC'18) ACIC'18 is a collection of binary-treatment datasets released for a data challenge. Following (Shi et al., 2019), we use a subset of the datasets with instances $N \in \{1, 5, 10\} \times 10^3$ (four datasets in each category). The covariates matrix for each dataset involves 177 features and is sub-sampled from a table of medical measurements taken from the Linked Birth and Infant Death Data (LBIDD) (MacDorman & Atkinson, 1998), that contains information corresponding to 100,000 subjects.

Fully Synthetic Datasets We generated a set of synthetic datasets according to the procedure described in (Hassanpour & Greiner, 2020); see Section A.2 for an overview. We considered all the viable datasets in a mesh generated by various sets of variables, of sizes $m_\Gamma, m_\Delta, m_\Upsilon \in \{0, 4, 8\}$ and $m_\Xi = 1$. This creates 24 scenarios[5] that consider all possible relative sizes of the factors Γ, Δ, and Υ. For each scenario, we synthesized multiple datasets with various initial random seeds in order to allow for statistical significance testing of the performance comparisons between various methods.

4.2 Evaluating Identification of the Underlying Factors

To evaluate the identification performance of the underlying factors, we use a fully synthetic dataset with $m_\Gamma = m_\Delta = m_\Upsilon = 8$ and $m_\Xi = 1$. We set x to be one of the four dummy vectors $V_{1..4}$ and input it to each trained representation network Z_j. Three of these vectors had "1" in the 8 positions associated with Γ, Δ, and Υ respectively, and the remaining 17 positions of each vector was filled with "0". The fourth vector was all "1" except for the last position (the noise) which was "0". This helps measure the maximum amount of information that is passed to the final layer of each representation network.[6]

[5] There are $3^3 = 27$ combinations in total; however, we removed three of these combinations that generate pure noise outcomes — *i.e.*, $\Delta = \Upsilon = \emptyset$: $(0, 0, 0)$, $(4, 0, 0)$, and $(8, 0, 0)$.

[6] Unlike the evaluation strategy presented in (Hassanpour & Greiner, 2020) that only looked at the first layer's weights of each representation network, we propagate the values through the entire network and check how much of each factor is exhibited in the final layer of every representation network.

H-VAE-CI	Z1	Z2	Z3	Z4	Z5	Z6	Z7
Γ	0.2262	0.2599	**1.2512**	1.1979	0.3109	0.2890	0.2959
Δ	0.6431	0.7704	0.5202	0.5602	**0.8229**	0.7073	**0.7170**
Υ	**0.8957**	0.7929	0.2622	0.2641	0.6551	0.5679	**0.7106**

S-VAE-CI	Z1	Z2
Γ	0.3407	0.3451
Δ	**0.7925**	0.7828
Υ	**0.8804**	0.8204

P-VAE-CI	Z1	Z2	Z3	Z4
Γ	0.2920	0.2922	**0.6374**	0.6676
Δ	**0.7799**	0.7563	0.1882	0.1888
Υ	**0.7686**	0.7735	0.4738	0.3737

DR-CFR	Γ_rep	Δ_rep	Υ_rep
Γ	**0.5289**	0.1733	0.8412
Δ	**0.5429**	**0.7038**	0.8402
Υ	0.0544	**0.6749**	1.0008

Figure 4: Performance analysis for decomposition of the underlying factors on the **synthetic** dataset with $m_{\Gamma,\Delta,\Upsilon}=8, m_{\Xi}=1$.

Table 1: **IHDP** (100 realizations) benchmark

METHOD	PEHE	ϵ_{ATE}
CFR	0.75 (0.57)	0.08 (0.10)
DR-CFR	0.65 (0.37)	0.03 (0.04)
DRAGON	NA	0.14 (0.15)
GANITE	2.81 (2.30)	0.24 (0.46)
CEVAE	2.50 (3.47)	0.18 (0.25)
S-VAE-CI	**0.51 (0.37)**	**0.01 (0.02)**
P-VAE-CI	**0.52 (0.36)**	**0.01 (0.03)**
H-VAE-CI	**0.49 (0.36)**	**0.01 (0.02)**

Table 2: **ACIC'18** ($N \leq 10K$) benchmark

METHOD	PEHE	ϵ_{ATE}
CFR	5.13 (5.59)	1.21 (1.81)
DR-CFR	3.86 (3.39)	0.80 (1.41)
DRAGON	NA	0.48 (0.77)
GANITE	3.55 (2.27)	0.69 (0.65)
CEVAE	5.30 (5.52)	3.29 (3.50)
S-VAE-CI	2.73 (2.39)	0.51 (0.82)
P-VAE-CI	2.62 (2.26)	0.37 (0.75)
H-VAE-CI	1.78 (1.27)	0.44 (0.77)

PEHE and ϵ_{ATE} measures (lower is better) represented in the form of "mean (standard deviation)".

We let $O_{i,j}$ be the `elu` output (here, $\in \mathbb{R}^{200}$) of the encoder network Z_j when $x=V_i$. The average of the 200 values of $O_{i,j}$ ($Avg(O_{i,j})$) represents the power of signal that was produced by the Z_j channel on the input V_i. The values shown in Figure 4's tables are the ratios of $Avg(O_{1,j})$, $Avg(O_{2,j})$, and $Avg(O_{3,j})$ divided by $Avg(O_{4,j})$ for each of the learned representation networks. Note that, a larger ratio indicates that the respective representation network Z_j has allowed more of the input signal V_i to pass through. Section A.3 includes more details on this procedure.

As expected, Z_3 and Z_4 capture Γ (*e.g.*, the Z_3 values for Γ in the {P, H}-VAE-CI tables are largest), and Z_1, Z_2, Z_5, Z_6, and Z_7 capture Δ and Υ. Note that decomposition of Δ from Υ has not been achieved by any of the methods except for H-VAE-CI, which successfully captures Υ by Z_1 and Δ by Z_5. This decomposition is quite vital for designing importance sampling weights that must be based on Δ only (Hassanpour & Greiner, 2020). Also observe that both {P, H}-VAE-CI are able to separate Γ from Δ. However, DR-CFR, which tried to disentangle all factors, failed not only to disentangle Δ from Υ, but also Γ from Δ.

4.3 EVALUATING TREATMENT EFFECT ESTIMATION

Evaluation of treatment effect estimation is often done with semi- or fully- synthetic datasets that include both factual and counterfactual outcomes. There are two categories of performance measures:

Individual-based: "Precision in Estimation of Heterogeneous Effect" $\text{PEHE}=\sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{\text{e}}_i-\text{e}_i)^2}$ uses $\hat{\text{e}}_i = \hat{y}_i^1 - \hat{y}_i^0$ as the estimated effect and $\text{e}_i = y_i^1 - y_i^0$ as the true effect (Hill, 2011); and

Population-based: "Bias of the Average Treatment Effect" $\epsilon_{\text{ATE}} = \left|\text{ATE} - \widehat{\text{ATE}}\right|$, where $\text{ATE} = \frac{1}{N}\sum_{i=1}^{N} y_i^1 - \frac{1}{N}\sum_{j=1}^{N} y_j^0$ and $\widehat{\text{ATE}}$ is calculated based on the estimated outcomes.

In this paper, we compare performances of our proposed methods **{S, P, H}-VAE-CI** versus the following treatment effect estimation methods: **CFR-Net** (Shalit et al., 2017), **DR-CFR** (Hassanpour & Greiner, 2020), **Dragon-Net** (Shi et al., 2019), **GANITE** (Yoon et al., 2018), and **CEVAE** (Louizos et al., 2017). The hyperparameters' search grid for the CFR-Net based algorithms (including our proposed methods) is available in Section A.4. For the other algorithms, we searched around their default hyperparameters' setting. We ran the experiments for the contender methods using their publicly available code-bases; note the following points:

- Since Dragon-Net is designed to estimate ATE only, we did not report its performance results for the PEHE measure (which, as expected, were significantly inaccurate).
- The original GANITE code-base was implemented for binary outcomes only. We modified the code (losses, etc.) such that it could process real-valued outcomes also.

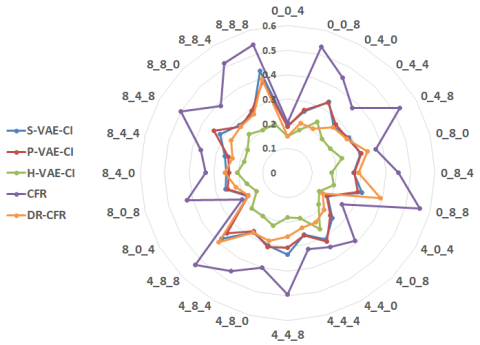

Figure 5: Radar graphs of PEHE (on the radii; lower is better) for the entire **synthetic** benchmark (24×3 with $N = 10{,}000$; each vertex denotes the respective dataset). Figure is best viewed in color.

Table 3: PEHE and ϵ_{ATE} measures (lower is better) represented in the form of "mean (standard deviation)" on the entire **synthetic** benchmark (average performance of 24×3 datasets, each with sample size of 10,000).

METHOD	PEHE	ϵ_{ATE}
CFR	0.41 (0.10)	0.027 (0.020)
DR-CFR	0.26 (0.07)	0.007 (0.004)
DRAGON	NA	0.007 (0.005)
GANITE	1.28 (0.43)	0.036 (0.015)
CEVAE	1.39 (0.32)	0.287 (0.217)
S-VAE-CI	**0.28 (0.05)**	**0.004 (0.003)**
P-VAE-CI	**0.28 (0.05)**	**0.004 (0.003)**
H-VAE-CI	**0.20 (0.03)**	**0.003 (0.002)**

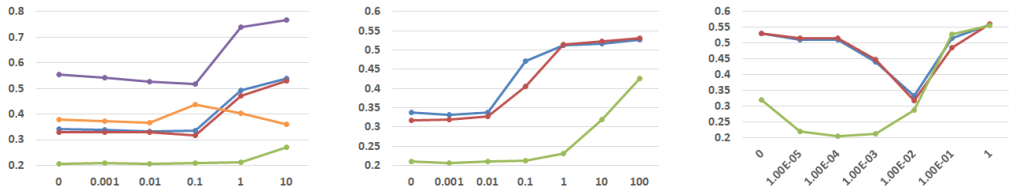

(a) α (b) β (c) γ

Figure 6: Hyperparameters' (x-axis) sensitivity analysis based on PEHE (y-axis) on the **synthetic** dataset with $m_{\Gamma,\Delta,\Upsilon} = 8$, $m_{\Xi} = 1$. Legend is the same as Figure 5. Plots are best viewed in color.

- We were surprised that CEVAE diverged when running on the ACIC'18 datasets. To avoid this, we had to run the ACIC'18 experiments on the binary covariates only.

Tables 1, 2, and 3 summarize the mean and standard deviation of the PEHE and ϵ_{ATE} measures (lower is better) on the IHDP, ACIC'18, and Synthetic benchmarks respectively. VAE-CI achieves the best performance among the contending methods. These results are statistically significant (in **bold**; based on the Welch's unpaired t-test with $\alpha = 0.05$) for the IHDP and Synthetic benchmarks. Although VAE-CI also achieves the best performance on the ACIC'18 benchmark, the results are not statistically significant due to the high standard deviation of the contending methods' performances.

Figure 5 visualizes the PEHE measures on the entire synthetic datasets with sample size of $N = 10{,}000$. We observe that H-VAE-CI's plot is inscribed by plots of all other methods, showcasing its superior performance under every possible selection bias scenario. We also performed hyperparameters' sensitivity analyses in terms of PEHE (see Figure 6). For the α hyperparameter (*i.e.*, the discrepancy penalty), DR-CFR and H-VAE-CI methods have the most robust performance. For β and γ, H-VAE-CI appears to have the most stable performance compared to the {S, P}-VAE-CI methods.

5 FUTURE WORKS AND CONCLUSION

Despite the success of the proposed methods in addressing causal inference, especially the Hybrid model, we are yet to invent algorithms that can learn to perfectly disentangle factors Δ and Υ. Our hypothesis is that isolating Δ could potentially enhance the quality of the importance sampling weights, which should improve the causal effect estimation performance. This is left to future work.

The goal of this paper was to estimate treatment effects (either for individuals or the entire population) from observational data. We designed three VAE-based (Kingma & Welling, 2014; Rezende et al., 2014) architectures (namely Series, Parallel, and Hybrid), that employed (Kingma et al., 2014)'s M1 and M2 models. The Hybrid model, as the best performing architecture, partially succeeded at decomposing the underlying factors Γ, Δ, and Υ; which helped in accurate estimation of treatment outcomes. Our empirical results demonstrated the superiority of the proposed methods, compared to both state-of-the-art discriminative as well as generative approaches in the literature.

REFERENCES

Christopher Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in β-VAE. *arXiv preprint:1804.03599*, 2018.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.

Jonathan Gordon and José Miguel Hernández-Lobato. Combining deep generative and discriminative models for bayesian semi-supervised learning. *Pattern Recognition*, 100:107156, 2020.

Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *JMLR*, 13(March), 2012.

Ruocheng Guo, Lu Cheng, Jundong Li, P. Richard Hahn, and Huan Liu. A survey of learning causality with data: Problems and methods. *arXiv preprint:1809.09337*, 2018.

Negar Hassanpour and Russell Greiner. Counterfactual regression with importance sampling weights. In *IJCAI*, 2019.

Negar Hassanpour and Russell Greiner. Learning disentangled representations for counterfactual regression. In *ICLR*, 2020.

Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. β-VAE: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2017.

Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1), 2011.

Matt Hoffman, Carlos Riquelme, and Matthew Johnson. The beta vae's implicit prior. 2017. URL `http://bayesiandeeplearning.org/2017/papers/66.pdf`.

Paul W Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81 (396):945–960, 1986.

Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *ICML*, 2016.

Diederik P Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.

Diederik P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *NeurIPS*, 2014.

Kun Kuang, Peng Cui, Bo Li, Meng Jiang, Shiqiang Yang, and Fei Wang. Treatment effect estimation with data-driven variable decomposition. In *AAAI*, 2017.

Christos Louizos, Kevin Swersky, Yujia Li, Max Welling, and Richard Zemel. The variational fair autoencoder. *arXiv preprint:1511.00830*, 2015.

Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *NeurIPS*. 2017.

Marian F MacDorman and Jonnae O Atkinson. Infant mortality statistics from the 1996 period linked birth/infant death dataset. *Monthly Vital Statistics Report*, 46(12), 1998.

Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint:0902.3430*, 2009.

Andrew McCallum, Chris Pal, Gregory Druck, and Xuerui Wang. Multi-conditional learning: Generative/discriminative training for clustering and classification. In *AAAI*, pp. 433–439, 2006.

Andrew Y Ng and Michael I Jordan. On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. In *NeurIPS*, pp. 841–848, 2002.

Judea Pearl. *Causality*. Cambridge University Press, 2009.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *ICML*, 2014.

Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 1983.

Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 1974.

Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *ICML*, 2017.

Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. In *NeurIPS*, 2019.

Liuyi Yao, Sheng Li, Yaliang Li, Mengdi Huai, Jing Gao, and Aidong Zhang. Representation learning for treatment effect estimation from observational data. In *NeurIPS*, 2018.

Jinsung Yoon, James Jordon, and Mihaela van der Schaar. GANITE: Estimation of individualized treatment effects using generative adversarial nets. In *ICLR*, 2018.