



# A Data-Centric Decomposition of Estimator Performance in Continuous Treatment Effect Estimation

Christopher Bockel-Rickermann  
Faculty of Economics and Business  
KU Leuven  
Leuven, Belgium  
christopher.rickermann@kuleuven.be

Daan Caljon  
Faculty of Economics and Business  
KU Leuven  
Leuven, Belgium  
daan.caljon@kuleuven.be

Toon Vanderschueren  
Faculty of Economics and Business  
KU Leuven  
Leuven, Belgium  
toon.vanderschueren@kuleuven.be

Tim Verdonck  
Department of Mathematics  
Universiteit Antwerpen - imec  
Antwerp, Belgium  
tim.verdonck@uantwerpen.be

Wouter Verbeke  
Faculty of Economics and Business  
KU Leuven  
Leuven, Belgium  
wouter.verbeke@kuleuven.be

## Abstract

Standard supervised learning algorithms prioritize predictive performance over causal inference, making them ill-suited for conditional average dose response (CADR) estimation. Although specialized CADR estimators have been proposed to address this, the field lacks clarity on which dataset characteristics truly challenge model performance, largely because current benchmarking practices fail to isolate different sources of estimation error. In this work, we analyze these current benchmarking practices and introduce a novel decomposition framework that disentangles the contribution of distinct data-generating components, such as confounding, dose distribution non-uniformity, and response surface complexity, to estimator performance. Applying this scheme to several established benchmarks, we uncover that widely used datasets do not primarily test for confounding robustness, as often assumed, but are instead dominated by challenges arising from non-uniform dose distributions. We further propose a new benchmark dataset with high CADR heterogeneity, where confounding does have a substantial impact. Our results call for a rethinking of current evaluation practices and advocate for more diagnostic, data-centric benchmarks to advance the development of robust CADR estimation methods.

## CCS Concepts

• **Computing methodologies** → **Causal reasoning and diagnostics; Machine learning.**

## Keywords

Dose response estimation, causal ML, performance decomposition

### ACM Reference Format:

Christopher Bockel-Rickermann, Daan Caljon, Toon Vanderschueren, Tim Verdonck, and Wouter Verbeke. 2026. A Data-Centric Decomposition of Estimator Performance in Continuous Treatment Effect Estimation. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26)*, August 09–13, 2026, Jeju Island, Republic of Korea. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3770855.3817731>



This work is licensed under a Creative Commons Attribution 4.0 International License. *KDD '26, Jeju Island, Republic of Korea*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2259-2/2026/08

<https://doi.org/10.1145/3770855.3817731>

## Resource Availability:

The source code of this paper has been made publicly available at <https://doi.org/10.5281/zenodo.20497700>.

## 1 Introduction

Continuous treatments—interventions with tunable intensity or dosage—are common in domains such as personalized medicine [15] and targeted marketing [37]. Understanding how individuals respond to varying treatment intensities is critical for optimal decision-making; yet, estimation of causal effects with continuous treatments from observational data remains underexplored in machine learning (ML). While a wide array of methods have been proposed for estimating heterogeneous treatment effects with binary treatments [24, 31, 32, 35, 56, 57, 63, 69], ML estimators for conditional average dose responses (CADR), also referred to as “individual” [53] or “heterogeneous” [71] dose responses, are confined to only a select number of recently proposed methods [3, 33, 39, 40, 53, 64, 70, 71]. Furthermore, these methods have not yet seen widespread usage in real-world applications.

Despite significant progress in the state of the art, CADR estimation research lacks alignment, focusing on different challenges and using different benchmarking datasets. Existing datasets often conflate multiple sources of difficulty, such as confounding, response complexity, and dose distribution imbalance, making it unclear what challenges truly drive estimator performance. To date, the field has relied on a selection of (semi-) synthetic benchmarking datasets created from manually defined data-generating processes (DGPs). These datasets claim to test estimators in the presence of certain challenges, most notably “confounding”.<sup>1</sup> Yet, as our experiments show, those DGPs expose estimators to more than just one challenge, among which confounding is not necessarily the most important one.

We believe that further progress in the field requires a deeper understanding of the nature of CADR estimation and the difficulties therein. To that end, we introduce a problem formulation for CADR estimation that unifies existing research. We conceptualize DGPs along this formulation and identify five components contributing to

<sup>1</sup> Some ML papers refer to confounding as “selection bias” which also relates to the out-of-sample generalizability of response estimates [20]. To reduce ambiguity, we adopt the terminology traditionally used in the causal inference literature [1, 10, 29, 42].




**Figure 1: Selected components of our decomposition scheme. To disentangle the effects of confounding from the effects of non-uniform distributions of doses, we evaluate estimators in three scenarios: 1) When doses are randomly sampled from a uniform distribution, 2) when those distributions are not uniform, but also not specific to a certain unit, and 3) when the data is confounded, so when dose assignment is specific to a certain unit. The distribution of doses across the total population is the same in steps 2) and 3). Our complete scheme includes two additional steps related to the distribution of interventions when there are multiple intervention options (cf. Section 4).**

CADR estimator performance. To facilitate future research, we propose a novel decomposition scheme that disentangles performance along these five components, allowing researchers to understand the sources of model performance (cf. Figure 1). Our framework lays the groundwork for robust, data-centric evaluation in continuous treatment effect estimation.

Our paper makes three important contributions: (1) We introduce a unifying problem formulation for CADR estimation and conceptualize synthetic DGPs for benchmarking dataset creation; (2) We propose a scheme for decomposing model performance along the mechanisms of a DGP and specific challenges in CADR estimation; (3) We test a selection of ML estimators and decompose their performance on four of the most popular benchmark datasets for CADR estimation. We elicit strengths and weaknesses and compare their performance against traditional supervised learning methods. As such, we aim to establish a standardized approach that facilitates an effective evaluation and comparison of methods. To this end, we propose a data-centric evaluation protocol that facilitates robust and interpretable benchmarking of continuous treatment effect estimators.

*Outline.* Section 2 introduces related work on evaluating ML estimators, and previous practices in decomposing model performance. We conceptualize CADR estimation in Section 3. In Section 4, we conceptualize DGPs and introduce our novel decomposition scheme. We illustrate our approach for a prominent dataset from Bica et al. [3] in Section 5, introducing established methods and discussing the findings. We provide results on different datasets in Section 6 and conclude in Section 7.

## 2 Context

*Evaluating conditional average intervention response estimators is fundamentally challenging.* Evaluating the performance of conditional average intervention response estimators is challenging, as per unit of interest we can only observe the response to a single “factual” intervention, leaving all other potential outcomes unobserved [27]. In conditional average treatment effect (CATE) estimation, there is only one counterfactual intervention: units have received either the treatment or the control. CADR estimation introduces

further complexity due to the dimensionality of the treatment space. With potentially distinct intervention types and a practically infinite range of dosages, the full individual dose-response curve is fundamentally unobservable. In consequence, using observational data limits evaluation to measuring accuracy in predicting factual responses, as in, e.g., supervised learning [22], which fails to capture counterfactual estimation accuracy.

*Lack of established benchmarking practices.* In ML research, CADR estimators are typically evaluated using semi- or fully synthetic datasets that are specified by a certain DGP. This DGP allows the calculation of counterfactual responses of any unit to any intervention, to overcome the challenges in evaluating using observational data. In CATE estimation, researchers have relied mostly on a single dataset for evaluating estimator performance [11]. For CADR estimation, there is no such established standard, and researchers have relied on several different datasets. Moreover, the creators of these datasets typically do not clarify the challenges present that might complicate CADR estimation. We will show in our experiments that a dataset often embodies several challenges (cf. Section 4), which further hinders the understanding of model performance. We aim to alleviate this issue by proposing a novel scheme to decompose performance along different aspects of benchmarking data.

*Decomposing model performance.* It is common practice in ML research to decompose (or to “ablate”) complex model architectures by systematically adding and removing components to measure their impact on performance. In CADR estimation, such studies have, e.g., been conducted by Schwab et al. [53] and Bica et al. [3]. While such ablations inform about which part of an architecture contributes to improvements in model performance, they do not allow us to understand the challenges in a certain dataset. This understanding is critical to effective methodology selection in real-life applications. We attempt to close this gap in CADR estimator research by providing a scheme to decompose datasets, not models. Such a decomposition enables us to understand the scenarios in which an estimator might work well or fail. This data-centric decomposition of model performance is in line with calls in other ML research domains [67, 68].




Figure 2: SWIG illustrating causal dependencies in observational data

### 3 Problem formulation

*Defining CADR estimation.* We find varying definitions of CADR estimation in the literature, typically differing in the availability of only one [26, 40] or several [3, 53] distinct interventions with an associated dose. In the following, we propose a unifying framework that subsumes any of these definitions.

We leverage the Neyman-Rubin potential outcomes (PO) framework [50, 59] and expect a unit of interest  $i$  to be specified by a realization  $\mathbf{x}_i$  of random variable  $\mathbf{X} \in \mathcal{X}$  with  $\mathcal{X} \subset \mathbb{R}^m$  being the  $m$ -dimensional feature space. The unit is exposed to a single intervention  $t_i$  sampled from  $T \in \mathcal{T}$ , with the intervention space  $\mathcal{T} = \{\omega_1, \dots, \omega_k\}$  being discrete with  $k$  different intervention options. Every unit is also exposed to a continuous intensity or dose  $d_i$  sampled from  $D \in \mathcal{D}$  with  $\mathcal{D} \subset \mathbb{R}$ . For the remainder of our paper and without loss of generality, we set  $\mathcal{D} = [0, 1]$ . For any realization of intervention and dose variables, there is a potential outcome  $Y(t, d) \in \mathcal{Y} \subset \mathbb{R}$ . In line with Schwab et al. [53] we define  $Y(t, d)$  as the “dose response”.

We are interested in finding an estimate of the conditional average dose response (CADR), defined as

$$\mu(t, d, \mathbf{x}) = \mathbb{E}[Y(t, d) | \mathbf{X} = \mathbf{x}] \quad (1)$$

for every  $t \in \mathcal{T}$ ,  $d \in \mathcal{D}$  and  $\mathbf{x} \in \mathcal{X}$ , which is in line with the definition by Bica et al. [3]. A detailed overview of the notation is presented in Appendix A.

*Understanding the causal structure of dose responses.* The challenges in estimating dose responses arise from the causal relationships between variables  $\mathbf{X}$ ,  $D$ ,  $T$ , and  $Y$ , which we illustrate in a single-world intervention graph (SWIG) [47] in Figure 2. In the SWIG, capital letters denote the natural random variables and lowercase letters denote the intervened-upon values. This makes it explicit how the potential outcomes downstream of an intervention are decoupled from the natural assignment mechanism. Typically, an intervention  $t$  is chosen based on the observed realization of the covariates  $\mathbf{x}$ . Given the intervention, a dose  $d$  is assigned. The observed potential outcome is subsequently dependent on the realization of all those variables  $\mathbf{x}$ ,  $t$ , and  $d$ . To find an unbiased estimate of  $\mu(\cdot)$  we must hence use an estimator that is flexible enough to capture the relationship between outcome and intervention variables, while correctly adjusting for the effects of  $\mathbf{X}$  on all  $Y$ ,  $T$ , and  $D$  [40], but also for the effect of  $T$  on  $D$ . The simultaneous influence

of  $\mathbf{X}$  on the intervention variables and the response is referred to as “confounding” [44].

*On the identifiability of CADR.* Estimating dose responses from observational data relies on a set of assumptions [60]. Specifically, the identifiability of intervention responses requires “strong ignorability” [49], subsuming “consistency”, “no hidden confounders”, and “overlap”:

ASSUMPTION 1. (Consistency) The observed outcome  $Y_i$  for a unit  $i$  that was assigned intervention  $t_i$  and dose  $d_i$  is the potential outcome  $Y_i(t_i, d_i)$ .

ASSUMPTION 2. (No hidden confounders) The assigned intervention  $T$  and dose  $D$  are conditionally independent of the potential outcome  $Y(t, d)$  given the covariates  $\mathbf{X}$ , so  $\{Y(t, d) | t \in \mathcal{T}, d \in \mathcal{D}\} \perp (T, D) | \mathbf{X}$

ASSUMPTION 3. (Overlap) Every unit has a greater-than-zero probability of receiving any possible combination of intervention and dose, so  $\forall t \in \mathcal{T} : \forall d \in \mathcal{D} : \forall \mathbf{x} \in \mathcal{X}$  with  $\mathbb{P}(\mathbf{x}) > 0 : 0 < \mathbb{P}((t, d) | \mathbf{x}) < 1$

We assume these assumptions to hold for the remainder of our work, as most previously proposed methods require them. On top of that, different from estimating treatment effects with binary treatments [30, 43], there is little research on the impacts of violations of these assumptions when estimating dose responses.

### 4 A decomposition scheme for performance of CADR estimators

*Abstracting DGPs.* When working with observational data, the underlying natural DGP is typically partially or fully unknown, adding to the challenge of evaluating CADR estimator performance (cf. Section 2). To counter this, ML research relies on synthetic DGPs for benchmarking intervention response estimators. Those DGPs assume a causal structure (in our case the SWIG in Figure 2) and define relationships between different variables, allowing for full control over challenges in the data, such as confounding. Generally, a DGP for CADR consists of four components: (1) the definition of observed units through a covariate vector  $\mathbf{X}$ , (2) an intervention assignment function  $a_t(\cdot)$  assigning an intervention to every unit, (3) a dose assignment function  $a_d(\cdot)$  assigning a dose, and (4) an outcome function  $\mu(\cdot)$ . Each of those components influences the resulting benchmarking dataset. These functions can be defined to take variable inputs. For example, by taking as input the covariate vector  $\mathbf{X}$ ,  $a_t(\cdot)$  and  $a_d(\cdot)$  introduce confounding in the resulting data.

We provide a list of the established CADR benchmarking datasets in Appendix D along with an overview of established CADR estimators, and the datasets used for their evaluation. The most widely used datasets are TCGA-2 [3], along with IHDP-1, News-3, and Synth-1, which were introduced in Nie et al. [40]. Typically, assignment functions are non-deterministic to comply with the strong ignorability assumption. For interventions, this is done by assigning non-zero probabilities to the various possible interventions before sampling a factual intervention per unit. Doses are sampled from a distribution with a strictly positive probability mass over  $\mathcal{D}$  and a mode conditional on the covariates of a unit, e.g., from a normal distribution [40] or a beta distribution [3]. The outcome function




**Figure 3: Dose distribution for different levels of confounding in data by Bica et al. [3]**

$\mu(\cdot)$  specifies a CADR by taking as input the resulting vectors  $\mathbf{X}$ ,  $D$ , and  $T$ . Individual responses are generated by adding random noise.

*Disentangling impacts of synthetic DGPs on estimator performance.* Confounding is the presence of a covariate vector that influences intervention and dose assignment, as well as the response of a unit. Different mechanisms have been proposed to introduce confounding in a benchmarking dataset: Bica et al. [3] set a modal intervention and dose that would maximize a unit’s CADR, whereas Nie et al. [40] use some polynomial non-linear function that maps a unit’s covariates to a modal dose. The level and complexity of confounding in a (semi-) synthetic dataset may vary, depending on the confounding mechanism in the synthetic DGP. However, only some of the established DGPs allow to vary the level of confounding. Therefore, it may be unclear what exactly is driving model performance: the ability of a method to model a complex CADR, as specified by  $\mu(\cdot)$ , or its ability to adjust for confounding.

Moreover, the assignment functions  $a_t(\cdot)$  and  $a_d(\cdot)$  introduce additional challenges to the dataset on top of confounding. As a certain assignment mechanism impacts the probability of being assigned different interventions and doses per unit, it simultaneously influences their distribution across the observed population (cf. Figure 1 and a detailed discussion in Appendix B). This is most evident upon analyzing benchmarking datasets with tunable levels of confounding, such as the TCGA-2 dataset proposed by Bica et al. [3]. When the level of dose confounding is set to  $\alpha = 1$ , so no confounding, doses are uniformly distributed across  $\mathcal{D}$ . When the level of confounding increases, the distribution of doses changes (cf. Figure 3). This leads to significantly fewer observations for some dose intervals, potentially impacting the performance of ML estimators [34].

*Decomposing performance by randomizing intervention assignment.* Building estimators robust to confounding poses a key challenge in CADR estimation [3, 40]. Yet, without further investigation, performance on a dataset could be attributed to any of the above-mentioned challenges in the data. Standard ML benchmarking practices do not reveal whether the performance of an estimator is impacted by non-linearity of responses, confounding factors, or the intervention and dose distributions across the population. To overcome this limitation, we propose a novel decomposition scheme for CADR estimator performance. Our scheme is performance metric agnostic, allowing for arbitrary and situation-specific evaluation criteria.

We consider two boundary scenarios: In the “randomized” scenario, interventions and doses are completely randomized and sampled uniformly. In the “non-randomized” scenario, the data aligns with its creators’ specifications, where interventions and doses adhere to assignment functions  $a_t(\cdot)$  and  $a_d(\cdot)$ . We then explore three intermediary scenarios that progressively move from the “randomized” to the “non-randomized” setup.

First, we generate scenario “ $t$  non-uniformity”. Per unit, an intervention is sampled from a joint distribution that is equivalent to the population distribution of interventions in the “non-randomized” scenario. This approach maintains the effects of  $a_t(\cdot)$  on the distribution of interventions across the total population, yet removes confounding, as  $\mathbb{P}(t|\mathbf{x}) = \mathbb{P}(t)$  for all  $t \in \mathcal{T}$  and  $\mathbf{x} \in \mathcal{X}$ . The doses are sampled at random. This scenario informs about the impact of population-level changes in intervention distributions.

Second, we generate scenario “ $t$  confounding”, in which we operate under random dose assignment, but confounded interventions as generated by  $a_t$ . This allows us to isolate the effects of intervention confounding from distributional effects investigated previously.

Third, and in addition to intervention confounding, we repeat the process for dose assignment, generating scenario “ $d$  non-uniformity”, in which  $\mathbb{P}(d|\mathbf{x}) = \mathbb{P}(d)$  for all  $d \in \mathcal{D}$  and  $\mathbf{x} \in \mathcal{X}$ .

The final scenario adds dose confounding, by generating doses according to  $a_d$ , yielding the “ $d$  confounding” or “non-randomized” scenario, which aligns with the original specifications of the data.

To be in line with the causal dependencies in observational data as outlined in Section 3, we chose to first decompose along the treatment assignment, yet our scheme is flexible and would allow for alterations. Importantly, our framework guarantees a *complete disentanglement between non-uniformity and confounding*: for the non-uniformity scenarios, we first generate the confounded intervention and dose assignments and then randomly re-assign them across units. This effectively decouples the covariates from the assignment mechanism through randomization while preserving the population-level distribution of interventions and doses. The full decomposition scheme is summarized in Algorithm 1. We refer to the Online Appendix<sup>2</sup> for the technical implementation. We summarize the resulting five scenarios below:

- (1) (*Randomized*) Random interventions and doses (sampled from uniform distributions)
- (2) ( *$t$  non-uniformity*) Non-uniformly distributed interventions and random doses
- (3) ( *$t$  confounding*) Confounded interventions and random doses
- (4) ( *$d$  non-uniformity*) Confounded interventions and non-uniformly distributed doses
- (5) (*Non-randomized /  $d$  confounding*) Confounded interventions and confounded doses

By iteratively changing key characteristics of the data, our decomposition scheme represents an experimental design to test for the effects of various contributing factors on CADR estimator performance. As such, we attempt to further bridge the gap between benchmarking practices in ML and experimental study design in the wider field of causal inference [51, 55].

<sup>2</sup><https://github.com/christopher-br/CADR-performance-deco>

## 5 Case study: decomposing performance on the TCGA-2 dataset

### 5.1 Experimental setup

*Dataset.* To demonstrate the workings of our decomposition scheme, we apply it to the TCGA-2 dataset proposed by Bica et al. [3].<sup>3</sup> Several other benchmarking datasets for CADR estimation have been proposed in previous studies, which we enlist in Appendix D. We selected the TCGA-2 dataset for several reasons. First, it comprises the assignment of one of three distinct interventions per unit and an associated dose, whereas most other datasets consider a single intervention. Second, the covariate matrix used in the TCGA data [5] is frequently used in other research. Third, the assignment mechanisms used in the DGP find use in several succeeding papers, such as in Nie et al. [40], Schweisthal et al. [54], and Vanderschueren et al. [62]. For technical details, we refer to the original paper [3].

*ML methods for CADR estimation.* We decompose the performance of selected established CADR estimators, as well as several supervised estimators, to provide a broad set of baseline methods. Hereby, we follow calls from related ML research to test novel algorithms against a wider array of methods [45]. Much of recent ML research has been devoted to developing neural network architectures to tackle analytical problems, with many papers ignoring classical methods such as regression or tree-based approaches. First, we apply three prominent ML estimators for CADR estimation, namely DRNet [53], SCIGAN [3], and VCNet [40]. Each of these is a neural network [17], providing several favorable characteristics to modeling data [28]. The estimators have been the first tailored ML methods for dose response estimation, and have been used as benchmark methodologies for several later-proposed methods. Each method uniquely tackles CADR estimation: SCIGAN uses generative adversarial networks (GANs) [18] to generate additional counterfactual outcomes per observed unit and effectively randomize intervention and dose assignment, as such removing confounding. DRNet trains separate models on a shared representation learner per combination of dose intervals and interventions, to reinforce the influence of the intervention variables in the model. VCNet follows a similar motivation, yet leverages a varying-coefficient architecture [21] to accomplish this. Appendix C provides a detailed description of the three methods. A complete overview of ML dose response estimators is provided in Appendix D. We further apply five supervised learning methods, which have not been sufficiently benchmarked in prior work (for a list of benchmark methodologies per established ML dose response estimator see again Appendix D). We apply a linear regression model, a regression tree [4], a generalized additive model (GAM) [66], xgboost [7] as a state-of-the-art implementation of a gradient-boosted decision tree [16], and a simple feed-forward multilayer perception (MLP). Inspired by kernel-based causal function estimators, we also include a kernel ridge regression with a product kernel design, i.e., multiplicatively combining separate kernels over covariates, intervention, and dose [46, 58]. In comparing ML CADR estimators with traditional supervised learning methods we aim to understand both the complexity

of benchmarking datasets and differences in the performance of methods concerning the challenges in dose response estimation. This facilitates practitioners and researchers in making a conscious tradeoff between interpretability and performance of estimators [2], as some of the targeted estimators introduce significant complexity over established techniques.

---

#### Algorithm 1: Performance decomposition

---

**Require:** Covariate matrix  $X$ , Intervention assignment function

$a_t(\cdot)$ , Dose assignment function  $a_d(\cdot)$ , Outcome function  $\mu(\cdot)$ , Machine learning method  $M$ , Performance metric  $P(\cdot)$

**Result:** Decomposed performance of  $M$  on DGP specified by  $X$ ,  $a_t$ ,  $a_d$  and  $\mu$

---

```
# 0) initialization
(1) sample  $ids_{train}, ids_{test}$ 
(2)  $T_{rand} \leftarrow$  sample at random from  $\{\omega_1, \dots, \omega_k\}$  # Get
    intervention vectors
(3)  $T_{conf} \leftarrow a_t(X)$ 
(4)  $T_{non-u} \leftarrow$  shuffle  $T_{conf}$ 
(5)  $D_{rand} \leftarrow$  sample at random from  $[0, 1]$  # Get dose
    vectors
(6)  $D_{conf} \leftarrow a_d(X, T_{conf})$ 
(7)  $D_{non-u} \leftarrow$  shuffle( $D_{conf}$ )

# 1) eval under random interventions and doses
(8)  $Y \leftarrow \mu(X, T_{rand}, D_{rand})$  # Get outcomes
(9) train  $M$  on  $(Y, X, T_{rand}, D_{rand})$  using  $ids_{train}$ 
(10) calculate  $P(M, Y, X, T_{rand}, D_{rand})$  using  $ids_{test}$ 
    # Calculate performance

# 2) eval under intervention non-uniformity
    and random doses
(11)  $Y \leftarrow \mu(X, T_{rand}, D_{rand})$  # Get outcomes
(12) train  $M$  on  $(Y, X, T_{non-u}, D_{rand})$  using  $ids_{train}$ 
(13) calculate  $P(M, Y, X, T_{non-u}, D_{rand})$  using  $ids_{test}$ 
    # Calculate performance

# 3) eval under intervention confounding and
    random doses
(14)  $Y \leftarrow \mu(X, T_{conf}, D_{rand})$  # Get outcomes
(15) train  $M$  on  $(Y, X, T_{conf}, D_{rand})$  using  $ids_{train}$ 
(16) calculate  $P(M, Y, X, T_{conf}, D_{rand})$  using  $ids_{test}$ 
    # Calculate performance

# 4) eval under intervention confounding and
    dose non-uniformity
(17)  $Y \leftarrow \mu(X, T_{conf}, D_{non-u})$  # Get outcomes
(18) train  $M$  on  $(Y, X, T_{conf}, D_{non-u})$  using  $ids_{train}$ 
(19) calculate  $P(M, Y, X, T_{conf}, D_{non-u})$  using  $ids_{test}$ 
    # Calculate performance

# 5) evaluate under intervention confounding
    and dose confounding
(20)  $Y \leftarrow \mu(X, T_{conf}, D_{non-u})$  # Get outcomes
(21) train  $M$  on  $(Y, X, T_{conf}, D_{conf})$  using  $ids_{train}$ 
(22) calculate  $P(M, Y, X, T_{conf}, D_{conf})$  using  $ids_{test}$ 
    # Calculate performance
```

---

<sup>3</sup>The level of intervention and dose confounding can be varied in the TCGA-2 dataset. We opt for the standard values proposed in the original paper [3] with the level of intervention confounding set to  $\kappa = 2$  and the level of dose confounding set to  $\alpha = 2$ .

**Table 1: Performance decomposition on TCGA-2 dataset [3].**

| Method       | Scenario               |                           |                        |                           |                        |
|--------------|------------------------|---------------------------|------------------------|---------------------------|------------------------|
|              | random.                | $\rightarrow t$ non-unif. | $\rightarrow t$ conf.  | $\rightarrow d$ non-unif. | $\rightarrow d$ conf.  |
| Lin. reg.    | 4.74 $\pm$ 0.02        | 4.84 $\pm$ 0.03           | 4.89 $\pm$ 0.02        | 5.41 $\pm$ 0.03           | 5.41 $\pm$ 0.03        |
| Reg. tree    | 0.40 $\pm$ 0.01        | 0.39 $\pm$ 0.01           | 0.39 $\pm$ 0.01        | <b>0.55</b> $\pm$ 0.04    | <b>0.56</b> $\pm$ 0.04 |
| GAM          | 3.17 $\pm$ 0.02        | 3.36 $\pm$ 0.04           | 3.30 $\pm$ 0.03        | 3.77 $\pm$ 0.23           | 3.77 $\pm$ 0.22        |
| Kernel ridge | 1.46 $\pm$ 0.02        | 1.39 $\pm$ 0.02           | 1.38 $\pm$ 0.01        | 1.34 $\pm$ 0.06           | 1.34 $\pm$ 0.06        |
| xgboost      | 0.98 $\pm$ 0.07        | 0.92 $\pm$ 0.10           | 0.88 $\pm$ 0.09        | 1.14 $\pm$ 0.07           | 1.13 $\pm$ 0.06        |
| MLP          | 3.14 $\pm$ 0.04        | 3.22 $\pm$ 0.07           | 3.20 $\pm$ 0.07        | 5.33 $\pm$ 0.14           | 5.30 $\pm$ 0.17        |
| SCIGAN       | 3.05 $\pm$ 1.17        | 2.34 $\pm$ 1.84           | 1.72 $\pm$ 0.48        | 4.40 $\pm$ 4.58           | 2.08 $\pm$ 0.77        |
| DRNet        | 0.97 $\pm$ 0.03        | 1.00 $\pm$ 0.04           | 1.02 $\pm$ 0.04        | 1.17 $\pm$ 0.05           | 1.16 $\pm$ 0.03        |
| VCNet        | <b>0.29</b> $\pm$ 0.03 | <b>0.38</b> $\pm$ 0.02    | <b>0.33</b> $\pm$ 0.02 | 0.60 $\pm$ 0.13           | 0.60 $\pm$ 0.13        |

random.: randomized; non-unif.: non-uniformity; conf.: confounding

**Performance evaluation.** Our decomposition scheme is agnostic to the selection of a performance metric and could be used across scenarios and use cases. Next to CADR estimators, it could similarly be used to evaluate “average” dose response estimators. For evaluating CADR estimator performance, Schwab et al. [53] propose three performance metrics: the “policy error” and “dose policy error”, which evaluate the capability of an estimator to identify the most effective interventions and doses by the magnitude of the response, as well as the “mean integrated squared error” (MISE), which evaluates the mean accuracy of the CADR estimates over the intervention and dose spaces. Since the MISE makes minimal assumptions about the domain of application or use of a CADR estimator, we adopt it in the experiments reported in this paper. For a number of test units  $N$ , the true CADR  $\mu(\cdot)$ , and the estimated CADR  $\hat{\mu}(\cdot)$  we calculate the MISE as

$$MISE = \frac{1}{N} \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \sum_{i=1}^N \int_{d \in \mathcal{D}} (\mu(t, d, \mathbf{x}_i) - \hat{\mu}(t, d, \mathbf{x}_i))^2 dd \quad (2)$$

In practice, the inner integral is approximated using Romberg integration with 64 equally spaced samples as in Schwab et al. [53]. We use 20% of the observations in the benchmarking dataset as a holdout test set to calculate the MISE. For completeness, we also show the results for Policy Error in the Online Appendix.

**Model selection.** Model selection in causal inference is challenging as the ground truth can never be observed in practice [12, 48, 52]. Several methodologies for tuning hyperparameters on observational data have been proposed, both stand-alone and accompanying an estimator. The choice of a selection procedure can have a large influence on model performance. To ensure each method performs (near-)optimally, models are selected based on the mean squared error in predicting the factual outcomes (factual selection criterion [12]) on a validation set (10% of the units in the covariate matrix), the remaining 70% of observations are used for training.

## 5.2 Results

**Insights into the dataset.** Upon analyzing the decomposed performance of all estimators as presented in Table 1, we conclude that the TCGA-2 dataset is challenging due to dose non-uniformity, rather than dose or intervention confounding, as it has the largest effect on performance. The non-uniformity of interventions only has a small detrimental effect on model performance. For this dataset, confounding seems to have very little effect on estimation performance, both for CADR estimators as well as supervised learning




**Figure 4: Distribution of errors per intervention and dose interval in the test set of the TCGA-2 dataset estimated by an MLP. A histogram of doses in the training set is added per plot in blue. Errors are correlated with dose non-uniformity, supporting that non-uniformity affects model performance.**

methods. The largest effect on model performance results from introducing dose non-uniformity, i.e., making some doses less likely to be observed across the population. This increases the MISE for all methods significantly when compared with scenario “ $t$  confounding”, in which doses are sampled from uniform distributions. The introduction of dose confounding in scenario “ $d$  confounding” again does not appear to significantly impact performance. This is surprising, given the proposition of the dataset to test methods for robustness against confounding biases. Further, the non-uniformity of doses and interventions is not a causal issue, but rather related to ML challenges outside of causal inference and treatment effect estimation, such as learning from imbalanced datasets [19, 23]. We also observe this behavior in analyzing performance for other datasets, as discussed in Section 6.

**Insights into model performance.** Comparing performance across the different estimators allows us to draw further conclusions. Estimators are only little affected by intervention non-uniformity, and not affected by their confounding. Conversely, performance might even be improved by confounding. This is counter-intuitive to the reasoning that confounding might adversely affect performance [3, 53]. A possible explanation for this result is that, under confounding, units have a higher probability of being assigned exactly those doses for which CADR heterogeneity is greatest. When comparing neural architectures, CADR estimators outperform the standard MLP. However, the best-performing model is a simple regression tree. We attribute these results to the overall low heterogeneity of the dose response space in the dataset (cf. Section 6).

**Understanding sources of performance gain.** The presented experimental results indicate that a decomposition of performance is imperative to evaluate the capabilities of CADR estimators. The observation that model performance does not degrade due to confounding, but due to non-uniform distributions of interventions




**Figure 5: MISE per method and dataset. Across datasets, confounding has little adverse effects on model performance. Full results including std. errors can be found in the Online Appendix.**

and doses indicates that the TCGA-2 dataset is not evaluating estimators for their robustness to confounding, but rather efficiency in learning from imbalanced or limited amounts of training data [14, 65]. This is surprising, as several ML papers use it to test estimators for robustness to confounding [3, 33, 64]. To confirm this hypothesis, we visualize the errors in CADR estimation made by an MLP per dose interval and intervention in Figure 4, next to a histogram plot of doses in the training data. The plots show that errors in predicting CADR increase with decreasing training observations for a specific dose. This is most notable for interventions  $\omega_2$  and  $\omega_3$ .

## 6 Insights from other datasets

*Datasets proposed by Nie et al. [40].* Other prominently used datasets for benchmarking CADR estimators were proposed by Nie et al. [40] (IHDP-1, News-3, and Synth-1). We decompose the performance of all methods from Section 5 on all three of these datasets and present results in Figure 5. Compared to the TCGA-2 dataset, the datasets only involve a single intervention, so our decomposition does not include Scenarios 2 and 3 ( $t$  non-uniformity and  $t$  confounding). The datasets are confounded as both the observed doses and outcomes are conditional on the covariate matrices. Yet, as in the TCGA-2 dataset, this confounding seems to have no additional effect on model performance. Moreover, also dose non-uniformity does not affect the models. We provide a population-level distribution of doses per dataset in Figure 6, which shows that distributions are less skewed compared to TCGA-2. This explains why supervised learning techniques, especially xgboost, are performing competitively on these datasets. Therefore, experiments using these data sets only enable limited insight into a method’s ability to tackle challenges inherent to CADR estimation.

*Decomposing performance under high CADR heterogeneity.* Working with synthetic datasets also allows visualizing the dose responses of individual units in the data (see Figure 6). All datasets discussed previously show little heterogeneity in the dose responses across different units. This might be a reason why supervised learning methods perform competitively against CADR estimators. We test this hypothesis by proposing a new benchmarking dataset, the “IHDP-3” dataset, which is distinct from the other datasets discussed in this study.

*IHDP-3 dataset.* Our novel dataset leverages the same covariate matrix as the IHDP-1 dataset [40], notably the covariates used in Hill

[24], but the assignment of intervention variables, as well as the outcome calculation, differ significantly. We consider a scenario with only one distinct intervention. Yet, different from previous datasets, the response to this intervention has higher heterogeneity in the covariates of a unit. Specifically, there are four different archetypes of units that respond distinctly to an intervention, e.g., units of one archetype might respond positively to an increased dose, while units from another archetype might respond negatively. Archetype assignment is not known ex-ante, but deterministic to comply with strong ignorability. We assign every unit in the covariate matrix to one of these archetypes based on their realization of binary covariates `b.marr` and `mom.lths`. The CADR per archetype is given in Table 2. Per unit, we generate the individual responses by adding normally distributed random noise.

**Table 2: CADR per archetype of units in the IHDP-3 dataset**

| AT | Dose response curve                                                                                                                           |
|----|-----------------------------------------------------------------------------------------------------------------------------------------------|
| A1 | $f_1(\mathbf{x}_i, d) = 10 * (\mathbf{x}_{i,0} + 12 * d * (d - \frac{3}{4} * (\mathbf{x}_{i,1} + \mathbf{x}_{i,2}))^2)$                       |
| A2 | $f_2(\mathbf{x}_i, d) = 10 * (\mathbf{x}_{i,1} + \sin(\pi * (\mathbf{x}_{i,2} + \mathbf{x}_{i,3}) * d))$                                      |
| A3 | $f_3(\mathbf{x}_i, d) = 10 * (\mathbf{x}_{i,2} + 12 * (\mathbf{x}_{i,3} * d - \mathbf{x}_{i,4} * d^2))$                                       |
| A4 | $f_4(\mathbf{x}_i, d) = \mathbf{x}_{i,0} * 3 * \sin(20 * \mathbf{x}_{i,2} * d) + 20 * \mathbf{x}_{i,3} * d - 20 * \mathbf{x}_{i,4} * d^2 + 5$ |

AT: Archetype,  $\mathbf{x}_{i,j}$ : Variable  $j$  in the covariate vector of unit  $i$

Next, we assume that doses are assigned to every observation based on their archetype, and assign every of the archetypes a modal dose in  $\{\frac{1}{8}, \frac{3}{8}, \frac{5}{8}, \frac{7}{8}\}$ . We then sample for every unit a factual dose from a beta distribution with the respective mode and tunable variance, following the approach first introduced by Bica et al. [3]. For a detailed DGP and the technical implementation, we refer to our source code in the Online Appendix.

The resulting data has a significantly higher heterogeneity in the dose responses across units as we visualize in Figure 6. The heterogeneity in CADR varies along the datasets. For the previously established ones, per intervention, responses are generated from a single, but differently parameterized function. This yields little heterogeneity in the CADR across units and might explain the good performance of supervised learning methods on these datasets. For IHDP-3, we can see higher degrees of heterogeneity with CADR depending on the archetype. In Figure 7, we visualize confounding in the datasets by generating t-SNE plots [25, 36] of their covariate spaces and color coding observations by their factual dose. Under confounding, we expect that units with different covariates would be assigned different factual doses. Yet, not all plots indicate this




**Figure 6: Dose response space in different datasets.** Per unit in the dataset, we visualize the CADR over different interventions and doses (top plots). Accompanying, we provide a histogram of the assigned doses in the data (bottom plots). We mark the conditional response to the factual dose with a dot. The factual dose also determines the color per curve, for which the bottom plot provides a mapping. For the clarification of parameters we refer to Table 5 and the sources therein.

behavior. Especially in the News-3 and IHDP-1 datasets, units with different factual doses cannot be separated.

We present the decomposed performance of CADR estimators on the IHDP-3 dataset in Table 3. Compared to the other datasets discussed in our study, we see a significant adverse effect of dose confounding on model performance, which is evident across all estimators. Similarly, results reveal the importance of model architecture, given the estimation problem. Training individual networks per dose strata, as implemented in DRNet, leads to errors that are only comparable to linear regression, indicating that the neural




**Figure 7: t-SNE plot of covariate space per dataset.** Color corresponds to the assigned dose. We see that in the IHDP-1 and News-3 dataset observations with different assigned doses cannot clearly be separated in the t-SNE plot, indicating that the confounding in the data might not be severe. In IHDP-3, TCGA-2, and Synth-1 the separation is clearer, which might indicate a stronger effect on model performance. For the clarification of parameters we refer to Table 5 and the sources therein.

architecture is not capable of handling high degrees of heterogeneity in CADR across units. Some neural architectures, especially VCNet and the standard MLP, outperform other methods in both the randomized scenario and under dose non-uniformity. The results further indicate that VCNet’s varying coefficient structure aids successful confounding adjustment.

These findings demonstrate that even widely used CADR benchmarks differ substantially in which challenges they pose. While IHDP-3 introduces heterogeneity and confounding, other datasets primarily test robustness to dose sparsity or non-linearity. This underscores the value of systematic, component-level decomposition of dataset characteristics, as facilitated by our proposed framework.

**Table 3: Performance decomposition on IHDP-3 dataset**

| Method       | Scenario           |                    |                    |
|--------------|--------------------|--------------------|--------------------|
|              | random.            | → $d$ non-unif. →  | $d$ conf.          |
| Lin. reg.    | 14.91 ± 0.21       | 14.79 ± 0.19       | 17.86 ± 0.52       |
| Reg. tree    | 6.92 ± 0.77        | 8.89 ± 2.23        | 10.82 ± 4.19       |
| GAM          | 15.22 ± 0.29       | 15.16 ± 0.32       | 17.30 ± 0.37       |
| Kernel ridge | 10.12 ± 0.22       | 10.85 ± 0.29       | 14.02 ± 0.35       |
| xgboost      | 7.06 ± 0.69        | 8.01 ± 0.57        | 10.92 ± 1.02       |
| MLP          | 3.47 ± 0.23        | 3.97 ± 0.44        | 10.14 ± 0.45       |
| SCIGAN       | 10.40 ± 2.24       | 12.33 ± 3.57       | 14.65 ± 6.24       |
| DRNet        | 14.92 ± 0.21       | 15.08 ± 0.26       | 16.17 ± 0.16       |
| VCNet        | <b>2.68</b> ± 0.35 | <b>3.76</b> ± 0.79 | <b>8.45</b> ± 0.76 |

random.: randomized; non-unif.: non-uniformity; conf.: confounding

## 7 Conclusions and future research

This paper aims to understand the nature of conditional average dose response (CADR) estimation and to provide tools to decompose estimator performance along challenges inherent to the field. We have provided a unifying problem formulation using a single-world intervention graph (SWIG) and conceptualized synthetic data-generating processes (DGPs) typically used to evaluate estimator performance. We have proposed a novel decomposition scheme that allows us to attribute estimator performance to five challenges: the non-linearity of response surfaces, confounding of interventions and doses, and their non-uniform distributions. From our experiments, we conclude that the inherent challenges are still poorly understood. Established benchmarks are challenging predominantly due to non-uniform distributions of interventions and doses, and non-linear response surfaces. In contrast, confounding seems to have little effect on model performance on these datasets. Additionally, by proposing a novel DGP and benchmarking dataset (IHDP-3), we show that confounding is a challenge for estimators only when the heterogeneity of CADR is high.

Our results show important limitations in existing ML research on CADR estimation and suggest that more research is needed to develop benchmarks that accurately test the capabilities of estimators. We hence encourage researchers to adopt our decomposition for any future research. Additionally, we provide three further takeaways for researchers and practitioners:

**(1) Confounding can materialize in various ways.** There is no clear-cut definition of how confounding might materialize in a DGP. Our experiments show that the confounding in most previously established datasets does not pose a challenge. Further research is needed to understand and quantify when tailored methods are needed.

**(2) Several challenges exist in CADR estimation.** In our experiments, the biggest impact on model performance was introduced by dose non-uniformity, not through confounding. This contrasts with claims from established datasets and is a novel finding, especially as dose non-uniformity is not a causal problem. A potential future research direction might therefore be the adoption of methodologies such as data-efficient ML [38, 41].

**(3) Supervised estimators might be appropriate to model CADR.** Finally, our results reveal that standard supervised learning methods might achieve competitive performance in estimating CADR. This supports takeaway (1) and calls for comparing any future CADR estimators against a complete set of established

benchmarking methods, such as gradient-boosted trees. Improving transparency might help practitioners gain trust in ML CADR estimators and aid future adoption.

In addition to guiding model development, our decomposition scheme can inform dataset construction. Researchers designing new benchmarks can use the scheme to ensure a balanced coverage of distinct estimation challenges. Likewise, practitioners can use it to diagnose which estimators align with the data characteristics in their application domain, providing a more principled basis for model selection than aggregate performance alone.

## Acknowledgments

This work was supported by the Research Foundation – Flanders (FWO) research projects G015020N, G038723N, and 1117322N.

## References

- [1] Joshua D. Angrist and Jörn-Steffen Pischke. 2009. *Mostly harmless econometrics*. Princeton Univ. Press, Princeton, NJ [u.a.]. Includes bibliographical references and index.
- [2] Andrew Bell, Ian Solano-Kamaiko, Oded Nov, and Julia Stoyanovich. 2022. It's Just Not That Simple: An Empirical Study of the Accuracy-Explainability Trade-off in Machine Learning for Public Policy. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 248–266. doi:10.1145/3531146.3533090
- [3] Ioana Bica, James Jordon, and Mihaela van der Schaar. 2020. Estimating the Effects of Continuous-valued Interventions using Generative Adversarial Networks. *Advances in Neural Information Processing Systems (2020)* (Feb. 2020). arXiv:2002.12326 [cs.LG] doi:10.48550/ARXIV.2002.12326
- [4] Leo Breiman, Jerome Friedman, R. A. Olshen, and Charles J. Stone. 2017. *Classification and Regression Trees*. Chapman and Hall/CRC, New York. doi:10.1201/9781315139470
- [5] Cancer Genome Atlas Research Network, John N. Weinstein, Eric A. Collisson, Gordon B. Mills, Kenna R. Mills Shaw, Brad A. Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, and Joshua M. Stuart. 2013. The Cancer Genome Atlas Pan-Cancer analysis project. *Nature Genetics* 45, 10 (Oct. 2013), 1113–1120. doi:10.1038/ng.2764
- [6] Rich Caruana. 1997. Multitask Learning. *Machine Learning* 28, 1 (1997), 41–75. doi:10.1023/a:1007379606734
- [7] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. arXiv:1603.02754 [cs.LG] doi:10.1145/2939672.2939785
- [8] Hugh A. Chipman, Edward I. George, and Robert E. McCulloch. 2010. BART: Bayesian additive regression trees. *The Annals of Applied Statistics* 4, 1 (March 2010), 266–298. doi:10.1214/09-AOAS285
- [9] Pádraig Cunningham and Sarah Jane Delany. 2021. k-Nearest Neighbour Classifiers - A Tutorial. *Comput. Surveys* 54, 6 (July 2021), 128:1–128:25. doi:10.1145/3459665
- [10] Scott Cunningham. 2021. *Causal inference*. Yale University Press, New Haven. Literaturverzeichnis: Seiten 541–553.
- [11] Alicia Curth, David Svensson, Jim Weatherall, and Mihaela van der Schaar. 2021. Really Doing Great at Estimating CATE? A Critical Look at ML Benchmarking Practices in Treatment Effect Estimation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. https://openreview.net/forum?id=FQLzQqGEAH
- [12] Alicia Curth and Mihaela van der Schaar. 2023. In Search of Insights, Not Magic Bullets: Towards Demystification of the Model Selection Dilemma in Heterogeneous Treatment Effect Estimation. arXiv:2302.02923 [stat.ML] doi:10.48550/ARXIV.2302.02923
- [13] Christian Fong, Chad Hazlett, and Kosuke Imai. 2018. Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements. *The Annals of Applied Statistics* 12, 1 (March 2018). doi:10.1214/17-aoas1101
- [14] George Forman and Ira Cohen. 2004. Learning from Little: Comparison of Classifiers Given Little Training. In *Knowledge Discovery in Databases: PKDD 2004*, Jean-François Boulicaut, Floriana Esposito, Fosca Giannotti, and Dino Pedreschi (Eds.). Springer, Berlin, Heidelberg, 161–172. doi:10.1007/978-3-540-30116-5\_17
- [15] Emil Frei and George P. Canellos. 1980. Dose: A critical factor in cancer chemotherapy. *The American Journal of Medicine* 69, 4 (Oct. 1980), 585–594. doi:10.1016/0002-9343(80)90472-6
- [16] Jerome H. Friedman. 2001. Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics* 29, 5 (2001), 1189–1232. https://www.jstor.org/stable/2699986

- [17] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (Oct. 2020), 139–144. doi:10.1145/3422622
- [19] Guo Haixiang, Li Yijing, Jennifer Shang, Gu Mingyun, Huang Yuan Yue, and Gong Bing. 2017. Learning from class-imbalanced data: Review of methods and applications. *Expert Systems with Applications* 73 (May 2017), 220–239. doi:10.1016/j.eswa.2016.12.035
- [20] Sebastian Haneuse. 2016. Distinguishing Selection Bias and Confounding Bias in Comparative Effectiveness Research. *Medical Care* 54, 4 (April 2016), e23–29. doi:10.1097/MLR.0000000000000011
- [21] Trevor Hastie and Robert Tibshirani. 1993. Varying-Coefficient Models. *Journal of the Royal Statistical Society: Series B (Methodological)* 55, 4 (Sept. 1993), 757–779. doi:10.1111/j.2517-6161.1993.tb01939.x
- [22] Trevor Hastie, Robert Tibshirani, and Jerome H. Friedman. 2017. *The elements of statistical learning* (second edition ed.). Springer, New York, NY. Description based on publisher supplied metadata and other sources..
- [23] Haibo He and Edwardo A. Garcia. 2009. Learning from Imbalanced Data. *IEEE Transactions on Knowledge and Data Engineering* 21, 9 (Sept. 2009), 1263–1284. doi:10.1109/TKDE.2008.239
- [24] Jennifer L. Hill. 2011. Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics* 20, 1 (Jan. 2011), 217–240. doi:10.1198/jcgs.2010.08162
- [25] Geoffrey E Hinton and Sam Roweis. 2002. Stochastic Neighbor Embedding. In *Advances in Neural Information Processing Systems*, Vol. 15. MIT Press.
- [26] Keisuke Hirano and Guido W. Imbens. 2004. The Propensity Score with Continuous Treatments. 73–84 pages. doi:10.1002/0470090456.ch7
- [27] Paul W. Holland. 1986. Statistics and Causal Inference. *J. Amer. Statist. Assoc.* 81, 396 (Dec. 1986), 945–960. doi:10.1080/01621459.1986.10478354
- [28] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. 1989. Multilayer feed-forward networks are universal approximators. *Neural Networks* 2, 5 (Jan. 1989), 359–366. doi:10.1016/0893-6080(89)90020-8
- [29] Guido Imbens and Donald Rubin. 2015. *Causal inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press, New York, NY.
- [30] Guido W. Imbens. 2004. Nonparametric Estimation of Average Treatment Effects Under Exogeneity: A Review. *Review of Economics and Statistics* 86, 1 (Feb. 2004), 4–29. doi:10.1162/003465304323023651
- [31] Fredrik D. Johansson, Uri Shalit, Nathan Kallus, and David Sontag. 2020. Generalization Bounds and Representation Learning for Estimation of Potential Outcomes and Causal Effects. arXiv:2001.07426 [cs.LG] doi:10.48550/ARXIV.2001.07426
- [32] Fredrik D. Johansson, Uri Shalit, and David Sontag. 2016. Learning Representations for Counterfactual Inference. arXiv:1605.03661 [stat.ML] doi:10.48550/ARXIV.1605.03661
- [33] Amirreza Kazemi and Martin Ester. 2024. Adversarially Balanced Representation for Continuous Treatment Effect Estimation. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 12 (March 2024), 13085–13093. doi:10.1609/aaai.v38i12.29207
- [34] Peter Kokol, Marko Kokol, and Sašo Zagoranski. 2022. Machine learning on small size samples: A synthetic knowledge synthesis. *Science Progress* 105, 1 (Jan. 2022), 00368504211029777. doi:10.1177/00368504211029777
- [35] Christos Louizos, Uri Shalit, Joris Mooij, David Sontag, Richard Zemel, and Max Welling. 2017. Causal Effect Inference with Deep Latent-Variable Models. arXiv:1705.08821 [stat.ML] doi:10.48550/ARXIV.1705.08821
- [36] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9, 86 (2008), 2579–2605. <http://jmlr.org/papers/v9/vandermaten08a.html>
- [37] Alex Miller and Kartik Hosanagar. 2020. Personalized discount targeting with causal machine learning. [https://web.archive.org/web/20220803140657id\\_/https://aisel.aisnet.org/cgi/viewcontent.cgi?article=1154&context=icis2020](https://web.archive.org/web/20220803140657id_/https://aisel.aisnet.org/cgi/viewcontent.cgi?article=1154&context=icis2020)
- [38] Baharan Mirzazoleiman, Jeff Bilmes, and Jure Leskovec. 2020. Coresets for Data-efficient Training of Machine Learning Models. 6950–6960 pages. <https://proceedings.mlr.press/v119/mirzazoleiman20a.html>
- [39] Lokesh Nagalapatti, Akshay Iyer, Abir De, and Sunita Sarawagi. 2024. Continuous Treatment Effect Estimation Using Gradient Interpolation and Kernel Smoothing. arXiv:2401.15447 [cs.LG] doi:10.48550/ARXIV.2401.15447
- [40] Lizhen Nie, Mao Ye, Dan Nicolae, et al. 2021. VCNet and Functional Targeted Regularization For Learning Causal Effects of Continuous Treatments. In *International Conference on Learning Representations*.
- [41] Matthew Olson, Abraham Wyner, and Richard Berk. 2018. Modern Neural Networks Generalize on Small Data Sets. In *Advances in Neural Information Processing Systems*, Vol. 31. Curran Associates, Inc.
- [42] Judea Pearl. 2022. *Causality* (second edition, reprinted with corrections ed.). Cambridge University Press, Cambridge.
- [43] Maya L Petersen, Kristin E Porter, Susan Gruber, Yue Wang, and Mark J van der Laan. 2012. Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research* 21, 1 (Feb. 2012), 31–54. doi:10.1177/0962280210386207
- [44] Miquel S. Porta (Ed.). 2014. *A dictionary of epidemiology* (sixth edition ed.). Oxford University Press, Oxford. Includes bibliographical references (pages 301-343). - Text in English. - Print version record.
- [45] Zhen Qin, Le Yan, Honglei Zhuang, Yi Tay, Rama Kumar Pasumarthi, Xuanhui Wang, Michael Bendersky, and Marc Najork. 2020. Are Neural Rankers still Outperformed by Gradient Boosted Decision Trees? [https://openreview.net/forum?id=UtiVf\\_q\\_vC](https://openreview.net/forum?id=UtiVf_q_vC)
- [46] Yordan P Raykov, Hengrui Luo, Justin D Strait, and Wasiur R KhudaBukhs. 2025. Kernel-based estimators for functional causal effects. *arXiv preprint arXiv:2503.05024* (2025).
- [47] Thomas S. Richardson and James M. Robins. 2013. Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper* 128, 30 (2013), 2013.
- [48] Craig A. Rolling and Yuhong Yang. 2013. Model Selection for Estimating Treatment Effects. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 76, 4 (Nov. 2013), 749–769. doi:10.1111/rssb.12043
- [49] Paul Rosenbaum and Donald Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, 1 (1983), 41–55. doi:10.1093/biomet/70.1.41
- [50] Donald B. Rubin. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66, 5 (oct 1974), 688–701. doi:10.1037/h0037350
- [51] Donald B. Rubin. 2008. For objective causal inference, design trumps analysis. *The Annals of Applied Statistics* 2, 3 (Sept. 2008). doi:10.1214/08-aos187
- [52] Alejandro Schuler, Michael Baiocchi, Robert Tibshirani, and Nigam Shah. 2018. A comparison of methods for model selection when estimating individual treatment effects. arXiv:1804.05146 [stat.ML] doi:10.48550/ARXIV.1804.05146
- [53] Patrick Schwab, Lorenz Linhardt, Stefan Bauer, Joachim M. Buhmann, and Walter Karlen. 2019. Learning Counterfactual Representations for Estimating Individual Dose-Response Curves. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 04 (apr 2019), 5612–5619. arXiv:1902.00981 [cs.LG] doi:10.1609/aaai.v34i04.6014
- [54] Jonas Schweisthal, Dennis Frauen, Valentyn Melnychuk, and Stefan Feuerriegel. 2023. Reliable Off-Policy Learning for Dosage Combinations. arXiv:2305.19742 [cs.LG] doi:10.48550/ARXIV.2305.19742
- [55] William R. Shadish, Thomas D. Cook, and Donald T. Campbell. 2015. *Experimental and quasi-experimental designs for generalized causal inference* ([nachdr.] ed.). Wadsworth Cengage Learning, Belmont, CA [u.a.].
- [56] Uri Shalit, Fredrik D. Johansson, and David Sontag. 2016. Estimating individual treatment effect: generalization bounds and algorithms. arXiv:1606.03976 [stat.ML] doi:10.48550/ARXIV.1606.03976
- [57] Claudia Shi, David M. Blei, and Victor Veitch. 2019. Adapting Neural Networks for the Estimation of Treatment Effects. arXiv:1906.02120 [stat.ML] doi:10.48550/ARXIV.1906.02120
- [58] R Singh, L Xu, and A Gretton. 2024. Kernel methods for causal functions: dose, heterogeneous and incremental response curves. *Biometrika* 111, 2 (06 2024), 497–516.
- [59] Jerzy Splawa-Neyman, D. M. Dabrowska, and T. P. Speed. 1990. On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9. *Statist. Sci.* 5, 4 (1990), 465–472. <http://www.jstor.org/stable/2245382>
- [60] Richard Stone. 1993. The Assumptions on Which Causal Inferences Rest. *Journal of the Royal Statistical Society: Series B (Methodological)* 55, 2 (Jan. 1993), 455–466. doi:10.1111/j.2517-6161.1993.tb01915.x
- [61] Stefan Tübbicke. 2021. Entropy Balancing for Continuous Treatments. *Journal of Econometric Methods* 11, 1 (Dec. 2021), 71–89. doi:10.1515/jem-2021-0002
- [62] Toon Vanderschueren, Robert Boute, Tim Verdonck, Bart Baesens, and Wouter Verbeke. 2023. Optimizing the preventive maintenance frequency with causal machine learning. *International Journal of Production Economics* 258 (apr 2023), 108798. doi:10.1016/j.ijpe.2023.108798
- [63] Stefan Wager and Susan Athey. 2018. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *J. Amer. Statist. Assoc.* 113, 523 (June 2018), 1228–1242. doi:10.1080/01621459.2017.1319839
- [64] Xin Wang, Shengfei Lyu, Xingyu Wu, Tianhao Wu, and Huanhuan Chen. 2022. Generalization Bounds for Estimating Causal Effects of Continuous Treatments. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 8605–8617.
- [65] Yaqing Wang, Quanming Yao, James T. Kwok, and Lionel M. Ni. 2020. Generalizing from a Few Examples: A Survey on Few-shot Learning. *Comput. Surveys* 53, 3 (June 2020), 63:1–63:34. doi:10.1145/3386252
- [66] Simon N. Wood. 2017. *Generalized Additive Models: An Introduction with R, Second Edition* (2 ed.). Chapman and Hall/CRC, Boca Raton. doi:10.1201/9781315370279
- [67] Jieqiang Yang, Pengyun Wang, Dejian Zou, Zitang Zhou, Kunyuan Ding, WENXUAN PENG, Haoqi Wang, Guangyao Chen, Bo Li, Yiyun Sun, Xuefeng Du, Kaiyang Zhou, Wayne Zhang, Dan Hendrycks, Yixuan Li, and Ziwei Liu. 2022.

- OpenOOD: Benchmarking Generalized Out-of-Distribution Detection. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 32598–32611.
- [68] Nanyang Ye, Kaican Li, Haoyue Bai, Runpeng Yu, Lanqing Hong, Fengwei Zhou, Zhenguo Li, and Jun Zhu. 2022. OoD-Bench: Quantifying and Understanding Two Dimensions of Out-of-Distribution Generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 7947–7958. [https://openaccess.thecvf.com/content/CVPR2022/papers/Ye\\_OoD-Bench\\_Quantifying\\_and\\_Understanding\\_Two\\_Dimensions\\_of\\_Out-of-Distribution\\_Generalization\\_CVPR\\_2022\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2022/papers/Ye_OoD-Bench_Quantifying_and_Understanding_Two_Dimensions_of_Out-of-Distribution_Generalization_CVPR_2022_paper.pdf)
- [69] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. 2018. GANITE: Estimation of Individualized Treatment Effects using Generative Adversarial Nets. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=ByKWUeWA->
- [70] Yi-Fan Zhang, Hanlin Zhang, Zachary C. Lipton, Li Erran Li, and Eric P. Xing. 2022. Exploring Transformer Backbones for Heterogeneous Treatment Effect Estimation. *arXiv:2202.01336 [cs.LG]* doi:10.48550/ARXIV.2202.01336
- [71] Minqin Zhu, Anpeng Wu, Haoxuan Li, Ruoxuan Xiong, Bo Li, Xiaoqing Yang, Xuan Qin, Peng Zhen, Jiecheng Guo, Fei Wu, and Kun Kuang. 2024. Contrastive Balancing Representation Learning for Heterogeneous Dose-Response Curves Estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

## A Notation

We summarize the relevant notation to our paper below. Random variables are denoted in capital letters, with realizations of such in lowercase. Matrices are denoted in boldface. Realizations of a variable at a certain position, are denoted with the position as subscript.

|                                                          |                                                                                       |
|----------------------------------------------------------|---------------------------------------------------------------------------------------|
| $\mathcal{X} \subset \mathbb{R}^m$                       | Covariate space                                                                       |
| $m$                                                      | Size of $\mathcal{X}$ (number of covariates)                                          |
| $\mathcal{T} \in \{\omega_1, \dots, \omega_k\}$          | Intervention space                                                                    |
| $k$                                                      | Number of possible interventions                                                      |
| $\mathcal{D} \subset \mathbb{R}$                         | Dose space                                                                            |
| $\mathcal{Y} \subset \mathbb{R}$                         | Outcome space                                                                         |
| $\mathbf{X} \in \mathcal{X}$                             | Covariates (random variable)                                                          |
| $T$                                                      | Interventions (random variable)                                                       |
| $D : \mathcal{T} \rightarrow \mathcal{D}$                | Potential dose function (function-valued random variable)                             |
| $Y : (\mathcal{T}, \mathcal{D}) \rightarrow \mathcal{Y}$ | Potential outcome function (function-valued random variable)                          |
| $\mu(t, d, \mathbf{x})$                                  | Conditional average dose response ( $\mathbb{E}[Y(t, d)   \mathbf{X} = \mathbf{x}]$ ) |

## B Methodology details

We illustrate the non-uniformity and confounding of doses in Figure 8, using a toy example with two individual units being assigned a dose, sampled from a probability distribution. This visualization is generalizable and serves to explain any effects on interventions as well.

In the base scenario, every dose in  $\mathcal{D} = [0, 1]$  is equally likely assigned to any of the units. Under non-uniformity, some doses are more likely to be assigned, specifically lower and higher ones, whereas medium doses around 0.5 are less likely. The individual distributions of the dose per unit are equivalent to the joint distribution. In the case of confounded doses, the joint distribution across the two units is the same as under non-uniformity but individual distributions differ. Specifically, Unit 1 is assigned lower doses on average, whereas Unit 2 is assigned higher doses.




**Figure 8: Impacts of non-uniformity and confoundedness on dose distributions per unit and per population in the case of two observed units and for doses only.**

## C ML estimators for conditional average dose responses

We will introduce the three CADR estimators considered in our paper in detail:

**DRNet [53].** DRNet uses a multitask learning approach [6] to estimate conditional average dose responses. The method trains a head network per individual intervention on a set of shared layers, as motivated by Shalit et al. [56] in the binary intervention setting. Per intervention, the method partitions the dose space  $\mathcal{D}$  into a set of strata. For each stratum, another individual network is trained to infer the dose response. The architecture can further be combined with regularization terms during training, to overcome potential covariate shifts between different interventions.

**SCIGAN [3].** SCIGAN assumes that there is a shift in covariates between different levels of intervention and dose. This shift leads to non-adjusted estimators overfitting the training data. At the core of SCIGAN is a specialized generative adversarial network (GAN) structure, which attempts to generate the outcomes of counterfactual interventions and doses per unit. Creating those counterfactual observations removes the covariate shift, and allows any estimator to learn an unbiased dose response model.

**VCNet [40].** VCNet proposes a varying-coefficient architecture [21]. The method trains a neural network that has a network structure varying in the assigned dose per unit, reinforcing the influence of the dose on the predicted outcome. Nie et al. [40] combine this architecture with estimating the generalized propensity score of the dose to calculate estimates of the *average* dose response through an approach proposed as “functional targeted regularization”, as motivated and inspired by Shi et al. [57]. In our experiments, we do not use this part of the architecture and focus on estimating the conditional average dose response.

## D Overview of established estimators and benchmarking datasets

When studying previously established dose response estimators, we find that methods have typically been evaluated on a small selection of benchmark datasets (cf. Table 4). Similarly, each paper does not compare performance against a complete set of benchmarking methods.

**Table 4: Overview of dose response estimators**

| Method             | Paper                   | Type | Multi-treat. | Benchmark datasets           | Benchmark methods                                                                          | Description                                                                                                                 |
|--------------------|-------------------------|------|--------------|------------------------------|--------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| DRNet              | Schwab et al. [53]      | CA   | ✓            | TCGA-1<br>MVICU-1<br>News-1  | BART [8]<br>Causal Forest [63]<br>GANITE [69]<br>HIE [26]<br>kNN [9]<br>MLP<br>TARNET [56] | Multi-head neural network with separate head network per treatment. Per treatment, separate head network per dose interval. |
| SCIGAN             | Bica et al. [3]         | CA   | ✓            | TCGA-2<br>MVICU-2<br>News-2  | DRNet<br>HIE<br>MLP                                                                        | GAN architecture to remove training data confounding.                                                                       |
| VCNet              | Nie et al. [40]         | A    | ✗            | IHDP-1<br>News-3<br>Synth-1  | BART<br>Causal forest<br>Dragonnet [57]                                                    | Varying coefficient network combined with targeted regularization.                                                          |
| ADMIT              | Wang et al. [64]        | A    | ✗            | TCGA-2<br>MVICU-2<br>Synth-1 | DRNet<br>EBCT [61]<br>HIE<br>SCIGAN<br>VCNet                                               | Representation balancing based on weighted populations.                                                                     |
| TransTEE           | Zhang et al. [70]       | CA   | ✓            | TCGA-2                       | DRNet<br>SCIGAN<br>TARNET<br>VCNet                                                         | General purpose transformer architecture for intervention response estimation.                                              |
| CRNet <sup>†</sup> | Zhu et al. [71]         | CA   | ✓            | IHDP-2<br>News-4<br>Synth-2  | Causal forest<br>CBGPS [13]<br>DRNet<br>HIE<br>SCIGAN<br>VCNet                             | Contrastive representation balancing.                                                                                       |
| GIKS <sup>†</sup>  | Nagalapatti et al. [39] | CA   | ✗            | IHDP-1<br>News-3<br>TCGA-2   | DRNet<br>TARNET<br>TransTEE                                                                | Adjusted learning objective based on weighted observations.                                                                 |
| ACFR               | Kazemi and Ester [33]   | CA   | ✗            | TCGA-2<br>News-2             | ADMIT<br>DRNet<br>HIE<br>MLP<br>SCIGAN<br>VCNet                                            | Representation balancing through adversarial learning.                                                                      |

<sup>†</sup>: No code base available, u.n.: unnamed, CA: Conditional average, A: Average, HIE: Hirano-Imbens estimator, MLP: multilayer perceptron

An overview of the proposed datasets for benchmarking dose response estimators can be found in Table 5 below, presenting the dimensionality of the covariate space, as well as the cardinality of the intervention space, and the presence of both intervention and dose confounding.

**Table 5: Benchmarking datasets for DR estimators**

| Dataset              | Paper              | Dim(x)        | T          | t confounding | d confounding |
|----------------------|--------------------|---------------|------------|---------------|---------------|
| MVICU-1              |                    | (8040, 49)    | 3          | ✓             | ✗             |
| News-1               | Schwab et al. [53] | (5000, 2870)  | {2,4,8,16} | ✓             | ✗             |
| TCGA-1               |                    | (9659, 20531) | 3          | ✓             | ✗             |
| MVICU-2              |                    | (8040, 49)    | 2          | ✓             | ✓             |
| News-2               | Bica et al. [3]    | (5000, 2870)  | 3          | ✓             | ✓             |
| TCGA-2*              |                    | (9659, 4000)  | 3          | ✓             | ✓             |
| News-3*              |                    | (2993, 498)   | 1          | n.a.          | ✓             |
| IHDP-1*              | Nie et al. [40]    | (747, 25)     | 1          | n.a.          | ✓             |
| Synth-1*             |                    | (700, 6)      | 1          | n.a.          | ✓             |
| IHDP-2 <sup>†</sup>  |                    | (2993, 498)   | 1          | ✓             | ✓             |
| News-4 <sup>†</sup>  | Zhu et al. [71]    | (747, 25)     | {2,4,8,16} | ✓             | ✓             |
| Synth-2 <sup>†</sup> |                    | (3000, 100)   | {1,2,5,10} | ✓             | ✓             |
| IHDP-3*              | <i>This paper</i>  | (747, 25)     | 1          | n.a.          | ✓             |

\*: Considered in this paper, <sup>†</sup>: No code base available, S: synthetic, R: real, n.a.: not applicable