
Learning Counterfactual Representations for Estimating Individual Dose-Response Curves

Patrick Schwab¹, Lorenz Linhardt², Stefan Bauer³, Joachim M. Buhmann², Walter Karlen¹

¹Institute of Robotics and Intelligent Systems, ETH Zurich, Switzerland

²Department of Computer Science, ETH Zurich, Switzerland

³MPI for Intelligent Systems, Tübingen, Germany

patrick.schwab@hest.ethz.ch

Abstract

Estimating what would be an individual’s potential response to varying levels of exposure to a treatment is of high practical relevance for several important fields, such as healthcare, economics and public policy. However, existing methods for learning to estimate counterfactual outcomes from observational data are either focused on estimating average dose-response curves, or limited to settings with only two treatments that do not have an associated dosage parameter. Here, we present a novel machine-learning approach towards learning counterfactual representations for estimating individual dose-response curves for any number of treatments with continuous dosage parameters with neural networks. Building on the established potential outcomes framework, we introduce performance metrics, model selection criteria, model architectures, and open benchmarks for estimating individual dose-response curves. Our experiments show that the methods developed in this work set a new state-of-the-art in estimating individual dose-response.

1 Introduction

Estimating dose-response curves from observational data is an important problem in many domains. In medicine, for example, we would be interested in using data of people that have been treated in the past to predict which treatments and associated dosages would lead to better outcomes for new patients [1]. This question is, at its core, a counterfactual one, i.e. we are interested in predicting what *would have happened* if we were to give a patient a specific treatment at a specific dosage in a given situation. Answering such counterfactual questions is a challenging task that requires either further assumptions about the underlying data-generating process or prospective interventional experiments, such as randomised controlled trials (RCTs) [2–4]. However, performing prospective experiments is expensive, time-consuming, and, in many cases, ethically not justifiable [5]. Two aspects make estimating counterfactual outcomes from observational data alone difficult [6, 7]: Firstly, we only observe the factual outcome and never the counterfactual outcomes that would potentially have happened had we chosen a different treatment option. In medicine, for example, we only observe the outcome of giving a patient a specific treatment at a specific dosage, but we never observe what would have happened if the patient was instead given a potential alternative treatment or a different dosage of the same treatment. Secondly, treatments are typically not assigned at random in observational data. In the medical setting, physicians take a range of factors, such as the patient’s expected response to the treatment, into account when choosing a treatment option. Due to this treatment assignment bias, the treated population may differ significantly from the general population. A supervised model naïvely trained to minimise the factual error would overfit to the properties of the treated group, and therefore not generalise to the entire population.

To address these problems, we introduce a novel methodology for training neural networks for counterfactual inference that extends to any number of treatments with continuous dosage parameters. In order to control for the biased assignment of treatments in observational data, we combine our

method with a variety of regularisation schemes originally developed for the discrete treatment setting, such as distribution matching [8, 9], propensity dropout (PD) [10], and matching on balancing scores [7, 11, 12]. In addition, we devise performance metrics, model selection criteria and open benchmarks for estimating individual dose-response curves. Our experiments demonstrate that the methods developed in this work set a new state-of-the-art in inferring individual dose-response curves. The source code for this work is available at <https://github.com/d909b/drnet>.

Contributions. We present the following contributions:

- We introduce a novel methodology for training neural networks for counterfactual inference that, in contrast to existing methods, is suitable for estimating counterfactual outcomes for any number of treatment options with associated exposure parameters.
- We develop performance metrics, model selection criteria, model architectures, and open benchmarks for estimating individual dose-response curves.
- We extend state-of-the-art methods for counterfactual inference for two non-parametric treatment options to the multiple parametric treatment options setting.
- We perform extensive experiments that show that our method sets a new state-of-the-art in inferring individual dose-response curves from observational data across several challenging datasets.

2 Related Work

Background. Causal analysis of treatment effects with rigorous experiments is, in many domains, an essential tool for validating interventions. In medicine, prospective experiments, such as RCTs, are the de facto gold standard to evaluate whether a given treatment is efficacious in treating a specific indication across a population [13, 14]. However, performing prospective experiments is expensive, time-consuming, and often not possible for ethical reasons [5]. Historically, there has therefore been considerable interest in developing methodologies for performing causal inference using readily available observational data [3, 11, 15–19]. The naïve approach of training supervised models to minimise the observed factual error is in general not a suitable choice for counterfactual inference tasks due to treatment assignment bias and the inability to observe counterfactual outcomes. To address the shortcomings of unsupervised and supervised learning in this setting, several adaptations to established machine-learning methods that aim to enable the estimation of counterfactual outcomes from observational data have recently been proposed [6–10, 20–22]. In this work, we build on several of these advances to develop a machine-learning approach for estimating individual dose-response with neural networks.

Estimating Individual Treatment Effects (ITE). ¹ Matching methods [12] are among the most widely used approaches to causal inference from observational data. Matching methods estimate the counterfactual outcome of a sample X to a treatment t using the observed factual outcome of its nearest neighbours that have received t . Propensity score matching (PSM) [11] combats the curse of dimensionality of matching directly on the covariates X by instead matching on the scalar probability $p(t|X)$ of receiving a treatment t given the covariates X . Another category of approaches uses adjusted regression models that receive both the covariates X and the treatment t as inputs. The simplest such model is Ordinary Least Squares (OLS), which may use either one model for all treatments, or a separate model for each treatment [23]. More complex models based on neural networks, like Treatment Agnostic Representation Networks (TARNETs), may be used to build non-linear regression models [9]. Estimators that combine a form of adjusted regression with a model for the exposure in a manner that makes them robust to misspecification of either are referred to as doubly robust [24]. In addition to OLS and neural networks, tree-based estimators, such as Bayesian Additive Regression Trees (BART) [25, 26] and Causal Forests (CF) [20], and distribution modelling methods, such as Causal Multi-task Gaussian Processes (CMGP) [21], Causal Effect Variational Autoencoders (CEVAEs) [22], and Generative Adversarial Nets for inference of Individualised Treatment Effects (GANITE) [6], have also been proposed for ITE estimation.² Other approaches, such as balancing neural networks (BNNs) [8] and counterfactual regression networks (CFRNET) [9], attempt to achieve balanced covariate distributions across treatment groups by explicitly minimising

¹The ITE is sometimes also referred to as the conditional average treatment effect (CATE).

²See [27] and [7] for empirical comparisons of large-numbers of machine-learning methods for ITE estimation for two and more available treatment options.

the empirical discrepancy distance between treatment groups using metrics such as the Wasserstein distance [28]. Most of the works mentioned above focus on the simplest setting with two available treatment options without associated dosage parameters. A notable exception is the generalised propensity score (GPS) [1] that extends the propensity score to treatments with continuous dosages.

In contrast to existing methods, we present the first machine-learning approach to learn to estimate individual dose-response curves for multiple available treatments with a continuous dosage parameter from observational data with neural networks. We additionally extend several known regularisation schemes for counterfactual inference to address the treatment assignment bias in observational data. To facilitate future research in this important area, we introduce performance metrics, model selection criteria, and open benchmarks. We believe this work could be particularly important for applications in precision medicine, where the current state-of-the-art of estimating the average dose response across the entire population does not take into account individual differences, even though large differences in dose-response between individuals are well-documented for many diseases [29–31].

3 Methodology

Problem Statement. We consider a setting in which we are given N observed samples X with p pre-treatment covariates x_i and $i \in [0 \dots p - 1]$. For each sample, the potential outcomes $y_{n,t}(s_t)$ are the response of the n th sample to a treatment t out of the set of k available treatment options $T = \{0, \dots, k - 1\}$ applied at a dosage $s_t \in \{s_t \in \mathbb{R}, a_t > 0 \mid a_t \leq s \leq b_t\}$, where a_t and b_t are the minimum and maximum dosage for treatment t , respectively. The set of treatments T can have two or more available treatment options. As training data, we receive factual samples X and their observed outcomes $y_{n,f}(s_f)$ after applying a specific observed treatment f at dosage s_f . Using the training data with factual outcomes, we wish to train a predictive model to produce accurate estimates $\hat{y}_t(n, s)$ of the potential outcomes across the entire range of s for all available treatment options t . We refer to the range of potential outcomes $y_{n,t}(s)$ across s as the *individual dose-response curve* of the n th sample. This setting is a direct extension of the Rubin-Neyman potential outcomes framework [32].

Assumptions. Following [1, 33], we assume unconfoundedness, which consists of three key parts: (1) Conditional Independence Assumption: The assignment to treatment t is independent of the outcome y_t given the pre-treatment covariates X , (2) Common Support Assumption: For all values of X , it must be possible to observe all treatment options with a probability greater than 0, and (3) Stable Unit Treatment Value Assumption: The observed outcome of any one unit must be unaffected by the assignments of treatments to other units. In addition, we assume smoothness, i.e. that units with similar covariates x_i have similar outcomes y , both for model training and selection.

Metrics. To enable a meaningful comparison of models in the presented setting, we use metrics that cover several desirable aspects of models trained for estimating individual dose-response curves. Our proposed metrics respectively aim to measure a predictive model’s ability (1) to recover the dose-response curve across the entire range of dosage values, (2) to determine the optimal dosage point for each treatment, and (3) to deduce the optimal treatment policy overall, including selection of the right treatment and dosage point, for each individual case. To measure to what degree a model covers the entire range of individual dose-response curves, we use the mean integrated square error³ (MISE) between the true dose-response y and the predicted dose-response \hat{y} as estimated by the model over N samples, all treatments T , and the entire range $[a_t, b_t]$ of dosages s .

$$\text{MISE} = \frac{1}{N} \frac{1}{|T|} \sum_{t \in T} \sum_{n=1}^N \int_{s=a_t}^{b_t} \left(y_{n,t}(s) - \hat{y}_{n,t}(s) \right)^2 ds \quad (1)$$

To further measure a model’s ability to determine the optimal dosage point for each individual case, we calculate the mean dosage policy error (DPE). The mean dosage policy error is the mean squared error in outcome y associated with using the estimated optimal dosage point \hat{s}_t^* according to the predictive model to determine the *true* optimal dosage point s_t^* over N samples and all treatments T .

$$\text{DPE} = \frac{1}{N} \frac{1}{|T|} \sum_{t \in T} \sum_{n=1}^N \left(y_{n,t}(s_t^*) - y_{n,t}(\hat{s}_t^*) \right)^2 \quad (2)$$

³A normalised version of this metric has been used in Silva [34].

where s_t^* and \hat{s}_t^* are the optimal dosage point according to the true dose-response curve and the estimated dose-response curve, respectively.

$$s_t^* = \arg \max_{s \in [a_t, b_t]} y_{n,t}(s) \quad (3) \quad \hat{s}_t^* = \arg \max_{s \in [a_t, b_t]} \hat{y}_{n,t}(s) \quad (4)$$

Finally, the policy error (PE) measures a model’s ability to determine the optimal treatment policy for individual cases, i.e. how much worse the outcome would be when using the estimated best optimal treatment option as opposed to the *true* optimal treatment option and dosage.

$$\text{PE} = \frac{1}{N} \sum_{n=1}^N \left(y_{n,t^*}(s_{t^*}^*) - y_{n,\hat{t}^*}(\hat{s}_{\hat{t}^*}^*) \right)^2 \quad (5)$$

where

$$t^* = \arg \max_{t \in \mathcal{T}} y_{n,t}(s_t^*) \quad (6) \quad \hat{t}^* = \arg \max_{t \in \mathcal{T}} \hat{y}_{n,t}(\hat{s}_t^*) \quad (7)$$

are the optimal treatment option according to the ground truth y and the predictive model, respectively. Considering the DPE and PE alongside the MISE is important to comprehensively evaluate models for counterfactual inference. For example, a model that accurately recovers dose response curves outside the regions containing the optimal response would achieve a respectable MISE but would not be a good model for determining the treatment and dosage choices that lead to the best outcome for a given unit. By considering multiple metrics, we can ensure that predictive models are a capable both in recovering the entire dose-response as well as in selecting the best treatment and dosage choices. We note that, in general, we can not calculate the MISE, DPE or PE without knowledge of the outcome-generating process, since the true dose-response function $y_{n,t}(s)$ is unknown.

Model Architecture. Model structure plays an important role in learning representations for counterfactual inference with neural networks [7, 9, 35]. A particularly challenging aspect of training neural networks for counterfactual inference is that the influence of the treatment indicator variable t may be lost in high-dimensional hidden representations [9]. To address this problem for the setting of two available treatments without dosage parameters, Shalit et al. [9] proposed the TARNET architecture that uses a shared base network and separate head networks for both treatment options. In TARNETs, the head networks are only trained on samples that received the respective treatment. Schwab et al. [7] extended the TARNET architecture to the multiple treatment setting by using k separate head networks, one for each treatment option. In the setting with multiple treatment options with associated dosage parameters, this problem is further compounded because we must maintain not only the influence of t on the hidden representations throughout the network, but also the influence of the continuous dosage parameter s . To ensure the influence of both t and s on hidden representations, we propose a hierarchical architecture for multiple treatments called dose response network (DRNet, Figure 1). DRNets ensure that the dosage parameter s maintains its influence by assigning a head to each of $E \in \mathbb{N}$ equally-sized dosage strata that subdivide the range of potential dosage parameters $[a_t, b_t]$. The hyperparameter E defines the trade-off between computational performance and the resolution $\frac{(b-a)}{E}$ at which the range of dosage values is partitioned. To further attenuate the influence of the dosage parameter s within the head layers, we additionally repeatedly append s to each hidden layer in the head layers. We motivate the proposed hierarchical structure with the effectiveness of the regress and compare approach to counterfactual inference [23], where one builds a separate estimator for each available treatment option. Separate models for each treatment option suffer from data-sparsity, since only units that received each respective treatment can be used to train a per-treatment model and there may not be a large number of samples available for each treatment. DRNets alleviate the issue of data-sparsity by enabling information to be shared both across the entire range of dosages through the treatment layers and across treatments through the base layers.

Model Selection. Given multiple models, it is not trivial to decide which model would perform better at counterfactual tasks, since we in general do not have access to the true dose-response to calculate error metrics like the ones given above. We therefore use a nearest neighbour approximation of the MISE to perform model selection using held-out factual data that has not been used for training. We calculate the nearest neighbour approximation NN-MISE of the MISE using:

$$\text{NN-MISE} = \frac{1}{N} \frac{1}{T} \sum_{t=1}^T \sum_{n=1}^N \int_{s=a_t}^{b_t} \left(y_{\text{NN}(n),t}(s) - \hat{y}_{n,t}(s) \right)^2 ds \quad (8)$$

where we substitute the true dose-response $y_{n,t}$ of the n th sample with the outcome $y_{\text{NN}(n),t}$ of an observed factual nearest neighbour of the n th sample at a dosage point s from the training set. Using

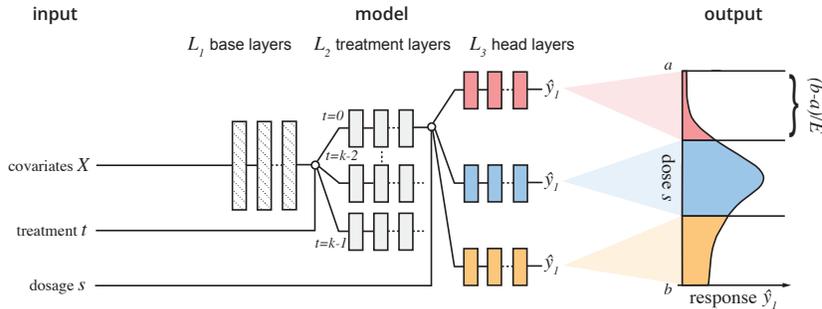


Figure 1: The dose response network (DRNet) architecture with shared base layers, k intermediary treatment layers, and $k * E$ heads for the multiple treatment setting with an associated dosage parameter s . The shared base layers are trained on all samples, and the treatment layers are only trained on samples from their respective treatment category. Each treatment layer is further subdivided into E head layers (only one set of $E = 3$ head layers for treatment $t = 0$ is shown above). Each head layer is assigned a dosage stratum that subdivides the range of potential dosages $[a_t, b_t]$ into E partitions of equal width $(b - a)/E$. The head layers each predict outcomes $\hat{y}_t(s)$ for a range of values of the dosage parameter s , and are only trained on samples that fall within their respective dosage stratum. The hierarchical structure of DRNets enables them to share common hidden representations across all samples (base layers), treatment options (treatment layers), and dosage strata (head layers) while maintaining the influence of both t and s on the hidden layers.

the nearest neighbour approximation of the MISE, we are able to perform model selection without access to the true counterfactual outcomes y . Among others, nearest neighbour methods have also been proposed for model selection in the setting with two available treatments without dosages [36].

Regularisation Schemes. DRNets can be combined with regularisation schemes developed to further address treatment assignment bias. To determine the utility of various regularisation schemes, we evaluated DRNets using distribution matching [9], propensity dropout [10], matching on the entire dataset [12], and on the batch level [7]. We naïvely extended these regularisation schemes since neither of these methods were originally developed for the dose-response setting (Appendix A).

4 Experiments

Our experiments aimed to answer the following questions:

- 1 How does the performance of our proposed approach compare to state-of-the-art methods for estimating individual dose-response?
- 2 How do varying choices of E influence counterfactual inference performance?
- 3 How does increasing treatment assignment bias affect the performance of dose-response estimators?

Datasets. Using real-world data, we performed experiments on three semi-synthetic datasets with two and more treatment options to gain a better understanding of the empirical properties of our proposed approach. To cover a broad range of settings, we chose datasets with different outcome and treatment assignment functions, and varying numbers of samples, features and treatments (Table 1). All three datasets were randomly split into training (63%), validation (27%) and test sets (10%).

News. The News benchmark consisted of 5000 randomly sampled news articles from the NY Times corpus⁴ and was originally introduced as a benchmark for counterfactual inference in the setting with two treatment options without an associated dosage parameter [8]. We extended the original dataset specification [7, 8] to enable the simulation of any number of treatments with associated dosage parameters. The samples X were news articles that consist of word counts $x_i \in \mathbb{N}$, outcomes $y_{s,t} \in \mathbb{R}$ that represent the reader’s opinion of the news item, and a normalised dosage parameter $s_t \in (0, 1]$ that represents the viewer’s reading time. There was a variable number of available treatment options t that corresponded to various devices that could be used to view the News items, e.g. smartphone, tablet, desktop, television or others [8]. We trained a topic model on the entire NY Times corpus to model that consumers prefer to read certain media items on specific viewing devices. We defined $z(X)$ as the topic distribution of news item X , and randomly picked k topic space centroids z_t and $2k$

⁴<https://archive.ics.uci.edu/ml/datasets/bag+of+words>

topic space centroids $z_{s_t,i}$ with $i \in 0, 1$ as prototypical news items. We assigned a random Gaussian outcome distribution with mean $\mu \sim \mathcal{N}(0.45, 0.15)$ and standard deviation $\sigma \sim \mathcal{N}(0.1, 0.05)$ to each centroid. For each sample, we drew ideal potential outcomes from that Gaussian outcome distribution $\tilde{y}_t \sim \mathcal{N}(\mu_t, \sigma_t) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, 0.15)$. The dose response \tilde{y}_s was drawn from a distance-weighted mixture of two Gaussians $\tilde{y}_s \sim d_0 \mathcal{N}(\mu_{s_t,0}, \sigma_{s_t,0}) + d_1 \mathcal{N}(\mu_{s_t,1}, \sigma_{s_t,1})$ using topic space distances $d = \text{softmax}(D(z(X), z_{s_t,i}))$ and the Euclidean distance as distance metric D . We assigned the observed treatment t using $t|x \sim \text{Bern}(\text{softmax}(\kappa \tilde{y}_t \tilde{y}_s))$ with a treatment assignment bias coefficient κ and an exponentially distributed observed dosage s_t using $s_t \sim \text{Exp}(\beta)$ with $\beta = 0.25$. The true potential outcomes $y_{s,t} = C \tilde{y}_t \tilde{y}_s$ were the product of \tilde{y}_t and \tilde{y}_s scaled by a coefficient $C = 50$. We used four different variants of this dataset with $k = 2, 4, 8,$ and 16 viewing devices, and $\kappa = 10, 10, 10,$ and $7,$ respectively. Higher values of κ indicate a higher expected treatment assignment bias depending on $\tilde{y}_t \tilde{y}_s$, with $\kappa = 0$ indicating no assignment bias.

Mechanical Ventilation in the Intensive Care Unit (MVICU). The MVICU benchmark models patients’ responses to different configurations of mechanical ventilation in the intensive care unit. The data was sourced from the publicly available MIMIC III database [37]. The samples X consisted of the last observed measurements x_i of various biosignals, including respiratory, cardiac and ventilation signals. The outcomes were arterial blood gas readings of the ratio of arterial oxygen partial pressure to fractional inspired oxygen PaO_2/FiO_2 which, at values lower than 300, are used as one of the clinical criteria for the diagnosis Acute Respiratory Distress Syndrome (ARDS) [38]. We modelled a mechanical ventilator with $k = 3$ adjustable treatment parameters: (1) the fraction of inspired oxygen, (2) the positive end-expiratory pressure in the lungs, and (3) tidal volume. To model the outcomes, we use the same procedure as for the News benchmark with a Gaussian outcome function and a mixture of Gaussian dose-response function, with the exception that we did not make use of topic models and instead performed the similarity comparisons D in covariate space. We used a treatment assignment bias $\kappa = 10$ and a scaling coefficient $C = 150$. Treatment dosages were drawn according to $s_t \sim \mathcal{N}(\mu_{\text{dose},t}, 0.1)$, where the distribution means were defined as $\mu_{\text{dose}} = (0.6, 0.65, 0.4)$ for each treatment.

The Cancer Genomic Atlas (TCGA). The TCGA project collected gene expression data from various types of cancers in 9659 individuals [39]. There were $k = 3$ available clinical treatment options: (1) medication, (2) chemotherapy, and (3) surgery. We used a synthetic outcome function that simulated the risk of cancer recurrence after receiving either of the treatment options based on the real-world gene expression data. We standardised the gene expression data using the mean and standard deviations of gene expression at each gene locus for normal tissue in the training set. To model the outcomes, we followed the same approach as in the MVICU benchmark with similarity comparisons done in covariate space using the cosine similarity as distance metric D , and parameterised with $\kappa = 10$ and $C = 50$. Treatment dosages in the TCGA benchmark were drawn according to $s_t \sim \mathcal{N}(0.65, 0.1)$.

Dataset	# Samples	# Features	# Treatments
News	5000	2870	2/4/8/16
MVICU	8040	49	3
TCGA	9659	20531	3

Table 1: Comparison of the benchmark datasets used in our experiments. We evaluate on three semi-synthetic datasets with varying numbers of treatments and samples.

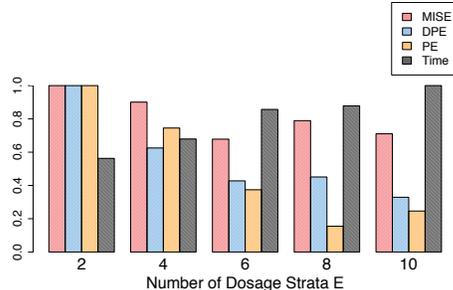


Figure 2: Analysis of the effect of choosing various numbers of dosage strata E (x-axis) on MISE (red), DPE (blue), PE (orange) and Time (black) as calculated on the MVICU benchmark. Metrics were normalised to the range $[0, 1]$. All other hyperparameters besides E were held equal.

Models. We evaluated DRNet, ablations, baselines, and all relevant state-of-the-art methods: k-nearest neighbours (kNN) [12], BART [25, 26], CF [20], GANITE [6], TARNET [9], and GPS [1] using the "causaldrf" package [40]. We evaluated which regularisation strategy for learning counterfactual representations is most effective by training DRNets using a Wasserstein regulariser between treatment group distributions (+ Wasserstein) [9], PD (+ PD) [10], batch matching (+ PM) [7], and matching the entire training set as a preprocessing step [41] using the PM algorithm (+ PSM_{PM}) [7]. To determine whether the DRNet architecture is more effective than its alternatives at learning representations for counterfactual inference in the presented setting, we also evaluated (1) a multi-layer perceptron (MLP) that received the treatment index t and dosage s as additional inputs, and (2) a TARNET for multiple treatments that received the dosage s as an extra input (TARNET) [7, 8] with all other hyperparameters beside the architecture held equal. As a final ablation of DRNet, we tested whether appending the dosage parameter s to each hidden layer in the head networks is effective by also training DRNets that only receive the dosage parameter once in the first hidden layer of the head network (- Repeat). We naïvely extended CF, GANITE and BART by adding the dosage as an additional input covariate, because they were not designed for treatments with dosages.

Hyperparameters. To ensure a fair comparison of the tested models, we took a systematic approach to hyperparameter search. Each model was given exactly the same number of hyperparameter optimisation runs with hyperparameters chosen at random from predefined hyperparameter ranges (Appendix B). We used 5 hyperparameter optimisation runs for each model on TCGA and 10 on all other benchmarks. Furthermore, we used the same random seed for each model, i.e. all models were evaluated on exactly the same sets of hyperparameter configurations. After computing the hyperparameter runs, we chose the best model based on the validation set NN-MISE. This setup ensures that each model received the same degree of hyperparameter optimisation. For all DRNets and ablations, we used $E = 5$ dosage strata with the exception of those presented in Figure 2.

Metrics. For each dataset and model, we calculated the $\sqrt{\text{MISE}}$, $\sqrt{\text{DPE}}$, and $\sqrt{\text{PE}}$. We used Romberg integration with 64 equally spaced samples from $y_{n,t}$ and $\hat{y}_{n,t}$ to compute the inner integral over the range of dosage parameters necessary for the MISE metric. To compute the optimal dosage points and treatment options in the DPE and PE, we used Sequential Least Squares Programming (SLSQP) to determine the respective maxima of $y_{n,t}(s)$ and $\hat{y}_{n,t}(s)$ numerically.

5 Results and Discussion

Counterfactual Inference. In order to evaluate the relative performances of the various methods across a wide range of settings, we compared the MISE of the listed models for counterfactual inference on the News-2/4/8/16, MVICU and TCGA benchmarks (Table 2; other metrics in Appendix D). Across the benchmarks, we found that DRNets outperformed all existing state-of-the-art methods in terms of MISE. We also found that DRNets that used additional regularisation strategies outperformed vanilla DRNets on News-2, News-4, News-8 and News-16. However, on MVICU and TCGA, DRNets that used additional regularisation performed similarly as standard DRNets. Where regularisation was effective, Wasserstein regularisation between treatment groups (+ Wasserstein) and batch matching (+ PM) were generally slightly more effective than PSM_{PM} and PD. In addition, not repeating the dosage parameter for each layer in the per-dosage range heads of a DRNet (- Repeat) performed worse than appending the dosage parameter on News-2, News-4 and News-8. Lastly, the results showed that DRNet improved upon both TARNET and the MLP baseline by a large margin across all datasets - demonstrating that the hierarchical dosage subdivision introduced by DRNets

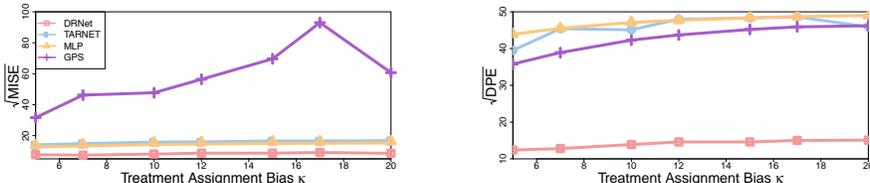


Figure 3: Comparison of DRNet (red), TARNET (blue), MLP (yellow) and GPS (purple) in terms of their $\sqrt{\text{MISE}}$ (top) and $\sqrt{\text{DPE}}$ (bottom) for varying levels of treatment assignment bias κ (x-axis) on News-2. DRNet performs better than other methods across the entire evaluated range of treatment assignment bias values, and is more robust to increasing levels of κ .

Table 2: Comparison of methods for counterfactual inference with multiple parametric treatments on News-2/4/8/16, MVICU and TCGA. We report the mean value \pm the standard deviation of $\sqrt{\text{MISE}}$ on the respective test sets over 5 repeat runs with new random seeds. n.r. = not reported for computational reasons (excessive runtime). \dagger = significantly different from DRNet ($\alpha < 0.05$).

Method	News-2	News-4	News-8	News-16	MVICU	TCGA
DRNet	8.0 \pm 0.1	11.6 \pm 0.1	10.2 \pm 0.1	10.3 \pm 0.0	31.1 \pm 0.4	9.6 \pm 0.0
- Repeat	\dagger 9.0 \pm 0.1	\dagger 11.9 \pm 0.2	10.3 \pm 0.1	10.4 \pm 0.1	31.0 \pm 0.3	10.2 \pm 0.2
+ Wasserstein	\dagger 7.7 \pm 0.2	11.5 \pm 0.0	\dagger 10.0 \pm 0.0	\dagger 10.2 \pm 0.0	32.9 \pm 2.9	10.2 \pm 0.9
+ PD	\dagger 9.0 \pm 0.2	\dagger 12.2 \pm 0.1	\dagger 10.6 \pm 0.2	10.3 \pm 0.1	\dagger 36.9 \pm 0.9	\dagger 11.9 \pm 1.4
+ PM	\dagger 8.4 \pm 0.3	\dagger 12.2 \pm 0.1	\dagger 11.4 \pm 0.3	\dagger 12.3 \pm 0.3	31.2 \pm 0.4	9.7 \pm 0.2
+ PSM _{PM}	\dagger 8.6 \pm 0.1	\dagger 12.2 \pm 0.2	\dagger 11.5 \pm 0.2	\dagger 12.2 \pm 0.3	\dagger 32.6 \pm 0.5	\dagger 11.4 \pm 0.6
MLP	\dagger 15.3 \pm 0.1	\dagger 14.5 \pm 0.0	\dagger 13.9 \pm 0.1	\dagger 14.0 \pm 0.0	\dagger 49.5 \pm 5.1	\dagger 15.3 \pm 0.2
TARNET	\dagger 15.5 \pm 0.1	\dagger 15.4 \pm 0.0	\dagger 14.7 \pm 0.1	\dagger 14.7 \pm 0.1	\dagger 58.0 \pm 4.8	\dagger 14.7 \pm 0.1
GANITE	\dagger 16.8 \pm 0.1	\dagger 15.6 \pm 0.1	\dagger 14.8 \pm 0.1	\dagger 14.8 \pm 0.0	\dagger 59.5 \pm 0.8	\dagger 15.4 \pm 0.2
kNN	\dagger 16.2 \pm 0.0	\dagger 14.7 \pm 0.0	\dagger 15.0 \pm 0.0	\dagger 14.5 \pm 0.0	\dagger 54.9 \pm 0.0	n.r.
GPS	\dagger 47.6 \pm 0.1	\dagger 24.7 \pm 0.1	\dagger 22.9 \pm 0.0	\dagger 15.5 \pm 0.1	\dagger 78.3 \pm 0.0	\dagger 26.3 \pm 0.0
CF	\dagger 26.0 \pm 0.0	\dagger 20.5 \pm 0.0	\dagger 19.6 \pm 0.0	\dagger 14.9 \pm 0.0	\dagger 57.5 \pm 0.0	\dagger 15.2 \pm 0.0
BART	\dagger 13.8 \pm 0.2	\dagger 14.0 \pm 0.1	\dagger 13.0 \pm 0.1	n.r.	\dagger 47.1 \pm 0.8	n.r.

is effective, and that an optimised model structure is paramount for learning representations for counterfactual inference.

Number of Dosage Strata E . To determine the impact of the choice of the number of dosage strata E on DRNet performance, we analysed the estimation performance and computation time of DRNets trained with various numbers of dosage strata E on the MVICU benchmark (Figure 2). With all other hyperparameters held equal, we found that a higher number of dosage strata in general improves estimation performance, because the resolution at which the dosage range is partitioned is increased. However, there is a trade-off between resolution and computational performance, as higher values of E consistently increased the computation time necessary for training and prediction.

Treatment Assignment Bias. To assess the robustness of DRNets and existing methods to increasing levels of treatment assignment bias in observational data, we compared the performance of DRNet to TARNET, MLP and GPS on the test set of News-2 with varying choices of treatment assignment bias $\kappa \in [5, 20]$ (Figure 3). We found that DRNet outperformed existing methods across the entire range of evaluated treatment assignment biases.

Limitations. A general limitation of methods that attempt to estimate causal effects from observational data is that they are based on untestable assumptions [2]. In this work, we assume unconfoundedness [1, 33], which implies that one must have reasonable certainty that the available covariate set X contains the most relevant variables for the problem setting being modelled. Making this judgement can be difficult in practice, particularly when one does not have much prior knowledge about the underlying causal process. Even without such certainty, this approach may nonetheless be a justifiable starting point to generate hypotheses when experimental data is not available [42].

6 Conclusion

We presented a deep-learning approach to learning to estimate individual dose-response to multiple treatments with continuous dosage parameters based on observational data. We extended several existing regularisation strategies to the setting with any number of treatment options with associated dosage parameters, and combined them with our approach in order to address treatment assignment bias inherent in observational data. In addition, we introduced performance metrics, model selection criteria, model architectures, and new open benchmarks for this setting. Our experiments demonstrated that model structure is paramount in learning neural representations for counterfactual inference of dose-response curves from observational data, and that there is a trade-off between model resolution and computational performance in DRNets. DRNets significantly outperform existing state-of-the-art methods in inferring individual dose-response curves across several benchmarks.

Acknowledgements

This work was partially funded by the Swiss National Science Foundation (SNSF) project No. 167302 within the National Research Program (NRP) 75 “Big Data”. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPUs used for this research. The results shown here are in whole or part based upon data generated by the TCGA Research Network: <http://cancergenome.nih.gov/>.

References

- [1] Guido W Imbens. The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3):706–710, 2000.
- [2] Richard Stone. The assumptions on which causal inferences rest. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 455–466, 1993.
- [3] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [4] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.
- [5] Arthur Schafer. The ethics of the randomized clinical trial. *New England Journal of Medicine*, 307(12):719–724, 1982.
- [6] 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.
- [7] Patrick Schwab, Lorenz Linhardt, and Walter Karlen. Perfect match: A simple method for learning representations for counterfactual inference with neural networks. *arXiv preprint arXiv:1810.00656*, 2018.
- [8] Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International Conference on Machine Learning*, pages 3020–3029, 2016.
- [9] Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *International Conference on Machine Learning*, 2017.
- [10] Ahmed M Alaa, Michael Weisz, and Mihaela van der Schaar. Deep counterfactual networks with propensity-dropout. *arXiv preprint arXiv:1706.05966*, 2017.
- [11] Paul R. Rosenbaum and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- [12] Daniel E Ho, Kosuke Imai, Gary King, and Elizabeth A Stuart. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political analysis*, 15(3):199–236, 2007.
- [13] Daniel Carpenter. *Reputation and power: organizational image and pharmaceutical regulation at the FDA*. Princeton University Press, 2014.
- [14] Laura E. Bothwell, Jeremy A. Greene, Scott H. Podolsky, and David S. Jones. Assessing the Gold Standard — Lessons from the History of RCTs. *New England Journal of Medicine*, 374(22):2175–2181, 2016.
- [15] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.
- [16] Joshua D Angrist, Guido W Imbens, and Donald B Rubin. Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434):444–455, 1996.
- [17] James M Robins, Miguel Angel Hernan, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.
- [18] Miguel A Hernán and James M Robins. Using big data to emulate a target trial when a randomized trial is not available. *American journal of epidemiology*, 183(8):758–764, 2016.
- [19] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.
- [20] Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 2017.

- [21] Ahmed M Alaa and Mihaela van der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. In *Advances in Neural Information Processing Systems*, pages 3424–3432, 2017.
- [22] 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*, pages 6446–6456, 2017.
- [23] Nathan Kallus. Recursive partitioning for personalization using observational data. In *International Conference on Machine Learning*, 2017.
- [24] Michele Jonsson Funk, Daniel Westreich, Chris Wiesen, Til Stürmer, M. Alan Brookhart, and Marie Davidian. Doubly robust estimation of causal effects. *American Journal of Epidemiology*, 173(7):761–767, 2011.
- [25] Hugh A Chipman, Edward I George, Robert E McCulloch, et al. BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.
- [26] Hugh Chipman and Robert McCulloch. BayesTree: Bayesian additive regression trees. *R package version 0.3-1.4*, 2016.
- [27] Michael C Knaus, Michael Lechner, and Anthony Strittmatter. Machine learning estimation of heterogeneous causal effects: Empirical monte carlo evidence. *arXiv preprint arXiv:1810.13237*, 2018.
- [28] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.
- [29] Darwin E Zaske, Patrick Irvine, Linda M Strand, Richard G Strate, Robert J Cipolle, and John Rotschafer. Wide interpatient variations in gentamicin dose requirements for geriatric patients. *JAMA*, 248(23):3122–3126, 1982.
- [30] Hans Oldenhof, Martin de Jong, Adri Steenhoek, and Rob Janknegt. Clinical pharmacokinetics of midazolam in intensive care patients, a wide interpatient variability? *Clinical Pharmacology & Therapeutics*, 43(3):263–269, 1988.
- [31] Charles L Campbell, Susan Smyth, Gilles Montalescot, and Steven R Steinhubl. Aspirin dose for the prevention of cardiovascular disease: a systematic review. *JAMA*, 297(18):2018–2024, 2007.
- [32] Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- [33] Michael Lechner. Identification and estimation of causal effects of multiple treatments under the conditional independence assumption. In *Econometric Evaluation of Labour Market Policies*, pages 43–58. Springer, 2001.
- [34] Ricardo Silva. Observational-interventional priors for dose-response learning. In *Advances in Neural Information Processing Systems*, pages 1561–1569, 2016.
- [35] Ahmed Alaa and Mihaela Schaar. Limits of estimating heterogeneous treatment effects: Guidelines for practical algorithm design. In *International Conference on Machine Learning*, pages 129–138, 2018.
- [36] Alejandro Schuler, Michael Baiocchi, Robert Tibshirani, and Nigam Shah. A comparison of methods for model selection when estimating individual treatment effects. *arXiv preprint arXiv:1804.05146*, 2018.
- [37] Mohammed Saeed, Mauricio Villarroel, Andrew T Reisner, Gari D Clifford, Li-Wei Lehman, George Moody, Thomas Heldt, Tin H Kyaw, Benjamin Moody, and Roger G Mark. Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): A public-access intensive care unit database. *Critical Care Medicine*, 39(5):952, 2011.
- [38] Niall D Ferguson, Eddy Fan, Luigi Camporota, Massimo Antonelli, Antonio Anzueto, Richard Beale, Laurent Brochard, Roy Brower, Andrés Esteban, Luciano Gattinoni, et al. The berlin definition of ards: an expanded rationale, justification, and supplementary material. *Intensive Care Medicine*, 38(10):1573–1582, 2012.
- [39] John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Mills Shaw, Brad A Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, Joshua M Stuart, Cancer Genome Atlas Research Network, et al. The cancer genome atlas pan-cancer analysis project. *Nature genetics*, 45(10):1113–1120, 2013.

- [40] Douglas Galagate. *Causal inference with a continuous treatment and outcome: Alternative estimators for parametric dose-response functions with applications*. PhD thesis, University of Maryland, 2016.
- [41] Daniel E Ho, Kosuke Imai, Gary King, Elizabeth A Stuart, et al. MatchIt: nonparametric preprocessing for parametric causal inference. *Journal of Statistical Software*, 42(8):1–28, 2011.
- [42] Guido W Imbens. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics*, 86(1):4–29, 2004.

Supplementary Material for: Learning Counterfactual Representations for Estimating Individual Dose-Response Curves

Patrick Schwab¹, Lorenz Linhardt², Stefan Bauer³, Joachim M. Buhmann², Walter Karlen¹

¹Institute of Robotics and Intelligent Systems, ETH Zurich, Switzerland

²Department of Computer Science, ETH Zurich, Switzerland

³MPI for Intelligent Systems, Tübingen, Germany

patrick.schwab@hest.ethz.ch

A Treatment Assignment Bias Regularisation

To address treatment assignment bias in DRNets, we evaluated four different regularisation strategies: (1) distribution matching between treatment groups using the Wasserstein regulariser (+ Wasserstein), (2) propensity dropout (+ PD), (3) matching on the entire dataset, and (4) matching on the batch level. Because neither of these regularisation strategies were originally developed for the setting with parametric treatment options, we implemented naïve extensions thereof for this setting. To extend (1) to this setting, we followed Schwab et al. (2018) and penalised pair-wise differences between treatment group distributions in the topmost shared hidden layer using the first treatment option as the control treatment. For (2), we applied PD for each treatment option to both the corresponding per-treatment layers and to the respective treatment option's head layers in each dosage stratum. To use (3) dataset-wide and (4) per-batch matching with parametric treatments, we followed the PM algorithm outlined in Schwab et al. (2018), and matched directly on the covariates X with the dosage parameter s added to the covariate set. For datasets of dimensionality higher than 200, we matched on a low-dimensional representation obtained via Principal Components Analysis (PCA) using 50 principal components in order to reduce the computational requirements.

B Hyperparameters

We used a standardised approach to hyperparameter optimisation for all methods. Each method was given exactly the same amount of hyperparameter optimisation runs (5 on TCGA, and 10 on all other benchmarks). We also fixed the random seed such that all methods were evaluated on exactly the same random hyperparameter configurations. For the methods based on neural network models (DRNet, + Repeat, + Wasserstein, + PD, + PM, + PSM_{PM}, MLP, TARNET, GANITE), we chose hyperparameters at random from predefined ranges (Table S1). For "+ Wasserstein", we additionally chose an imbalance penalty weight at random from 0.1, 1.0, or 10.0. For GANITE, we also randomly chose the supervised loss weights α and β from 0.1, 1.0, and 10.0 (Yoon et al., 2018) as an additional hyperparameter. For BART and CF, we used the default hyperparameters from their respective implementations in the R-packages "bartMachine" (Kapelner & Bleich, 2013) and "grf" (Athey et al., 2016). Because CF was designed for estimating the difference in treatment effect between two treatment options and not for directly estimating treatment outcomes \hat{y}_t , we used a baseline ridge regression model with regularisation strength $\alpha = 0.5$ to estimate a control outcome \hat{y}_0 for the first treatment and one CF model to estimate the difference in treatment effect between that control treatment and all other treatment options (Schwab et al., 2018). For KNN, we used 5 nearest neighbours to compute the potential outcomes matching on the covariates X with the dosage parameter s added as an additional covariate. For GPS, we used the implementation in the "causaldrf" R-package (Galagate, 2016) with a normal treatment model, a linear treatment formula,

Table S1: Hyperparameter ranges used in our experiments.

Hyperparameter	Values
Batch size B	32, 64, 128
Number of units per hidden layer M	24, 48, 96
Number of hidden layers L	2, 3
Dropout percentage p_{dropout}	[0.0, 0.2]

and a polynomial of degree 2 as the outcome formula. To reduce the computational requirements for GPS to a manageable level, we additionally preprocessed the covariates X using PCA dimensionality reduction with 16 principal components for benchmarks with a covariate-space dimensionality higher than 200.

C Computing Infrastructure

We used computing infrastructure consisting of compute nodes with Intel i5 and Xeon CPUs to run the experiments described in this paper.

D Dosage Policy Error and Policy Error Results

We report the Dosage Policy Error (DPE) and Policy Error (PE) results in Tables S2 and S3, respectively. We found that DRNets outperformed all existing state-of-the-art methods in terms of DPE and PE with the exception of the PE on News-8, where DRNets achieved the second-best result after GPS.

Table S2: Comparison of methods for counterfactual inference with multiple parametric treatments on News-2/4/8/16, MVICU and TCGA. We report the mean value \pm the standard deviation of $\sqrt{\text{DPE}}$ on the respective test sets over 5 repeat runs with new random seeds. n.r. = not reported for computational reasons (excessive runtime). \dagger = significantly different from DRNet ($\alpha < 0.05$).

Method	News-2	News-4	News-8	News-16	MVICU	TCGA
DRNet	14.0 \pm 0.1	14.3 \pm 0.2	17.2 \pm 0.6	19.1 \pm 2.1	5.0 \pm 1.5	2.1 \pm 0.4
- Repeat	17.5 \pm 7.1	\dagger 18.5 \pm 3.8	16.8 \pm 2.9	16.1 \pm 3.0	5.4 \pm 3.1	5.4 \pm 3.9
+ Wasserstein	\dagger 13.8 \pm 0.1	14.2 \pm 0.3	16.7 \pm 0.6	18.7 \pm 2.4	20.3 \pm 19.0	2.9 \pm 1.3
+ PD	\dagger 13.7 \pm 0.1	14.3 \pm 0.2	17.6 \pm 0.0	16.5 \pm 2.0	\dagger 38.9 \pm 19.6	\dagger 12.7 \pm 5.4
+ PM	\dagger 13.7 \pm 0.1	13.9 \pm 0.5	\dagger 21.3 \pm 2.7	21.0 \pm 0.7	\dagger 12.5 \pm 3.5	1.8 \pm 0.1
+ PSM _{PM}	14.0 \pm 0.1	14.3 \pm 0.1	\dagger 19.5 \pm 1.3	19.2 \pm 1.6	\dagger 28.1 \pm 1.4	2.5 \pm 1.1
MLP	\dagger 47.1 \pm 0.0	\dagger 43.2 \pm 0.1	\dagger 39.8 \pm 0.4	\dagger 40.7 \pm 0.0	\dagger 42.8 \pm 22.8	\dagger 31.1 \pm 0.1
TARNET	\dagger 45.0 \pm 1.1	\dagger 41.5 \pm 1.1	\dagger 38.6 \pm 1.1	\dagger 38.1 \pm 1.4	\dagger 106. \pm 48.8	\dagger 38.6 \pm 1.1
GANITE	\dagger 39.4 \pm 0.1	\dagger 35.5 \pm 0.1	\dagger 32.6 \pm 0.2	\dagger 32.0 \pm 0.3	\dagger 103. \pm 9.1	\dagger 26.5 \pm 0.6
kNN	\dagger 44.6 \pm 0.0	\dagger 41.4 \pm 0.0	\dagger 39.2 \pm 0.0	\dagger 38.1 \pm 0.0	\dagger 60.5 \pm 0.0	n.r.
GPS	\dagger 42.3 \pm 0.0	\dagger 34.7 \pm 0.0	\dagger 22.5 \pm 0.0	17.3 \pm 0.1	\dagger 81.6 \pm 0.0	\dagger 23.8 \pm 0.0
CF	\dagger 47.1 \pm 0.0	\dagger 43.2 \pm 0.0	\dagger 38.3 \pm 0.0	\dagger 35.2 \pm 0.0	\dagger 116. \pm 0.0	\dagger 30.9 \pm 0.0
BART	\dagger 31.6 \pm 1.4	\dagger 25.0 \pm 1.5	\dagger 23.1 \pm 0.4	n.r.	\dagger 39.9 \pm 2.4	n.r.

Table S3: Comparison of methods for counterfactual inference with multiple parametric treatments on News-2/4/8/16, MVICU and TCGA. We report the mean value \pm the standard deviation of $\sqrt{\text{PE}}$ on the respective test sets over 5 repeat runs with new random seeds. n.r. = not reported for computational reasons (excessive runtime). \dagger = significantly different from DRNet ($\alpha < 0.05$).

Method	News-2	News-4	News-8	News-16	MVICU	TCGA
DRNet	15.7 \pm 0.2	15.5 \pm 0.4	15.1 \pm 0.2	19.2 \pm 14.9	12.9 \pm 1.2	2.1 \pm 0.1
- Repeat	20.4 \pm 9.6	20.0 \pm 10.	24.3 \pm 12.	5.3 \pm 8.8	13.7 \pm 2.5	2.9 \pm 1.3
+ Wasserstein	\dagger 15.2 \pm 0.1	15.4 \pm 0.5	15.6 \pm 1.1	19.2 \pm 14.9	13.6 \pm 1.3	2.2 \pm 0.0
+ PD	15.3 \pm 0.6	\dagger 32.8 \pm 0.0	14.9 \pm 0.0	\dagger 0.9 \pm 0.0	\dagger 48.1 \pm 27.1	\dagger 21.8 \pm 9.2
+ PM	\dagger 15.1 \pm 0.1	12.4 \pm 5.3	\dagger 29.3 \pm 7.2	33.0 \pm 11.7	12.1 \pm 1.9	2.3 \pm 0.4
+ PSM _{PM}	\dagger 16.2 \pm 0.4	15.3 \pm 0.3	\dagger 23.7 \pm 7.3	\dagger 1.3 \pm 0.8	\dagger 23.3 \pm 4.7	\dagger 2.3 \pm 0.1
MLP	\dagger 49.5 \pm 0.1	\dagger 49.6 \pm 0.0	\dagger 48.5 \pm 0.7	\dagger 48.3 \pm 0.3	\dagger 23.7 \pm 9.8	\dagger 37.1 \pm 2.3
TARNET	\dagger 47.2 \pm 2.1	\dagger 48.0 \pm 1.8	\dagger 47.7 \pm 1.1	\dagger 44.8 \pm 3.2	\dagger 102. \pm 44.0	\dagger 47.7 \pm 1.1
GANITE	\dagger 42.6 \pm 0.3	\dagger 40.3 \pm 0.3	\dagger 42.7 \pm 0.4	34.4 \pm 0.6	\dagger 96.7 \pm 9.9	\dagger 25.9 \pm 1.1
kNN	\dagger 45.2 \pm 0.0	\dagger 42.1 \pm 0.0	\dagger 45.5 \pm 0.0	\dagger 46.4 \pm 0.0	\dagger 59.1 \pm 0.0	n.r.
GPS	\dagger 44.6 \pm 0.0	\dagger 13.3 \pm 0.0	\dagger 13.3 \pm 0.0	\dagger 1.6 \pm 0.0	\dagger 140. \pm 0.0	\dagger 20.0 \pm 0.0
CF	\dagger 48.9 \pm 0.0	\dagger 49.6 \pm 0.0	\dagger 49.6 \pm 0.0	\dagger 48.3 \pm 0.0	\dagger 108. \pm 0.0	\dagger 35.3 \pm 0.0
BART	\dagger 35.5 \pm 14.	\dagger 34.6 \pm 4.2	\dagger 44.5 \pm 1.2	n.r.	13.2 \pm 1.0	n.r.