

---

# Estimating treatment effects from single-arm trials via latent-variable modeling

---

Manuel Haußmann<sup>1</sup>

Tran Minh Son Le<sup>1</sup>

Viivi Halla-aho<sup>3</sup>

Samu Kurki<sup>4</sup>

Jussi V. Leinonen<sup>4</sup>

Miika Koskinen<sup>3</sup>

Samuel Kaski<sup>1,2</sup>

Harri Lähdesmäki<sup>1</sup>

<sup>1</sup>Department of Computer Science  
Aalto University  
Espoo, Finland

<sup>2</sup>Department of Computer Science  
University of Manchester  
Manchester, United Kingdom

<sup>3</sup>ICT Management  
Helsinki University Hospital  
Helsinki, Finland

<sup>4</sup>Bayer Oy  
Espoo, Finland

## Abstract

Randomized controlled trials (RCTs) are the accepted standard for treatment effect estimation but they can be infeasible due to ethical reasons and prohibitive costs. Single-arm trials, where all patients belong to the treatment group, can be a viable alternative but require access to an external control group. We propose an identifiable deep latent-variable model for this scenario that can also account for missing covariate observations by modeling their structured missingness patterns. Our method uses amortized variational inference to learn both group-specific and identifiable shared latent representations, which can subsequently be used for (i) patient matching if treatment outcomes are not available for the treatment group, or for (ii) direct treatment effect estimation assuming outcomes are available for both groups. We evaluate the model on a public benchmark as well as on a data set consisting of a published RCT study and real-world electronic health records. Compared to previous methods, our results show improved performance both for direct treatment effect estimation as well as for effect estimation via patient matching.


## 1 INTRODUCTION

Randomized controlled trials (RCTs) are the ‘gold standard’ in medical research and social sciences for the estima-

tion of treatment effects. When conducted properly, RCTs provide control over the treatment assignment and therefore the removal of confounding factors. However, this advantage comes with the price of several shortcomings. Recruiting a sufficient number of patients and collecting their data is usually a costly endeavor both in terms of time as well as financial investments, limiting us to smaller sample sizes in general. The study size constraint may further be exacerbated in the case of rare diseases, and RCTs also have ethical challenges in case of serious diseases when no one can be left untreated. Single-arm trials offer a way out of these predicaments at the price of requiring external control data, thus no longer being randomized.

These control observations are usually taken from the control arms of historical RCTs, or from so-called real-world data (RWD) in the form of electronic health records (EHR) are a promising source for these control observations. In the medical domain, EHR refers to records collected during regular healthcare and hospital visits and not as part of a specific RCT study. The degree to which RWD and RCT can replace and complement each other is an open research question, both theoretical (Collins et al., 2020; Eichler et al., 2021) as well as practical (Franklin and Schneeweiss, 2017). Large-scale projects evaluating trial replicability with RWD data are currently ongoing to better understand their respective strengths and limitations (Dahabreh et al., 2020; Franklin et al., 2021).

In this work, we consider the task of augmenting single-arm trials with external controls (Gray et al., 2020; Schmidli



**Figure 1: Overview.** For the task of **treatment effect estimation** from **single-arm trial data**, patient information for a control group has to be extracted from **electronic health records** which have been collected, e.g., during hospital visits. Due to their different sources the two covariate distributions only partially overlap. Our model maps them into group-specific latent spaces and a shared, identifiable predictive space. This low-dimensional representation can then be subsequently used for treatment effect estimation. If outcome information is available for both groups, we can obtain a *direct estimate* of the effect from the potential outcomes ( $Y(0)$ ,  $Y(1)$ ). We introduce an additional task where treatment outcome information is available only for the control group. In this scenario our method estimates the treatment effect via *patient matching*.

et al., 2020; Chen et al., 2021). To estimate treatment effects we have to assume that these two groups have some intrinsic similarity. Yet, given their different sources, their covariate specifics will vary, e.g., concerning the availability of medical history records, breadth of lab measurements, measurement precision, variance in the demographics, etc. We focus specifically on three problems that arise in this setting. First, due to the differences in the covariate distributions between data collected in RCT and RWD, there is limited overlap between the two sets of observations which has to be overcome to provide a reliable treatment effect estimate (D’Amour et al., 2021). Second, most prior work in the machine learning literature (Johansson et al., 2016; Shalit et al., 2017; Shi et al., 2019; Curth and van der Schaar, 2021b; Bica and van der Schaar, 2022; Jiang et al., 2023) has focused on the assumption of having access to treatment outcome information from both the single-arm trial patients as well as the external controls. We extend this assumption by considering the additional task of having to find a suitable set of external controls pre-treatment (i.e. before the treatment has been administered to the single-arm group). Here, we have access to pre-treatment covariates for both groups as well as the outcome values for the external control group. This setup is highly relevant as it avoids leaking any information from the trial results to the model during the model learning and matching of the patients. Without such new techniques, the inclusion of trial results while inferring the latent variable model could bias further statistical estimates that rely on the learned latent representation. Finally, we have to account for the fact that real-world data will usually include missing measurements whose non-random patterns have to be accounted for and modeled properly.

We introduce a latent variable model to infer group-specific as well as shared identifiable latent representations for subsequent treatment effect estimation. The former allows us to explain away specifics that are unique to the treatment and the control subsets, respectively, while a shared representation provides a compressed latent space for treatment effect estimation as well as patient matching, i.e., the selection of a subset of control patients that are most similar to the treated group. This differs from the popular approach in the medical literature that relies on various forms of propensity score estimators (Stuart, 2010), which do not infer such a space prior to estimating a matching score. The problem setting and our approach are conceptualized in Figure 1.

**Contributions.** In this work, we consider the task of estimating treatment effects from single-arm trials with external controls and also introduce an additional scenario where outcome information is not available for the treatment group during inference. We contribute

- (i) a principled way of handling these tasks via amortized latent-variable models with identifiability guarantees that can infer a predictive latent space between two different covariate distributions for subsequent treatment effect prediction, and simultaneously model structured missingness patterns;
- (ii) an extensive ablation study demonstrating that our method is competitive if outcome information is available for both treated and control groups, as well as if it is available only for the latter. The method improves upon prior work on several variations of a semi-synthetic benchmark as well as on a curated large-scale data set that combines data from a published RCT study (Bakris et al., 2020) with real-world EHRs.

## 2 RELATED WORK

**Treatment Effect Estimation.** Treatment effect estimation from observational studies (OS) has a long history. See, e.g., [Imbens \(2004\)](#) for a review on *average treatment effect (ATE)* estimation approaches used within the statistics and econometrics literature, which rely on a variety of machine learning approaches, e.g., random forests ([Hill, 2011](#); [Athey et al., 2019](#)). Given the increasing amount of observational data available, a recent trend in the machine learning literature has been to focus on models for *individual treatment effect* or *conditional average treatment effect (CATE)* estimations (see, e.g., [Bica et al. \(2021\)](#) for an introduction). A broad spectrum of approaches exists to construct deterministic CATE estimators ([Künzel et al., 2019](#)). See [Curth and van der Schaar \(2021a\)](#) for a recent overview and unifying framework to classify different neural network-based approaches. A common approach is to build upon the theory of representation learning ([Bengio et al., 2013](#)) to construct deterministic representations that are then used by task-specific mappings ([Johansson et al., 2016](#); [Shalit et al., 2017](#); [Shi et al., 2019](#); [Jiang et al., 2023](#)). In parallel, there is an increasing interest in using generative models for the task of CATE estimation, relying, e.g., on variational autoencoders ([Kingma and Welling, 2014](#); [Louizos et al., 2017](#); [Lu et al., 2020](#); [Zhang et al., 2021](#)), generative adversarial networks ([Goodfellow et al., 2014](#); [Yoon et al., 2018](#)), or energy-based approaches ([Lecun et al., 2006](#); [Zhang et al., 2022](#)). Our method belongs to this second group of generative approaches, focuses on amortized variational inference, and extends prior work to the specific task of learning with external controls in settings that can have strongly divergent groups of samples.

**Combining RCT with OS.** Combining small randomized control trial data sets with observational data has been the focus of several studies in recent years. It can be seen as a form of domain adaptation ([Ben-David et al., 2010](#)) as the aim is to combine different data sources with similar yet distinct characteristics. [Kallus et al. \(2018\)](#) consider combining the two to get rid of hidden confounders requiring the very restrictive assumption of linear biases. [Cheng and Cai \(2021\)](#) propose to learn separate estimators, combining them in a second step via a weighting scheme. However, finding these weights requires a separate validation set. [Hatt et al. \(2022\)](#) propose a two-step approach in which they first rely on learning a preliminary representation solely from observational study data, which they aim to fine-tune by learning data-specific structures via the randomized trial data. While these approaches focus on the combination of two complete studies, our setting focuses on single-arm information from each of our groups. Closest to our formulation are [Bica and van der Schaar \(2022\)](#), who consider combining multiple heterogeneous data sources into a joint model for CATE estimation, requiring full treatment outcome information for all data sources.

**External Controls and Patient Matching.** Matching consists of pairing each patient in the treatment group with a suitable patient from the set of external controls, based on a similarity score, and has long been a topic of ongoing research with a wide spectrum of approaches ([Stuart, 2010](#)). E.g., [Li and Fu \(2017\)](#) rely on kernel methods and maximum mean discrepancy regularization to construct latent spaces for subsequent matching. [Luo and Zhu \(2020\)](#), in turn, rely on central subspaces. [Athey et al. \(2018\)](#) can approximately ensure balancing between the covariates in high-dimensional spaces under the assumption of only linear biases between them.

**Unconfoundedness and Identifiability.** A common assumption for causal inference is the absence of hidden confounders. To avoid this assumption prior works rely, e.g., on having multiple causes ([Wang and Blei, 2019](#)), access to multiple treatments over time ([Bica et al., 2020](#)), or by relying on invariant risk minimization ([Arjovsky et al., 2019](#); [Shi et al., 2019](#)). Obtaining guarantees for identifiability of learned latent spaces is similarly an area of ongoing research ([Xi and Bloem-Reddy, 2023](#); [Moran et al., 2022](#)). Our approach relies on using auxiliary variables following the work by [Khemakhem et al. \(2020\)](#).

**Estimation under Missingness.** Modeling with missing treatment outcome information has been pursued primarily with the goal of average treatment effect estimation ([Williamson et al., 2012](#); [Zhang et al., 2016](#); [Kennedy, 2020](#)). [Kuzmanovic et al. \(2022\)](#) generalize such prior approaches by focusing on building CATE estimators with missing treatment information. We differ from these as we do not assume randomly missing outcome data but are restricted to *systematically* missing treatment outcome information from the treatment group. This is necessary when we want to match a set of external control patients with a group of patients recruited specifically for the study at hand. Additionally, in real-world applications, we are often faced with missing covariate information ([Perez-Lebel et al., 2022](#)). Modeling missing not at random (MNAR) covariates ([Rubin, 1976](#)) has received a lot of focus in the generative modeling literature recently ([Collier et al., 2020](#); [Ipsen et al., 2021](#); [Ghalebikesabi et al., 2021](#)). Throughout this work we build on the work by [Collier et al. \(2020\)](#).

## 3 BACKGROUND

**Problem Specification.** We have access to a sample of  $N$  patients with covariates  $\mathbf{x} \in \mathcal{X}$ , outcomes  $y \in \mathcal{Y}$  and treatment assignments  $t \in \{0, 1\}$ , such that  $\mathcal{D} = \{(\mathbf{x}_1, y_1, t_1), \dots, (\mathbf{x}_N, y_N, t_N)\}$ . Following [Rubin \(2005\)](#)'s potential outcomes formulation we assume potential outcomes  $Y_n(0), Y_n(1) \in \mathcal{Y}$  of which due to the fundamental problem of causality ([Pearl, 2009](#)) only one can be observed, such that  $y_n = t_n Y_n(1) + (1 - t_n) Y_n(0)$ .

Throughout, we have to assume three standard assumptions necessary for estimating causal effects from observational

data: (i) *consistency*, i.e., if patient  $n$  has received treatment  $t_n$  we observe potential outcome  $y_n = Y_n(t_n)$ ; (ii) *unconfoundedness*,  $Y(0), Y(1) \perp\!\!\!\perp T|X$ , i.e., there are no unobserved confounders; (iii) *overlap*,  $0 < \pi(x) < 1 \forall x$ , i.e., the treatment assignment is not deterministic, where  $\pi(x) \triangleq \mathbb{P}(T = 1|X = x)$ , is known as the *propensity score*. Assumptions (ii) and (iii) provide us with tension as they can be difficult to simultaneously fulfill (D’Amour et al., 2021). To ensure unconfoundedness, including an increasing number of covariates is usually considered to be helpful. However, to ensure overlap, a low-dimensional space is desired. We see below how this tension is handled within our model. Given the diverse nature of the treatment and control groups in our setup, this third constraint is the most critical and requires careful consideration.

Observing outcomes of both arms, we estimate the *conditional average treatment effect (CATE)*,

$$\tau_C(x) \triangleq \mathbb{E}[Y(1) - Y(0)|X = x] = \mu_1(x) - \mu_0(x), \quad (1)$$

where  $\mu_t(x) \triangleq \mathbb{E}[Y(t)|X = x]$ . The *average treatment effect (ATE)* is then defined as  $\mathbb{E}[\tau_C(X)]$ .

The main focus of our work is the case where outcomes are only available for the control group, in which case we consider the *average treatment effect for the treated (ATT)*,

$$\tau_A \triangleq \mathbb{E}[Y(1) - Y(0)|T = 1], \quad (2)$$

and estimate it with a suitable control group by matching external control patients with the treated.


## 4 OUR MODEL

We introduce a latent-variable model for the task of treatment effect estimation from single-arm trials with external controls focusing also on the task where outcome information is only available for the control group.

**A Generative Model.** Previous models (Louizos et al., 2017; Lu et al., 2020) have so far solely focused on learning latent embeddings within the setting of having observed outcome information for both the treatment and the control groups. From their joint

$$p(t, \mathbf{x}, \mathbf{y}, \mathbf{z}) = p(t|\mathbf{z})p(\mathbf{x}|\mathbf{z})p(\mathbf{y}|t, \mathbf{z})p(\mathbf{z}), \quad (3)$$

we notice the dual role played by the latent representation  $\mathbf{z}$ . In (3)  $\mathbf{z}$  is assumed to be predictive of  $\mathbf{y}$  and, at the same time, to offer a reliable encoding of  $\mathbf{x}$ . These tasks counteract each other. Accurately modeling  $p(\mathbf{y}|t, \mathbf{z})$  requires extracting predictive information from the covariates  $\mathbf{x}$  (as  $\mathbf{z}$  is latent), which we assume to sufficiently overlap between the treatment and the control groups. Modeling  $p(\mathbf{x}|\mathbf{z})$ , however, forces the model to encode a complete representation, i.e., to take covariate information into consideration that is irrelevant to the predictive task, and to further model group-specific variations existing in the data. This was observed



**Figure 2: Plate Diagram of Our Model.** Black solid arrows denote the generative model, red dashed arrows the inferential dependency. Empty, partially filled, and filled circles refer to latent, partially, and fully observed variables.

to be detrimental, e.g., by Lu et al. (2020), who dropped it from their final objective.

We take these different tasks into account and reduce the tension by extending the original latent space representation by adding an additional latent variable  $\mathbf{u} \in \mathcal{U}$  to the joint distribution. This variable aims to explain away group-specific characteristics, allowing  $\mathbf{z}$  to be predictive of the treatment, yet still be guided by the reconstructive task it performs jointly with  $\mathbf{u}$  through the resulting likelihood  $p(\mathbf{x}|\mathbf{u}, \mathbf{z})$ . However, as deep unconstrained latent-variable models are not identifiable (Khemakhem et al., 2020)  $\mathcal{Z}$  is also not identifiable due to  $p(\mathbf{z})$  being unconditional. We instead have to rely on a conditional prior  $p(\mathbf{z}|\mathbf{c})$ , where  $\mathbf{c}$  is an additional observed variable, separate from the covariates  $\mathbf{x}$ . In practice,  $\mathbf{c}$  contains covariates that are best considered as non-random variables, such as gender, age, country of origin, etc. With a normal prior for  $\mathbf{u}$ , our model (see Figure 2) is given as

$$\begin{aligned} \mathbf{u} &\sim p(\mathbf{u}) = \mathcal{N}(\mathbf{u}|0, \sigma_u^2 \mathbb{1}), \\ \mathbf{z} &\sim p(\mathbf{z}|\mathbf{c}) = \mathcal{N}(\mathbf{z}|\lambda(\mathbf{c}), \sigma_z^2 \mathbb{1}), \\ \mathbf{x} &\sim p(\mathbf{x}|\mathbf{u}, \mathbf{z}), \\ t &\sim p(t|\mathbf{z}) = \text{Ber}(t|\sigma(f(\mathbf{z}))), \\ y &\sim p(y|t, \mathbf{z}) = \mathcal{N}(y|\mu_t^y(\mathbf{z}), \alpha^{-1}), \end{aligned}$$

where  $\lambda(\cdot)$ ,  $f(\cdot)$  and  $\mu_t^y(\cdot)$  are neural nets,  $\alpha$  is a precision parameter,  $\sigma(x) \triangleq 1/(1 + \exp(-x))$  the logistic sigmoid, and  $\mathbb{1}$  an identity matrix. We assume the likelihood  $p(\mathbf{x}|\mathbf{u}, \mathbf{z})$  to factorize over the covariates, i.e.,  $p(\mathbf{x}|\mathbf{u}, \mathbf{z}) = \prod_{i=1}^{d_x} p(x_i|\mathbf{u}, \mathbf{z})$ , where  $d_x \triangleq \dim(\mathbf{x})$ . In the experiments, we model the individual covariates  $x_i$  via Bernoulli, Categorical, or Normal distributions, depending on their domain. For continuous covariates, we assume a homoscedastic noise model, i.e.,  $\mathcal{N}(x_i|g(\mathbf{z})_i, \beta_i^{-1})$  with precision parameters  $\beta_i$ . The precisions  $\alpha$ , and  $\{\beta_i\}_i$  are optimized together with the remaining model parameters. Assuming  $\mathcal{U}$  is high-dimensional enough, we can rely on it to model the group-specific latent spaces of both the treatment and control groups. Alternatively, they can be explicitly separated into group-specific latent spaces, such that  $\mathcal{U} = [\mathcal{U}_0, \mathcal{U}_1]$ . This allows for an easy extension of the current two-group setup to multiple groups which may only

become available at a later stage. We explore this variant in our experiments but do not find the additional separation to be necessary in most cases.

**Inference.** We rely on amortized variational inference (Kingma and Welling, 2014; Rezende et al., 2014) to infer an approximation to the intractable posterior  $p(\mathbf{U}, \mathbf{Z}|\mathcal{D})$ , where  $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_N)$ ,  $\mathbf{C}$ ,  $\mathbf{X}$ ,  $\mathbf{Z}$  analogously. We assume a mean-field variational posterior

$$q(\mathbf{U}, \mathbf{Z}|\mathbf{C}, \mathbf{X}) = \prod_{n=1}^N q(\mathbf{u}_n|\mathbf{x}_n)q(\mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n), \quad \text{where}$$

$$q(\mathbf{u}_n|\mathbf{x}_n) = \mathcal{N}(\mathbf{u}_n|\mu^u(\mathbf{x}_n), \sigma^u(\mathbf{x}_n)) \quad \text{and}$$

$$q(\mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n) = \mathcal{N}(\mathbf{z}_n|\mu^z(\mathbf{c}_n, \mathbf{x}_n), \sigma^z(\mathbf{c}_n, \mathbf{x}_n)),$$

with neural networks  $\mu^u(\cdot)$ ,  $\mu^z(\cdot)$ ,  $\sigma^u(\cdot)$ , and  $\sigma^z(\cdot)$ . We minimize the Kullback-Leibler divergence between the two, i.e.,  $\text{KL}(q(\mathbf{U}, \mathbf{Z}|\mathbf{C}, \mathbf{X}) \| p(\mathbf{U}, \mathbf{Z}|\mathcal{D}))$ , by maximizing the corresponding *evidence lower bound (ELBO)* with respect to all parameters. If outcome values are available for both groups, the ELBO is given as

$$\begin{aligned} \log p(\mathcal{D}) &\geq \sum_{n=1}^N \mathbb{E}_{q(\mathbf{u}_n, \mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n)} [\log p(\mathbf{x}_n|\mathbf{u}_n, \mathbf{z}_n)] \\ &\quad + \mathbb{E}_{q(\mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n)} \left[ \underbrace{\log p(y_n|t_n, \mathbf{z}_n)}_{(*)} + \log p(t_n|\mathbf{z}_n) \right] \\ &\quad - \text{KL}(q(\mathbf{u}_n|\mathbf{x}_n) \| p(\mathbf{u}_n)) \\ &\quad - \text{KL}(q(\mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n) \| p(\mathbf{z}_n|\mathbf{c}_n)) \triangleq \mathcal{L}_{\text{elbo}}^{\text{full}}. \end{aligned} \quad (4)$$

If outcome values  $y_n$  are available only for the control group, we optimize a modified ELBO. The missing outcome for the single-arm treatment group is masked out, i.e., we replace  $(*)$  in the equation above with

$$\mathbb{E}_{q(\mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n)} [(1 - t_n) \log p(y_n|t_n, \mathbf{z}_n)].$$

See Appendix B for further details.

We follow prior work (Johansson et al., 2016; Shalit et al., 2017; Lu et al., 2020) and rely on tools from domain adaptation (Ben-David et al., 2010) to further constrain  $\mathbf{z}$ . Specifically, we rely on the gradient-reversal layer approach by Ganin and Lempitsky (2015) providing us with a generative adversarial network-based regularizer (Goodfellow et al., 2014). Our model tries to fool a discriminator whose task is to distinguish between the encoded samples from the two groups, i.e., whose aim is to maximize  $\log p(t_n|\mathbf{z}_n)$ . See Appendix C for details on the modified objective.

**Identifiability.** As mentioned above, fully unconstrained latent-variable models,  $p_\theta(\mathbf{x}, \mathbf{z}) = p_\theta(\mathbf{x}|\mathbf{z})p_\theta(\mathbf{z})$ , are non-identifiable in the sense that the implication

$$\forall(\theta, \theta') : p_\theta(\mathbf{x}) = p_{\theta'}(\mathbf{x}) \Rightarrow \theta = \theta',$$

does not necessarily hold (Khemakhem et al., 2020), i.e., different model parameterizations can lead to the same

marginal. Khemakhem et al. (2020) propose to constrain the model via an additional set of variables  $\mathbf{c}$  and assume a joint distribution that factorizes as

$$p_\theta(\mathbf{x}, \mathbf{z}|\mathbf{c}) = p_\theta(\mathbf{x}|\mathbf{z})p_\theta(\mathbf{z}|\mathbf{c}),$$

where  $\mathbf{x} = f_\theta(\mathbf{z}) + \varepsilon$ , with  $\varepsilon \sim p(\varepsilon)$  being an independent noise variable. The prior  $p(\mathbf{z}|\mathbf{c})$  is assumed to be parameterizable as a factorizing exponential family,

$$p(\mathbf{z}|\mathbf{c}) = \prod_{i=1}^{d_z} m(z_i)/Z(\mathbf{c}) \exp(T(z_i)^\top \lambda(\mathbf{c})). \quad (5)$$

Assuming a sufficiently diverse set of variables  $\mathbf{c} \in \mathcal{C}$ , there need to be  $d_z k + 1$  distinct values, where  $k$  is the number of sufficient statistics, then  $\theta \triangleq (f, T, \lambda)$  is identifiable up to permutations and translations. Xi and Bloem-Reddy (2023) later showed that the assumption on a factorizing exponential family is not necessary and that, due to the requirement on the independent noise variable  $\varepsilon$  in  $p(\mathbf{x}|\mathbf{z})$  theoretical guarantees only hold for continuous covariates  $\mathbf{x}$ . Empirical results by Khemakhem et al. (2020) indicate that identifiability still tends to be achievable for discrete covariates. Throughout our experiments, whenever we encounter a mixture of discrete and continuous covariates, we still rely on a prior as specified in (6) and observe improved results, while theoretical guarantees could be maintained by ignoring discrete covariates within the likelihood.

Within our setup, we are interested in identifying the predictive latent space of  $\mathbf{z}$ . We keep the prior  $p(\mathbf{u})$  unconstrained and model  $\mathbf{z}$ 's conditional prior  $p(\mathbf{z}|\mathbf{c})$  as a factorized normal distribution, whose  $\lambda(\cdot)$  are parameterized by a neural net. The argument by Khemakhem et al. (2020) then guarantees identifiability of  $\mathbf{z}$  up to the constraints mentioned above. See Appendix B.1 for a detailed discussion on the required conditions and derivations.

**Patient Matching.** Matching consists of extracting a subset of observations from the control group that is similar to the observations in the single-arm treatment group with respect to a feature  $s$ . Such a feature is often created via the propensity score.<sup>1</sup> Given  $s$ , a distance measure estimates the similarity between two observations. That is, the matching control  $\mathbf{x}_{\text{match}}^i$  for an observation  $\mathbf{x}_i$  with  $s_i$  is

$$\mathbf{x}_{\text{match}}^i = \mathbf{x}_j, \quad \text{with } j = \arg \min_{j \in J_c} d(s_i, s_j),$$

where  $d(\cdot, \cdot)$  is a distance metric and  $J_c = \{j|t_j = 0\}$  is the index set over the control observations.

As  $\mathcal{Z}$  is designed to extract a predictive encoding, we perform patient matching within this compressed, low-dimensional latent space, and match in  $\mathcal{X}$  only for ablation

<sup>1</sup>A classical result by Rosenbaum and Rubin (1983) proves that if the average treatment effect is identifiable from observational data after adjusting for  $X$ , then adjusting for the propensity score  $\pi(X)$  is sufficient.

purposes. If we consider the inferred posterior means in  $\mathcal{Z}$  as point estimates that encode the covariates, we either (i) estimate a propensity score given the encodings  $\mu^z(\mathbf{x}_i)$ , i.e.,  $s_i = \hat{\pi}(T = 1 | \mu^z(\mathbf{x}_i))$ , (ii) use the posterior mean directly, i.e.,  $s_i = \mu^z(\mathbf{x}_i)$ , or (iii) the full variational posterior, i.e.,  $s_i = (\mu^z(\mathbf{x}_i), \sigma^z(\mathbf{x}_i))$ . Given similarity distances, we rely on nearest-neighbor matching with replacement. This straightforward approach could be extended to further improve the overall overlap by taking additional measures into account, such as, e.g., calipers (Austin, 2011). Throughout the experiments, we use the Euclidean distance for  $d(\cdot, \cdot)$ . See Appendix C.6 for matching without replacement and an ablation using distributional distance measures.

**Missingness.** So far, our model assumes fully observed covariates  $\mathbf{x}$  — an assumption that is often violated in practice. We generally cannot assume that EHRs have measurements of all the features we require in a specific study for every patient. Additionally, we cannot assume that missing measurements are completely random but rather have to assume dependency structures within the missingness, i.e., missing not at random (MNAR) (Rubin, 1976).<sup>2</sup> To account for this we build upon the proposal by Collier et al. (2020). Given masking variables  $\mathbf{m} \in \mathcal{M} = \{0, 1\}^{d_x}$ , where  $\mathbf{m}_{ni} = 1$  indicates that the  $i$ -th covariate has been observed in the  $n$ -th sample, we observe masked covariates  $\tilde{\mathbf{x}} = \mathbf{m} \circ \mathbf{x} + (1 - \mathbf{m}) \circ \boldsymbol{\eta}$ , for some imputation  $\boldsymbol{\eta}$  and  $\circ$  the element-wise product. To account for the MNAR structure, we model  $\mathbf{m}$  to depend on an additional latent variable  $\mathbf{z}^m \sim p(\mathbf{z}^m)$  as well as  $(\mathbf{u}, \mathbf{z})$  such that

$$\mathbf{m} | \mathbf{z}^m, \mathbf{u}, \mathbf{z} \sim \prod_{i=1}^{d_x} \text{Ber}(m_i | \sigma(g(\mathbf{z}^m, \mathbf{u}, \mathbf{z})_i)),$$

$$\mathbf{x} | \mathbf{m}, \mathbf{u}, \mathbf{z}^x \sim p(\mathbf{x} | \mathbf{m}, \mathbf{u}, \mathbf{z}^x).$$

See Appendix B.2 for the complete model.

## 5 EVALUATION

**Data Sets.** Given the fundamental problem of causality, we have to rely on semi-synthetic data sets to evaluate the efficacy of the proposed method. We first consider adaptations of the commonly used semi-synthetic IHDP benchmark (Hill, 2011) consisting of 25 real covariates (one categorical, 19 binary, five continuous) and synthetic continuous treatment outcomes  $y$ . Additionally, we create a survival analysis setup where  $y$  represents time-to-event information, with the aim of estimating hazard ratios. As a second real-world data set, we combine 140, partially missing, covariates (100 binary, 40 continuous) of 833 participants of a published RCT study (Bakris et al., 2020) with real-world electronic health record (EHR) data from 2646 patients. Details on these data sets as well as further experimental details not mentioned in the subsequent subsections

<sup>2</sup>While MNAR is our focus in this work, the generative model could easily be adapted to missing (completely) at random assumptions as well.

are discussed in Appendix C. As far as possible, all variants share the same architectures and hyperparameters. See Appendix D for further evaluations. An implementation of our proposed approach is available at [https://github.com/manuelhaussmann/lvm\\_singlearm](https://github.com/manuelhaussmann/lvm_singlearm).

**Experimental Assumptions.** We consider two cases. Either (a) *outcome information is available for both groups*, i.e., the single-arm trial group and the external control group, which is the common assumption used in the deep learning literature. In this case, estimators can be inferred directly. Or (b) *outcome information is only available for the control group*. Here, inference of a representation space is required, which can be used in a second step to select a suitable subset of control patients via matching. We evaluate scenarios where either *all* covariates are potentially predictive, or only a *subset* of them, and where the covariate overlap between the two groups is either *high*, where we keep the original covariates, which are already only partially overlapping, or *low*, where we additionally shift them to decrease their overlap even further.

**Baselines & Ablations.** We compare our method to several baselines. *SingleNet* serves as our simplest deterministic baseline and learns a single estimator  $\mu(x, t)$ . *TNet* generalizes this to two separate neural net-based estimators  $\mu_t(x)$ . *TarNet* and *CFRNet* in turn learn a shared representation space from which they predict treatment outcomes (Shalit et al., 2017). *SNet* further generalizes this by learning a mixture of unique and shared representations (Curth and van der Schaar, 2021a). VAE serves as our vanilla generative baseline following the standard architecture proposed by Shalit et al. (2017). *CEVAE* extends it with an additional treatment outcome estimator (Louizos et al., 2017). Finally, *TEDVAE* generalizes CEVAE by splitting the latent space into separate predictive parts similar to the deterministic SNet generalization (Zhang et al., 2021). We further compare with *CFor*, a random forest-based approach for causal estimation (Wager and Athey, 2018), and three propensity score estimators: *PScov* uses the observed covariates, *PSpca* maps them into a  $d_z$ -dimensional feature vector, and *PSlat* learns an estimator on the latent encoding space  $\mathcal{Z}$  our approach infers. All three use an ElasticNet (Zou and Hastie, 2005) to infer the estimator.

We report the following ablations of our model. (*Our*) optimizes (4) with an unconditional, i.e., non-identifiable prior  $p(\mathbf{z})$ , (*+I*) specifies the conditional  $p(\mathbf{z} | \mathbf{c})$ . (*+sep*) infers separate representations  $\mathcal{U}$  for each group. For simplicity, all variants follow the TarNet architecture as closely as possible. Finally, (*+snet*) and (*+tedvae*) demonstrate that our proposal can easily be implemented with modern deep models by simply replacing the TarNet architecture backbone with SNet or TEDVAE (see Appendix C.3 for a more detailed discussion of such replacements).

**Table 1: Full & Partial Outcome Observation.** Variations of modified IHDP covariates and simulated treatment outcomes. Our proposal improves on both deterministic and probabilistic baselines. (‘-’ refers to a method not being applicable)

METHOD	(a) full outcome observation (RMSE of CATE)				(b) partial outcome observation (AE of ATT)			
	all+high		subset+low		all+high		subset+low	
	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample
CFor	0.556±0.000	0.5967±0.001	3.449±0.052	2.597±0.031	-	-	-	-
PScov	-	-	-	-	0.115±0.005	0.336±0.017	0.691±0.059	0.702±0.063
PSpca	-	-	-	-	0.096±0.005	0.300±0.014	0.496±0.036	0.613±0.051
PSlat	-	-	-	-	0.084±0.004	0.187±0.013	0.276±0.025	0.272±0.023
SingleNet	0.241±0.005	0.328±0.009	0.924±0.016	0.878±0.027	-	-	-	-
TNet	0.175±0.003	0.275±0.008	0.278±0.009	0.328±0.017	-	-	-	-
TARNet	0.177±0.004	0.280±0.008	0.272±0.009	0.333±0.019	0.043±0.003	0.130±0.008	0.140±0.011	0.239±0.016
CFRNet	0.171±0.003	0.279±0.008	0.279±0.008	0.338±0.017	0.042±0.002	0.131±0.007	0.345±0.016	0.370±0.017
SNet	0.168±0.003	0.264±0.008	0.211±0.008	<b>0.243±0.016</b>	0.043±0.003	0.137±0.007	0.221±0.015	0.366±0.026
VAE	-	-	-	-	0.085±0.004	0.266±0.012	0.264±0.014	0.427±0.022
CEVAE	0.182±0.002	0.287±0.007	0.288±0.009	0.350±0.017	0.052±0.002	0.174±0.009	0.150±0.009	0.259±0.015
TEDVAE	0.176±0.002	0.283±0.007	0.293±0.009	0.336±0.018	0.052±0.003	0.179±0.009	0.144±0.009	0.272±0.016
Ours	0.152±0.002	0.267±0.008	0.201±0.007	0.272±0.016	<b>0.037±0.002</b>	<b>0.113±0.007</b>	<b>0.114±0.008</b>	<b>0.190±0.016</b>
+I	0.141±0.002	0.251±0.008	0.194±0.007	0.262±0.016	<b>0.038±0.002</b>	<b>0.122±0.007</b>	<b>0.110±0.007</b>	<b>0.176±0.014</b>
+sep	0.151±0.002	0.267±0.008	0.209±0.007	0.278±0.016	<b>0.037±0.002</b>	<b>0.111±0.007</b>	0.117±0.009	<b>0.193±0.016</b>
+sep+I	0.143±0.002	0.255±0.007	0.197±0.007	0.266±0.017	<b>0.035±0.002</b>	<b>0.117±0.008</b>	<b>0.104±0.007</b>	<b>0.185±0.015</b>
+snet	0.141±0.002	0.256±0.008	<b>0.185±0.007</b>	<b>0.253±0.017</b>	<b>0.038±0.002</b>	<b>0.117±0.007</b>	0.194±0.014	0.322±0.020
+snet+I	<b>0.130±0.002</b>	<b>0.242±0.008</b>	<b>0.181±0.006</b>	<b>0.250±0.017</b>	<b>0.038±0.002</b>	<b>0.120±0.008</b>	0.190±0.012	0.322±0.023
+snet+sep	0.143±0.002	0.254±0.008	0.191±0.007	0.258±0.017	0.040±0.003	<b>0.109±0.007</b>	0.179±0.013	0.316±0.020
+snet+sep+I	<b>0.131±0.002</b>	<b>0.247±0.008</b>	<b>0.180±0.007</b>	<b>0.251±0.017</b>	0.039±0.003	<b>0.115±0.007</b>	0.196±0.015	0.292±0.020
+tedvae	0.143±0.002	0.257±0.008	<b>0.187±0.007</b>	0.258±0.019	<b>0.039±0.003</b>	<b>0.119±0.007</b>	<b>0.103±0.008</b>	<b>0.189±0.015</b>
+tedvae+I	0.142±0.002	0.255±0.008	0.191±0.006	<b>0.253±0.016</b>	<b>0.036±0.002</b>	0.124±0.007	<b>0.102±0.007</b>	<b>0.191±0.017</b>

mean ± standard error over 300 random replications; statistically significant best models marked **bold**; lowest mean underlined

**Performance Metrics.** If outcome information is available for both groups (case (a)) we evaluate the root mean squared error (RMSE) between true and estimated CATE values, also known as the *precision in estimation of heterogeneous effects (PEHE)* (Hill, 2011). Otherwise (case (b)) we compute the ATT between the treated and the matched subset of the control and report the absolute error (AE) of the estimate. We consider *within-sample* and *out-of-sample* performance. Within-sample considers covariates  $\mathbf{x}$  and outcomes  $y = Y(t)$  already observed during training and requires inferring the counterfactual  $Y(1 - t)$ . Out-of-sample performance considers previously unseen covariates  $\mathbf{x}$  and potential outcomes for the prediction.


### 5.1 Full Outcome Observation

We summarize the performance of CATE estimation on IHDP covariates assuming outcomes are available for both groups (treatment and control) in Table 1 (a). Comparing the performance of deterministic and latent-variable approaches, we observe that switching from a sufficiently advanced deterministic model (*TNet*, *TARNet*, *CFRNet*, *SNet*) to a simple generative model (*CEVAE*) tends to reduce performance throughout the different scenarios. As was argued already in prior work by Lu et al. (2020), forcing the latent encoding  $\mathcal{Z}$  to fully encode the covariates instead of focusing on being predictive constraints too much. However, splitting the latent space into  $\mathcal{U}$  and  $\mathcal{Z}$  clearly improves upon the deterministic baselines (*Ours*). If  $\mathcal{U}$  is sufficiently high-dimensional, further splitting it into group-specific latent spaces provides little to no benefit in this

setup (*+sep*). Switching from an unconditional prior ( $p(\mathbf{z})$ ) to a conditional one ( $p(\mathbf{z}|\mathbf{c})$ ) improves the predictive performance even further (*+I*). Finally, switching to a more modern architecture (*+snet*) gives an additional reduction in RMSE. The tree-based baseline (*CFor*) performs significantly worse than all neural net-based methods on this setting, as was already observed by Curth and van der Schaar (2021a), and fails completely if there is only limited overlap between the two sets of covariates. Further results are provided in Appendix D.2.

### 5.2 Partial Outcome Observation

We summarize the performance of ATT estimation on IHDP covariates assuming outcomes are available only for the control group in Table 1 (b). In this setting, as before, switching from a deterministic baseline to a generative variant requires the added structure provided by a separate latent space  $\mathcal{U}$  for significant improvements (*Ours*). While relying on an identifiable version of our approach (*+I*) tends to improve performance, it is not significant, while relying on an SNet structure may even decrease performance depending on the setup. Comparing all baselines with a VAE that simply learns a joint representation for both groups without any outcome information (*VAE*) shows that some guidance, even if it is only accessible from a single group, is necessary to infer a representation suitable for subsequent matching (*Ours*). While matching based on a propensity score estimator is not competitive with the deep models, it improves significantly when it is trained on an encoded set of covariates (*PSlat*) compared to the observed ones (*PScov*).



**Figure 3: Missingness.** Moving from a scenario where all covariates are observed ( $\blacklozenge$ ) to one with structured missingness that is not modeled ( $\blacklozenge$ ) reduces performance, as expected. Modeling the MNAR pattern as part of our generative model ( $\blacklozenge$ ) reliably improves performance in all of our variants. Visualized is the RMSE of within-sample CATE estimation in the *all+high* scenario. Shown are 100 random replications along with their respective means ( $\blacklozenge$ ).

Further results are provided in [Appendix D.3](#).

### 5.3 Missingness

We summarize the performance of CATE estimation on IHDP covariate in a missing not at random (MNAR) scenario in [Figure 3](#). To simulate this scenario we randomly drop covariates, with the probability depending on the observed values for  $\mathbf{x}_n$ . Throughout, we assume that missingness only occurs in  $\mathbf{x}$ , not in  $\mathbf{c}$  or  $y$ . Increasing the amount of missingness in the covariates decreases performance in all models (see [Appendix D.4](#)). Proper modeling of the MNAR structure consistently gives significant performance improvements for all variations of our proposal. We discuss the precise mechanism of how we create this data set in [Appendix C.5.1](#). See [Appendix D.4](#) for an ATT estimation scenario and further results.

### 5.4 Survival

So far our evaluation has assumed a scalar outcome  $y$ , i.e.,  $\mathcal{Y} = \mathbb{R}$ . However, our method is not limited to this case. We now demonstrate its usefulness in the area of survival analysis ([Clark et al., 2003](#)), where  $y$  specifies a (potentially censored) survival time. Given access to the survival times of a control group via, e.g., their electronic health records, the goal is to select a suitable subgroup whose survival curve mimics the true, but unknown, counterfactual survival curve of the single-arm treatment group. We create synthetic survival data based on IHDP covariates and assume a Weibull distribution for  $p(y|t, \mathbf{z})$ , both for our model as well as the baselines. To evaluate the matching performance, we evaluate the squared error between the estimated hazard ratios relying either on the true counterfactual or the matched subset. Our proposal method improves upon the deterministic as well as generative baselines. See [Figure 4](#) for an exemplary visualization on a subset of the methods and [Table 2](#) for a summary of the results.

**Table 2: Survival.** Median squared error between the estimated and true hazard ratios for survival curve estimation. Our approach improves both upon its deterministic as well as its generative baselines in most variants.

METHOD	Median Error
Naive	0.394 [0.093,0.898]
TARNet	0.162 [0.018,0.501]
CFRNet	0.513 [0.166,1.059]
SNet	0.486 [0.173,1.120]
VAE	0.288 [0.064,0.663]
CEVAE	0.173 [0.022,0.572]
<b>Ours</b>	<b>0.135</b> [0.053,0.526]
+I	0.180 [0.020,0.448]
+sep	0.155 [0.036,0.509]
+sep+I	0.156 [0.033,0.473]
+snet	0.331 [0.060,1.425]
+snet+I	0.201 [0.316,1.120]

median over 100 random replications with [lower, upper quartile];  
lowest median marked bold


### 5.5 Real-World Experimental Data

We summarize the performance of ATT estimation in [Table 3](#). In our clinical data setup, as was the case for IHDP, we cannot evaluate performance against a true treatment outcome, given that the counterfactual reality is necessarily unknowable. We therefore have to again rely on synthetic outcomes on top of real covariates. See [Appendix C.5](#) for details on their creation. Evaluating the absolute error of ATT estimation, we observe that switching to our proposal improves performance over the baselines. Overall, a conditional prior  $p(\mathbf{z}|\mathbf{c})$  tends to improve upon their non-identifiable partner methods while splitting  $\mathcal{U}$  into two separate sub-spaces tends to hurt overall performance.

## 6 CONCLUSION

We contributed with this work a principled latent-variable model with the ability to estimate treatment effects from single-arm trials with external controls. The model learns a predictive low-dimensional latent space as well as separate (group-specific) representations to account for the different characteristics between treated and untreated patients. We performed an extensive ablation study on both synthetic benchmark data as well as a real-world combination of RCT and EHR data. Our model improved upon an extensive set of baselines across the benchmarks. In this work, we made technical choices regarding architectures and hyperparameters such that we can conclude on the importance of the different modeling choices proposed in our approach, without introducing experimental biases. The next step is to fine-tune technical choices, e.g., by incorporating further advanced architectures such as the FlexTENet ([Curth and van der Schaar, 2021b](#)), or estimators, such as the R-





**Figure 4: Survival.** Kaplan-Meier estimates of survival curves for three of the models for a single random seed. During training survival information for the treatment group is unavailable and we observe survival times only for the **control group**. After patient matching, the **matched subset** is compared against the unknown, **counterfactual survival curve** of the single-arm group. The shaded areas are the 95% confidence intervals on the respective Kaplan-Meier estimators. While the match found by *CFRNet* overlaps completely with the unmatched control group, *TARNet* can get close to the desired survival curve. *Ours* largely overlaps with it.

**Table 3: Real-World Medical Data.** Combining single-arm trial data with electronic health records. We observe the same pattern as before, i.e., switching to our generative proposal improves upon the respective baselines.

AE OF ATT	within-sample	out-of-sample
PScov	2.984±1.444	2.477±0.740
TARNet	2.514±1.007	2.236±0.892
CFRNet	2.271±0.884	2.230±0.865
SNet	3.032±1.141	2.319±0.678
VAE	2.445±0.828	2.123±0.886
CEVAE	2.587±1.051	2.221±0.954
Ours	2.143±0.918	2.032±0.874
+I	2.277±0.935	<u>1.997±0.899</u>
+sep	2.893±1.332	3.163±1.381
+sep+I	2.775±1.266	2.520±1.109
+snet	2.451±0.901	2.277±0.892
+snet+I	<u>2.121±0.824</u>	2.247±0.892

mean ± standard error over 5 random replications;  
lowest mean underlined

Learner (Lu et al., 2020; Nie and Wager, 2020) in the same way as we already demonstrated.

## ACKNOWLEDGEMENTS

This work was supported by the Academy of Finland Flagship programme: Finnish Center for Artificial Intelligence FCAI. Samuel Kaski was supported by the UKRI Turing AI World-Leading Researcher Fellowship, [EP/W002973/1]. This work was supported by the Research Council of Finland (decision number: 359135). We acknowledge the computational resources provided by the Aalto Science-IT project. Additional computations have been performed using resources at HUS Academic analytics platform.

## References

- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Susan Athey, Guido W. Imbens, and Stefan Wager. Approximate residual balancing: debiased inference of average treatment effects in high dimensions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(4):597–623, 2018.
- Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.
- Peter C. Austin. Optimal caliper widths for propensity-score matching when estimating differences in means and differences in proportions in observational studies. *Pharmaceutical Statistics*, 10(2):150–161, 2011.
- George L. Bakris, Rajiv Agarwal, Stefan D. Anker, Bertram Pitt, Luis M. Ruilope, Peter Rossing, Peter Kolkhof, Christina Nowack, Patrick Schloemer, Amer Joseph, and Gerasimos Filippatos. Effect of finerenone on chronic kidney disease outcomes in type 2 diabetes. *New England Journal of Medicine*, 383:2219–2229, 12 2020.
- Keith Battocchi, Eleanor Dillon, Maggie Hei, Greg Lewis, Paul Oka, Miruna Oprescu, and Vasilis Syrgkanis. EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation, 2019. URL <https://github.com/py-why/EconML>. Version 0.14.1.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010.
- 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 and Mihaela van der Schaar. Transfer learning on heterogeneous feature spaces for treatment effects estimation. In *Advances in Neural Information Processing Systems*, 2022.
- Ioana Bica, Ahmed Alaa, and Mihaela Van Der Schaar. Time series deconfounder: Estimating treatment effects over time in the presence of hidden confounders. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 884–895. PMLR, 13–18 Jul 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.
- Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- Jie Chen, Martin Ho, Kwan Lee, Yang Song, Yixin Fang, Benjamin A Goldstein, Weili He, Telba Irony, Qi Jiang, Mark van der Laan, Hana Lee, Xiwu Lin, Zhaoling Meng, Pallavi Mishra-Kalyani, Frank Rockhold, Hongwei Wang, and Roseann White. The current landscape in biostatistics of real-world data and evidence: Clinical study design and analysis. *Statistics in Biopharmaceutical Research*, 0(0):1–14, 2021.
- David Cheng and Tianxi Cai. Adaptive combination of randomized and observational data. *arXiv preprint arXiv:2111.15012*, 2021.
- T G Clark, M J Bradburn, S B Love, and D G Altman. Survival analysis part i: Basic concepts and first analyses. *British Journal of Cancer*, 89:232–238, 7 2003.
- Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- Mark Collier, Alfredo Nazabal, and Chris Williams. Vae’s in the presence of missing data. In *ICML Workshop on the Art of Learning with Missing Values (Artemiss)*, 2020.
- Rory Collins, Louise Bowman, Martin Landray, and Richard Peto. The magic of randomization versus the myth of real-world evidence. *New England Journal of Medicine*, 382(7):674–678, 2020.
- Alicia Curth and Mihaela van der Schaar. Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 1810–1818. PMLR, 13–15 Apr 2021a.
- Alicia Curth and Mihaela van der Schaar. On inductive biases for heterogeneous treatment effect estimation. In *Advances in Neural Information Processing Systems*, volume 34, pages 15883–15894. Curran Associates, Inc., 2021b.
- Issa J Dahabreh, James M Robins, and Miguel A Hernán. Benchmarking observational methods by comparing randomized trials and their emulations. *Epidemiology*, 31(5):614–619, 2020.
- Alexander D’Amour, Peng Ding, Avi Feller, Lihua Lei, and Jasjeet Sekhon. Overlap in observational studies with high-dimensional covariates. *Journal of Econometrics*, 221(2):644–654, 2021. ISSN 0304-4076.
- Hans-Georg Eichler, Francesco Pignatti, Brigitte Schwarzer-Daum, Ana Hidalgo-Simon, Irmgard Eichler,

- Peter Arlett, Anthony Humphreys, Spiros Vamvakas, Nikolai Brun, and Guido Rasi. Randomized controlled trials versus real world evidence: Neither magic nor myth. *Clinical Pharmacology & Therapeutics*, 109(5):1212–1218, 2021.
- Jessica M. Franklin and Sebastian Schneeweiss. When and how can real world data analyses substitute for randomized controlled trials? *Clinical Pharmacology & Therapeutics*, 102(6):924–933, 2017.
- Jessica M. Franklin, Elisabetta Patorno, Rishi J. Desai, Robert J. Glynn, David Martin, Kenneth Quinto, Ajinkya Pawar, Lily G. Bessette, Hemin Lee, Elizabeth M. Garry, Nileesa Gautam, and Sebastian Schneeweiss. Emulating randomized clinical trials with nonrandomized real-world evidence studies. *Circulation*, 143(10):1002–1013, 2021.
- Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1180–1189, Lille, France, 07–09 Jul 2015. PMLR.
- Sahra Ghalebikesabi, Rob Cornish, Chris Holmes, and Luke Kelly. Deep generative missingness pattern-set mixture models. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 3727–3735. PMLR, 13–15 Apr 2021.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- Christen M Gray, Fiona Grimson, Deborah Layton, Stuart Pocock, and Joseph Kim. A framework for methodological choice and evidence assessment for studies using external comparators from real-world data. *Drug safety*, 43(7):623–633, 2020.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.
- Tobias Hatt, Jeroen Berrevoets, Alicia Curth, Stefan Feuerriegel, and Mihaela van der Schaar. Combining observational and randomized data for estimating heterogeneous treatment effects. *arXiv preprint arXiv:2202.12891*, 2022.
- Jennifer L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Aapo Hyvarinen, Hiroaki Sasaki, and Richard Turner. Non-linear ica using auxiliary variables and generalized contrastive learning. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 859–868. PMLR, 16–18 Apr 2019.
- Guido W. Imbens. Nonparametric Estimation of Average Treatment Effects Under Exogeneity: A Review. *The Review of Economics and Statistics*, 86(1):4–29, 02 2004. ISSN 0034-6535.
- Niels Bruun Ipsen, Pierre-Alexandre Mattei, and Jes Frellsen. not-{miwae}: Deep generative modelling with missing not at random data. In *International Conference on Learning Representations*, 2021.
- Ziyang Jiang, Zhuoran Hou, Yiling Liu, Yiman Ren, Keyu Li, and David Carlson. Estimating causal effects using a multi-task deep ensemble. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 15023–15040. PMLR, 23–29 Jul 2023.
- Fredrik D. Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 3020–3029, New York, New York, USA, 20–22 Jun 2016. PMLR.
- Tom Joy, Yuge Shi, Philip Torr, Tom Rainforth, Sebastian M Schmon, and Siddharth N. Learning multimodal VAEs through mutual supervision. In *International Conference on Learning Representations*, 2022.
- Nathan Kallus, Aahlad Manas Puli, and Uri Shalit. Removing hidden confounding by experimental grounding. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Edward H. Kennedy. Efficient nonparametric causal inference with missing exposure information. *The International Journal of Biostatistics*, 16(1):20190087, 2020.
- Ilyes Khemakhem, Diederik Kingma, Ricardo Monti, and Aapo Hyvarinen. Variational autoencoders and non-linear ica: A unifying framework. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2207–2217. PMLR, 26–28 Aug 2020.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *International Conference on Learning Representations*, 2014.

- 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*, 116:4156–4165, 3 2019.
- Samu Kurki, Viivi Halla-aho, Manuel Haussmann, Harri Lähdesmäki, Jussi V. Leinonen, and Miika Koskinen. A comparative study of clinical trial and real-world data in patients with diabetic kidney disease. *Scientific Reports*, 14, 1 2024.
- Milan Kuzmanovic, Tobias Hatt, and Stefan Feuerriegel. Estimating conditional average treatment effects with missing treatment information. *arXiv preprint arXiv:2203.01422*, 2022.
- Yann Lecun, Sumit Chopra, Raia Hadsell, Marc Aurelio Ranzato, and Fu Jie Huang. *A tutorial on energy-based learning*. MIT Press, 2006.
- Sheng Li and Yun Fu. Matching on balanced nonlinear representations for treatment effects estimation. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Chaochao Lu, Yuhuai Wu, José Miguel Hernández-Lobato, and Bernhard Schölkopf. Invariant causal representation learning for out-of-distribution generalization. In *International Conference on Learning Representations*, 2022.
- Danni Lu, Chenyang Tao, Junya Chen, Fan Li, Feng Guo, and Lawrence Carin. Reconsidering generative objectives for counterfactual reasoning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21539–21553. Curran Associates, Inc., 2020.
- Wei Luo and Yeying Zhu. Matching using sufficient dimension reduction for causal inference. *Journal of Business & Economic Statistics*, 38(4):888–900, 2020.
- Taro Makino, Krzysztof J. Geras, and Kyunghyun Cho. Generative multitask learning mitigates target-causing confounding. *CoRR*, abs/2202.04136, 2022.
- Laura Manduchi, Ricards Marcinkevics, Michela C. Massi, Thomas Weikert, Alexander Sauter, Verena Gotta, Timothy Müller, Flavio Vasella, Marian C. Neidert, Marc Pfister, Bram Stieltjes, and Julia E Vogt. A deep variational approach to clustering survival data. In *International Conference on Learning Representations*, 2022.
- Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32, 2016.
- Gemma Elyse Moran, Dhanya Sridhar, Yixin Wang, and David Blei. Identifiable deep generative models via sparse decoding. *Transactions on Machine Learning Research*, 2022.
- Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2): 299–319, 09 2020. ISSN 0006-3444.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- Judea Pearl. *Causality*. Cambridge University Press, 2 edition, 2009.
- Alexandre Perez-Lebel, Gaël Varoquaux, Marine Le Morvan, Julie Josse, and Jean-Baptiste Poline. Benchmarking missing-values approaches for predictive models on health databases. *GigaScience*, 11, 4 2022.
- Sebastian Pölsterl, 2019. URL <https://k-d-w.org/blog/2019/07/survival-analysis-for-deep-learning/>.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1278–1286, Beijing, China, 22–24 Jun 2014. PMLR.
- 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, 04 1983. ISSN 0006-3444.
- Donald B. Rubin. Inference and missing data. *Biometrika*, 63:581–592, 1976.
- Donald B Rubin. Causal inference using potential outcomes. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- Heinz Schmidli, Dieter A Häring, Marius Thomas, Adrian Cassidy, Sebastian Weber, and Frank Bretz. Beyond randomized clinical trials: use of external controls. *Clinical Pharmacology & Therapeutics*, 107(4):806–816, 2020.
- Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70

- of *Proceedings of Machine Learning Research*, pages 3076–3085. PMLR, 06–11 Aug 2017.
- Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Elizabeth A. Stuart. Matching Methods for Causal Inference: A Review and a Look Forward. *Statistical Science*, 25(1):1 – 21, 2010.
- Terry M Therneau. *A Package for Survival Analysis in R*, 2023. URL <https://CRAN.R-project.org/package=survival>. R package version 3.5-7.
- Nilesh Tripuraneni, Michael Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. In *Advances in Neural Information Processing Systems*, volume 33, pages 7852–7862. Curran Associates, Inc., 2020.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113: 1228–1242, 7 2018.
- Yixin Wang and David M. Blei. The blessings of multiple causes. *Journal of the American Statistical Association*, 114(528):1574–1596, 2019.
- E.J. Williamson, A. Forbes, and R. Wolfe. Doubly robust estimators of causal exposure effects with missing data in the outcome, exposure or a confounder. *Statistics in Medicine*, 31(30):4382–4400, 2012.
- Quanhan Xi and Benjamin Bloem-Reddy. Indeterminacy in generative models: Characterization and strong identifiability. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 6912–6939. PMLR, 25–27 Apr 2023.
- Jinsung Yoon, James Jordon, and Mihaela van der Schaar. GANITE: Estimation of individualized treatment effects using generative adversarial nets. In *International Conference on Learning Representations*, 2018.
- Weijia Zhang, Lin Liu, and Jiuyong Li. Treatment effect estimation with disentangled latent factors. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35: 10923–10930, 5 2021.
- Yao Zhang, Jeroen Berrevoets, and Mihaela Van Der Schaar. Identifiable energy-based representations: An application to estimating heterogeneous causal effects. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 4158–4177. PMLR, 28–30 Mar 2022.
- Zhiwei Zhang, Wei Liu, Bo Zhang, Li Tang, and Jun Zhang. Causal inference with missing exposure information: Methods and applications to an obstetric study. *Statistical Methods in Medical Research*, 25(5):2053–2066, 2016. PMID: 24318273.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.

## Checklist

1. For all models and algorithms presented, check if you include:
  - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]  
*See Appendix B.*
  - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]  
*See Appendix D.1 for an empirical evaluation of the runtime.*
  - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [No]  
*An implementation of all models included in the experiments will be available upon acceptance.*
2. For any theoretical claim, check if you include:
  - (a) Statements of the full set of assumptions of all theoretical results. [Yes]  
*See Appendix B.*
  - (b) Complete proofs of all theoretical results. [Yes]  
*See Appendix B.*
  - (c) Clear explanations of any assumptions. [Yes]  
*See Appendix B.*
3. For all figures and tables that present empirical results, check if you include:
  - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]  
*See Appendix C.*
  - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]  
*See Appendix C.*
  - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]  
*See Appendix C.*
  - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]  
*See Appendix C.*
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
  - (a) Citations of the creator If your work uses existing assets. [Yes]
  - (b) The license information of the assets, if applicable. [Yes]
  - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
  - (d) Information about consent from data providers/curators. [Yes]
  - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Yes]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
  - (a) The full text of instructions given to participants and screenshots. [No]  
*This work did not collect new data from human subjects. The real-world data deals with prior data collected from patients. The specifics of the data collection are currently not included to keep the anonymity, but will be available upon publication.*
  - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [No]  
*See above.*
  - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]  
*Patients were not payed for their treatment.*

# Latent-variable Modeling of Treatment Effects from Single-arm Trials

## Supplementary Materials

### A FURTHER RELATED WORK

**Multi-task, Multi-modal, and Transfer Learning.** How to combine information from multiple data sources is an ongoing research question. Multi-task learning considers the problem of solving multiple related tasks from a single data set (Caruana, 1997; Maurer et al., 2016; Makino et al., 2022). Multi-modal learning is closely related but focuses instead on using heterogeneous data sources for a single target task (Joy et al., 2022). Transfer learning considers the task of transferring a learned model from one data set to another (Pan and Yang, 2010; Tripuraneni et al., 2020). Our task is related to this field via domain adaptation (Ben-David et al., 2010). We assume two different data sources with similar yet distinct characteristics and extract a predictive representation by combining the two.

### B THEORY

#### B.1 Identifiability

##### B.1.1 Prior Work

Throughout this work we follow the approach of Khemakhem et al. (2020), who rely on auxiliary covariates for identifiability guarantees. See also the recent work by Xi and Bloem-Reddy (2023), who extend this work by providing additional results and guarantees.

Khemakhem et al. (2020) observe that a general unconstrained latent-variable model of the form

$$p(\mathbf{x}, \mathbf{z}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z}),$$

where  $\mathbf{x}$  is assumed to be observed and  $\mathbf{z}$  to be latent is, in general, not identifiable in the sense of the implication

$$\forall(\theta, \theta') : p_{\theta}(\mathbf{x}) = p_{\theta'}(\mathbf{x}) \Rightarrow \theta = \theta'.$$

To solve this problem, they follow prior work by Hyvarinen et al. (2019) and introduce an additional observed variable  $\mathbf{c}$

$$p_{\theta}(\mathbf{x}, \mathbf{z}|\mathbf{c}) = p_f(\mathbf{x}|\mathbf{z})p_{\mathbf{T},\lambda}(\mathbf{z}|\mathbf{c}),$$

with  $\theta = (f, \mathbf{T}, \lambda)$  and  $f$  an injective function, e.g., a neural net. The likelihood  $p_f(\mathbf{x}|\mathbf{z})$  is assumed to be of the form

$$p_f(\mathbf{x}|\mathbf{z}) = f(\mathbf{z}) + \varepsilon,$$

with  $\varepsilon \sim p(\varepsilon)$ , i.e., composed of an additive independent noise source. The covariates  $\mathbf{x}$  are assumed to be continuous.<sup>3</sup> The conditional prior in turn is assumed to be conditionally factorial into a product of exponential family distributions,

$$p_{\mathbf{T},\lambda}(\mathbf{z}|\mathbf{c}) = \prod_{i=1}^{d_z} \frac{m(z_i)}{Z(\mathbf{c})} \exp(\mathbf{T}_i(z_i)^{\top} \lambda(\mathbf{c})),$$

where  $\mathbf{T}_i$  are the sufficient statistics with  $\dim(\mathbf{T}_i) = k$ , and  $\lambda(\cdot)$  some function of  $\mathbf{c}$ , e.g., a neural net. Although this factorization assumption is sufficient within our current setup, it has since been shown that it is not necessary (Lu et al., 2022; Xi and Bloem-Reddy, 2023).

We quote the following two definitions and theorems by Khemakhem et al. (2020) following their numbering scheme and adapting them to our notation.

<sup>3</sup>Khemakhem et al. (2020) speculated that this constraint on a continuous covariate could also be extended to discrete covariates, but Xi and Bloem-Reddy (2023) show, that this is not provable in its current form. Nevertheless, both Khemakhem et al. (2020) and we do observe good empirical performance nevertheless also for discrete covariates  $\mathbf{x}$ .

**Definition 1.** Let  $\sim$  be an equivalence relation on  $\Theta$ . We say that  $p(\mathbf{x}|\mathbf{z})p(\mathbf{z})$  is identifiable up to  $\sim$  if

$$p_{\theta}(\mathbf{x}) = p_{\theta'}(\mathbf{x}) \Rightarrow \theta \sim \theta'.$$

The elements of the quotient space  $\Theta/\sim$  are called the identifiability classes.

**Definition 2.** Let  $\sim$  be the equivalence relation on  $\Theta$  defined as follows:

$$(f, \mathbf{T}, \lambda) \sim (\tilde{f}, \tilde{\mathbf{T}}, \tilde{\lambda}) \Leftrightarrow \exists A, \mathbf{b} \quad \mathbf{T}(f^{-1}(\mathbf{x})) = A\tilde{\mathbf{T}}(\tilde{f}^{-1}(\mathbf{x})) + \mathbf{b} \quad \forall \mathbf{x} \in \mathcal{X}$$

where  $A$  is an  $d_z k \times d_z k$  matrix and  $\mathbf{b}$  is a vector. If  $A$  is invertible, we denote this relation by  $\sim_A$ . If  $A$  is a block permutation matrix, we denote it by  $\sim_P$ .

**Theorem 1.** Assume that we observe data sampled from a generative model defined according to the model specified above with parameters  $(f, \mathbf{T}, \lambda)$ . Assume the following holds:

- (i) The set  $\{\mathbf{x} \in \mathcal{X} | \varphi_{\varepsilon}(\mathbf{x}) = 0\}$  has measure zero, where  $\varphi_{\varepsilon}$  is the characteristic function of  $p(\varepsilon)$ ;
- (ii) the mixing function  $f$  is injective;
- (iii) the sufficient statistics  $\mathbf{T}$  are differentiable almost everywhere and  $(T_{i,j})_{1 \leq j \leq k}$  are linearly independent on any subset of  $\mathcal{X}$  of measure greater than zero;
- (iv) there exist  $d_z k + 1$  distinct points  $\mathbf{c}^0, \dots, \mathbf{c}^{d_z k}$  such that the matrix

$$L = (\lambda(\mathbf{c}^1) - \lambda(\mathbf{c}^0), \dots, \lambda(\mathbf{c}^{d_z k}) - \lambda(\mathbf{c}^0)),$$

of size  $d_z k \times d_z k$  is invertible,

then the parameters  $(f, \mathbf{T}, \lambda)$  are  $\sim_A$ -identifiable.

**Theorem 4** Assume the following:

- (i) The family of distributions  $q_{\phi}(\mathbf{z}|\mathbf{x}, \mathbf{c})$  contains  $p_{\theta}(\mathbf{z}|\mathbf{x}, \mathbf{c})$ ,
- (ii) we maximize

$$\mathbb{E}_{\mathcal{D}} [\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x}, \mathbf{u})} [\log p_{\theta}(\mathbf{x}, \mathbf{z}|\mathbf{c}) - \log q_{\phi}(\mathbf{z}|\mathbf{x}, \mathbf{c})]],$$

with respect to both  $\theta$  and  $\phi$ ,

then in the limit of infinite data  $\mathcal{D}$ , the VAE learns the true parameters  $\theta^* = (f^*, \mathbf{T}^*, \lambda^*)$  up to an equivalence class  $\sim_A$ .

### B.1.2 Our setup

The joint probability distribution we use throughout most part of the paper is

$$p(\mathbf{x}, y, t, \mathbf{u}, \mathbf{z}|\mathbf{c}) = p(\mathbf{x}|\mathbf{u}, \mathbf{z})p(t|\mathbf{z})p(y|\mathbf{z})p(\mathbf{u})p(\mathbf{z}|\mathbf{c}),$$

where for this discussion we assume  $\mathbf{x} \in \mathbb{R}^{d_x}$ , i.e., all covariates are continuous and normally distributed. As discussed above, the guarantees do not hold for discrete covariates. Throughout our experiments, we still include them, and empirically still observe good results, but to guarantee identifiability they could also be dropped from the likelihood, and, e.g., just be used as additional inputs to the variational posterior. Similarly, we focus only on the latent-variable model with respect to  $\mathbf{x}$ , instead of  $(\mathbf{x}, y)$  as (i)  $y$  may be partially unobserved and is not part of the variational posterior violating the assumptions of Theorem 4, and (ii)  $p(y|t, \mathbf{z})$  is, e.g., in the survival experiment a Weibull distribution, i.e., does not fulfill the constraint of an independent noise source. Note also that Theorem 4 requires us to optimize the original ELBO, i.e., without any regularization terms. We only include these in the two-arm setup but drop them from the one-arm setup to fulfill this constraint.

We keep the prior  $p(\mathbf{u})$  unconstrained and only constrain  $p(\mathbf{z}|\mathbf{c})$  as a factorizing normal distribution

$$p(\mathbf{z}|\mathbf{c}) = \prod_{i=1}^{d_z} m(z_i)/Z(\mathbf{c}) \exp(T(z_i)^{\top} \lambda(\mathbf{c})), \quad (6)$$



with  $\lambda(\cdot)$  parameterized by a neural network. Throughout the experiments, we fix the prior variance of  $p(\mathbf{z}|\mathbf{c})$  and only parameterize the mean, that is  $k = 1$ . As we require  $d_z k + 1$  distinct values, this allows us to use a lower variety in  $\mathbf{c}$ , at the price of only being able to guarantee  $\sim_A$  identifiability, instead of  $\sim_P$  (Khemakhem et al., 2020, Theorem 3).

This gives us, after marginalizing over  $y, t$

$$\begin{aligned} p(\mathbf{x}, \mathbf{z}|\mathbf{c}) &= \int p(\mathbf{x}|\mathbf{u}, \mathbf{z})p(\mathbf{u})p(\mathbf{z}|\mathbf{c}) \, d\mathbf{u} \\ &= \mathbb{E}_{p(\mathbf{u})} [p(\mathbf{x}|\mathbf{u}, \mathbf{z})] p(\mathbf{z}|\mathbf{c}) \\ &= \mathbb{E}_{p(\mathbf{u})} [p_\varepsilon(\mathbf{x} - f(\mathbf{u}, \mathbf{z}))] p(\mathbf{z}|\mathbf{c}) \\ &= p_\varepsilon\left(\mathbf{x} - \underbrace{\mathbb{E}_{p(\mathbf{u})} [f(\mathbf{u}, \mathbf{z})]}_{\triangleq h(\mathbf{z})}\right) p(\mathbf{z}|\mathbf{c}), \end{aligned}$$

where  $p_\varepsilon(\mathbf{x} - h(\mathbf{z})) \triangleq p_\varepsilon(\varepsilon) = \mathcal{N}(\varepsilon|0, \mathbf{1})$ . Assuming the remaining conditions of Theorem 1 hold, this gives us the desired  $\sim_A$  identifiability for  $(h, \mathbf{T}, \lambda)$ .

## B.2 Missingness

Not only can missing covariate information be assumed to be present in most real-world data sets, e.g., electronic health records, but one should also expect it to be missing not at random (MNAR) (Rubin, 2005), i.e., the missingness pattern depends on the missing covariates themselves. To properly account for that we extend our model with binary masking variables  $\mathbf{m} \in \mathcal{M} = \{0, 1\}^{d_x}$ , where  $m_{ni} = 0$  indicates that the  $i$ -th covariate of observation  $n$  is missing. We observe

$$\tilde{\mathbf{x}} = \mathbf{m} \circ \mathbf{x} + (1 - \mathbf{m}) \circ \boldsymbol{\eta},$$

where  $\circ$  is elementwise multiplication, and  $\boldsymbol{\eta}$  some imputation scheme, e.g.,  $\boldsymbol{\eta} = \mathbf{0}$  for zero imputation.


We model MNAR by assuming a prior over  $\mathbf{m}$  that depends on  $\mathbf{u}$  and  $\mathbf{z}$ , to account for the not-at-random structure, and assume an additional latent  $\mathbf{z}^m$  to model any remaining pattern. Our joint model is given as

$$\begin{aligned} \mathbf{u} &\sim p(\mathbf{u}) = \mathcal{N}(\mathbf{u}|0, \sigma_u^2 \mathbf{1}), \\ \mathbf{z}|\mathbf{c} &\sim p(\mathbf{z}|\mathbf{c}) = \mathcal{N}(\mathbf{z}|\lambda(\mathbf{c}), \sigma_z^2 \mathbf{1}), \\ \mathbf{z}^m &\sim p(\mathbf{z}^m) = \mathcal{N}(\mathbf{z}^m|0, \sigma_{z^m}^2 \mathbf{1}), \\ \mathbf{m}|\mathbf{u}, \mathbf{z}, \mathbf{z}^m &\sim \prod_{i=1}^{d_x} \text{Ber}(m_i|\sigma(g(\mathbf{z}^m, \mathbf{u}, \mathbf{z})_i)), \\ \mathbf{x}|\mathbf{m}, \mathbf{u}, \mathbf{z} &\sim p(\mathbf{x}|\mathbf{m}, \mathbf{u}, \mathbf{z}), \\ t &\sim p(t|\mathbf{z}) = \text{Ber}(t|\sigma(f(\mathbf{z}))), \\ y &\sim p(y|t, \mathbf{z}) = \mathcal{N}(y|\mu_t^y(\mathbf{z}), \alpha^{-1}), \end{aligned}$$

where  $f(\cdot)$ ,  $g(\cdot)$ ,  $\mu_t^y(\cdot)$  are neural nets, and  $\sigma(x) \triangleq 1/(1 + \exp(-x))$  is the logistic sigmoid. As in the main text, we only focus on identifiability with respect to  $\mathbf{z}$ . We assume  $p(\mathbf{x}|\mathbf{m}, \mathbf{u}, \mathbf{z})$  to factorize over the covariates, with each covariate likelihood being modeled via normal, Bernoulli, or categorical distributions depending on their respective domains. See Figure 5 for a summarizing plate diagram.

We define  $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_N)$  and  $\mathbf{Z}^m, \mathbf{Z}, \mathbf{X}$  are defined analogously. We again rely on a mean-field variational posterior, given as

$$\begin{aligned} q(\mathbf{U}, \mathbf{Z}^m, \mathbf{Z}|\mathbf{C}, \mathbf{X}) &= \prod_{n=1}^N q(\mathbf{u}_n|\mathbf{m}_n, \tilde{\mathbf{x}}_n)q(\mathbf{z}_n|\mathbf{c}_n, \mathbf{m}_n, \tilde{\mathbf{x}}_n)q(\mathbf{z}_n^m|\mathbf{m}_n, \tilde{\mathbf{x}}_n) \\ &= \prod_{n=1}^N \mathcal{N}(\mathbf{u}_n|\mu^u(\mathbf{m}_n, \tilde{\mathbf{x}}_n), \sigma^u(\mathbf{m}_n, \tilde{\mathbf{x}}_n))\mathcal{N}(\mathbf{z}_n|\mu^z(\mathbf{c}_n, \mathbf{m}_n, \tilde{\mathbf{x}}_n), \sigma^z(\mathbf{c}_n, \mathbf{m}_n, \tilde{\mathbf{x}}_n)) \\ &\quad \cdot \mathcal{N}(\mathbf{z}_n^m|\mu^{z^m}(\mathbf{m}_n, \tilde{\mathbf{x}}_n), \sigma^{z^m}(\mathbf{m}_n, \tilde{\mathbf{x}}_n)), \end{aligned}$$



**Figure 5:** *Missing not at random model.* Empty, partially filled, and filled circles represent latent, partially observed and observed variables.  $\tilde{\mathbf{x}}$  is deterministic as described in the text.

The full ELBO is given as

$$\begin{aligned} \log p(\mathcal{D}) \geq & \sum_{n=1}^N \mathbb{E}_{q(\mathbf{u}_n, \mathbf{z}_n, \mathbf{z}_n^m | \mathbf{c}_n, \mathbf{x}_n)} [\log p(\mathbf{x}_n | \mathbf{m}_n, \mathbf{u}_n, \mathbf{z}_n) + \log p(\mathbf{m}_n | \mathbf{u}_n, \mathbf{z}_n^m, \mathbf{z}_n)] \\ & + \mathbb{E}_{q(\mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n)} [\log p(y_n | t_n, \mathbf{z}_n) + \log p(t_n | \mathbf{z}_n)] \\ & - \text{KL}(q(\mathbf{u}_n | \tilde{\mathbf{x}}_n) \parallel p(\mathbf{u}_n)) - \text{KL}(q(\mathbf{z}_n | \mathbf{c}_n, \tilde{\mathbf{x}}_n) \parallel p(\mathbf{z}_n | \mathbf{c}_n)) \\ & - \text{KL}(q(\mathbf{z}_n^m | \mathbf{m}_n, \tilde{\mathbf{x}}_n) \parallel p(\mathbf{z}_n^m)) \end{aligned}$$

See [Appendix D.4](#) for related experiments.

### B.3 A split $\mathcal{U}$

The main paper discusses a separate latent space  $\mathcal{U}$  and formulates the joint model as

$$\begin{aligned} \mathbf{u} & \sim p(\mathbf{u}) = \mathcal{N}(\mathbf{u} | 0, \sigma_u^2 \mathbf{1}), \\ \mathbf{z} & \sim p(\mathbf{z} | \mathbf{c}) = \mathcal{N}(\mathbf{z} | \lambda(\mathbf{c}), \sigma_z^2 \mathbf{1}), \\ \mathbf{x} & \sim p(\mathbf{x} | \mathbf{u}, \mathbf{z}), \\ t & \sim p(t | \mathbf{z}) = \text{Ber}(t | \sigma(f(\mathbf{z}))), \\ y & \sim p(y | t, \mathbf{z}) = \mathcal{N}(y | \mu_t^y(\mathbf{z}), \alpha^{-1}). \end{aligned}$$

Assuming data from  $n_{\text{data}}$  sources, another approach is to use a separate  $\mathcal{U}$  for each of them, i.e., we have  $\mathcal{U} = \mathcal{U}_1 \times \dots \times \mathcal{U}_{n_{\text{data}}}$ . The joint then becomes

$$\begin{aligned} \mathbf{u}^g & \sim p(\mathbf{u}^g) = \mathcal{N}(\mathbf{u}^g | 0, \sigma_u^2 \mathbf{1}), \quad \text{for } g = 1, \dots, n_{\text{data}}, \\ \mathbf{z} & \sim p(\mathbf{z} | \mathbf{c}) = \mathcal{N}(\mathbf{z} | \lambda(\mathbf{c}), \sigma_z^2 \mathbf{1}), \\ \mathbf{x} & \sim p(\mathbf{x} | \mathbf{u}, \mathbf{z}), \\ t & \sim p(t | \mathbf{z}) = \text{Ber}(t | \sigma(f(\mathbf{z}))), \\ y & \sim p(y | t, \mathbf{z}) = \mathcal{N}(y | \mu_t^y(\mathbf{z}), \alpha^{-1}). \end{aligned}$$

In our specific setup of one treatment and one control group, we have  $n_{\text{data}} = 2$ . However, this can directly be generalized to multiple control sources and separate treatment trials.

### B.4 Outcome only for the control group

A common assumption in the machine learning literature is to predict treatment effects given pre- and post-treatment information, i.e., covariates as well as outcomes for all treatment and control groups. The additional alternative setting that we consider is that the model is blinded to post-treatment information from the treatment group, i.e., the model has to infer

a suitable latent embedding solely from pre-treatment covariates of the single-arm trial as well as covariate and outcome information from the control electronic health records.

This blinding ensures, that before having selected a fixed set of controls to compare against no post-treatment information is leaked that could potentially bias future tests.

Assuming that we have  $N_0$  control observations and  $N_1$  treated observations, our data set becomes  $\mathcal{D} = \{(x_1, y_1, t_1), \dots, (x_{N_0}, y_{N_0}, t_{N_0}), (x_{N_0+1}, t_{N_0+1}), \dots, (x_{N_0+N_1}, t_{N_0+N_1})\}$  and the ELBO becomes

$$\begin{aligned} \log p(\mathcal{D}) &\geq \sum_{n=1}^{N_0} \mathbb{E}_{q(\mathbf{u}_n, \mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n)} [\log p(\mathbf{x}_n | \mathbf{u}_n, \mathbf{z}_n)] + \mathbb{E}_{q(\mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n)} [\log p(y_n | t_n, \mathbf{z}_n) + \log p(t_n | \mathbf{z}_n)] \\ &\quad + \sum_{n=N_0+1}^{N_0+N_1} \mathbb{E}_{q(\mathbf{u}_n, \mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n)} [\log p(\mathbf{x}_n | \mathbf{u}_n, \mathbf{z}_n)] + \mathbb{E}_{q(\mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n)} [\log p(t_n | \mathbf{z}_n)] \\ &\quad - \text{KL}(q(\mathbf{u}_n | \mathbf{x}_n) \parallel p(\mathbf{u}_n)) - \text{KL}(q(\mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n) \parallel p(\mathbf{z}_n | \mathbf{c}_n)) \\ &= \sum_{n=1}^N \mathbb{E}_{q(\mathbf{u}_n, \mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n)} [\log p(\mathbf{x}_n | \mathbf{u}_n, \mathbf{z}_n)] + \mathbb{E}_{q(\mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n)} [(1 - t_n) \log p(y_n | t_n, \mathbf{z}_n) + \log p(t_n | \mathbf{z}_n)] \\ &\quad - \text{KL}(q(\mathbf{u}_n | \mathbf{x}_n) \parallel p(\mathbf{u}_n)) - \text{KL}(q(\mathbf{z}_n | \mathbf{c}_n, \mathbf{x}_n) \parallel p(\mathbf{z}_n | \mathbf{c}_n)), \end{aligned}$$

where we use  $N = N_0 + N_1$  and recover the original ELBO in (4) up to the correction factor for notational simplicity.

## C EXPERIMENTAL DETAILS

This section goes through various details on the experimental setups. A reference implementation is provided at [anonymus](#).

### C.1 Significance Testing

Throughout our experiments, we rely on paired Student t-tests to evaluate which models have a comparable performance. The t-tests are computed relatively to the model with the best empirical mean and are one-sided, i.e., whether the alternative has a lower mean. We mark all methods that do not lead to a rejection of the null hypothesis at a significance level of  $p < 0.05$  in bold. This applies to all experiments throughout this paper. The lowest empirical mean is additionally underlined. The number of repetitions for each experiment is specified below the respective table of results.

### C.2 Models Under Consideration

Within this paper, we compare against a range of deterministic and probabilistic baselines as well as several baselines. We discuss their relation to our method within this section. Unless noted otherwise, we reimplemented all baselines in PyTorch ([Paszke et al., 2019](#)).

#### Non neural net-based baselines.

- *CFor* is a popular causal forest-based approach for causal estimation introduced by [Wager and Athey \(2018\)](#) and serves as our main non-neural network-based baseline. Given its lack of a representation space to perform matching on, it can only be used in the scenario where outcome information is available for both groups. We rely on the implementation provided by the EconML package ([Battocchi et al., 2019](#)).
- *PScov*, *PSpca*, and *PSlat* serve as propensity score (PS)-based matching methods that are constructed on various sets of inputs. *PScov* uses the features provided by the original covariates in  $\mathcal{X}$ . *PSpca* first maps them into a  $\text{dim}(z)$  space via a principal component analysis (PCA). *PSlat* finally relies on  $\mathcal{Z}$  space encodings that were provided by one of our neural network-based approaches. Throughout the experiments, we always use the encoding of the best-performing neural net as input to *PSlat*. Independent of their input domain, all three construct their estimator of  $\pi(x) = \mathbb{P}(T = 1 | X = x)$  via an ElasticNet ([Zou and Hastie, 2005](#)). As these propensity score methods do not infer estimators  $\mu_t(\cdot)$  we only compare against them in the matching experiments. We rely on the ElasticNet implementation provided by scikit-learn ([Virtanen et al., 2020](#)).

#### Deterministic neural nets.

- *SingleNet* is a simple neural net, that learns a simple estimator  $\mu(\mathbf{x}, t)$  by concatenating  $t$  to the features  $\mathbf{x}$ , such that  $\mu_t(\mathbf{x}) = \mu(\mathbf{x}, t)$ . Its objective is then given as a simple MSE loss,

$$\min \sum_{n=1}^N (\mu(\mathbf{x}_n, t_n) - y_n)^2.$$

The estimator  $\mu(\cdot, \cdot)$  is parameterized by a neural net. As it does not infer a shared latent space we only use it in scenario (a) where outcome information is available for both arms.

- *TNet* slightly generalizes upon *SingleNet* by learning two separate neural net estimators  $\mu_0(\mathbf{x})$  and  $\mu_1(\mathbf{x})$ . Its objective is again an MSE loss

$$\min \sum_{n=1}^N (\mu_{t_n}(\mathbf{x}_n) - y_n)^2.$$

We parameterize  $\mu_0(\cdot)$  and  $\mu_1(\cdot)$  by neural nets. Similar to the *SingleNet* it does not learn a shared latent space to perform matching on and is therefore not applicable to the matching scenario.

- *TARNet* (Shalit et al., 2017) is a generalization of the *TNet* that consists of three neural nets.  $h(\cdot)$  maps the observed covariates into a shared representation space.  $\mu_0(\cdot)$  and  $\mu_1(\cdot)$  are then estimated based on this representation. The MSE loss is given as

$$\min \sum_{n=1}^N (\mu_{t_n}(h(\mathbf{x}_n)) - y_n)^2.$$

- *CFRNet* (Shalit et al., 2017) extends *TARNet* with an additional MMD (see Appendix C.7 for details) term in the loss for further regularization.
- *SNet* (Curth and van der Schaar, 2021a) further generalizes upon these by replacing the single  $h(\cdot)$  into five sub-nets  $h_i(\cdot), i \in \{1, \dots, 5\}$  that separate the representation space into various sub-spaces, which serve as input to  $\mu_t(\cdot)$  as well as a separate classifier  $g(\cdot)$  that predicts  $t$  given these representations which is optimized via a cross-entropy loss. These subspaces in turn serve as (i) input only to  $\mu_0$ , (ii) input to both  $\mu_0, \mu_1$ , (iii) input only to  $\mu_1$ , (iv) input to all estimators  $\mu_0, \mu_1, g$ , or (v) input only to  $g$ . In addition, SNet replaces the MMD regularizer with an orthogonalization regularizer on the representation space mappings  $h_i(\cdot)$ . See Curth and van der Schaar (2021a) for details.

Comparing the loss terms of these deterministic models to our ELBO objective (4) they can be interpreted as optimizing (up to variations in the model structure and additional loss terms),

$$\sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}_n|\mathbf{x}_n)} [\log p(y_n|\mathbf{z}_n, t_n)],$$

for a deterministic delta distribution  $q(\mathbf{z}|\mathbf{x}) = \delta(\mathbf{z} - h(\mathbf{x}))$ , and  $p(y|\mathbf{z}, t) = \mathcal{N}(y|\mu_t(\mathbf{z}), 0.5)$ .<sup>4</sup> *SNet* comes even closer by adding an additional  $\log p(t|\mathbf{z})$  to the objective.

**Generative baselines.** While our deterministic baselines only consider the covariates as input features, the generative approaches also rely on learning representations that can (approximately) reconstruct them. All generative models rely on mean-field variational posteriors and use the same prior and likelihood assumptions as far as possible.

- *VAE* is a simple variational auto-encoder (Kingma and Welling, 2014), whose ELBO is given as

$$\sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}_n|\mathbf{x}_n)} [\log p(\mathbf{x}_n|\mathbf{z}_n)] - \text{KL}(q(\mathbf{z}_n|\mathbf{x}_n) \parallel p(\mathbf{z}_n)).$$

As it lacks a model for the outcome likelihood  $y$ , we only compare against it in the matching scenarios.

- *CEVAE* (Louizos et al., 2017) provides an output likelihood for the outcome  $y$  together with an additional likelihood for  $t$ , i.e., optimizes an ELBO given as

$$\sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}_n|\mathbf{x}_n)} [\log p(\mathbf{x}_n|\mathbf{z}_n) + \log p(y_n|\mathbf{z}_n, t_n) + \log p(t_n|\mathbf{z}_n)] - \text{KL}(q(\mathbf{z}_n|\mathbf{x}_n) \parallel p(\mathbf{z}_n)).$$

<sup>4</sup>The 0.5 is due to recover the MSE loss without the scaling factor 1/2 the normal distribution introduces in its exponential.

Note that Louizos et al. (2017) include additional reconstructive terms in their objective which we leave out to keep a valid ELBO.

- *TEDVAE* (Zhang et al., 2021) generalizes *CEVAE* similar to how *SNet* generalized the *CFRNet*. It splits the latent  $\mathcal{Z}$  into three parts such that the joint is given as

$$p(\mathbf{x}, \mathbf{z}^1, \mathbf{z}^2, \mathbf{z}^3, y, t) = p(\mathbf{x}|\mathbf{z}^1, \mathbf{z}^2, \mathbf{z}^3)p(t|\mathbf{z}^1, \mathbf{z}^2)p(y|\mathbf{z}^2, \mathbf{z}^3)p(\mathbf{z}^1)p(\mathbf{z}^2)p(\mathbf{z}^3).$$

The corresponding ELBO objective is then given as

$$\sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}_n|\mathbf{x}_n)} [\log p(\mathbf{x}_n|\mathbf{z}_n) + \log p(y_n|\mathbf{z}_n^2, \mathbf{z}_n^3, t_n) + \log p(t_n|\mathbf{z}_n^1, \mathbf{z}_n^2)] - \text{KL}(q(\mathbf{z}_n|\mathbf{x}_n) \| p(\mathbf{z}_n)).$$

where we use  $\mathbf{z}_n = (\mathbf{z}_n^1, \mathbf{z}_n^2, \mathbf{z}_n^3)$  in the notation.

**Variations of our proposal.** Throughout the experiments, we evaluate a wide range of variations on our model. We describe each of the building blocks in turn.

- *Ours* is the basic model we rely on with a joint given as  $p(t, \mathbf{u}, \mathbf{x}, y, \mathbf{z}) = p(t|\mathbf{z})p(\mathbf{u})p(\mathbf{x}|\mathbf{u}, \mathbf{z})p(y|\mathbf{z}, t)p(\mathbf{z})$ . The ELBO is given as

$$\begin{aligned} & \sum_{n=1}^N \mathbb{E}_{q(\mathbf{u}_n, \mathbf{z}_n|\mathbf{x}_n)} [\log p(\mathbf{x}_n|\mathbf{u}_n, \mathbf{z}_n)] + \mathbb{E}_{q(\mathbf{z}_n|\mathbf{x}_n)} [\log p(y_n|t_n, \mathbf{z}_n) + \log p(t_n|\mathbf{z}_n)] \\ & - \text{KL}(q(\mathbf{u}_n|\mathbf{x}_n) \| p(\mathbf{u}_n)) - \text{KL}(q(\mathbf{z}_n|\mathbf{x}_n) \| p(\mathbf{z}_n)) \end{aligned}$$

- “+I” indicates a variation that is identifiable in  $\mathbf{z}$ . Its ELBO differs from *Ours* in the conditional prior and the additional covariates  $\mathbf{c}$ , highlighted in red,

$$\begin{aligned} & \sum_{n=1}^N \mathbb{E}_{q(\mathbf{u}_n, \mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n)} [\log p(\mathbf{x}_n|\mathbf{u}_n, \mathbf{z}_n)] + \mathbb{E}_{q(\mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n)} [\log p(y_n|t_n, \mathbf{z}_n) + \log p(t_n|\mathbf{z}_n)] \\ & - \text{KL}(q(\mathbf{u}_n|\mathbf{x}_n) \| p(\mathbf{u}_n)) - \text{KL}(q(\mathbf{z}_n|\mathbf{c}_n, \mathbf{x}_n) \| p(\mathbf{z}_n|\mathbf{c}_n)) \end{aligned}$$

- “+sep” indicates a variation that splits  $\mathcal{U}$  into two parts, each responsible for one of the two groups. The differences in its ELBO compared to *Ours* are highlighted in red

$$\begin{aligned} & \sum_{n=1}^N \mathbb{E}_{q(\mathbf{u}_n^t, \mathbf{z}_n|\mathbf{x}_n)} [\log p(\mathbf{x}_n|\mathbf{u}_n^t, \mathbf{z}_n)] + \mathbb{E}_{q(\mathbf{z}_n|\mathbf{x}_n)} [\log p(y_n|t_n, \mathbf{z}_n) + \log p(t_n|\mathbf{z}_n)] \\ & - \text{KL}(q(\mathbf{u}_n^t|\mathbf{x}_n) \| p(\mathbf{u}_n)) - \text{KL}(q(\mathbf{z}_n|\mathbf{x}_n) \| p(\mathbf{z}_n)). \end{aligned}$$

- “+snet” uses the deterministic *SNet* structure and regularization, adding a generator  $p(\mathbf{x}|\mathbf{u}, \mathbf{z})$  on top of the representation  $\mathbf{z}$ , where  $\mathbf{z}$  is the combination of all five subspaces.
- “+tedvae” uses the generative *TEDVAE* structure, adds a separate  $p(\mathbf{u})$ , and adapts the covariate log-likelihood to  $p(\mathbf{x}|\mathbf{u}, \mathbf{z}^1, \mathbf{z}^2, \mathbf{z}^3)$ .

We parameterize the probability densities for the model variations as

$$\begin{aligned}
 p(\mathbf{u}), p(\mathbf{u}^t), p(\mathbf{z}), p(\mathbf{z}^m) &= \mathcal{N}(0, \sigma^2 \mathbf{1}), \quad \sigma = 1 \text{ for all experiments,} \\
 p(\mathbf{z}|\mathbf{c}) &= \mathcal{N}(\lambda(\mathbf{c}), \sigma_z^2 \mathbf{1}), \quad \sigma_z = 1 \text{ for all experiments,} \\
 p(x|\mathbf{z}) &= \begin{cases} \mathcal{N}(f_{z \rightarrow x}(\mathbf{z}), \sigma_x), & \text{if } x \text{ is continuous} \\ \text{Ber}(\sigma(f_{z \rightarrow x}(\mathbf{z}))), & \text{if } x \text{ is binary} \\ \text{Cat}(\zeta(f_{z \rightarrow x}(\mathbf{z}))), & \text{if } x \text{ is categorical} \end{cases}, \\
 p(x|\mathbf{u}, \mathbf{z}) &\text{ is defined analogously but with } f_{uz \rightarrow x}(f_{u \rightarrow uz}(\mathbf{u}), f_{z \rightarrow uz}(\mathbf{z})), \\
 p(t|\mathbf{z}) &= \text{Ber}(\sigma(f_{z \rightarrow t}(\mathbf{z}))), \\
 p(y|\mathbf{z}) &= \begin{cases} \mathcal{N}(f_{z \rightarrow y}(\mathbf{z}), \sigma_y) & \text{if continuous} \\ \text{see Appendix C.8} & \text{if time to event} \end{cases}, \\
 q(\mathbf{u}|\mathbf{x}) &= \mathcal{N}(\mu^u(\mathbf{x}), \sigma^u(\mathbf{x})), \\
 q(\mathbf{u}^t|\mathbf{x}) &= \mathcal{N}(\mu_t^u(\mathbf{x}), \sigma_t^u(\mathbf{x})), \\
 q(\mathbf{z}|\mathbf{x}) &= \mathcal{N}(\mu^z(\mathbf{x}), \sigma^z(\mathbf{x})), \\
 q(\mathbf{z}^m|\mathbf{m}) &= \mathcal{N}(\mu^{z^m}(\mathbf{x}), \sigma^{z^m}(\mathbf{x})), \\
 q(\mathbf{z}|\mathbf{c}, \mathbf{x}) &\text{ is defined analogously with an additional input } \mathbf{c}, \\
 &\text{the split } \mathbf{z} \text{ cases, i.e., } +snet, +tedvae \text{ split the output of } \mu^z, \sigma^z \text{ accordingly,}
 \end{aligned}$$

where  $\sigma(\cdot)$  is the logistic sigmoid,  $\zeta(\cdot)$  the softmax, and all remaining functions are parameterized via neural networks. Standard deviations  $\sigma_x$  and  $\sigma_y$  are treated as free parameters that are optimized via gradient descent as part of the training routine.

### C.3 Modern

*TEDVAE* (Zhang et al., 2021) and *SNet* (Curth and van der Schaar, 2021a) serve as two examples of more modern generative and deterministic models and how they can easily be adapted to our setup. *SNet* stands in as an example of a deterministic neural net with a more advanced architectural structure and regularization method. Simply interpreting its mapping from  $\mathcal{X}$  to the representation space  $\mathcal{Z}$  as the parameterization of an amortized variational posterior and adding a generator that mirrors it by learning a mapping from  $\mathcal{Z}$  to  $\mathcal{X}$ , as well as adding a separate representation  $\mathcal{U}$  are all that is needed. *TEDVAE* provides the same illustration for generative models, whose adaptation is even simpler. Here an encoder-decoder structure is already provided lacking simply a separate  $\mathcal{U}$  representation. Other novel approaches can be adapted to our setup in the same manner. For comparison, we replicate the corresponding rows from Table 1 below. Our adaptations improve upon these baselines in all but one setting (highlighted in the table via underlining).

METHOD	(a) full outcome observation (RMSE of CATE)				(b) partial outcome observation (AE of ATT)			
	all+high		subset+low		all+high		subset+low	
	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample
SNet	0.168±0.003	0.264±0.008	0.211±0.008	<u>0.243±0.016</u>	0.043±0.003	0.137±0.007	0.221±0.015	0.366±0.026
Ours+snet	0.141±0.002	0.256±0.008	0.185±0.007	0.253±0.017	0.038±0.002	0.117±0.007	0.194±0.014	0.322±0.020
TEDVAE	0.176±0.002	0.283±0.007	0.293±0.009	0.336±0.018	0.052±0.003	0.179±0.009	0.144±0.009	0.272±0.016
Ours+tedvae	0.143±0.002	0.257±0.008	0.187±0.007	0.258±0.019	0.039±0.003	0.119±0.007	0.103±0.008	0.189±0.015

### C.4 Architectures and Hyperparameters

We rely on the same set of architectures and hyperparameters throughout the various experiments and methods, adapting them only to the specific dimensionalities. Please refer to Appendix C.2 for a description of where the respective nets appear.

**Neural architectures.** We use the notation  $A, B, C$ , to indicate fully connected layers of  $A, B, C$  neurons respectively. All models use exponentiated linear units (Clevert et al., 2016) as activation functions between layers. They follow the layer width and depth of Shalit et al. (2017) as closely as possible.

Function	Mapping	Comments
$\mu(\cdot, \cdot)$	$\dim(\mathbf{x}) + 1, 200, 200, \dim(z), 100, 100, 1$	
$\mu_i(\cdot), i \in \{0, 1\}$	$\dim(\mathbf{x}), 200, 200, \dim(z), 100, 100, 1$	<i>TNet</i>
$\mu_i(\cdot), i \in \{0, 1\}$	$\dim(z), 100, 100, 1$	otherwise
$h(\cdot)$	$\dim(\mathbf{x}), 200, 200, \dim(z)$	
$h_i(\cdot), i \in \{1, \dots, 5\}$	$\dim(\mathbf{x}), 200, 200, s_i$	$s_i$ defined in the hyperparameters
$g(\cdot), f_{z \rightarrow t}$	$\dim(\mathbf{z}), 1$	
$\lambda(\cdot)$	$\dim(\mathbf{c}), 10, 10, \dim(\mathbf{z})$	
$f_{z \rightarrow uz}(\cdot)$	$\dim(\mathbf{z}), 50$	
$f_{u \rightarrow uz}(\cdot)$	$\dim(\mathbf{z}), 150$	
$f_{uz \rightarrow x}(\cdot)$	$200, 200, \dim(\mathbf{x})$	
$f_{z \rightarrow x}(\cdot)$	$\dim(\mathbf{z}), 200, 200, \dim(\mathbf{x})$	
$f_{z \rightarrow y}(\cdot)$	$\dim(\mathbf{z}), 100, 100, 1$	
$\mu_t^u(\cdot), \sigma_t^u(\cdot)$	$\dim(\mathbf{x}), 200, 200, \dim(\mathbf{u})$	
$\mu_t^z(\cdot), \sigma_t^z(\cdot)$	$\dim(\mathbf{x}), 200, 200, \dim(\mathbf{z})$	
$\mu_t^{zm}(\cdot), \sigma_t^{zm}(\cdot)$	$\dim(\mathbf{m}), 200, 200, \dim(\mathbf{z}^m)$	

**Hyperparameters and further settings.** In the fully observed setup, the representation space  $\mathcal{Z}$  is five-dimensional, with  $s = (1, 1, 2, 1)$  to specify the split in the SNet representation space. The dimensionality of  $\mathcal{U}$  is fixed at 50 throughout all experiments.

Discrete covariates are treated as Bernoulli or categorical variables turned into one-hot encodings for all models in the case of categorical covariates. Continuous covariates are treated as normal variables with a homoscedastic noise model, with a variance parameter per covariate inferred via gradient descent. These are modeled as normal variables with a global precision parameter inferred via gradient descent.

We perform gradient descent via the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 0.001 and weight decay of 0.0001. A random subset of 10% of the training data in each of the replications is used as a validation set for early stopping. The upper limit for the maximal number of epochs was set to 500, which was never reached.

For baselines that do not model missing values, we impute the empirical mean for missing continuous covariates and the median for missing binary ones. (Categorical covariates were assumed to always be observed for simplicity.)

## C.5 Semisynthetic Experimental Data

Due to the fundamental problem of causality (Pearl, 2009), the true treatment effect is not known for real data as the counterfactual outcome can never be observed. Any experimental evaluation is therefore constrained to rely on either completely synthetic or semi-synthetic data sets. In the main paper, we focus on the latter case and rely on real covariates with simulated treatment outcomes. This allows us to stay realistic with respect to covariate distributions, while still being able to generate factual and counterfactual outcomes to simulate various treatment effect scenarios.

We rely on two sets of covariates for this task, which we describe below in greater detail. The first is based on covariates from the popular IHDP benchmark (Hill, 2011). The second is a newly curated real-world data set combining RCT study data with electronic health records. We describe each of them in turn.

### C.5.1 IHDP

**Data.** The Infant Health and Development Program (IHDP) data set as used by Hill (2011) consists of 25 covariates. 19 of them are binary, five of them continuous and one categorical.<sup>5</sup> This observational data set contains 139 treated children and 608 untreated ones after preprocessing, where Hill (2011) removes all children with nonwhite mothers from the treated group to create an artificial reduction in overlap between the two groups.

**Synthetic outcome creation.** As is common, we follow Hill (2011). Given the covariates  $\mathbf{x}$ , an offset vector  $\mathbf{w}$ , whose entries all equal 0.5, and a regression vector  $\beta$ , whose coefficients are randomly sampled from (0, 0.1, 0.2, 0.3, 0.4) with probabilities (0.6, 0.1, 0.1, 0.1, 0.1), the potential outcomes are generated as

$$Y(0) \sim \mathcal{N}(\exp((\mathbf{x} + \mathbf{w})^\top \beta), 1) \quad \text{and} \quad Y(1) \sim \mathcal{N}(\mathbf{x}^\top \beta - \omega, 1),$$

where  $\omega$  is an offset, chosen so that the true treatment effect on the treated is given as 4. For the different scenarios described below,  $\mathbf{x}$  is further modified before creating the synthetic outcomes  $Y(0)$  and  $Y(1)$ . We rely on the precomputed

<sup>5</sup>The categorical covariate is modelled as continuous in most prior work.

train/val/test splits provided by [Shalit et al. \(2017\)](#). The number of random data sets used in each experiment is marked in the respective result tables or figure captions.

We generate  $\mathbf{c}$  by picking three random covariates from the 19 binary covariates for each seed. Note that given the randomness in which covariates are predictive, this does not guarantee that the chosen  $\mathbf{c}$  provides good guidance to  $\mathbf{z}$ . Similar to how the researcher cannot know a priori which covariates are useful.

**All vs Subset.** The *all* scenario relies on all 25 covariates  $\mathbf{x}$  to create the synthetic outcomes. As the  $\beta$  vector is randomly sampled and can contain zeros, this still leaves us with potentially fewer than 25 predictive covariates.

Compared to this, *subset* explicitly constraints to fewer than 25 covariates by choosing a random subset of  $n_{\text{pred}}$  covariates for each seed before following the procedure described above. As before,  $n_{\text{pred}}$  only provides us with an upper bound of predictive covariates. Throughout our experiments, we use  $n_{\text{pred}} = 15$ .

**High vs Low.** The *high* setup refers to using the original  $\mathbf{x}$  as defined by [Hill \(2011\)](#). We keep the overlap as is, which has already been perturbed to a certain degree as discussed above.

For *low* we try to mimic real-world scenarios where the overlap tends to be even lower. To increase the divergence between the two data sources we modify the control group as follows. We shift the five continuous covariates as

$$X_c \leftarrow X_c + s + \varepsilon, \quad \text{where } s \sim \text{Ber}(0.5) \cdot 6 - 3, \quad \varepsilon \sim \mathcal{N}(0, 9),$$

for the  $c$ -th continuous covariate. That is,  $s \in \{-3, 3\}$  with equal probability, and  $\varepsilon$  adds a large amount of variation to still keep some overlap. Of the 19 binary covariates, we pick a random subset of five increasing the probability of three of them to be one by 80% and the probability of the other two to be 0 by 80%.

**Missingness.** We compare three settings of increasing missing not at random structure. In each of the three settings, the covariates for each observation are missing based on the following:

- *Strong.* For discrete covariates: If  $x_i = 1$  it has a probability of  $p = 0.2$  of being observed, if  $x_i = 0$  it has a probability of  $p = 0.3$  of being observed. A continuous covariate with  $x_i > 0$  has a probability of  $p = 0.1$  of being observed.
- *Medium.* For discrete covariates: If  $x_i = 1$  it has a probability of  $p = 0.5$  of being observed, if  $x_i = 0$  it has a probability of  $p = 0.7$  of being observed. A continuous covariate with  $x_i > 0$  has a probability of  $p = 0.6$  of being observed.
- *Weak.* For discrete covariates: If  $x_i = 1$  it has a probability of  $p = 0.9$  of being observed, if  $x_i = 0$  it has a probability of  $p = 0.9$  of being observed. A continuous covariate with  $x_i > 0$  has a probability of  $p = 0.9$  of being observed.

Missing discrete covariates are imputed by the median of the observed data set, missing continuous by zeros.

**Survival analysis.** Data generation details for the survival analysis experiments are given in [Appendix C.8](#).

## C.5.2 Real-world data

**Data.** This data set consists of two parts. The single-arm trial data is taken from [Bakris et al. \(2020\)](#), a published RCT study, evaluating the treatment of patients suffering from chronic kidney disease and type 2 diabetes. For our purpose, only a subset of the treated group of patients is selected for the model to train on ( $N=833$ ). The control group is constructed from electronic health records provided by Helsinki University Hospital, whose ethics committee gave ethical approval and a study permit. A set of 2646 patients was selected who roughly fulfilled the inclusion criteria of the RCT study. Although both data sets are originally longitudinal, the measurements closest to the index date are collected into a single set of covariates, leaving proper longitudinal modeling for future work. Note that while the index date is properly defined in the RCT it remains a noisy choice for the EHR data. For each patient, we pick the time a subset of the original inclusion criteria is fulfilled. A subset of 100 binary covariates of the most prevalent diagnosis and medical history is selected (from approximately 4000), as well as 40 continuous covariates (lab measurements age, weight, height, ...) giving us a total of 140 covariates. Due to their wide range of potential values, we standardize them by subtracting the median and scaling them by the interquartile range.<sup>6</sup> As the EHR data contains a lot of missing covariates, we impute missing binary covariates via their median, while missing continuous ones are imputed with zeros (justified due to the prior standardization of the data). To create synthetic treatment outcomes, we require fully observed covariates, i.e., cannot model the true missingness pattern inherent in the model. We provide a longer discussion on the differences and similarities in the respective covariate distributions in [Kurki et al. \(2024\)](#).

<sup>6</sup>See `sklearn.preprocessing.RobustScaler` for an implementation.



**Synthetic outcome creation.** We generate synthetic outcomes as follows. First, a random subset of 30 covariates is selected from the original 140 and mapped with a single hidden layer neural network  $(30,10,\dim(\mathbf{z}))^7$  into a five-dimensional representation space. The outcomes are then given as

$$Y(0) \sim \mathcal{N}(\mathbf{a}^\top \mathbf{v}, 0.3^2) \quad \text{and} \quad Y(1) \sim \mathcal{N}((\mathbf{a} + 0.5)^\top \mathbf{w} - \omega, 0.3^2),$$

where  $\mathbf{v}, \mathbf{w} \sim \mathcal{N}(0, \mathbb{1})$ ,  $\mathbf{a}$  the five dimensional representation, and  $\omega$  such that the ATT is four.

## C.6 Alternative Matching Approaches

**Matching without replacement.** Throughout all of our experiments, we conduct matching with replacement, i.e., we match control patients to treated patients independently of whether they have been matched before. While this ensures that every treated patient gets a control that is closest to them, it effectively reduces the number of used patients and might induce statistical problems due to multiple samples being identical. While this is the most common approach, some alternatives can be considered. The first one would be to pick patients greedily for a random permutation of the treated patients. While this is simple, it will depend on the order and thus won't be optimal. Another approach is to consider the matching as a task that aims to find the optimal minimal weight solution for a bipartite graph. Considering each group as one set of vertices with the edges between the groups given by the respective distances, the optimal matching is given by solving a straightforward linear program.<sup>8</sup>

It should be noted that these two approaches solve the task of reusing patients, but suffer from the problem that the chosen subset is now heavily interdependent, i.e., the choice of one patient is no longer independent of the other choices. We evaluate this latter approach on a subset of methods in Table 4.

**Table 4: Alternative matching.** We compare matching with replacement (*with*) to matching without replacement (*w/out*). Matching without replacement performs similar, or even improves upon the default matching with replacement approach.

METHOD	all+high				subset+low			
	within sample		out-of-sample		within sample		out-of-sample	
	with	w/out	with	w/out	with	w/out	with	w/out
VAE	0.085±0.004	<u>0.083±0.004</u>	0.266±0.012	<u>0.255±0.011</u>	<u>0.264±0.014</u>	0.282±0.014	0.427±0.022	<u>0.368±0.018</u>
CEVAE	<u>0.052±0.002</u>	<u>0.052±0.002</u>	0.174±0.009	<u>0.173±0.008</u>	<u>0.150±0.009</u>	0.192±0.011	<u>0.259±0.015</u>	<u>0.263±0.014</u>
TEDVAE	0.052±0.003	<u>0.050±0.002</u>	<u>0.179±0.009</u>	0.180±0.009	<u>0.144±0.009</u>	0.198±0.010	0.272±0.016	<u>0.271±0.014</u>
Ours	<u>0.037±0.002</u>	<u>0.037±0.002</u>	0.113±0.007	<u>0.106±0.006</u>	0.114±0.008	<u>0.107±0.008</u>	0.190±0.016	<u>0.180±0.013</u>
+I	0.038±0.002	<u>0.036±0.002</u>	0.122±0.007	<u>0.116±0.007</u>	0.110±0.007	<u>0.108±0.008</u>	0.176±0.014	<u>0.167±0.011</u>
+sep	0.037±0.002	<u>0.035±0.002</u>	0.111±0.007	<u>0.109±0.007</u>	0.144±0.009	<u>0.110±0.008</u>	0.272±0.016	<u>0.174±0.012</u>

*mean ± standard error over 300 random replications; the lower of the two means between each variation is underlined*

**Varying metrics.** Within the main paper, we always rely on the Euclidian distance

$$d(\mathbf{a}, \mathbf{b}) = \sqrt{\sum_{i=1}^d (a_i - b_i)^2},$$

for  $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$ . However, given that our variational posteriors consist of distributions instead of point estimates, we can also rely on distributional distance metrics that take posterior variances into account.

The *Wasserstein distance* between two univariate Normal distributions is analytically tractable and given as

$$W(\mathcal{N}(\mu_0, \sigma_0), \mathcal{N}(\mu_1, \sigma_1))^2 = (\mu_0 - \mu_1)^2 + (\sigma_0 - \sigma_1)^2.$$

The squared *Hellinger distance* between two univariate Normal distributions is given as

$$H(\mathcal{N}(\mu_0, \sigma_0), \mathcal{N}(\mu_1, \sigma_1))^2 = 1 - \sqrt{\frac{2\sigma_0\sigma_1}{\sigma_0^2 + \sigma_1^2}} \exp\left(-\frac{(\mu_0 - \mu_1)^2}{4(\sigma_0^2 + \sigma_1^2)}\right).$$

<sup>7</sup>See Appendix C.4 for details on this notation.

<sup>8</sup>Our implementation relies on a SciPy (Virtanen et al., 2020) routine.

Both Wasserstein and Hellinger factorize for mean-field normal distributions which is why we only state their univariate formulae.

For  $q(\mathbf{z}_i) = \mathcal{N}(\mathbf{z}_i, \mu_i, \sigma_i^2)$ , and we  $q(\mathbf{z}_j) = \mathcal{N}(\mathbf{z}_j | \mu_j, \sigma_j^2)$ , we compare the ATT performance for matching in the reduced setting via (i) the Euclidean distance  $d(\mu_i, \mu_j)$ , (ii) the Wasserstein distance  $W(q(\mathbf{z}_i), q(\mathbf{z}_j))$ , (iii) the Hellinger distance  $H(q(\mathbf{z}_i), q(\mathbf{z}_j))$ , and (iv) via the Euclidean distance on a propensity estimator  $\hat{\pi}$ , i.e.,  $d(\hat{\pi}(\mathbf{z}_i), \hat{\pi}(\mathbf{z}_j))$ .

As in the remainder of the paper, we stick to mean-field normal posteriors. However, the low-dimensional structure of  $\mathbf{z}$  allows us to infer more complex posteriors without too many constraints due to computational or memory costs. E.g., by switching to a multivariate Normal, for which the two distances remain analytically tractable, while at the same time allowing for more structured variance estimation. We leave such extensions to future work.

The results are summarized in [Appendix D.5](#).

### C.7 Predictive Latent Space $\mathcal{Z}$ Regularization

As discussed in our main paper, one crucial assumption for causal inference from observational data is overlap, i.e., that  $0 < \pi(x) < 1, \forall x$  where  $\pi(x) = \mathbb{P}(T = 1 | X = x)$  is the propensity score. It is therefore common practice ([Shalit et al., 2017](#); [Johansson et al., 2016](#); [Lu et al., 2020](#)) to further regularize the latent space, in our case  $\mathcal{Z}$ , on top of the existing regularization through the loss. In our case it is already implicitly regularized by the KL terms within our ELBO (4), i.e., in the identifiable setup  $\text{KL}(q(\mathbf{z}) \parallel p(\mathbf{z}|\mathbf{c}))$  and  $\text{KL}(q(\mathbf{z}) \parallel p(\mathbf{z}))$  in the unidentifiable one. To further encourage overlap through additional regularization we can rely on several approaches. One could rely on distributional approximation methods, such as the maximum mean discrepancy method (MMD) ([Gretton et al., 2012](#)), as was proposed, e.g., by [Johansson et al. \(2016\)](#); [Shalit et al. \(2017\)](#). MMD aims to minimize the distance between two aggregate variational posteriors

$$q_t(\mathbf{z}) \triangleq \frac{1}{N_t} \sum_{n:t_n=t} q(\mathbf{z}_n | \mathbf{x}_n), \quad t \in \{0, 1\},$$

by minimizing  $\text{MMD}^2(q_0, q_1)$  defined as

$$\text{MMD}^2(q_0, q_1) = \mathbb{E}_{q_0(z)q_0(z')} [k(z, z')] - 2\mathbb{E}_{q_0(z)q_1(z')} [k(z, z')] + \mathbb{E}_{q_1(z)q_1(z')} [k(z, z')],$$

where we use the exponentiated quadratic kernel as  $k(\cdot, \cdot)$  and approximate the expectations via sampling. The loss to be optimized is then given as

$$\mathcal{L}_{\text{elbo}} + \kappa \text{MMD}^2(q_0, q_1),$$

where  $\kappa$  serves as a balancing factor. Throughout our experiments, we use a fixed  $\kappa = 1$

Another approach is to use critic-based approaches (e.g., [Lu et al., 2020](#)), i.e., learning to fool a critic whose task is to differentiate between treated and control observations, as is commonly done in generative adversarial networks (GANs) ([Goodfellow et al., 2014](#)). Within this work, we rely primarily on the gradient reversal layer formulation introduced by [Ganin and Lempitsky \(2015\)](#) for GAN-based regularization. An additional discriminator tries to classify instances  $\mathbf{z}_n$  correctly as treated/control, i.e., to maximize the log-likelihood  $\log p(t_n | \mathbf{z}_n)$ . The main model tries to maximize its increasing overlap in the process. These two diverging objectives can be incorporated in our existing objective (4) by relying on a pseudo-function

$$R_\kappa(\mathbf{z}) \triangleq \mathbf{z} \quad \text{with} \quad \frac{dR_\kappa}{d\mathbf{z}} \triangleq -\kappa \mathbf{1},$$

and using  $\tilde{f}(\mathbf{z}) \triangleq (f \circ R_\kappa)(\mathbf{z})$  for the likelihood  $p(t|\mathbf{z}) = \text{Ber}(t|\sigma(\tilde{f}(\mathbf{z})))$ , and  $f$  a neural network. The inherent objective of the discriminator  $f$  is then to maximize  $\log p(t|\mathbf{z})$ , while the objective of our generative model becomes a modification of the ELBO, i.e., its objective is to

$$\text{maximize} \quad \mathcal{L}_{\text{elbo}} - 2 \log p(t|\mathbf{z}),$$

where  $\mathcal{L}_{\text{elbo}}$  is given by (4). We use  $\kappa = 1$  in the experiments.

These regularization approaches apply to the case where we want full overlap in the latent space  $\mathcal{Z}$ . This is the case for having access to outcome data from both groups, i.e., the single-arm patients as well as the control. For our estimators  $\mu_t$  to be predictive of the counterfactual of the other group, we want the overlap between these groups to be maximal. However, if we only have access to outcome data  $y$  for the control group, we only require the encodings of the treated group to overlap with the encodings of the control group, but not vice versa. I.e., as long as we can select a suitable subset of control patients, it is irrelevant how the remainder of the controls are encoded. It is therefore sometimes already sufficient to rely on the indirect regularization provided by the KL terms and the generative constraints to ensure a sufficient overlap.

We do this in the one-arm experiments for IHDP. Note, that this is also necessary to fulfill the assumptions of [Khemakhem et al. \(2020\)](#)'s Theorem 4 (see above).

### C.8 Survival

Our survival analysis experiment uses the IHDP covariates as its starting point after transforming them following the *subset+low* strategy. We only experiment with right-censored observations.

Synthetic outcomes are then generated via the following steps, where we roughly follow [Pölsterl \(2019\)](#); [Manduchi et al. \(2022\)](#).

1. Sample  $\beta$  as in the default IHDP setting from  $(0, 0.1, 0.2, 0.3, 0.4)$  with probabilities  $(0.6, 0.1, 0.1, 0.1, 0.1)$
2. For every patient  $n$  compute a risk score assuming they are treated  $r_n^t$  or control  $r_n^c$  as

$$\begin{aligned} r_n^c &= \text{clip}(\text{soft}(\mathbf{x}_n^\top \beta + ((\mathbf{x} - 0.5)^\top \beta)^2)) \\ r_n^t &= \text{clip}(\text{soft}(\mathbf{x}_n^\top \beta + |\mathbf{x}^\top \beta|)) \end{aligned}$$

where  $\text{soft}(x) = \log(1 + \exp(x))$  is the softplus function, and we clip them to be within  $[0.5, 15]$ .

3. For each  $\mathbf{r}^c = (r_1^c, \dots, r_N^c)$  and  $\mathbf{r}^t = (r_1^t, \dots, r_N^t)$ , we, relying on a generic  $\mathbf{r}$  for notational simplicity in the following list, create survival times  $y$  and censoring indicators  $\delta$  as follows
  - (a) Define mean survival time  $T_0$  ( $= 365$  in our experiments) and  $p_{\text{cens}}$  the probability of being censored ( $= 0.3$  in our experiments).
  - (b) Let  $\lambda_n = \exp(r_n)/T_0$ , for  $n = 1, \dots, N$ .
  - (c) Let  $u_n = -\frac{\log a_n}{\lambda_n}$ , for  $a_n \sim \mathcal{U}(0, 1)$  for  $n = 1, \dots, N$ .
  - (d) Let  $q_{\text{cens}} = q_{1-p_{\text{cens}}}(\mathbf{u})$ , where  $\mathbf{u} = (u_1, \dots, u_N)$  and  $q_\alpha(\cdot)$  is the  $\alpha$ -th quantile.
  - (e) Sample  $t_{\text{cens}} \sim \mathcal{U}(\min_n u_n, q_{\text{cens}})$ .
  - (f) Let  $\delta_n = \begin{cases} 1, & \text{if } u_n \leq t_{\text{cens}} \\ 0, & \text{else} \end{cases}$ , for  $n \in \{1, \dots, N\}$ .
  - (g) Let  $y_n = \begin{cases} u_n, & \text{if } \delta = 1 \\ t_{\text{cens}}, & \text{else} \end{cases}$ , for  $n \in \{1, \dots, N\}$ .
4. Normalize survival times  $y$  to the interval  $[0.001, 1.001]$ .

During inference, each model has access to pre-treatment covariates for the treated group, i.e.  $\mathbf{x}$ , with post-treatment outcome data being hidden, and to pre-treatment covariates  $\mathbf{x}$ , as well as, (censored) survival times  $y$  and event indicators  $\delta$  for the control group.

We model the likelihood  $p(y|t, \mathbf{z})$  for each model as a Weibull distribution,  $\mathcal{W}(\lambda, k)$ , whose density is given as

$$f(y; \lambda, k) = \frac{k}{\lambda} \left(\frac{y}{\lambda}\right)^{k-1} \exp\left(-\left(\frac{y}{\lambda}\right)^k\right) I(y \geq 0),$$

with an indicator function  $I$ . For a survival function  $S(y|t, \mathbf{z}) = \int_y^\infty f(\tau; \lambda, k) d\tau$ , the likelihood is given as

$$\begin{aligned} p(y|t, \mathbf{z}) &= f(y; \lambda(\mathbf{z}), k)^\delta S(y|t, \mathbf{z})^{1-\delta} \\ &= \left(\frac{k}{\lambda_t(\mathbf{z})} \left(\frac{y}{\lambda_t(\mathbf{z})}\right)^{k-1} \exp\left(-\left(\frac{y}{\lambda_t(\mathbf{z})}\right)^k\right)\right)^\delta \left(\exp\left(-\left(\frac{y}{\lambda_t(\mathbf{z})}\right)^k\right)\right)^{1-\delta}, \end{aligned}$$

where the shape parameter  $k$  is a fixed hyperparameter ( $k = 1$  throughout our experiments), and  $\lambda_t(\mathbf{z}) = \log(1 + \exp(\mathbf{z}^\top \beta_t))$ .  $\beta_0$  and  $\beta_1$  are optimized via gradient descent together with the remaining parameters.

## C.9 Computation of Performance Metrics

**RMSE of CATE.** As specified in the main paper, the conditional average treatment effect is defined as

$$\mathbb{E}[Y(1) - Y(0)|X = x] = \mu_1(x) - \mu_0(x).$$

Given outcome information for treated ( $t = 1$ ) and untreated ( $t = 0$ ) observations, all our methods and baselines can infer estimators  $\hat{\mu}_0(x)$  and  $\hat{\mu}_1(x)$  to get CATE estimates at test time. Given that the outcomes are synthetically generated, we know the true CATE for each observation and can directly compute and report the RMSE.

We report two RMSEs. First the *within sample* RMSE. During training time, for observation  $(\mathbf{x}_n, t_n, y_n)$ , we only observe the factual outcome  $y_n = Y(t_n)$ . To estimate the CATE at test time, a model needs to correctly predict the unseen counterfactual  $Y(1 - t_n)$  as well. The *out-of-sample* RMSE refers to new covariates  $\mathbf{x}_m$ , i.e., the model has never seen them, nor their factual outcome.

**AE of ATT.** The average treatment effect for the treated is defined as

$$\mathbb{E}[Y(1) - Y(0)|T = 1] = \mathbb{E}[Y(1)|T = 1] - \mathbb{E}[Y(0)|T = 1].$$

In our specific setting, we assume that outcomes  $y$  are only observed for the control group, i.e., for  $t = 0$ .

Once we are unblinded, we can estimate  $\mathbb{E}[Y(1)|T = 1]$  by the sample average  $\frac{1}{N_1} \sum_{n=1:t_n=1}^N y_n$ .

Building an estimator for  $\mathbb{E}[Y(0)|T = 1]$  requires the selection of a suitable subset of control observations that match the characteristics of the treated observations. To avoid any statistical bias, we require this matching to be based only on pre-treatment information, i.e., only on observed covariates for the single-arm trial data. After matching, we estimate  $\mathbb{E}[Y(0)|T = 1]$  via the empirical average  $\frac{1}{N_J} \sum_{j \in J} y_j$ , where  $J$  are the indices of the matched patients, and  $N_J$  the number of matched samples. After matching the treatment outcome is unblinded and the ATT can be estimated and compared with the synthetic true ATT.

**Squared error for Time-to-event data.** For our semi-synthetic survival experiment, the goal is to compare estimated hazard ratios. The evaluation proceeds in four steps: (i) The true hazard ratio is estimated by fitting a Cox proportional hazards regression model on the factual and counterfactual survival curves of the single-arm group; (ii) each model is fit using pre-treatment covariates for both groups and post-treatment survival outcomes for the control group; (iii) a subset of control patients is selected and a second cox proportional hazards regression model is fit on the survival curve for the selected subset as well as the, now unblinded, observed factual outcomes of the treated group; (iv) the squared difference between the two hazard ratio estimates is computed and reported.

Cox models are fitted in R using the `survival` package by [Therneau \(2023\)](#).

## D FURTHER EVALUATION

In this section we provide extended results on the experiments provided in the main paper.

### D.1 Runtime

**IHDP.** Given the small size of the models and the small size of the IHDP data set, models can be trained efficiently and fast on a modern CPU. Training a deterministic model takes about 30 seconds and about twice as much for a generative one. Differences between the individual deterministic/generative approaches are too minuscule to be relevant. These numbers apply to all three experimental setups, CATE estimation, ATT estimation, and survival analysis. Depending on the experimental setup we ran 100–300 replications.

**Real-world data.** Training on the real-world set takes about five minutes for a deterministic model and about twice as much for a generative one.

### D.2 Full Results For Both Arms

We report results for all four settings on the IHDP data set in [Table 5](#).

### D.3 Full Results For Single-arms

We report results for all four settings on the IHDP data set in [Table 6](#).

### D.4 Missingness

We report results for various degrees of missingness in [Table 7](#).

**Table 5:** *Outcome for both groups.* Extended results related to Table 1 in the main paper.

RMSE of CATE	all+high		all+reduced		subset+high		subset+reduced	
	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample
CFor	0.556±0.000	0.5967±0.001	4.571±0.043	4.074±0.047	0.508±0.001	0.507±0.001	3.449±0.052	2.597±0.031
PScov	–	–	–	–	–	–	–	–
PSpca	–	–	–	–	–	–	–	–
PSlat	–	–	–	–	–	–	–	–
SingleNet	0.241±0.005	0.328±0.009	0.863±0.013	0.839±0.030	0.250±0.005	0.274±0.006	0.924±0.016	0.878±0.027
TNet	0.175±0.003	0.275±0.008	0.210±0.008	0.346±0.020	0.217±0.004	0.246±0.005	0.278±0.009	0.328±0.017
TARNet	0.177±0.004	0.280±0.008	0.220±0.009	0.343±0.019	0.221±0.004	0.251±0.005	0.272±0.009	0.333±0.019
CFRNet	0.171±0.003	0.279±0.008	0.297±0.009	0.397±0.020	0.179±0.003	0.210±0.005	0.279±0.008	0.338±0.017
SNet	0.168±0.003	0.264±0.008	0.228±0.009	<b>0.299±0.018</b>	0.159±0.003	0.194±0.005	0.211±0.008	<b>0.243±0.016</b>
VAE	–	–	–	–	–	–	–	–
CEVAE	0.182±0.002	0.287±0.007	0.224±0.008	0.348±0.019	0.223±0.003	0.256±0.004	0.288±0.009	0.350±0.017
TEDVAE	0.176±0.002	0.283±0.007	0.246±0.008	0.344±0.019	0.202±0.002	0.238±0.004	0.293±0.009	0.336±0.018
Ours	0.152±0.002	0.267±0.008	0.174±0.006	0.304±0.019	0.161±0.003	0.202±0.005	0.201±0.007	0.272±0.016
+I	0.141±0.002	0.251±0.008	<b>0.159±0.005</b>	<b>0.293±0.017</b>	0.157±0.003	0.195±0.005	0.194±0.007	0.262±0.016
+sep	0.151±0.002	0.267±0.008	0.184±0.007	0.310±0.019	0.162±0.002	0.204±0.005	0.209±0.007	0.278±0.016
+sep+I	0.143±0.002	0.255±0.007	<b>0.167±0.006</b>	<b>0.294±0.017</b>	0.162±0.003	0.202±0.005	0.197±0.007	0.266±0.017
+snet	0.141±0.002	0.256±0.008	0.183±0.007	0.308±0.018	0.145±0.002	<b>0.184±0.005</b>	<b>0.185±0.007</b>	<b>0.253±0.017</b>
+snet+I	<b>0.130±0.002</b>	<b>0.242±0.008</b>	<b>0.161±0.005</b>	<b>0.288±0.017</b>	<b>0.139±0.002</b>	<b>0.182±0.005</b>	<b>0.181±0.006</b>	<b>0.250±0.017</b>
+snet+sep	0.143±0.002	0.254±0.008	0.168±0.006	0.297±0.018	<b>0.143±0.002</b>	<b>0.186±0.005</b>	0.191±0.007	0.258±0.017
+snet+sep+I	<b>0.131±0.002</b>	<b>0.247±0.008</b>	<b>0.160±0.005</b>	<b>0.290±0.017</b>	<b>0.142±0.002</b>	<b>0.182±0.005</b>	<b>0.180±0.007</b>	<b>0.251±0.017</b>
+tedvae	0.143±0.002	0.257±0.008	0.186±0.007	0.310±0.018	0.146±0.002	<b>0.184±0.005</b>	<b>0.187±0.007</b>	0.258±0.019
+tedvae+I	0.142±0.002	0.255±0.008	0.174±0.007	0.305±0.019	<b>0.142±0.002</b>	<b>0.182±0.005</b>	0.191±0.006	<b>0.253±0.016</b>
+tedvae+sep	0.143±0.002	0.257±0.008	0.192±0.007	0.317±0.019	0.146±0.002	0.188±0.005	0.200±0.007	0.263±0.017
+tedvae+sep+I	0.138±0.002	0.250±0.008	0.181±0.006	0.308±0.018	<b>0.142±0.002</b>	<b>0.185±0.005</b>	0.190±0.007	0.261±0.016

mean ± standard error over 300 random replications; statistically significant best models marked **bold**; lowest mean underlined

## D.5 Matching Metrics

We evaluate two different scenarios on a subset of the methods discussed. ATT estimation in (i) the *all+high* scenario, and (ii) in the *subset+low* scenario. Adding distributional information within the mean-field assumption of  $q(\mathbf{Z})$  provides little benefit and tends to even hurt performance. Matching via a propensity score estimator in the latent space is never competitive in our setting. The results following the current mean-field assumption are summarized in Table 8.

**Table 6:** Outcome only for the control group. Extended results related to Table 1 in the main paper.

METHOD	all+high		all+reduced		subset+high		subset+reduced	
	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample
CFor	–	–	–	–	–	–	–	–
PScov	0.115±0.005	0.336±0.017	0.648±0.082	0.804±0.265	0.111±0.005	0.373±0.019	0.691±0.059	0.702±0.063
PSpca	0.096±0.005	0.300±0.014	0.556±0.056	0.811±0.265	0.115±0.005	0.310±0.014	0.496±0.036	0.613±0.051
PSlat	0.084±0.004	0.187±0.013	0.501±0.039	0.290±0.024	0.096±0.006	0.260±0.016	0.276±0.025	0.272±0.023
SingleNet	–	–	–	–	–	–	–	–
TNet	–	–	–	–	–	–	–	–
TARNet	0.043±0.003	0.130±0.008	0.087±0.008	0.108±0.007	<b>0.075±0.004</b>	0.223±0.012	0.140±0.011	0.239±0.016
CFRNet	0.042±0.002	0.131±0.007	0.321±0.015	0.327±0.016	0.080±0.004	0.224±0.012	0.345±0.016	0.370±0.017
SNet	0.043±0.003	0.137±0.007	0.111±0.010	0.178±0.013	<b>0.075±0.004</b>	0.224±0.012	0.221±0.015	0.366±0.026
VAE	0.085±0.004	0.266±0.012	0.280±0.015	0.360±0.028	0.094±0.004	0.271±0.014	0.264±0.014	0.427±0.022
CEVAE	0.052±0.002	0.174±0.009	0.096±0.006	0.178±0.011	0.086±0.004	0.234±0.012	0.150±0.009	0.259±0.015
TEDVAE	0.052±0.003	0.179±0.009	0.108±0.007	0.188±0.012	0.082±0.004	0.233±0.012	0.144±0.009	0.272±0.016
Ours	<b>0.037±0.002</b>	<b>0.113±0.007</b>	0.074±0.006	<b>0.091±0.008</b>	<b>0.074±0.004</b>	<b>0.212±0.011</b>	<b>0.114±0.008</b>	<b>0.190±0.016</b>
+I	<b>0.038±0.002</b>	<b>0.122±0.007</b>	<b>0.055±0.005</b>	<b>0.080±0.006</b>	0.079±0.004	0.229±0.011	<b>0.110±0.007</b>	<b>0.176±0.014</b>
+sep	<b>0.037±0.002</b>	<b>0.111±0.007</b>	<b>0.063±0.005</b>	<b>0.083±0.006</b>	<b>0.074±0.004</b>	0.222±0.012	0.117±0.009	<b>0.193±0.016</b>
+sep+I	<b>0.035±0.002</b>	<b>0.117±0.008</b>	<b>0.055±0.005</b>	<b>0.078±0.006</b>	<b>0.074±0.004</b>	0.226±0.012	<b>0.104±0.007</b>	<b>0.185±0.015</b>
+snet	<b>0.038±0.002</b>	<b>0.117±0.007</b>	0.091±0.008	0.169±0.012	<b>0.076±0.004</b>	<b>0.217±0.013</b>	0.194±0.014	0.322±0.020
+snet+I	<b>0.038±0.002</b>	<b>0.120±0.008</b>	0.094±0.007	0.159±0.012	<b>0.077±0.004</b>	<b>0.212±0.011</b>	0.190±0.012	0.322±0.023
+snet+sep	0.040±0.003	<b>0.109±0.007</b>	0.095±0.007	0.173±0.012	<b>0.071±0.004</b>	<b>0.211±0.011</b>	0.179±0.013	0.316±0.020
+snet+sep+I	0.039±0.003	<b>0.115±0.007</b>	0.085±0.006	0.156±0.011	<b>0.076±0.004</b>	<b>0.211±0.011</b>	0.196±0.015	0.292±0.020
+tedvae	<b>0.039±0.003</b>	<b>0.119±0.007</b>	0.076±0.006	0.094±0.007	<b>0.074±0.004</b>	<b>0.196±0.011</b>	<b>0.103±0.008</b>	<b>0.189±0.015</b>
+tedvae+I	<b>0.036±0.002</b>	0.124±0.007	<b>0.054±0.004</b>	<b>0.085±0.006</b>	<b>0.072±0.004</b>	0.234±0.013	<b>0.102±0.007</b>	<b>0.191±0.017</b>
+tedvae+sep	0.039±0.002	0.125±0.007	0.067±0.005	<b>0.089±0.007</b>	<b>0.075±0.004</b>	<b>0.209±0.011</b>	<b>0.112±0.009</b>	0.197±0.014
+tedvae+sep+I	0.039±0.002	0.124±0.007	<b>0.061±0.004</b>	<b>0.091±0.007</b>	<b>0.074±0.004</b>	0.218±0.012	<b>0.104±0.008</b>	<b>0.173±0.012</b>

mean ± standard error over 300 random replications; statistically significant best models marked **bold**; lowest mean underlined

**Table 7: Missingness.** Increasing the amount of missing covariates decreases performance as expected. Modeling this in our proposed approach recovers this to a certain degree.

*(a) outcome for both groups (all + high)*

<i>RMSE of CATE</i>	none		weak		medium		strong	
	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample
TARNet	0.176±0.006	0.263±0.012	0.304±0.006	0.314±0.010	0.462±0.011	0.463±0.019	0.531±0.014	0.607±0.025
CFRNet	0.170±0.005	0.258±0.012	0.302±0.007	0.307±0.011	0.457±0.011	0.452±0.019	0.527±0.014	0.600±0.025
SNet	0.172±0.005	0.254±0.013	0.301±0.007	0.293±0.012	0.465±0.012	0.444±0.018	0.536±0.015	0.644±0.027
CEVAE	0.185±0.005	0.277±0.012	0.295±0.005	0.322±0.010	0.438±0.009	0.482±0.017	0.511±0.012	0.635±0.022
TEDVAE	0.176±0.004	0.270±0.012	0.297±0.005	0.318±0.010	0.448±0.011	0.464±0.018	0.514±0.013	0.629±0.023
Our	0.152±0.004	0.249±0.012	0.297±0.005	0.318±0.010	0.412±0.010	0.464±0.019	0.492±0.013	0.632±0.026
+I	0.144±0.004	0.236±0.013	0.268±0.006	0.301±0.010	0.412±0.011	0.453±0.017	0.489±0.013	0.621±0.026
+sep	0.149±0.004	0.245±0.013	0.267±0.005	0.298±0.011	0.412±0.011	0.459±0.019	0.490±0.013	0.622±0.025
+I+sep	0.144±0.004	0.243±0.013	0.268±0.006	0.300±0.011	0.411±0.011	0.452±0.018	0.488±0.013	0.629±0.026
Our+mask	–	–	0.224±0.004	0.275±0.012	0.331±0.006	0.348±0.013	0.430±0.010	0.518±0.019
+I	–	–	0.221±0.004	0.264±0.011	0.327±0.007	0.356±0.014	0.415±0.010	0.506±0.021
+sep	–	–	0.223±0.005	0.267±0.011	0.329±0.006	0.348±0.014	0.428±0.010	0.495±0.017
+I+sep	–	–	0.217±0.004	0.257±0.011	0.326±0.007	0.345±0.013	0.426±0.010	0.492±0.018

*mean ± standard error over 100 random replications*

*(b) outcome only for the control group (subset + low)*

<i>AE of ATT</i>	none		weak		medium		strong	
	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample	within sample	out-of-sample
TARNet	0.138±0.016	0.172±0.018	0.216±0.016	0.238±0.024	0.289±0.019	0.196±0.020	0.448±0.030	0.266±0.027
CFRNet	0.415±0.031	0.414±0.036	0.405±0.029	0.428±0.031	0.318±0.023	0.335±0.029	0.413±0.027	0.313±0.030
SNet	0.177±0.020	0.240±0.024	0.248±0.021	0.294±0.027	0.305±0.021	0.291±0.026	0.434±0.026	0.390±0.035
CEVAE	0.241±0.021	0.274±0.024	0.268±0.021	0.353±0.028	0.289±0.021	0.302±0.025	0.397±0.028	0.312±0.033
TEDVAE	0.238±0.020	0.286±0.025	0.267±0.020	0.331±0.027	0.289±0.022	0.274±0.023	0.409±0.027	0.303±0.034
Our	0.106±0.015	0.160±0.018	0.192±0.014	0.215±0.024	0.271±0.021	0.246±0.020	0.442±0.028	0.291±0.027
+I	0.109±0.016	0.185±0.020	0.174±0.013	0.214±0.020	0.280±0.021	0.236±0.021	0.418±0.027	0.365±0.059
+sep	0.104±0.016	0.181±0.019	0.193±0.015	0.229±0.023	0.288±0.021	0.245±0.021	0.445±0.029	0.268±0.026
+I+sep	0.115±0.017	0.162±0.019	0.191±0.015	0.331±0.027	0.260±0.020	0.215±0.021	0.433±0.030	0.332±0.056
Our+mask	–	–	0.150±0.013	0.200±0.021	0.197±0.014	0.171±0.019	0.378±0.023	0.218±0.022
+I	–	–	0.137±0.014	0.171±0.020	0.176±0.014	0.178±0.018	0.337±0.023	0.239±0.025
+sep	–	–	0.143±0.012	0.176±0.019	0.203±0.015	0.165±0.018	0.375±0.024	0.230±0.022
+I+sep	–	–	0.138±0.013	0.175±0.019	0.179±0.014	0.186±0.019	0.337±0.021	0.228±0.027

*mean ± standard error over 100 random replications*

**Table 8:** *Varying metrics for matching.* We compare relying on the Euclidean metric using the posterior mean, to using distributional metrics or propensity score (PS) matching in the latent space.

*(i) ATT estimation for all + high*

<i>AE of ATT</i>	within sample				out-of-sample			
	METHOD	Euclidian	Wasserstein	Hellinger	PS	Euclidian	Wasserstein	Hellinger
VAE	<u>0.085±0.004</u>	<u>0.085±0.004</u>	0.112±0.006	0.116±0.005	<u>0.266±0.012</u>	<u>0.266±0.012</u>	0.388±0.021	0.315±0.015
CEVAE	<u>0.052±0.002</u>	<u>0.052±0.002</u>	0.109±0.006	0.086±0.004	<u>0.174±0.009</u>	<u>0.174±0.009</u>	0.443±0.025	0.255±0.011
TEDVAE	<u>0.052±0.003</u>	<u>0.053±0.003</u>	0.168±0.012	0.087±0.004	<u>0.179±0.009</u>	<u>0.179±0.009</u>	0.495±0.028	0.258±0.015
Ours	<u>0.037±0.002</u>	<u>0.040±0.003</u>	<u>0.037±0.003</u>	0.085±0.004	<u>0.113±0.007</u>	<u>0.112±0.007</u>	0.128±0.008	0.177±0.015
+I	<u>0.038±0.002</u>	<u>0.039±0.002</u>	0.040±0.002	0.079±0.004	<u>0.122±0.007</u>	<u>0.124±0.007</u>	0.119±0.007	0.182±0.013
+sep	<u>0.037±0.002</u>	<u>0.036±0.002</u>	0.039±0.002	0.079±0.004	<u>0.111±0.007</u>	<u>0.115±0.007</u>	0.116±0.007	0.167±0.012

mean ± standard error over 300 random replications; lowest mean underlined

*(ii) ATT estimation for subset + low*

<i>AE of ATT</i>	within sample				out-of-sample			
	METHOD	Euclidian	Wasserstein	Hellinger	PS	Euclidian	Wasserstein	Hellinger
VAE	<u>0.264±0.014</u>	<u>0.264±0.014</u>	0.288±0.015	0.579±0.071	<u>0.427±0.022</u>	<u>0.427±0.022</u>	0.534±0.034	0.697±0.055
CEVAE	<u>0.150±0.009</u>	<u>0.150±0.009</u>	0.277±0.020	0.554±0.072	<u>0.259±0.015</u>	<u>0.258±0.015</u>	0.539±0.040	0.582±0.045
TEDVAE	<u>0.144±0.009</u>	<u>0.144±0.009</u>	0.277±0.022	0.548±0.055	<u>0.272±0.016</u>	<u>0.271±0.016</u>	0.574±0.048	0.624±0.052
Ours	<u>0.114±0.008</u>	<u>0.113±0.008</u>	0.120±0.008	0.271±0.023	<u>0.190±0.016</u>	<u>0.200±0.017</u>	0.205±0.016	0.252±0.019
+I	<u>0.110±0.007</u>	<u>0.106±0.007</u>	0.108±0.007	0.260±0.025	<u>0.176±0.014</u>	<u>0.177±0.014</u>	0.179±0.015	0.279±0.029
+sep	<u>0.144±0.009</u>	<u>0.144±0.009</u>	0.277±0.022	0.310±0.039	<u>0.272±0.016</u>	<u>0.271±0.016</u>	0.574±0.048	0.275±0.028

mean ± standard error over 300 random replications; lowest mean underlined