

# From Observational Data to Clinical Recommendations: A Causal Framework for Estimating Patient-level Treatment Effects and Learning Policies

Rom Gutman<sup>1</sup>, Shimon Sheiba<sup>1</sup>, Omer Noy Klein<sup>2</sup>, Naama Dekel Bird<sup>1</sup>, Amit Gruber<sup>3,4</sup>, Doron Aronson<sup>3,4</sup>, Oren Caspi<sup>3,4</sup>, and Uri Shalit<sup>1,5</sup>

<sup>1</sup>Faculty of Data and Decision Sciences, Technion, Israel

<sup>2</sup>Blavatnik School of Computer Science, Tel-Aviv University, Israel

<sup>3</sup>Department of Cardiology, Rambam Health Care Campus, Israel

<sup>4</sup>The Bruce Rappaport Faculty of Medicine, Technion, Haifa, Israel

<sup>5</sup>Department of Statistics and Operations Research, School of Mathematical Sciences, Tel Aviv University, Israel

## Abstract

We propose a framework for building patient-specific treatment recommendation models, building on the large recent literature on learning patient-level causal models and inspired by the target trial paradigm of Hernán and Robins [1]. We focus on safety and validity, including the crucial issue of causal identification when using observational data. We do not provide a specific model, but rather a way to integrate existing methods and know-how into a practical pipeline. We further provide a real world use-case of treatment optimization for patients with heart failure who develop acute kidney injury during hospitalization. The results suggest our pipeline can improve patient outcomes over the current treatment regime.

## 1 Introduction

The rapid accumulation of patient health data is driving an increased interest in using the data to learn models which recommend patient-specific treatments [2–7]. The goal is to use patient health data, such as electronic medical records (EMRs), to learn individualized treatment rules to improve patient care compared to existing treatment policies. Ideally, the learned individualized treatment rules will help physicians make better decisions about the best treatment for each patient. This idea holds a particular interest in scenarios where clinicians experience high uncertainty in treatment decision-making. For example, consider the case of acute heart failure (AHF) patients who have developed acute kidney injury (AKI). Caregivers face a challenging trade-off when deciding the optimal level of diuretics to administer, having to decide whether to target the congestion problem or the organ perfusion problem (intra-vascular volume) [8–10]. The estimation of optimal patient-level treatments is a causal problem [11–13], as it requires evaluating and comparing the causal effect of each treatment under consideration. Thus, learning treatment recommendation models

from patient health data requires extra care, both due to the biases inherent in causal data analysis [14, 15] and the high stakes of the task.

However, despite their wide use, there are real concerns about the possible biases of observational data, chief among them the possibility of unobserved confounding between treatment and outcome leading to bias in the effect estimates. Furthermore, while observational health data is widely used to estimate average treatment effects (ATEs), much less work has been done on estimating patient-level effects using such data [16].

Indeed, estimating patient-level treatment effects is a daunting task – we cannot hope to correctly estimate the true effect for each and every patient. However, we claim that in many cases estimating the accurate effect for every patient is not necessary for having a positive impact. Rather, in cases where there are no clear treatment guidelines, a reasonable goal for a treatment recommendation system would be to make recommendations that, if followed, would improve the overall average patient outcomes compared to the current practice [15, 17]. This idea is formalized in the notion of *policy value* (see Section 2.5). Importantly, recommendations need not be given for all patients – in order to ensure safety, recommendations may be deferred for patients for whom the recommendation model shows high uncertainty regarding which treatment is best.

We propose the Target Recommendation System, a framework for safely learning and rigorously evaluating patient-level treatment recommendation models from patient health data.

Our goal is to enable researchers to responsibly develop treatment recommendation models that will lead to overall better patient outcomes. The framework includes suggested guidelines for assessing whether learning such models is feasible for a given clinical task and dataset, and best practices for applying and validating the many patient-level causal estimation methods currently available.

Given a clinical question of interest (“How should we manage diuretics for heart failure patients with kidney injury?”), the goal of the proposed framework is to help practitioners address two questions: (1) Is the clinical question answerable given the available data? And, if positive, (2) How to safely estimate a treatment recommendation model from data, and how to evaluate its potential value?

Toward these goals, we first give a set of sufficient conditions regarding the clinical task and the available data. We propose guidelines for addressing the questions of hidden confounding, and more generally addressing the problem of *causal identification* [18], i.e. the conditions under which using the available data for estimating the desired causal effects is possible in the first place. We then propose a workflow for learning and validating a treatment recommendation model, integrating a large selection of recent work in statistical machine learning and causal inference. Our focus here is not on any specific algorithm, but rather on how to best make use of the rapidly growing literature in the field.

We apply our framework to a real-world case study involving patients with a severe and difficult-to-treat condition: patients hospitalized with AHF, who developed AKI during their hospital stay. We estimate and validate treatment recommendation policies, showing that they could improve patient welfare compared to current practice.

## 2 Preliminaries

In this section, we give the basic definitions needed for developing the framework. A full formal introduction to causal inference is beyond the scope of this paper, and we refer the readers to [18, 13, 19, 20] References on more specific concepts are given throughout the section.

## 2.1 The potential outcomes framework and the conditional average treatment effect

Our goal is to estimate the effect of a treatment, or intervention,  $T$  on an outcome  $Y$  given covariates  $X$ . We consider a binary treatment, although our framework can be generalized to multiple treatments. We assume for each unit (i.e., patient) there exist two *potential outcomes* [11]  $Y^1, Y^0$ , where  $Y^t$  is the outcome that would have occurred had the patient received treatment  $t$ .

Under the potential outcomes framework, the Average Treatment Effect (ATE) is defined as  $\mathbb{E}[Y^1 - Y^0]$ . The ATE can be interpreted as the difference between the outcome of treating all the population with  $T = 1$  versus treating them all with  $T = 0$ . The Conditional Average Treatment Effect (CATE) [21] is defined as the treatment effect conditioned on  $X = x$ :

$$\tau(x) := \mathbb{E}[Y^1 - Y^0 \mid X = x]. \quad (1)$$

The CATE represents the average gain or loss from changing the treatment for the sub-population with covariates  $X = x$ . In this paper,  $X$  is assumed to be a relatively high-dimensional vector, and in practice is usually unique for each patient.

We further assume we observe a sample of  $n$  patients  $\{(x_i, t_i, y_i)\}_{i=1}^n$ , where  $x_i \in \mathcal{X} \subset \mathbb{R}^d$  represents the baseline (pre-treatment) covariates of the  $i$ -th patient,  $t_i \in \{0, 1\}$  is a treatment given in a single point in time, and  $y_i \in \mathcal{Y}$  corresponds to a continuous or binary outcome of interest.

Only one of the potential outcomes  $Y^t$  can ever be observed for each patient [11], a major challenge known as “The Fundamental Problem of Causal Inference” [12]. Thus, estimation of either of the above causal quantities using the observed outcomes  $y_i$  requires that some set of causal identification conditions be satisfied, conditions which we outline in Section 2.2 below. If these conditions are indeed satisfied, unbiased estimates of the ATE and CATE can be obtained from the patient data using estimation methods, as described in Section 2.3.

## 2.2 Conditions for causal effect identifiability

Treatment effects can be estimated from observational data under a set of four jointly sufficient conditions [22, 23], commonly used in causal inference: **stable unit treatment value assumption (SUTVA)** [23], **consistency** [23], **common support (overlap, positivity)** [22] and **ignorability (conditional independence)** [22].

**Assumption 1 (SUTVA)** *The potential outcome of any unit will not be affected by the treatment assignment of another unit, and, there are no different forms or versions of each treatment level, which lead to different potential outcomes.*

**Assumption 2 (Consistency)**  $Y = Y^1 \cdot T + Y^0 \cdot (1 - T)$

*The observed outcome  $Y = y_i$  for each patient is in fact the potential outcome of the patient under the treatment  $T = t_i$  that they had received.*

**Assumption 3 (Common Support)**  $p(T = t \mid X = x) > 0 \quad \forall x \in \mathcal{X}, t \in \{0, 1\}$

*All the patients have a non-zero probability of receiving each possible treatment  $T$ .*

**Assumption 4 (Ignorability)**  $\{Y^1, Y^0\} \perp\!\!\!\perp T \mid X$

*This conditional independence statement implies that there are no unmeasured variables that affect both the treatment assignment and the outcome (a.k.a. confounders).*

While all causal effect estimation methods must rely on identification conditions, some of these conditions are provably untestable from data. Prominently, it is well-known that there is no data-dependent test for the validity of ignorability (Assumption 4) [24, 13, 18]. This condition is essential because hidden confounding can lead to non-vanishing bias in causal estimates, even with an unbounded sample size [13]. As a mitigation strategy, methods of *sensitivity analysis* have been developed to test the stability of the treatment recommendations under varying levels of hidden confounding, including for CATE [25–31]. See Section 2.4 for more details.

The quantity in Assumption 3 is commonly denoted by  $e(x) := p(T = t \mid X = x)$  and is known as the *propensity score*. This is the treatment assignment probability for a patient with observed covariates  $x$ , as reflected in the observed data, i.e. it encodes the actual clinical practice. The propensity score is often estimated from data and used for causal identification and estimation [32], and we discuss it in detail in Section 3.3.

### 2.3 Methods for CATE estimation

Under the identification assumptions specified in Section 2.2, the CATE function  $\tau(x)$  can be estimated from observed data as:

$$\tau(x) = \mathbb{E}[Y \mid T = 1, x] - \mathbb{E}[Y \mid T = 0, x].$$

This is also known in the literature as heterogeneous treatment effect (HTE) (e.g., [5]). In recent years, many methods have been proposed to generate estimates  $\hat{\tau}$  of  $\tau$ .

A commonly used set of approaches for estimating the CATE function is known as meta-learners [33], where the CATE estimation is decomposed into several supervised learning subproblems. The simplest approach, named *S-Learner* [33], fits a single model  $f(x, t)$  using the entire sample, with  $T$  acting as a feature:  $f(x, t) \approx \mathbb{E}[Y \mid T = t, x]$ . CATE is then estimated as:  $\hat{\tau}(x_i) = f(x_i, 1) - f(x_i, 0)$ . The so-called *T-learner* [33] method fits two separate models  $f_t(x) \approx \mathbb{E}[Y \mid T = t, x]$ , each to the population receiving the corresponding treatment  $t$ , and estimate the CATE as  $\hat{\tau}(x_i) = f_1(x_i) - f_0(x_i)$ . In both cases, the model(s)  $f$  can be learned using any supervised learning model. As the S-Learner fits a single model for both treatment arms, the predicted outcomes of both arms share the same noise, thereby potentially reducing variance. However, the treatment covariate might get washed out and result in a zero-biased estimation of CATE. The T-learner approach offers more flexibility by employing separate models for each treatment arm, but it does not leverage the common components of the response, which can introduce both bias and unneeded variance, particularly when the sizes of the treatment groups differ.

Several works have expanded the meta-learners approach with the goal of overcoming some of the shortcomings of both the S- and T-Learners [33–35]. These methods utilize the learner’s estimation with additional knowledge, such as propensity scores, to establish what are known as pseudo-outcomes. For instance, the *X-learner* [33] overcomes imbalanced treatment arm population sizes by establishing pseudo-outcomes using both arms of the T-learner with a weighting function. The *R-learner* [34] and *DR-learner* [35] combine the propensity score with the outcome predictions to establish pseudo-labels which are in turn fitted using yet another model. See [36] for a detailed overview of these methods.

Other approaches aim to estimate the CATE function directly. A well-known approach is the Causal Forest [37], a tree-based non-parametric method that directly models the treatment effect similar to the well-known random forest [38]. Additional examples are causal boosting [39], the doubly-robust targeted maximum likelihood estimation (TMLE) [40, 41], Gaussian process-based



methods [42], and deep learning-based approaches [43–46]. A comprehensive description of existing methods is provided by Bica et al. [2].

## 2.4 Causal and statistical uncertainty

In Section 2.3 we described methods for estimating treatment effects. Building on these, incorporating uncertainty estimation is important for enhancing the reliability of these models as they work in tandem with clinicians.

There are two primary sources of uncertainty in CATE estimates: statistical and causal. The statistical uncertainty, common in supervised learning, includes finite sample variance, inherent noise, and model misspecification, among others. This uncertainty has been extensively studied with both Bayesian methods [47–49] and frequentist approaches [37, 50, 51]. The causal uncertainty stems from the reliance of CATE estimators on unverifiable assumptions, as mentioned above, notably the assumption of the absence of hidden confounders. This uncertainty is typically addressed by methods of sensitivity analysis which assume some bounded level of hidden confounding. Examples of such approaches tailored for CATE estimation can be found in [25, 52, 27–31]. Some of these models account for both sources of uncertainty together [53, 28, 31]. Notably, causal uncertainty does not diminish with sample size, but might diminish by measuring previously unmeasured confounders.

Methods for estimating both types of uncertainty typically have predefined parameters that represent the allowed uncertainty levels. In this work we will denote these as  $\alpha_{\text{stat}}$  and  $\alpha_{\text{causal}}$ , respectively. For example, the parameters can respectively represent the expected confidence level and assumed level of unobserved confounding.

These uncertainty parameters in turn yield a range of possible estimates for the CATE, establishing bounds for the possible values of  $\tau(x)$ . We denote these bounds as  $[\hat{\tau}_{\theta}(x), \overline{\hat{\tau}}_{\theta}(x)]$ , where  $\theta$  represents the set of used uncertainty parameters, for example  $\theta = (\alpha_{\text{stat}}, \alpha_{\text{causal}})$ .

Importantly, when the range  $[\hat{\tau}_{\theta}(x), \overline{\hat{\tau}}_{\theta}(x)]$  includes both positive and negative values, it implies that it is uncertain whether the treatment is beneficial or harmful for patients with covariates  $x$  for the defined uncertainty parameters  $\theta$ . As we will see below, in such cases we might want to defer making a treatment recommendation.

## 2.5 Learning and evaluating treatment policies

In this work, the main motivation for estimating the CATE function  $\hat{\tau}$  is to inform a treatment policy for achieving optimal outcomes. Formally, let  $\pi : X \rightarrow \mathcal{T}$  be a treatment policy that maps a patient’s feature vector  $x \in \mathcal{X}$  to a treatment assignment  $t \in \mathcal{T}$ . This is also known as an individual treatment rule (ITR). Here we will consider the more general problem of learning a policy that either makes a treatment recommendation or alternately *defers* the recommendation due to insufficient certainty [53].

In this paper, we use CATE estimates to construct a policy. Given a CATE estimate  $\hat{\tau}$ , we apply a decision rule  $\psi$  producing a policy  $\pi$  such that

$$\pi(x_i) = \psi(\tau(x_i)).$$

A standard approach for binary treatments is to assign treatment according to the sign of a CATE model  $\hat{\tau}$ , such that  $\psi_{\text{sign}}(x_i) = \mathbb{1}_{\hat{\tau}(x_i) \geq 0}$ .

Given a policy  $\pi$ , we are interested in evaluating the expected patient outcomes if treatments were assigned according to  $\pi$ , and comparing these expected outcomes with outcomes under other

treatment policies, e.g., the actual treatment assignment observed in the data. Generally, a good policy would be one that is better than current practice. The expected outcome under a certain policy is known as the *policy value*, defined as:

$$V(\pi) := \mathbb{E} \left[ Y^{\pi(X)} \right].$$

Similarly to ATE and CATE, this estimand also suffers from the fundamental problem of causal inference, as for any given  $x_i$ , the quantity  $Y^{\pi(X)}$  is a counterfactual whenever  $\pi(x_i) \neq t_i$ . Thus, the estimation of  $V(\pi)$  requires causal identification conditions, and the same ones given in Section 2.2 are commonly used. More generally the problem of estimating  $V(\pi)$  resembles that of estimating ATE, and similar approaches are often employed in this context. Most commonly, these methods rely either on outcome modeling (“plug-in”, “direct methods”), weighted methods (inverse propensity weighting/scoring – IPW) [54], or a combination of both in a doubly-robust manner [55, 56].

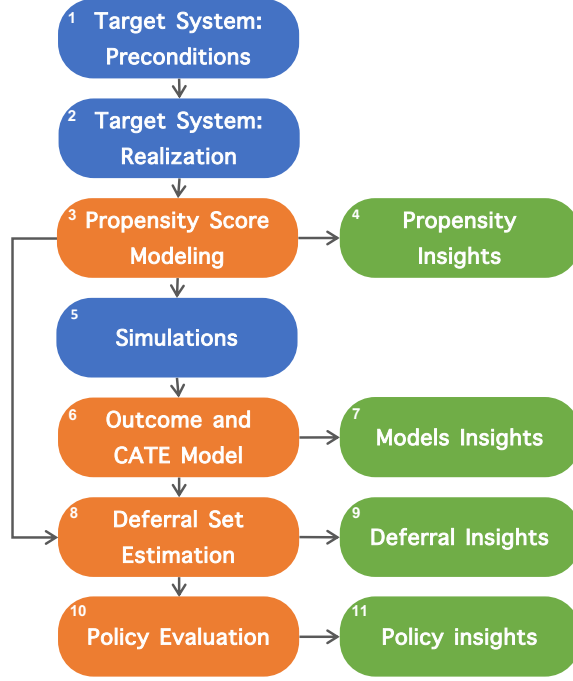
While below we mostly focus on learning a treatment policy via the CATE function, there are many approaches for optimizing  $\pi(x)$  directly from data, with the aim of learning a policy attaining high values of  $V(\pi)$  [54, 57, 55, 58–61]. Some approaches focus on learning a more complex decision rule  $\psi$  [59], combining propensity information with outcome models in a doubly-robust manner [59, 60] or directly selecting the policy based on maximizing the potential outcome [57]. Several works have expanded this direction using semi-parametric methods [61, 59] or using other optimization methods [60].

### 3 Modeling Framework

We introduce a framework for developing a treatment recommendation system based on patient health data. Given patient covariates  $x$ , such a system either outputs a recommended treatment  $t \in \mathcal{T}$ , or a “defer” option ( $\perp$ ), meaning that it has no recommendation for the specific case. Fig. 1 provides a high-level map of the framework’s steps, where we distinguish between three types of steps:

- **Identification:** The identification steps are designed to help practitioners specify the desired treatment recommendation system and assess the feasibility of constructing it with the available data. This entails a close discussion with the clinical partners to define the problem setting, followed by a preliminary examination of the framework using a semi-synthetic simulation.
- **Estimation:** The estimating are meant to guide practitioners on how to build the target system models, including methods for estimating propensity scores, outcome and CATE models, and estimating the policy value.
- **Validation:** The validation steps are aimed at building confidence in the system’s output and gaining clinical insights. This includes proposed evaluation practices to assess the outcomes of the estimation steps.

The following subsections provide a detailed description of the framework.



**Figure 1:** The outline of the target system, including **identification**, **estimation** and **validation** steps. Steps 1-2 refer to defining the clinical and causal question of interests and step 5 is designed to validate its feasibility. Steps 3,6,8,10 are for estimating the causal quantities of interest such as the CATE, and establishing a recommendation policy (step 10). Steps 4,7,9,11 are for validating and evaluating the corresponding estimation steps, where we recommend best practices to ensure each step’s validity.

### 3.1 Target Recommendation System: Preconditions

The “Target Trial” is a framework for estimating average treatment effects from observational data, which has had many successes in recent years [62, 14, 1, 63, 3]. The framework postulates that one should formulate an ideal clinical trial which would be used to estimate the average effect of the intervention in question and that the analysis of the observational data should emulate this ideal trial. Here, we propose a similar idea for the clinical decision support, framed as the “Target Recommendation System” (TRS). TRS helps define the estimands, when will the system be called upon, with what input, and how will it be integrated within the clinical workflow.

Therefore, we start by formulating the basic questions needed to define the system. These should be discussed between the modeling and clinical experts on the team, including a close understanding of the clinical staff’s routines, to gain firsthand insights into their decision-making processes when faced with treatment choices. Ideally, the discussions should take place before one embarks on collecting and analyzing patient data [64], as they bear on the causal assumptions needed for obtaining valid estimates from observational data (Section 2.2), as well as the general applicability of the proposed system. The basic points to address are as follows:

**Single treatment decision at a well-defined time-point.** The TRS should focus on a single clinical decision which occurs at a well-defined time point in the clinical workflow. Examples include

selecting among a set of possible medications, determining whether to embark on a pre-specified procedure (e.g., a surgery), whether to take a diagnostic action (e.g., a specific type of biopsy). This also includes preliminary inclusion and exclusion criteria, giving a preliminary definition of the target population for which the decision and the system are relevant.

The importance of a well-specified decision time-point (called time-zero by Hernán and Robins [1]) is double. First, this will likely be the point in the clinical workflow where the system will be called upon to provide a recommendation. Thus, any input to the TRS must be based on data that is available before the decision time-point. Additionally, pre-specifying the decision time-point helps us grapple with questions of confounding: since confounders are factors that affect the clinical decision (and the outcome), understanding what information was available to the decision maker at decision time-point allows us to narrow down the set of potential confounders, addressing the ignorability assumption (Section 2.2).

**Small action space.** We focus on treatment decisions with a relatively small set of treatment options (actions). Furthermore, these actions should in principle be applicable to most patients within the target population. Large action spaces usually imply a small sample size for at least some of the actions, and are much more likely to lead to violations of the common support assumption (Section 2.2).

**Ambiguity or lack of clinical guidelines.** If the decision about the assignment of the treatment in question follows clear clinical guidelines, then it is very hard to generate evidence for the treatment recommendation system, as the overlap assumption (Section 2.2) is likely violated: Patients of certain characteristics would (nearly) always receive the treatment specified by the guidelines, and there will be no data about the response of these patients to the treatment which goes against the guidelines. On the other hand, the absence of clear clinical guidelines grants clinicians flexibility in selecting a treatment, potentially leading to variability between clinicians, which in turn could lead to similar patients receiving different treatments.


Moreover, cases where there are no clinical guidelines are often the cases where clinicians are open to, and in need of, assistance in determining the best treatment for a given patient.

**Well-defined and meaningful outcomes.** The treatment decision should have a clinical outcome that is well-defined and acknowledged by the domain experts as clinically meaningful when deciding on treatment. This is the outcome which the TRS will aim to optimize.

Assuming the above preconditions have been discussed and affirmatively addressed, the output of this stage should be the following: (1) **who are the clinical practitioners** expected to use the system; (2) **who are the patients** for which the recommendation is needed (3) the **time point** within the clinical course when a recommendation will either be solicited by the clinician or offered by the system; (4) what information do we expect to be available as **input** to the system at this time point; (5) the exact types of recommendations the system is expected to present to the clinical staff, i.e. the system’s **outputs**.

In cases where the above conditions are not met, one can try and redefine the treatment recommendation system. This includes, but is not limited to, redefining the treatment decision, restricting the recommendation to a smaller sets of treatments, or redefining the target population. One might also conclude that the task is unfeasible, which we refer to as an *exit point*.

In Appendix A we present suggested questions that can aid in defining exactly the setting and context of the treatment recommendation system.



**Figure 2:** A schematic illustration of the different types of causal variables for the question of the causal effect of the treatment  $T$  on the outcome  $Y$ ; this is a modified version of Figure 1 from Tennant et al. [65].  $Z$  is an instrumental variable,  $M$  is a mediating variable,  $W$  is an effect modifier,  $C$  is a *measured* confounder,  $U$  is an *unmeasured* confounder, and  $P$  is a proxy observation of some part of  $U$  (e.g., a blood test might be a proxy for the state of a physiological system).

### 3.2 Target Recommendation System: Realization

This stage aims to assess whether the assumptions for causal effect identifiability discussed in Section 2.2 hold in the available data.

We begin by evaluating the SUTVA condition (Assumption 1). This involves confirming that treatments are abundantly available, preventing a scenario where one patient’s treatment assignment compromises another patient’s access to treatment. In addition, treatments for contagious diseases might also violate SUTVA, as treating one patient might cause or prevent other patients from falling ill. Also, treatments that are nominally considered the same should in fact agree across different settings in either their mechanism of effect or their medical protocols; for example, a procedure must not be done in substantially different ways across two hospital units participating in the same study.

For the consistency assumption (Assumption 2) to hold, we must ensure that treatment allocations are accurately recorded in the dataset.

We discuss the common support condition (Assumption 3) in detail in Section 3.3.

In the rest of this subsection, we focus on the ignorability assumption (Assumption 4). Notably, we are dealing here with potential confounders for actions taken by trained human decision-makers. Since humans typically take into account a limited number of factors in their decision-making process [66–68], this reduces the possible scope of hidden confounders, although non-conscious factors should be taken into account as well.

We propose two alternative approaches to evaluating whether the ignorability assumption holds. The first is to build a causal graph [13] of all the covariates in the data and in addition any other

factors that might influence the decision making process of the clinicians. This should be done in close consultation with the clinical experts. Notably, the causal graph could include variables that are not represented as-is in the dataset. The causal graph allows researchers to map different variables and their relationships, and potentially use the *backdoor criterion* [69], a widely used condition for causal identifiability. Tennant et al. [65] offer a practical guideline on how to include and report such graphs.

Following Tennant et al. [65], in Fig. 2 we provide a schematic illustration of the different types of variables and how they affect each other. A simplified approach to the analysis is to make sure that all confounders ( $C$ ) are represented in the data, or, when that is impossible, that relevant proxies ( $P$ ) exist for confounders that are not in the data. For example, one can use hematocrit and hemoglobin biomarkers as proxies for congestion in AHF patients [70]. In addition, including effect modifiers ( $W$ ) is beneficial in terms of reducing variance for outcome modeling. On the other hand, including instrumental variables ( $Z$ ) in the analysis could increase both bias and variances and should be avoided [71–73]. Mediating variables ( $M$ ) occur after the treatment and thus should also be discarded.

However, we note that constructing a causal graph that comprehensively captures the multitude of potential variables involved in a clinical problem is often challenging and time-consuming. Thus, in Appendix B we provide an alternative to the causal graph approach: a step-by-step protocol addressing the question of confounding and which covariates should be included and excluded from the analysis.

As in the last step (Section 3.1), in cases where the data ultimately does not meet the causal identification requirements, the researcher should re-consider the settings, data acquisition setup, and the research question itself.

### 3.3 Propensity Score Modeling


In this step, we focus on the estimation of propensity score  $e(x)$ , as defined in Section 2.2. The purpose of estimating it in our framework is twofold. First, to evaluate the population for which the common support assumption holds (Section 2.2), and for whom we can plausibly estimate the CATE function. Second, propensity scores are used both in some CATE estimators (e.g. X-learner [33] and R-learner [34]) and in policy value estimators and baselines, as we will see in Section 3.10. We elaborate on the use of propensity score as an overlap evaluator in the validation step presented in Section 3.4.

From a modeling perspective, estimating the propensity score is a relatively straightforward task of estimating the conditional probability of the binary variable  $T$  conditioned on a covariate vector  $X$ . We recommend using and comparing several models for modeling the propensity score, such as (regularized) logistic regression, XGBoost [74], and outcome adaptive Lasso [75].

As the propensity score should ideally represent the true probability of treatment conditioned on the confounders, in-sample and out-of-sample calibration is an important metric for propensity model evaluation [76]. Calibration can be improved for any given model by using post-processing methods [77, 78].

We note that standard discrimination metrics for classifiers such as accuracy or the area under the receiver-operator characteristics curve (AUROC) should be treated here more as descriptive measures of the propensity score. For instance, when a model achieves an AUROC of 1, it signifies that the model assigns a higher score ( $\hat{p}(T = 1 \mid X = x)$ ) to each unit  $x$  with treatment  $T = 1$  compared to units with treatment  $T = 0$ . Although this behavior can be desirable in standard

classification problems and would indicate a strong predictor, in the context of propensity estimation such an AUROC score suggests a possible lack of overlap.



**Figure 3:** Schematic illustration of propensity overlap, presenting the distribution of propensity scores (x-axis) of patients received treatment  $T = 1$  (blue) or treatment  $T = 0$  (red). **A** Strong overlap, where minor trimming could be considered. **B** Mild overlap, where trimming for extreme values is advised. **C** Non-overlap case, where the researchers should re-consider the research question (i.e. using this data as-is is not advised).

### 3.4 Propensity score model insights

After fitting and evaluating the propensity score models, the researcher should validate their clinical face validity and use them to evaluate whether common support (Assumption 3) holds for a substantial sub-population.

As a clinical face validity check, we recommend evaluating the models using standard interpretability tools such as examining the weights of linear models, or using SHAP values [79]. The results can then be evaluated with the clinical team. The idea is to ensure that the important factors for clinical decision-making are reflected in the propensity score model. If concerns are raised (e.g. the model indicates some variable is highly predictive of treatment, but the clinicians find this unreasonable based on their experience), further analysis of the models and data is required.

If the models are deemed clinically reasonable, our focus shifts to a further objective of the propensity score model, namely assessing the extent of common support: to what degree could each unit plausibly receive each of the treatments. Towards that end, we recommend visualizing the distribution of propensity scores across the treatment arms. Fig. 3 presents a stylized demonstration of different levels of overlap: Fig. 3A exhibits near-perfect common support, Fig. 3B shows good common support, while Fig. 3C demonstrates no common support. The first two cases illustrate adequate representation for each treatment arm, although trimming might be necessary, as we discuss in the next paragraph. However, the last case (Fig. 3C) raises a warning sign, indicating a complete lack of common support. In such instances, the research objectives should be reevaluated, as there is likely no reliable statistical method for estimating meaningful treatment recommendations.

Even if common support is satisfied for most patients (as in Fig. 3A and Fig. 3B), some patients might still be outside the common support region. As part of our framework, we propose that recommendations for patients such as these, i.e. those with extreme propensity scores, be *deferred*. The reasons are twofold: First, patients with extreme propensity scores are those who typically



receive only one type of treatment, and therefore we cannot reliably estimate their counterfactual outcomes [80, 81, 18, 46]. Furthermore, patients with extreme propensity scores are typically those for whom clinicians do not require recommendations from a data-driven model, as they are consistently treated with a single treatment.

Deferral by propensity score is performed by settings lower and upper bounds (denoted  $\eta_l$  and  $\eta_h$ , respectively), such that the overlap set is defined as:

$$S_{\text{overlap}} := \{x \in X \mid \eta_l \leq \hat{e}(x) \leq \eta_h\}. \quad (2)$$

When determining the thresholds of the overlap region ( $\eta_l$  and  $\eta_h$ ), factors such as the feature space dimensionality and the number of patients under each treatment arm in the low overlap regions should be taken into consideration [82, 83, 81]. The goal is for the downstream CATE models to have sufficient support for estimating counterfactuals.

An additional validity check is to characterize the above populations for which common support does not hold, i.e. patients who according to the propensity model are treated nearly deterministically with a single treatment option. We propose consulting with the clinical partners whether in their experience treatment for this population is indeed chosen unambiguously. The goal is to further test whether the model makes clinical sense and to highlight possible omissions in the data or model.

We further propose testing whether vulnerable populations tend to have more extreme, or less extreme, propensity score values compared to the rest of the population, indicating possible systemic biases in the clinical population. Such biases should be highlighted and examined down the pipeline, as a possible target for improved recommendations.

Finally, we note that the notion of deferral is related to, but distinct from, what is commonly known as *trimming* [82, 83] in the observational study literature: the practice of removing from the study units with extreme propensity scores. When estimating an ATE, trimming raises issues regarding the target population [82, 83], as the ATE is only estimated for the non-trimmed subpopulation. However, in our case, we do not suffer from this problem, as we are aiming for individual recommendations rather than ATE estimates.

### 3.4.1 Sub-population characterization

Throughout this framework, we come up with a need to characterize in an interpretable way certain patient sub-populations. For example, here we wish to characterize the patients outside of the overlap (common support) zone, i.e. those with extreme propensity scores. We thus give some general recommendations for the task of characterizing a sub-population.

We first recommend generating some descriptive statistics of the patients included and excluded from the sub-population, and presenting them in simple tables or histograms.

We then recommend fitting a linear model or a single decision-tree  $f : X \rightarrow Z$ , where  $Z$  is the sub-population label (e.g. “deferred” and “not deferred”), and visualizing the top coefficients that contribute the most to the model’s classification. For a linear model, this can be measured by absolute coefficient value. One can further use SHAP values visualized over linear or non-linear models [79].

## 3.5 Semi-synthetic outcome simulation

The previous steps are meant to establish whether the data is suitable for the task, mainly from a causal perspective. In this step, the goal is to gain a better understanding of whether the data

suffices from the statistical perspective, i.e. in terms of sample size and the ability to fit machine learning models which are adequate to the task.

Toward that end, we propose using a semi-synthetic outcome simulation. The simulation serves two main purposes: (1) Debugging and testing whether, under relatively favorable conditions, the proposed framework can use data to uncover a treatment policy that is reasonably close to the optimal one. (2) Estimating whether the sample size is sufficient to discover treatment effects and plausibly estimate a policy that is superior to current practice, if such a policy exists.

We propose creating a simulation based on the existing covariates  $X$  and observed treatment assignments  $T$ , where the only simulated component are the potential outcomes for each unit,  $(Y_i^1, Y_i^0)$ . Thus, one can obtain the “true” treatment effect (within the context of the simulation) for each unit by computing  $Y_i^1 - Y_i^0$ . Once the potential outcomes are simulated, one can estimate the optimal policy  $\pi^*$  and its optimal policy value  $V(\pi^*)$  under the simulation, as well as the policy value of the current policy,  $\pi_{\text{current}}(x_i) = t_i$ , and the policy value of any other policy of interest.

We propose using several different functional and parametric forms for the simulated outcomes, e.g. linear, tree-based etc. The simulated potential outcomes should meet the following criteria: (1) they should have the same scale as the true outcomes, (2) exhibit similar variance as the true outcomes, (3) possess a clinically reasonable distribution of CATE values, and (4) allow for explicitly controlling varying levels of divergence between current practice and optimal practice.

Next we follow the **estimation** phase steps of the framework using the simulated data: The first of these steps is estimating CATE models (denoted  $\mathcal{A}$ ) based on the simulated data (see Section 3.6), derive a corresponding policy  $\pi_{\mathcal{A}}$  (Section 2.5), and estimate its policy value  $\hat{V}(\pi_{\mathcal{A}})$  (see Section 3.10).

We recommend running the simulation multiple times with varying settings and random seeds. The analysis should then focus on the degree to which policy value estimators (Section 3.10) agree with the ground-truth policy values, using measures such as Pearson correlation, scatter plots etc., and the degree to which CATE models managed to lead to a beneficial policy, by comparing their policy values to policy values of baseline policies (Section 3.11) on the one hand and the optimal policy on the other hand. Most importantly, we recommend checking (1) whether the estimated policy value for the policy  $\pi_{\mathcal{A}}$  is close to its true policy value under the simulation, i.e. is  $\hat{V}(\pi_{\mathcal{A}}) \approx V(\pi_{\mathcal{A}})$ ; (2) whether the learned policy  $\pi_{\mathcal{A}}$  has a better policy value than the current treatment policy, i.e. is  $V(\pi_{\mathcal{A}}) > V(\pi_{\text{current}})$ ; and (3) whether the policy value is reasonably close to the optimal policy,  $\hat{V}(\pi_{\mathcal{A}}) \approx V(\pi^*)$ .

In Appendix C we provide the simulation we created for our case study with detailed explanations of how we aimed to achieve the simulation goals.

### 3.6 Outcome and CATE modeling

Having established an overlap set (Eq. (2)), the next step is estimating the CATE function for members of the set  $S_{\text{overlap}}$ . The field of CATE estimation has seen significant growth in recent years, with numerous models and methods being developed and actively researched, as discussed in Section 2.3. In this step, our focus is on addressing the major difficulties of CATE estimation, rather than advocating for any specific method.

A major challenge in CATE estimation is the absence of a held-out set, and more generally the difficulty of evaluating and comparing the accuracy of CATE estimates. This is unlike regression or classification models, where held-out error can be used for evaluation and model selection. Thus, practitioners might be at loss with regard to which of the many possible options for CATE modelling

should they use [84].

To address this challenge, we propose a two-pronged approach that can be applied to any of the aforementioned methods. First, considering that most CATE models involve regression or classification models as components (e.g. all the meta-learner approaches), we recommend evaluating these “component models” in a manner consistent with the evaluation of traditional models. Although such evaluations are not sufficient for selecting between different families of CATE meta-learners (e.g., T-learner vs. X-learner), they can still be utilized to compare and reject poor-performing models within a given family. For instance, if we intend to use a random forest regression model as a component in a meta-learner, and we find that its held-out mean squared error is exceedingly high (for example it is approximately equal to the overall variance of the regression target), it is advisable to exclude this model from the CATE estimation process. Having said this, it is important to note that a CATE model may use an underlying regression or classification model that demonstrates superior performance in terms of held-out accuracy, yet still leads to inferior accuracy in CATE estimation and the subsequent derivation of a treatment policy [85]. Thus, the goal of this first step is merely to reject the worst-performing models, not to choose the best one.

Second, once distinctly under-performing models have been weeded out, we suggest using *held-out policy value* as the primary metric for model evaluation and selection. While not a perfect metric, it serves as the closest proxy to the ultimate goal of the treatment recommendation model, which is to improve patient welfare. We discuss the policy value in detail in Section 3.11 below.

### 3.7 Outcome and CATE model insights

For CATE models that have reasonably performing components and competitive policy value performance, we propose a series of validity tests.

First, we recommend conducting an error analysis of the component regression or classification models. This analysis aims to identify the strengths and weaknesses of the component models, detect sub-populations where the models might perform poorly, and highlight significant discrepancies among them. These insights, to some extent, can also feed back into the previous step and guide the selection of component models.

We further advise employing standard interpretability methods for the component models and involving clinical experts in examining the findings. If the models heavily rely on covariates that appear to be incongruous, it is important to investigate this phenomenon as it could indicate fragility in the models. However, it is worth noting that highly predictive models may rely on administrative data, such as the timing of tests, and these should not be dismissed outright [86, 87].

We then suggest estimating standard correlation metrics, such as Pearson and Spearman correlation, between the estimates of the selected CATE models. Strong agreement among models with different parametric assumptions provides more confidence in their results. On the other hand, if the correlations between models are low it may indicate the presence of unobserved confounding or suggest that each model explains different aspects of the treatment effect variance, with no clear determination of which is accurate. Therefore, a low correlation between the outputs of different CATE models serves as a warning sign regarding the validity of the estimates.

Next, we propose visualizing the CATE calibration graph as outlined in [88] to illustrate the data heterogeneity. This method divides the dataset into segments based on quantiles of the estimated CATE values. In each segment, we calculate the ATE using a method such as Augmented Inverse Probability Weighting (AIPW) [89] or some other ATE estimation technique. Then, we plot the estimated ATE within each segment versus the average predicted CATE within this segment. This

may indicate how well the predictor identified homogeneous groups in the heterogeneous population, by evaluating how aligned those two estimates are.

Finally, the CATE estimates should be used to infer the ATE and compare it to any known effects from the existing literature, if such exist. A wide divergence from any existing ATE estimates is not necessarily indicative of a mistake, but should encourage a deeper dive into potentially substantial differences between the setting and population used in the study and the settings and populations used in previous studies.

### 3.8 Deferral set estimation

As we focus on high-stakes clinical recommendations, we wish to create models that refrain from unfounded recommendations. This requires assessing the uncertainty of CATE estimates and deferring decisions where this uncertainty is high.

As discussed in Section 2.4, we recall that for CATE models there exist both statistical and causal uncertainty. We propose estimating both sources of uncertainty jointly (e.g. using the method proposed by Jesson et al. [28], Yadlowsky et al. [26], Oprescu et al. [31]), and use the estimated uncertainties to establish a deferral rule  $\text{Rej}_\theta$  parameterized by uncertainty parameters  $\theta$ :

$$\text{Rej}_\theta : \mathcal{X} \rightarrow \{0, 1\}, \quad (3)$$

where  $\text{Rej}_\theta(x) = 1$  implies that, given uncertainty parameters  $\theta$ , for a patient with features  $X = x$  we defer making a treatment recommendation. A common rule, as detailed in Section 2.4, is to defer the decision for a sample if its joint uncertainty interval includes zero, i.e.  $0 \in [\hat{\tau}_\theta(x), \overline{\hat{\tau}}_\theta(x)]$ .

For each of the methods jointly modeling statistical and causal uncertainty the user must specify the uncertainty parameters; the statistical parameter is typically a prediction interval or a Bayesian credible interval, while the causal parameter is often a bound on the level of divergence between the true and observed propensity scores [25]. Importantly, we note that we are not restricted to standard prediction interval levels such as setting  $\alpha_{\text{stat}} = 0.95$ . We might be willing to tolerate more uncertainty than that in order to defer fewer recommendations. We encourage varying the degrees of accepted statistical and causal uncertainties ( $\alpha_{\text{stat}}, \alpha_{\text{causal}}$ ) to understand their impact on the proportion of deferrals across the population, and on the held-out policy value of the non-deferred patients.

Importantly, the final deferral rule  $\text{Rej}'_\theta$  should also include the patients who were “trimmed” during the propensity evaluation (i.e.,  $x_i \notin S_{\text{overlap}}$  (Section 3.4)). Therefore, we recommend a combined deferral rule. For example, given a propensity estimator  $\hat{e}(x)$  and using uncertainty intervals with uncertainty parameter  $\theta$ , we obtain a deferral rule:

$$\text{Rej}'_\theta(x) = \begin{cases} 1, & \text{if } \eta_l > \hat{e}(x) \vee \hat{e}(x) > \eta_h \\ 1 & \text{if } 0 \in [\hat{\tau}_\theta(x), \overline{\hat{\tau}}_\theta(x)] \\ 0 & \text{otherwise} \end{cases}. \quad (4)$$

### 3.9 Deferral set insights

Once a deferral rule  $\text{Rej}'_\theta$  has been determined, we propose examining and characterizing the deferred patient population, understanding which populations tend to be deferred more than others. A special emphasis should be put on the role of vulnerable populations who might end up being under-served or over-served by the recommendation system.

To this end, we recommend using similar methods as described in Section 3.4.1 and consulting with the domain experts whether the characterization of the deferred sub-population makes clinical sense.

### 3.10 Policy

The main purpose of our proposed framework is to generate a *policy*: a recommendation for the treatment of eligible patients. These recommendations are generated based on a model’s CATE estimate, coupled with a decision rule converting the CATE estimate to a recommendation. Formally, given a CATE estimation method  $\mathcal{A}$  yielding CATE estimates  $\hat{\tau}_{\mathcal{A}}$  (Section 3.6), a deferral rule  $\text{Rej}'_{\theta}$  (Section 3.8), and a decision rule  $\psi : \mathbb{R} \rightarrow \mathcal{T}$ , the policy for a patient with features  $x$  is defined as:

$$\pi_{\mathcal{A}}(x) = \begin{cases} \perp & \text{if } \text{Rej}'_{\theta}(x) = 1 \\ \psi(\hat{\tau}_{\mathcal{A}}(x)) & \text{otherwise.} \end{cases}$$

As mentioned in Section 2.5, there are several possible choices for the decision rule  $\psi$ , where the most commonly used is simply the sign of  $\hat{\tau}$ :  $\psi(\hat{\tau}_{\mathcal{A}}(x)) = \mathbb{I}_{\{\hat{\tau}_{\mathcal{A}}(x) \geq 0\}}$  (or some threshold other than 0, depending on the context).

The quality of a policy is assessed via the policy value  $V(\pi)$  (Section 2.5), evaluated on the held-out data. For policies with a defer option such as the ones we propose, the policy value estimate has two components. When the policy defers, we assume the action and its outcome would match the policy observed historically in the data (i.e. current assigned treatment); thus, for a unit  $x_i$  with observed treatment  $t_i$  and outcomes  $y_i$  we assume the outcome under “defer” would simply be  $y_i$ , i.e. we assume the clinicians would treat this patient as they always do. For emphasis we denote this outcome  $Y_{\text{factual}}$ . When the policy does not defer and gives a treatment recommendation, we estimate the policy value using standard methods. Formally, given a policy function  $\pi : \mathbb{X} \rightarrow \{0, 1, \perp\}$ , where  $\perp$  is deferral, the policy value of  $\pi$  is given by:

$$V(\pi) = \mathbb{E}[Y(\pi(x)) \mid \pi(x) \neq \perp] p(\pi(x) \neq \perp) + \mathbb{E}[Y_{\text{factual}} \mid \pi(x) = \perp] p(\pi(x) = \perp), \quad (5)$$

where  $\mathbb{E}[Y(\pi(x)) \mid \pi(x) \neq \perp]$  describe what will be the average outcome  $Y$  using policy  $\pi$ . As this quantity is a causal quantity which requires estimating unobserved potential outcomes, causal identifying assumptions are required to hold for this quantity to be estimated. In Section 2.5 we discuss some of the methods that have been suggested to estimate policy value [55, 58–60, 90, 91], where common approaches include *Inverse Propensity Weighting* (IPW) as well as *Doubly-Robust* (DR) approaches which combine the propensity and outcome models [55]. We denote them as  $V^{IPW}$  and  $V^{DR}$ , respectively. All of these should be used on held-out data, since the policy is learned and thus naive estimates might suffer from overfitting. In this work we suggest using the DR methods as given by [55, 92], with bootstrap sampling. See Appendix D for details.

### 3.11 Provide insights on chosen policy

Once several policies  $\pi$  (derived from different CATE estimators, rejection rules, etc.) and their corresponding estimated policy values  $\hat{V}(\pi)$  are obtained, we suggest a series of validation steps to aid in policy selection. These steps focus on two main aspects: (1) Understanding the value of the policies and (2) Assessing the clinical sensibility of the policies.

For the first aim, we wish to ensure the learned policies outperform several baseline approaches: “Doctors”, which is the treatment as it is assigned in the historical data (i.e., the correct practice),

“Random”, which randomly assigns treatment in the same proportion as the historical policy, “Propensity”, which assigns treatment based on the estimated propensity scores (e.g. it assign  $T = 1$  if  $\hat{e}(x) > 0.5$ ), and “Treat-all-with- $t$ ”, which assigns all patients to treatment  $T = t$ ; i.e. there is one such policy for each treatment arm.

We propose the following approaches for comparing policies: (1) Direct and visual comparison, and (2) rank graph comparison.

1. **Direct and visual comparison:** First, we suggest comparing the policies’ performance in terms of held-out policy value within each bootstrap round, counting overall “wins”. See Table 4 for an example.

Additionally, it is useful to visually examine the overall performance of various policies graphically, e.g. using box plots. This comparison may provide insight into the distribution of policy values; see Fig. 7 for an illustration.

2. **Rank graph comparison:** We further propose measuring the policy values as function of treatment proportion [43, 91]. This analysis assumes each policy can provide a ranking of patients by treatment benefit. The graph then presents the policy’s value under a varying threshold  $q \in [0, 1]$ , which represents assigning  $T = 1$  only to a proportion  $q$  of all the patients. Thus, the extreme thresholds ( $q = 0, q = 1$ ) for each policy have the same policy values as “Treat-all-with- $t$ ” ( $t = 0, t = 1$ ) policies. See Fig. 8 for an example of such a graph, and Appendix E for technical details.

Following the above analysis, some conclusions might be made: If “Treat-all-with- $t$ ” policies are the best policies, it might suggest a fixed treatment decision is better than using personalized policy. If the policy “Doctors” has the highest value, it suggests current practice would not be improved by the model. Finally, if the “propensity” baseline is the best performing, it suggests that a sort of “clinical wisdom of the crowds” phenomenon might be the best option, as the propensity score estimator aggregates the clinicians’ decision making process. Unlike the previous two cases, here one might move on with the propensity policy as a the basis for the treatment recommendation model going forward.

Following the selection of policy model, we recommend characterizing the different sub-populations affected by the policy recommendations. Specifically, we suggest examining: (1) patients in each treatment recommendation arm, and (2) patients whose recommendations are in agreement or disagreement with the actual care as reflected in the data. The characterization of treatment assignment policies is also known as policy summarization, which is a challenging task that is still an open area of research [93, 94]. We recommend characterizing the above-mentioned sets with the approaches described in Section 3.4.1.

Finally, we suggest plotting the *Average Potential Outcome* ( $\mathbb{E}[Y \mid i \in S_{sub}]$ ) of the above sub-populations  $S_{sub}$  and present it in a tree-shape known as *Outcome Tree*. See Fig. 6 for an example.

## 4 Case study: acute care

In this case study, we explore a clinical dilemma arising when treating hospitalized patients suffering from acute kidney injury (AKI) as a consequence of acute decompensated heart failure (ADHF), also known as type 1 cardiorenal syndrome (CRS-1). When treating patients suffering from CRS-1, clinicians grapple with a difficult choice: whether to prioritize decongestion or intravascular volume

and end-organ perfusion (very simplistically: whether to prioritize treating the heart condition or the kidney failure) [8–10]. This conundrum manifests itself in decisions around volume optimization therapies, such as adjusting the dosage of loop diuretics for patients. As of now, clear guidelines for managing this situation remain elusive [8, 95]. For additional information on ADHF and AKI, please see Appendix F. In the subsequent sections, we trace the steps of our proposed framework (Section 3) as applied to this clinical challenge.

#### 4.1 Target Recommendation System: Preconditions

The target recommendation system was formulated over several months of in-depth discussions with the clinical team, including on-site visits to observe patient treatment in practice. Consequently, in alignment with the preconditions specified in Section 3.1, we have framed the target system as follows: Our objective is to aid in determining the optimal diuretics dosage for patients hospitalized with ADHF who develop AKI during their hospital stay. We identify the decision point as the first time a laboratory result indicating a rise of  $> 0.3 \frac{mg}{dl}$  in serum creatinine from the baseline (admission) value [8–10], providing a 48-hour window for the treatment decision.

Two options for diuretic dosage are considered: (1) **“Decrease”**: reduction ( $T = 0$ ), or (2) **“Increase”**: maintenance or increase ( $T = 1$ ). The outcome we evaluated relates to renal function [95]: Specifically, renal function is evaluated by the percent return to levels of creatinine at baseline ( $RTB$ ), where 100% denotes complete return to baseline and 0% indicates no change [8]. In general higher RTB values are desirable, though other clinical factors could be taken into consideration. The RTB measure is calculated as follows:

$$RTB := \frac{crea_d - crea_o}{crea_d - crea_b}, \quad (6)$$

where  $crea_d$  is the lab result at the treatment decision point,  $crea_o$  is the lab result at the outcome point (the last measurement within a week from the decision point), and  $crea_b$  is the patient’s baseline creatinine value taken at time of hospital admission.

We note that as currently defined, our formulation violates SUTVA as there are two versions of treatment for  $T = 1$ : increasing or maintaining the dosage are not the same thing. Nonetheless, the two versions represent similar clinical intent of focusing on prioritizing decongestion by either maintaining or increasing the diuretic dosage. Thus, we can think of this clinical decision to focus on cardiac function as representing a closely aligned, if not the identical, treatment arm. In upcoming work, we look into a three-treatment-arm analysis of this decision.

#### 4.2 Target Recommendation System: Realization

During our on-site visits and discussions with clinicians, we have made efforts to identify any possible factors that might affect treatment decisions and outcomes. We believe that following this process we have identified a comprehensive set of confounders for the model. All factors mentioned by our clinical collaborators as potentially impacting their treatment decisions were documented in the hospital electronic health record system<sup>1</sup>.

Additionally, we identified covariates affecting the outcome for improved prediction accuracy, based on insights from a study on a related patient cohort [87] and based on domain expertise from our clinical partners (Section 3.2).

---

<sup>1</sup>The system is in-house and maintained by the hospital, and includes forms created in the past by the clinical team in order to collect data about ADHF and AKI patients



Consequently, we included approximately 200 covariates, covering all lab test results and timings, patient demographics, prior hospitalizations, other medications at admission and before the treatment decision, and the time elapsed from admission to the treatment decision. Missing data was imputed using the median of the training set, along with a missingness indicator. Time series data such as lab tests were used with summary statistics: mean, median, 10th and 90th percentiles, standard deviation, and first and last measurement for each signal from the admission time until the decision point. See a full list in Table S.1 in the appendix.

### 4.3 Data

We used data from a cohort of 6,940 patients, encompassing 12,027 patient visits, gathered between 2004 and 2015 at a large university medical center [87]. The eligibility criteria were patients admitted primarily due to heart failure, as per the European Society of Cardiology criteria [96–98], who during hospitalization experienced AKI. As stated above AKI was defined as any instance of serum creatinine increase exceeding  $> 0.3 \frac{mg}{dl}$  from the value at admission. Applying this criterion yielded a study cohort of 2,157 stays. This study follows the Declaration of Helsinki and has been approved by the institutional review committee on human research. See detailed description of the cohort in Table S.1 in the appendix.

We note that while we made the utmost effort to include all possible confounders, we believe that going into a prospective trial we will be able to further improve the quality of our dataset, especially by applying definitions of treatments which more precisely align with current care. We will explore these refinements in future work.

### 4.4 Estimation

In the following subsections, we describe the results following the outline in Section 3. Some further results are detailed in Appendix F.6. The data were split into train, validation, and test sets, containing 1305, 322, and 530 stays, respectively.

#### 4.4.1 Propensity model

For propensity score and overlap estimation (Section 3.3), we fit a model using XGBoost [74]. Table S.3 includes all the relevant metrics for the propensity model; in particular, we found the model to be well-calibrated, see Fig. 4.


#### 4.4.2 Propensity model insights

To gain a better understanding of the propensity model (following Section 3.4), we used SHAP [99] to extract the features most correlated with the predictions, see Fig. S.1. We then consulted with domain experts who checked the validity of our findings, approving that the top features are indeed likely to correlate with the clinical decision. Common support was selected based on the training data, and it contains subjects with propensity scores ranging from 0.21 to 0.9, see Fig. S.2.

Out of the 530 patients in the test set, 461 were in the region of common support.

#### 4.4.3 Semi-synthetic outcome simulation

As discussed in Section 3.5, we performed a simulation study whose goals are to ascertain whether causal effects can be uncovered from the data under reasonably favorable assumptions. Given the



**Figure 4:** Calibration curve of propensity estimator, using XGBoost, on the validation set.

data of patients in the common support region (Section 4.4.1), we simulated potential outcomes for both treatment arms. The generated outcomes were chosen such that the CATE is a linear function of the covariates; see Appendix C. Specifically, the linear CATE was designed to be a weighted average of the clinical decision maker’s implied assessment of the best treatment as given by a log-linear propensity score estimator, and an arbitrary vector. We run 6 simulation rounds with different seeds.

For the analysis, we used XGBoost, Ridge & Lasso Regression, and BART [47, 48], and fitted models separately on both treatment arms, a practice known as “T-Learner” (see Section 2.3). We also used causal forest [37] as a direct estimator of the CATE function. Moreover, we used three ensemble-based policies: “Average”, which takes the average outcome prediction of all ML models as the estimated potential outcome and determines the treatment based on it; “Majority”, which casts a vote where each ML model decides the treatment recommendation, and the majority vote determines the recommendation; and “Consensus”, which provides a recommendation only for patients on which all ML models agree regarding the treatment, deferring all other recommendations.

These policies were compared against the baselines “Doctors”, “Random”, “Propensity”, and “Treatment-all-with-t” (Section 3.11), as well as an “Optimal” policy (which we can derive since this is a simulation), as per Section 3.5.

In Section 3.5, we suggested 3 indications that are required from the simulation to test whether

Policy	$\hat{V}^{\text{DR}}(\pi)$	$\hat{V}^{\text{IPW}}(\pi)$	$V(\pi)$
XGBoost [74]	-0.378 (0.082)	-0.417 (0.069)	-0.392 (0.082)
Causal Forest [37]	-0.325 (0.098)	-0.365 (0.081)	-0.334 (0.096)
Ridge	-0.442 (0.083)	-0.472 (0.080)	-0.432 (0.074)
Lasso	-0.447 (0.098)	-0.488 (0.093)	-0.450 (0.095)
BART [47]	-0.405 (0.088)	-0.430 (0.082)	-0.397 (0.073)
Average	-0.446 (0.094)	-0.469 (0.088)	-0.451 (0.087)
Majority	-0.372 (0.107)	-0.795 (0.049)	-0.339 (0.087)
Consensus	-0.444 (0.091)	-0.478 (0.085)	-0.442 (0.084)
Propensity	0.079 (0.104)	0.063 (0.130)	0.067 (0.101)
Increase all	0.1 (0.115)	0.096 (0.105)	0.116 (0.116)
Decrease all	-0.16 (0.074)	-0.192 (0.072)	-0.136 (0.063)
Doctors	-0.081 (0.077)	-0.081 (0.077)	-0.081 (0.077)
<b>Optimal</b>	-	-	<b>-0.524 (0.074)</b>

**Table 1:** Mean and SEM of simulation results over 6 randomization runs. For each policy we provided the Doubly Robust (DR)  $\hat{V}^{\text{DR}}(\pi)$  and Inverse Propensity Weighting (IPW)  $\hat{V}^{\text{IPW}}(\pi)$  estimated policy value, along with the real policy values  $V(\pi)$ . We also provide the optimal policy value, based on the simulated potential outcomes values. Lower value is better.

the data is suitable for the task at hand. In the rest of the sub-section, we analyze the results, presented in Table 1 and Fig. 5, in light of those goals.

First, the results indicate that the estimated policy value is close to its real policy value. This suggests that, under our current framework and evaluation method, the estimated policy value is similar to the real policy value. In Fig. 5 we can further see that DR policy estimation is better aligned with the actual policy.


We note that using the same CATE estimator as a plug-in estimator for policy value estimation can cause a bias in estimation, known as *congeniality bias* [84]. The results in Fig. 5 further indicate that the DR estimator has limited bias and gives a good estimation of the policy value.

We further observe that most policies outperform the current practice (“Doctors”), with the exceptions of causal forest, propensity, and majority vote. Unsurprisingly, policies that have linear underlying estimators (Ridge, Lasso) perform well, in tune with the linear nature of the outcome generation process. Along with XGBoost and BART, these policies yielded results that approximate the optimal policy.

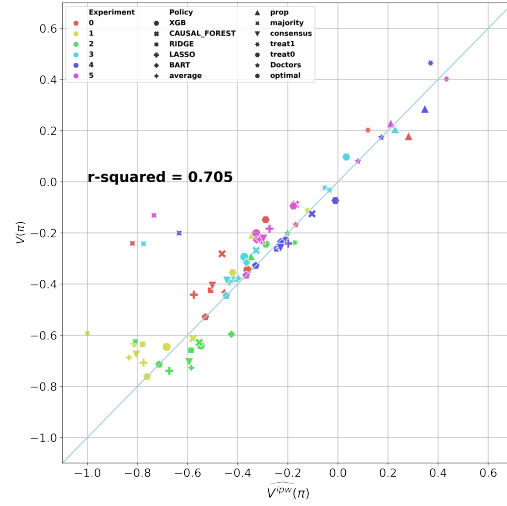
Based on these results, we conclude that our data is sufficiently robust and our modeling approach is appropriate for the task. Further details can be found in Appendix F.3.

#### 4.4.4 Outcome and CATE model

For the CATE estimator models (Section 3.6), we used the same methods as in Section 4.4.3, and added DragonNet [46] as a direct estimator of the CATE function. Although our outcome of interest is *RTB* of creatinine, we note it is a noisy estimand. Looking into the definition of *RTB* (Eq. (6)), we see that “outcome creatinine” ( $crea_o$ ) is the only unknown measurement at decision time. Therefore, we set  $crea_o$  to be the outcome of interest for models at this stage, and transform it using Eq. (6) for downstream policy value estimation (Section 4.4.7). In Table 2, we present the prediction models’ performance metrics. The results suggest satisfactory performance, with slight



(a) Simulation: Estimated policy values using doubly robust (DR) estimation vs. the true policy values.



(b) Simulation: Estimated policy values using inverse propensity weighting (IPW) estimation vs. the true policy values.

**Figure 5:** Scatter plot of the true policy value versus the estimated policy values, using DR (Fig. 5a) and IPW (Fig. 5b) methods. The graph represents 6 simulation runs, using multiple policies: T-learner of XGboost (“XGB”), Ridge (“RIDGE”), Lasso (“LASSO”), and BART (“BART”), Causal Forest as a direct estimator (“CAUSAL\_FOREST”), average (“average”), majority (“majority”) and consensus (“consensus”) as Ensemble-based policies, and the baseline policies – propensity based policy (“prop”), “Increase” to all (“treat1”), “Decrease” to all (“treat0”), Current Policy (“Doctors”) and the optimal policy.

over-fitting for the BART model. It is notable from the results that “Increase” prediction task was more difficult. This finding may stem from the diverse nature of the “Increase” arm, which includes both increase and stay with the same dosage. In future work, we intend to model them as 2 separate arms.

Models			RMSE	MAE	$R^2$	std
$T = 0$ : <b>Decrease</b>	Train	BART	0.533	0.346	0.824	1.272
		XGB	0.593	0.348	0.782	
		RIDGE	0.744	0.455	0.657	
		LASSO	0.754	0.433	0.648	
	Validation	BART	0.808	0.538	0.58	1.251
		XGB	0.771	0.504	0.618	
		RIDGE	0.772	0.491	0.617	
		LASSO	0.773	0.48	0.615	
$T = 1$ : <b>Increase</b>	Train	BART	0.6	0.411	0.781	1.282
		XGB	0.659	0.44	0.735	
		RIDGE	0.688	0.477	0.712	
		LASSO	0.808	0.533	0.602	
	Validation	BART	0.988	0.573	0.588	1.545
		XGB	0.99	0.556	0.587	
		RIDGE	1.059	0.646	0.527	
		LASSO	1.021	0.581	0.56	

**Table 2:** Performance metrics for outcome prediction models, in terms of Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and the correlation coefficient ( $R^2$ ). We present the results for the T-learner models: BART, XGBoost (“XGB”), Ridge, and Lasso.

#### 4.4.5 Outcome and CATE model insights

In Appendix F.4 we present several analysis results (Following Section 3.7). First, we show the model’s error distribution along with an analysis of the important features of each model. We find these closely align with current medical knowledge. Furthermore, we outline the correlations between the CATE models, which indicates we have a positive, while small, correlation between the policies. See Tables S.4 and S.5 in the appendix. Additionally, we show the CATE calibration graph in Fig. S.7 in the appendix, which indicates our CATE estimates are calibrated in the sense presented in Athey and Wager [88].

#### 4.4.6 Deferral set

As explained in Section 3.8, our framework suggests only providing recommendations for patients for whom we have higher certainty that the recommended treatment will be beneficial. For patients for whom there is high uncertainty, we defer making a recommendation. The initial set of deferrals includes patients that were not part of the propensity-overlap region (see Section 4.4.1). Another set of patients to be deferred are patients that have high – statistical or causal – uncertainty for their recommendation. We obtain this using Quince [28], with Dragonnet [46] as an underlying

CATE estimator. Using a causal sensitivity parameter of  $\exp(0.1)$  and a statistical point estimate we end up with an additional 139 (out of 481) deferred patients.

Therefore, we have two sets of deferrals: “*Inclusive*”, which is based only on the common support set, and “*Conservative*”, which includes both common support and high uncertainty deferrals. In Appendix F.5 we present insights on the trimmed patients, following Section 3.9.

#### 4.4.7 Policy

Following the CATE model results we established treatment policies following Section 3.10 using multiple different CATE estimators. For policy value estimation, we used XGBoost (XGB) as our plug-in estimator (see Appendix D), and the outcome is given in RTB terms, as described in Section 4.1.

In the next section, we present the results based on the “*Inclusive*” set. In Fig. S.10 and Fig. S.11 we further present results for the “*Conservative*” set.

#### 4.4.8 Policy Insights

Having established our candidate policies, we evaluate them on the actual data and compare them with current policy as performed by the clinicians at the hospital. The results given in this section are on held-out data, taken over 10K bootstrap rounds.

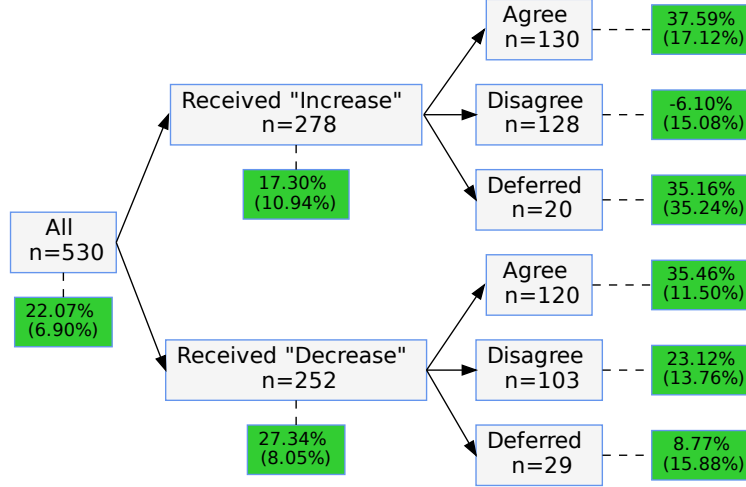
As a first evaluation phase, we investigated the breakdown of our policy recommendation versus the current treatment. Each of the treatment arms was divided into three bins – whether our policy agrees with the observed treatment, disagrees, or was deferred.

Fig. 6 shows the results in the form of an *Outcome Tree*, as suggested in Section 3.11, with an XGBoost T-learner as the basis for the policy. The analysis suggests that in cases where the algorithm agrees with the current treatment, the patients have a mean RTB of 37% in “*Increase*” and 35.46% “*Decrease*” arms. Those results are better than the cases where the policy disagrees, and better than the overall results in their respective arms.

Next, we evaluate the policy value. The results are presented in Table 3. They indicate that the policy value of most models is better than the current treatment, as seen when using both DR and IPW policy value estimators. For example, it seems from the evaluation that using either Ridge (45.8%) or XGB (40.9%) as outcome models would almost double the RTB ratio compared to the current treatment policy value (22.1%).

We further present the policy values in box-plots in Fig. 7. For clarity, we show only two best-performing policies (Ridge and XGBoost), the full results are in Appendix F.6 (Fig. S.9 and Fig. S.10). The results indicate that our recommended policies yield better value to the patients compared to current care, where there is no intersection between their respective inter-quartile ranges. Interestingly, the results suggest that treating all patients with *decrease* will result in better outcomes than the current treatment policy, yielding a value similar to both ridge and XGB. We discuss this finding in the end of the section.

In Fig. 8 we show the rank graph; for clarity we again show only the two best performing policies (Ridge and XGBoost). We can see that for most given thresholds, applying a ridge- or XGBoost-based policy will yield better performance than the current treatment policy across most quantiles, where ridge-based is statistically better with 95% confident interval. The analysis suggests, similar to Fig. 7, that we cannot reject the possibility that using a model based policy has the same value as simply decreasing dosages for all patients – we discuss this point below.



**Figure 6:** Outcome Tree: A graph representing the mean RTB value for each policy group, using XGBoost T-learner as the basis for the policy. The first split (left) is by actual treatment received. The second split (right) is by whether the policy agreed or disagreed with the actual treatment, or whether it deferred. The gray boxes represent the number of patients in each subgroup, and the green boxes represent their mean (sem) RTB value.

In Table 4, we present the result of a direct comparison of policy performances across bootstrap rounds. This analysis reveals that the Ridge and XGBoost policies outperform the existing treatment in 9972 and 9850 out of 10000 rounds, respectively. Additionally, these policies surpass most others in performance. Notably, however, we see here too that these policies advantage over the “Decrease” policy are not statistically significant.

Overall, our analysis suggests that a policy that recommends decreasing the dosage for all patients would yield a similar performance to our model-based policy, in terms of patient RTB. While the performance is similar, we note that our recommendations decrease the dosage only for about 43% - 55% of the patients which allows more flexibility with respect to other outcomes, which we present next.

**Patient Re-Hospitalization** As outlined in the opening of Section 4, the clinician’s treatment decision is based on trying to maintain a delicate balance between kidney and cardiac function. Thus, we further evaluate our policies on cardiac function-related outcomes. The cardiac function is estimated by checking if the patient was re-hospitalized due to AHF within 30 days of the treatment decision, which would indicate a failure in managing the patient’s heart condition.

We evaluated the policy values using the same approaches described above, where for the DR policy value estimation we used an  $L_2$  regularized Logistic-Regression trained to predict 30-day re-hospitalization on the train split, as a plug-in estimator.

We present the re-hospitalization policy-value box-plot in Fig. 9. The results indicate that the



Model	$\hat{V}^{\text{DR}}(\pi)$	$\hat{V}^{\text{IPW}}(\pi)$
Ridge	45.8%	43.8%
XGBoost	40.9%	40.9%
DragonNet	37.0%	38.3%
Lasso	31.2%	31.5%
Causal Forest	28.8%	28.4%
BART	20.2%	21.5%
Average	35.2%	35.4%
Majority	33.4%	31.4%
Consensus	33.9%	33.5%
Propensity	26.4%	32.1%
Random	22.8%	26.5%
Increase all	13.4%	13.3%
Decrease all	37.6%	41.6%
Doctors	22.1%	22.1%

**Table 3:** The policy values of using different policies, using both Doubly Robust (DR) and Inversed Propensity Weighting (IPW) estimators. The results are separated by their type – first, the estimated policies, second the ensemble policies, followed by the simple baselines.

XGBoost-based policy we obtained earlier achieves a lower re-hospitalization rate than current care, and similar to the “Increase” policy. Furthermore, unlike the analysis of the main outcome, here the “Decrease” policy performs worse than current care and is similar to the Ridge-based policy. This, combined with the results on the main outcome (Fig. 7), suggests that XGBoost-based policy might balance the two outcomes better than current care and other baselines, improving both kidney and cardiac function.


## 5 Related work

In recent years, several authors sought to give guidelines on how to learn treatment effect and produce a treatment policy from data. In this section, we will provide a short overview of those works and our added contribution.

**Observational studies** Powell et al. [3] and Dang et al. [100] suggested guidelines for using observational data for estimation of average treatment effects. Those works share a similar perspective to ours regarding causal identification and the importance of cross-discipline collaboration. However, our focuses on establishing a policy for treatment recommendation on the individual level, rather than estimating a single overall average effect.

As mentioned in Section 2.3, estimation of *Heterogeneous Treatment Effect* (HTE) from both observational and experimental data is an active area of research [5, 101–108]. It mainly focuses on identifying subgroups with differing causal effects. As in the ATE case, our work is using tools from the HTE literature, but our focus on is on estimating a useful individualized policy rather than estimating effects for sub-groups.

Sharma and Kiciman [109] developed a Python package and a framework to estimate causal effects, including in observational cases, in 4 stages closely resembling our suggested framework. While this framework suggests many evaluation methods, they are mostly related to ATE and do




**Figure 7:** Policy value box-plot, the result of running 10K bootstraps evaluation on held-out data, showing 6 policies: *Current*: current treatment, *Random*: randomly assigning treatment at the same proportion as current treatment, *Keep/Increase*: all patients given “increase”, *Decrease*: all patient given “decrease”, *XGboost*: T-learner of XGBoost models, *Ridge*: T-learner of Ridge models.

not take into consideration the deferral option in policy, and lack guidelines for how to conduct the identification phase. Their framework is not focused on using healthcare data and recommending actions to clinicians.

**Individual Treatment Rules (ITR)** As explained in Section 2.5, ITR is a common name for a learned function  $\pi(X)$  that assigns a treatment to each individual based on their characteristics.

Boominathan et al. [110] have suggested a method for forming a policy from data in cases where we know all the outcomes (whether based on biological knowledge, advanced lab tests, etc.), but they are not available to the decision maker during treatment assignment. This setting does not apply to our case. Moreover, the authors do not tackle the identification and evaluation steps we suggest in this work. Meid et al. [4] suggested a general framework to learn ITR in cases where HTE has been found for a given treatment, based on effect modifiers. While their work has similar settings to ours, we go further from the simulation study, and provided the additional required steps of identification and validation for this task. Moreover, we provided steps in cases of high uncertainty and deferral.

Meid et al. [111] presented a real-word case study, using CATE to recommend personalized treatment. Following similar ideas to our suggested framework, their work explored the estimation and some of the evaluation steps. Our work expands on this by introducing analysis and validation



**Figure 8:** Policy value rank graph, result of running 10K bootstraps evaluation on held-out data, showing two T-learner models: XGBoost (“XGB”) and Ridge (“RIDGE”), compared with “Doctors” treatment policy (“doctors”). Each model “\_rank” markers (e.g., “RIDGE Rank”) represent the policy value of providing “decrease” recommendation to exactly N% of the population, based on a ranking of the estimated CATE values given by the model. We used a step-size of 2%. We further compared it to randomly ranking patients policy (“random\_rank”). The left and right endpoints represent the policy value of “Treat-all-with-t” policies, respectively. See explanation in Section 3.11.

phases for each stage, and our focus on identification. Moreover, our framework includes a deferral option for high uncertainty or weak overlap cases.

Petersen and van der Laan [112] suggested a theoretical framework (“Causal Roadmap”) on how to think about estimating causal parameters, as an alternative to the “Target Trial” [62, 14, 1, 63] discussed in Section 3.1. Their framework focuses on selecting the most suitable model in reference to the presumed data-generating process and causal assumptions.

In companion papers, Montoya et al. [61] and Montoya et al. [56] demonstrate an application of this theoretical framework to establish ITR, in RCT settings. In Montoya et al. [61] they consider different treatment estimation models and policy learning, again under RCT settings. Following that, Montoya et al. [56] explore policy value estimation. We differ from this work as we suggest a framework for learning a policy in the observational settings with a focus on *causal identification*, and in addition our discussion of validation steps and the role of the deferral option.

	BART	CF	Doc	LASSO	Prop	RIDGE	XGB	Avg	Con	Rand	Maj	DN	Dec	Inc
BART	-	2803	2434	332	1752	21	21	38	14	2124	538	138	343	5820
CF	7197	-	4956	2408	3448	273	538	1454	1408	4342	2232	1349	479	7775
Doc	7566	5044	-	1356	2213	28	150	701	559	4132	1064	383	365	8513
LASSO	9668	7592	8644	-	7054	456	772	2025	2414	7559	6818	2310	2798	9679
Prop	8248	6552	7787	2946	-	141	567	1679	1692	5855	4271	1259	1040	8780
RIDGE	9979	9727	9972	9544	9859	-	7036	9007	9418	9835	9905	8839	8058	9988
XGB	9979	9462	9850	9228	9433	2964	-	8261	8816	9528	9735	7239	6339	9962
Avg	9962	8546	9299	7975	8321	993	1739	-	6170	8702	8525	4447	4325	9896
Con	9986	8592	9441	7586	8308	582	1184	3830	-	8678	8668	3654	3846	9918
Rand	7876	5658	5868	2441	4145	165	472	1298	1322	-	3186	1041	1113	8554
Maj	9462	7768	8936	3182	5729	95	265	1475	1332	6814	-	1177	945	9691
DN	9862	8651	9617	7690	8741	1161	2761	5553	6346	8959	8823	-	4692	9920
Dec	9657	9521	9635	7202	8960	1942	3661	5675	6154	8887	9055	5308	-	9487
Inc	4180	2225	1487	321	1220	12	38	104	82	1446	309	80	513	-

**Table 4:** This table represents the direct comparison of policies, where columns and rows represent the different policies, and each cell represents how many bootstrap rounds out of 10,000 was the policy value of the “row” policy better than the “column” policy. For example, the T-learner policy with XGB (“XGB”) had a higher policy value 9,850 times (out of 10k rounds) than the current policy (“Doc”), which is 98.5% of the rounds. The models are: T-learners with backbone ML models – BART, Lasso, Ridge, and XGBoost (“XGB”), Direct methods – Causal Forest (“CF”), DragonNet (“DN”), baseline methods – average (“Avg”), consensus (“Con”), majority (“Maj”), Random (“Rand”), and policies representing providing “Increase” and “Decrease” to all (“Inc” and “Dec” respectively), alongside model based on the propensity score (“Propensity”) and current treatment (“Doc”).

## 6 Discussion


In this work we propose a framework for generating patient-level treatment recommendation models based on patient health data, including observational data. We provide detailed steps and propose how to validate them, with the goal of developing a reliable and robust treatment recommendation policy.

Rather than prescribing a specific causal effect estimation method, we give a general guideline for developing treatment policies. We recognize that various estimation methods may be more suitable depending on the target system’s characteristics and complexity. We applied our framework to a real-world treatment recommendation challenge in acute healthcare settings. As outlined in Section 4.4.7, our framework yields promising results, suggesting that the learned treatment policy may outperform current care and lead both to better renal functions and lower re-hospitalization rates.

A notable challenge in healthcare today is the gap between the rapid development of models and their actual implementation in bedside settings [6, 113–119]. Although our work does not directly tackle this gap, it is motivated by it. Our framework aims to help taking a step towards closing the gap in two key ways. First, we explicitly focus on actionable recommendations as part of the treatment process. Second, given the high stakes of deploying systems in healthcare settings, our framework emphasizes ways of mitigating risks by prioritizing responsible, expert-in-the-loop model development and testing.

While those steps are not exhaustive, we believe they can contribute to addressing the challenges inherent in this gap. We recognize that our rigorous approach may not be necessary in every scenario. For example, in lower-stakes environments such as e-commerce recommendation systems where speed is crucial, some steps of our framework could be omitted.

While the framework provides promising results, it is essential to check its usefulness “in the



**Figure 9:** Policy value box-plot, a result of running 10K bootstraps evaluation on held-out data, in terms of re-hospitalization 30 days from decision-point. The DR policy value where estimated using  $L_2$ -regularized logistic regression estimator. The policies are: *Current*: current treatment, *Random*: randomly assigning treatment at the same proportion as current treatment, *Keep/Increase*: all patients given “increase”, *Decrease*: all patient given “decrease”, *XGboost*: T-learner of XGBoost models, *Ridge*: T-learner of Ridge models.

wild” [120], where conditions may differ. Testing such systems in real-world settings remains an active area of research [121, 122]. The ultimate test for the type of systems we discuss here would be a clinical trial in the field, comparing clinical outcomes when recommendations are available vs. when they are not. This requires real-time integration with the hospital EHR, as the system will need access to the full set of covariates in order to make real-time recommendations. Moreover, questions of generalization among different hospitals and health settings will also come into play here [123].

Looking forward, we believe that incorporating more advanced estimation methods should improve the accuracy of our policy. Additionally, improvements such as leveraging multi-modal data, better utilization of time-series information, accounting for survival bias, and leveraging mechanistic knowledge whenever possible can contribute to refining the model’s performance.

For instance, we note that the analysis of the patients’ re-hospitalization might suffer from competing risks or survival bias: some patients might have died during the 30 days from the decision point, and thus their potential outcome of re-hospitalization is invalid [124]. Moreover, when accounting for secondary outcomes, such as re-hospitalization, the confounders might differ

from those pertaining to the main outcome.

More broadly, there are some crucial areas which our framework does not directly address. First, multiple or continuous treatments, while sharing many similarities with binary treatments, still require special care especially when it comes to propensity score modeling and the accompanying deferral rules. A more difficult challenge is balancing multiple outcomes; for example, in our cardio-renal case, clinicians wish to optimize both renal function and cardiac functions, goals which might be at odds, and a balance which is difficult to quantify. Notions such as the win-ratio [125, 126] might be relevant in such cases

Another extension to the framework would be developing dynamic treatment recommendation policies [127, 118]. These would support physicians beyond the initial decision-making stage. In cases such as the cardio-renal scenario we discuss, physicians could benefit from guidance on managing the patient’s entire treatment plan, especially after the onset of Acute Kidney Injury (AKI). This is a challenging task, as each decision increases the complexity of possible treatments and reduces the overlap between the patients’ journeys [128].

A further promising research effort is modeling directly the deferral decision, as opposed to our current uncertainty based approach, which might be sub-optimal in terms of maximizing the combined expert-model decision making [129–132]. Downstream from our recommendation system is the question of human-computer interaction, and how to present recommendations to clinicians effectively [133–135].

Finally, beyond the distribution shift due to the difference between retrospective and prospective environments and across hospitals mentioned above, we must also consider shifts caused by the presence of the treatment recommendation system itself [136, 137]. This might change the entire treatment regime, beyond the decision point, and thus may impact the estimation accuracy and validity of policy estimation.

## References

- [1] 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. ISSN 14766256. doi: 10.1093/aje/kwv254.
- [2] Ioana Bica, Ahmed M. Alaa, Craig Lambert, and Mihaela van der Schaar. From Real-World Patient Data to Individualized Treatment Effects Using Machine Learning: Current and Future Methods to Address Underlying Challenges. *Clinical Pharmacology & Therapeutics*, 109(1):87–100, 1 2021. ISSN 1532-6535. doi: 10.1002/CPT.1907.
- [3] Michael Powell, Allison Koenecke, James Brian Byrd, Akihiko Nishimura, Maximilian F. Konig, Ruoxuan Xiong, Sadiqa Mahmood, Vera Mucaj, Chetan Bettgowda, Liam Rose, Suzanne Tamang, Adam Sacarny, Brian Caffo, Susan Athey, Elizabeth A. Stuart, and Joshua T. Vogelstein. Ten Rules for Conducting Retrospective Pharmacoepidemiological Analyses: Example COVID-19 Study. *Frontiers in Pharmacology*, 12:1799, 7 2021. ISSN 16639812. doi: 10.3389/fphar.2021.700776.
- [4] Andreas D. Meid, Carmen Ruff, Lucas Wirbka, Felicitas Stoll, Hanna M. Seidling, Andreas Groll, and Walter E. Haefeli. Using the causal inference framework to support individualized drug treatment decisions based on observational healthcare data. *Clinical Epidemiology*, 12: 1223–1234, 2020. ISSN 11791349. doi: 10.2147/CLEP.S274466.
- [5] David M. Kent, Ewout Steyerberg, and David Van Klaveren. Personalized evidence based medicine: Predictive approaches to heterogeneous treatment effects. *BMJ (Online)*, 363, 2018. ISSN 17561833. doi: 10.1136/bmj.k4245.
- [6] Bar Eini-Porat, Ofra Amir, Danny Eytan, and Uri Shalit. Tell me something interesting: Clinical utility of machine learning prediction models in the icu. *Journal of Biomedical Informatics*, page 104107, 2022.
- [7] Mattia Proserpi, Yi Guo, Matt Sperrin, James S. Koopman, Jae S. Min, Xing He, Shan-nan Rich, Mo Wang, Iain E. Buchan, and Jiang Bian. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nature Machine Intelligence*, 2(7): 369–375, 7 2020. ISSN 25225839. doi: 10.1038/s42256-020-0197-y.
- [8] Jubran Boulos, Wisam Darawsha, Zaid A. Abassi, Zaher S. Azzam, and Doron Aronson. Treatment patterns of patients with acute heart failure who develop acute kidney injury. *ESC Heart Failure*, 6(1):45–52, 2 2019. ISSN 20555822. doi: 10.1002/ehf2.12364.
- [9] Kevin Damman, Mattia A.E. Valente, Adriaan A. Voors, Christopher M. O’Connor, Dirk J. Van Veldhuisen, and Hans L. Hillege. Renal impairment, worsening renal function, and outcome in patients with heart failure: An updated meta-analysis. *European Heart Journal*, 35(7):455–469, 2 2014. ISSN 0195668X. doi: 10.1093/eurheartj/eh386.
- [10] WH Wilson Tang and Wilfried Mullens. Cardiorenal syndrome in decompensated heart failure. *Heart*, 96(4):255–260, 2010.
- [11] Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974. ISSN 00220663. doi: 10.1037/h0037350.



- [12] Paul W. Holland and Donald B. Rubin. Causal inference in retrospective studies. *ETS Research Report Series*, 1987(1):203–231, 6 1987. ISSN 2330-8516. doi: 10.1002/j.2330-8516.1987.tb00211.x.
- [13] Judea Pearl. *Causality*. Cambridge university press, illustrate edition, 2009. ISBN 9780521895606.
- [14] Miguel A Hernán. With great data comes great responsibility: Publishing comparative effectiveness research in epidemiology. *Epidemiology*, 22(3):290–291, 2011. ISSN 10443983. doi: 10.1097/EDE.0b013e3182114039.
- [15] Uri Shalit. Can we learn individual-level treatment policies from clinical data? *Biostatistics*, 21(2):359–362, 11 2019. ISSN 1465-4644. doi: 10.1093/biostatistics/kxz043.
- [16] Jingpu Shi and Beau Norgeot. Learning causal effects from observational data in healthcare: A review and summary. *Frontiers in Medicine*, page 2027, 2022.
- [17] Carlos Fernández-Loría and Foster Provost. Causal Decision Making and Causal Effect Estimation Are Not the Same... and Why It Matters. *arXiv preprint arXiv:2104.04103*, 2021.
- [18] MA Hernan and J Robins. *Causal Inference: What if*. Chapman & Hill/CRC, Boca Raton, 2020.
- [19] Victor Chernozhukov, Christian Hansen, Nathan Kallus, Martin Spindler, and Vasilis Syrgkanis. Applied causal inference powered by ml and ai. *arXiv preprint arXiv:2403.02467*, 2024.
- [20] Stefan Wager. Causal inference: A statistical learning approach, 2024.
- [21] Guido W Imbens. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and Statistics*, 86(1):4–29, 2004. ISSN 00346535. doi: 10.1162/003465304323023651.
- [22] 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, 4 1983. ISSN 00063444. doi: 10.1093/biomet/70.1.41.
- [23] Donald B Rubin. Randomization Analysis of Experimental Data: The Fisher Randomization Test Comment. *Journal of the American Statistical Association*, 75(371):591–593, 1980.
- [24] Stephen R. Cole and Miguel A. Hernán. Constructing inverse probability weights for marginal structural models. *American Journal of Epidemiology*, 168(6):656–664, 7 2008. ISSN 00029262. doi: 10.1093/aje/kwn164.
- [25] Zhiqiang Tan. A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 101(476):1619–1637, dec 2006. doi: 10.1198/016214506000000023.
- [26] Steve Yadowlsky, Hongseok Namkoong, Sanjay Basu, John Duchi, and Lu Tian. Bounds on the conditional and average treatment effect with unobserved confounding factors. *The Annals of Statistics*, 50(5):2587–2615, 2022.

- [27] Nathan Kallus, Xiaojie Mao, and Angela Zhou. Interval estimation of individual-level causal effects under unobserved confounding. In *The 22nd international conference on artificial intelligence and statistics*, pages 2281–2290. PMLR, 2019.
- [28] Andrew Jesson, Sören Mindermann, Yarin Gal, and Uri Shalit. Quantifying ignorance in individual-level causal-effect estimates under hidden confounding. In *International Conference on Machine Learning*, pages 4829–4838. PMLR, 2021.
- [29] Ying Jin, Zhimei Ren, and Emmanuel J Candès. Sensitivity analysis of individual treatment effects: A robust conformal inference approach. *arXiv preprint arXiv:2111.12161*, 2021.
- [30] Mingzhang Yin, Claudia Shi, Yixin Wang, and David M Blei. Conformal sensitivity analysis for individual treatment effects. *Journal of the American Statistical Association*, pages 1–14, 2022.
- [31] Miruna Oprescu, Jacob Dorn, Marah Ghoummaid, Andrew Jesson, Nathan Kallus, and Uri Shalit. B-learner: Quasi-oracle bounds on heterogeneous causal effects under hidden confounding. *arXiv preprint arXiv:2304.10577*, 2023.
- [32] Peter C. Austin. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate Behavioral Research*, 46(3):399–424, 5 2011. ISSN 00273171. doi: 10.1080/00273171.2011.568786.
- [33] Sören R. Künnel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences of the United States of America*, 116(10):4156–4165, 2019. ISSN 10916490. doi: 10.1073/pnas.1804597116.
- [34] Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319, 5 2021. ISSN 0006-3444.
- [35] Edward H Kennedy. Towards optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint arXiv:2004.14497*, 2020.
- [36] Gabriel Okasa. Meta-learners for estimation of causal effects: Finite sample cross-fit performance. *arXiv preprint arXiv:2201.12692*, 2022.
- [37] Stefan Wager and Susan Athey. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 7 2018. ISSN 1537274X. doi: 10.1080/01621459.2017.1319839.
- [38] Leo Breiman. Random Forests. *Machine Learning*, 45(1):5–32, 2001. ISSN 08856125. doi: 10.1023/A:1010933404324.
- [39] Scott Powers, Junyang Qian, Kenneth Jung, Alejandro Schuler, Nigam H. Shah, Trevor Hastie, and Robert Tibshirani. Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in Medicine*, 37(11):1767–1787, 5 2018. ISSN 10970258. doi: 10.1002/sim.7623.
- [40] Susan Gruber and Mark J van der Laan. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *The international journal of biostatistics*, 6 (1):Article 26, 2010. ISSN 1557-4679. doi: 10.2202/1557-4679.1260.

- [41] Mark J. van der Laan and Alexander R. Luedtke. Targeted Learning of the Mean Outcome under an Optimal Dynamic Treatment Rule. *Journal of Causal Inference*, 3(1), 10 2015. ISSN 2193-3677. doi: 10.1515/jci-2013-0022.
- [42] Li-Fang Cheng, Bianca Dumitrascu, Michael Zhang, Corey Chivers, Michael Draugelis, Kai Li, and Barbara E. Engelhardt. Patient-Specific Effects of Medication Using Latent Force Models with Gaussian Processes. *arXiv preprint arXiv:1906.00226*, 2019.
- [43] Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *34th International Conference on Machine Learning, ICML 2017*, volume 6, pages 4709–4718, 2017. ISBN 9781510855144.
- [44] Fredrik D. Johansson, Nathan Kallus, Uri Shalit, and David Sontag. Learning Weighted Representations for Generalization Across Designs. *arXiv preprint arXiv:1802.08598*, 2 2018.
- [45] Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal Effect Inference with Deep Latent-Variable Models. In I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 6446–6456. Curran Associates, Inc., 5 2017.
- [46] Claudia Shi, David M Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. *arXiv preprint arXiv:1906.02120*, 2019. ISSN 23318422.
- [47] Hugh A Chipman, Edward I George, and Robert E. McCulloch. BART: Bayesian additive regression trees. *Annals of Applied Statistics*, 4(1):266–298, 2010. ISSN 19326157. doi: 10.1214/09-AOAS285.
- [48] Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011. ISSN 10618600. doi: 10.1198/jcgs.2010.08162.
- [49] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.
- [50] Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.
- [51] Lihua Lei and Emmanuel J Candès. Conformal inference of counterfactuals and individual treatment effects. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 83(5):911–938, 2021.
- [52] Steve Yadlowsky, Hongseok Namkoong, Sanjay Basu, John Duchi, and Lu Tian. Bounds on the conditional and average treatment effect with unobserved confounding factors. *arXiv preprint arXiv:1808.09521*, 8 2018.
- [53] Andrew Jesson, Sören Mindermann, Uri Shalit, and Yarin Gal. Identifying causal-effect inference failure with uncertainty-aware models. *Advances in Neural Information Processing Systems*, 33, 2020.

- [54] Min Qian and Susan A. Murphy. Performance guarantees for individualized treatment rules. *The Annals of Statistics*, 39(2):1180–1210, 4 2011. ISSN 0090-5364. doi: 10.1214/10-aos864.
- [55] Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4):485–511, 2014.
- [56] Lina M Montoya, Mark J van der Laan, Jennifer L Skeem, and Maya L Petersen. Estimators for the value of the optimal dynamic treatment rule with application to criminal justice interventions. *The International Journal of Biostatistics*, 2022.
- [57] Dimitris Bertsimas, Agni Orfanoudaki, and Rory B. Weiner. Personalized Treatment for Coronary Artery Disease Patients: A Machine Learning Approach. *arXiv preprint arXiv:1910.08483*, 10 2019.
- [58] Susan Athey and Stefan Wager. Efficient Policy Learning. *arXiv preprint arXiv:1702.02896*, 2 2017.
- [59] Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1):133–161, 2021.
- [60] Nathan Kallus. Balanced policy evaluation and learning. In *Advances in Neural Information Processing Systems*, volume 2018-Decem, pages 8895–8906, 2018.
- [61] Lina M Montoya, Mark J van der Laan, Alexander R Luedtke, Jennifer L Skeem, Jeremy R Coyle, and Maya L Petersen. The optimal dynamic treatment rule superlearner: considerations, performance, and application to criminal justice interventions. *The International Journal of Biostatistics*, 2022.
- [62] Donald B Rubin. Teaching statistical inference for causal effects in experiments and observational studies. *Journal of Educational and Behavioral Statistics*, 29(3):343–367, 2004.
- [63] Barbra A Dickerman, Xabier García-Albéniz, Roger W Logan, Spiros Denaxas, and Miguel A Hernán. Avoidable flaws in observational analyses: an application to statins and cancer. *Nature medicine*, 25(10):1601–1606, 2019.
- [64] Mark P. Sendak, Joshua D’Arcy, Sehj Kashyap, Michael Gao, Marshall Nichols, Kristin Corey, William Ratliff, and Suresh Balu. A Path for Translation of Machine Learning Products into Healthcare Delivery. *EMJ Innovations*, 1 2020. ISSN 2513-8634. doi: 10.33590/emjinnov/19-00172.
- [65] Peter WG Tennant, Eleanor J Murray, Kellyn F Arnold, Laurie Berrie, Matthew P Fox, Sarah C Gadd, Wendy J Harrison, Claire Keeble, Lysie R Ranker, Johannes Textor, et al. Use of directed acyclic graphs (dags) to identify confounders in applied health research: review and recommendations. *International journal of epidemiology*, 50(2):620–632, 2021.
- [66] Herbert A Simon. A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118, 1955.
- [67] Andre W Kushniruk. Analysis of complex decision-making processes in health care: cognitive approaches to health informatics. *Journal of biomedical informatics*, 34(5):365–376, 2001.

- [68] Vimla L Patel, David R Kaufman, and Jose F Arocha. Emerging paradigms of cognition in medical decision-making. *Journal of biomedical informatics*, 35(1):52–75, 2002.
- [69] Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- [70] Kévin Duarte, Jean-Marie Monnez, Eliane Albuissou, Bertram Pitt, Faiez Zannad, and Patrick Rossignol. Prognostic value of estimated plasma volume in heart failure. *JACC: Heart Failure*, 3(11):886–893, 2015.
- [71] Peng Ding, TJ VanderWeele, and James M Robins. Instrumental variables as bias amplifiers with general outcome and confounding. *Biometrika*, 104(2):291–302, 2017.
- [72] Brian Sauer, M Alan Brookhart, Jason A Roy, and Tyler J VanderWeele. Covariate selection. In *Developing a protocol for observational comparative effectiveness research: a user’s guide*. Agency for Healthcare Research and Quality (US), 2013.
- [73] Paul Huenermund, Beyers Louw, and Mikko Rönkkö. The choice of control variables: How causal graphs can inform the decision. In *Academy of Management Proceedings*, page 294. Academy of Management Briarcliff Manor, NY 10510, 2022. doi: 10.5465/AMBPP.2022.294.
- [74] Tianqi Chen and Carlos Guestrin. XGBoost. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD ’16*, pages 785–794, 3 2016. ISBN 9781450342322. doi: 10.1145/2939672.2939785.
- [75] Ashkan Ertefaie Susan M Shortreed. Outcome adaptive lasso: Variable selection for causal inference. *Biometrics*, 73(4):1111–1122, dec 2017. ISSN 0006-341X. doi: 10.1111/biom.12679.
- [76] Rom Gutman, Ehud Karavani, and Yishai Shimoni. Propensity score models are better when post-calibrated. *arXiv preprint arXiv:2211.01221*, 2022.
- [77] John Platt et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74, 1999.
- [78] Bianca Zadrozny and Charles Elkan. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *Icml*, volume 1, pages 609–616, 2001.
- [79] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.
- [80] Rajeev H Dehejia and Sadek Wahba. Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs. *Journal of the American statistical Association*, 94(448): 1053–1062, 1999.
- [81] Colin B Fogarty, Mark E Mikkelsen, David F Gaieski, and Dylan S Small. Discrete optimization for interpretable study populations and randomization inference in an observational study of severe sepsis mortality. *Journal of the American Statistical Association*, 111(514): 447–458, 2016.
- [82] Richard K Crump, V Joseph Hotz, Guido W Imbens, and Oscar A Mitnik. Dealing with limited overlap in estimation of average treatment effects. *Biometrika*, 96(1):187–199, 2009.

- [83] Alexander D’Amour, Peng Ding, Avi Feller, Lihua Lei, and Jasjeet Sekhon. Overlap in observational studies with high-dimensional covariates. *Journal of Econometrics*, 221(2):644–654, 2021.
- [84] Alicia Curth and Mihaela Van Der Schaar. In search of insights, not magic bullets: Towards demystification of the model selection dilemma in heterogeneous treatment effect estimation. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 6623–6642. PMLR, 23–29 Jul 2023.
- [85] Alicia Curth, David Svensson, Jim Weatherall, and Mihaela van der Schaar. 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)*, 2021.
- [86] Denis Agniel, Isaac S Kohane, and Griffin M Weber. Biases in electronic health record data due to processes within the healthcare system: retrospective observational study. *Bmj*, 361, 2018.
- [87] Rom Gutman, Doron Aronson, Oren Caspi, and Uri Shalit. What drives performance in machine learning models for predicting heart failure outcome? *European Heart Journal - Digital Health*, 9 2022. doi: 10.1093/EHJDH/ZTAC054.
- [88] Susan Athey and Stefan Wager. Estimating treatment effects with causal forests: An application. *Observational Studies*, 5(2):37–51, 2019.
- [89] Adam N Glynn and Kevin M Quinn. An introduction to the augmented inverse propensity weighted estimator. *Political analysis*, pages 36–56, 2010.
- [90] Owen E Leete, Nathan Kallus, Michael G Hudgens, Sonia Napravnik, and Michael R Kosorok. Balanced policy evaluation and learning for right censored data. *arXiv preprint arXiv:1911.05728*, 2019.
- [91] Kosuke Imai and Michael Lingzhi Li. Experimental evaluation of individualized treatment rules. *Journal of the American Statistical Association*, 118(541):242–256, 2023.
- [92] Jared K Lunceford and Marie Davidian. Stratification and weighting via the propensity score in estimation of causal treatment effects: A comparative study. *Statistics in Medicine*, 23 (19):2937–2960, 2004. ISSN 02776715. doi: 10.1002/sim.1903.
- [93] Isaac Lage, Daphna Lifschitz, Finale Doshi-Velez, and Ofra Amir. Toward robust policy summarization. *Autonomous agents and multi-agent systems*, 2019:2081, 2019.
- [94] Anton Matsson and Fredrik D Johansson. Case-based off-policy evaluation using prototype learning. In *Uncertainty in Artificial Intelligence*, pages 1339–1349. PMLR, 2022.
- [95] Theresa A McDonagh, Marco Metra, Marianna Adamo, Roy S Gardner, Andreas Baumbach, Michael Böhm, Haran Burri, Javed Butler, Jelena Čelutkienė, Ovidiu Chioncel, John G F Cleland, Andrew J S Coats, Maria G Crespo-Leiro, Dimitrios Farmakis, Martine Gilard, Stephane Heymans, Arno W Hoes, Tiny Jaarsma, Ewa A Jankowska, Mitja Lainscak,

- Carolyn S P Lam, Alexander R Lyon, John J V McMurray, Alexandre Mebazaa, Richard Mindham, Claudio Muneretto, Massimo Francesco Piepoli, Susanna Price, Giuseppe M C Rosano, Frank Ruschitzka, Anne Kathrine Skibelund, and ESC Scientific Document Group. 2021 ESC Guidelines for the diagnosis and treatment of acute and chronic heart failure: Developed by the Task Force for the diagnosis and treatment of acute and chronic heart failure of the European Society of Cardiology (ESC) With the special contribution of the Heart Failure Association (HFA) of the ESC. *European Heart Journal*, 42(36):3599–3726, 08 2021. ISSN 0195-668X. doi: 10.1093/eurheartj/ehab368.
- [96] Diab Mutlak, Jonathan Lessick, Shehrban Khalil, Sergey Yalonetsky, Yoram Agmon, and Doron Aronson. Tricuspid regurgitation in acute heart failure: is there any incremental risk? *European Heart Journal - Cardiovascular Imaging*, 19(9):993–1001, 9 2018. ISSN 2047-2404. doi: 10.1093/ehjci/jex343.
- [97] Doron Aronson, Wisam Darawsha, Aula Atamna, Marielle Kaplan, Badira F. Makhoul, Diab Mutlak, Jonathan Lessick, Shemy Carasso, Shimon Reisner, Yoram Agmon, Robert Dragu, and Zaher S. Azzam. Pulmonary Hypertension, Right Ventricular Function, and Clinical Outcome in Acute Decompensated Heart Failure. *Journal of Cardiac Failure*, 19(10):665–671, 10 2013. ISSN 10719164. doi: 10.1016/j.cardfail.2013.08.007.
- [98] Piotr Ponikowski, Adriaan A. Voors, Stefan D. Anker, Héctor Bueno, John G. F. Cleland, Andrew J. S. Coats, Volkmar Falk, José Ramón González-Juanatey, Veli-Pekka Harjola, Ewa A. Jankowska, Mariell Jessup, Cecilia Linde, Petros Nihoyannopoulos, John T. Parissis, Burkert Pieske, Jillian P. Riley, Giuseppe M. C. Rosano, Luis M. Ruilope, Frank Ruschitzka, Frans H. Rutten, and Peter van der Meer. 2016 ESC Guidelines for the diagnosis and treatment of acute and chronic heart failure. *European Journal of Heart Failure*, 18(8):891–975, 8 2016. ISSN 13889842. doi: 10.1002/ejhf.592.
- [99] Scott M Lundberg and Su In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 2017-Decem, pages 4766–4775, 5 2017.
- [100] Lauren E Dang, Susan Gruber, Hana Lee, Issa Dahabreh, Elizabeth A Stuart, Brian D Williamson, Richard Wyss, Iván Díaz, Debashis Ghosh, Emre Kıcıman, et al. A causal roadmap for generating high-quality real-world evidence. *arXiv preprint arXiv:2305.06850*, 2023.
- [101] Richard L. Kravitz, Naihua Duan, and Joel Braslow. Evidence-based medicine, heterogeneity of treatment effects, and the trouble with averages. *Milbank Quarterly*, 82(4):661–687, 12 2004. ISSN 0887378X. doi: 10.1111/j.0887-378X.2004.00327.x.
- [102] Rodney A Hayward, David M Kent, Sandeep Vijan, and Timothy P Hofer. Reporting clinical trial results to inform providers, payers, and consumers. *Health Affairs*, 24(6):1571–1581, 2005.
- [103] Issa J. Dahabreh, Rodney Hayward, and David M. Kent. Using group data to treat individuals: Understanding heterogeneous treatment effects in the age of precision medicine and patient-centred evidence. *International Journal of Epidemiology*, 45(6):2184–2193, 12 2016. ISSN 14643685. doi: 10.1093/ije/dyw125.



- [104] Shuai Chen, Lu Tian, Tianxi Cai, and Menggang Yu. A general statistical framework for subgroup identification and comparative treatment scoring. *Biometrics*, 73(4):1199–1209, 2017. ISSN 15410420. doi: 10.1111/biom.12676.
- [105] Sarah C. Anoke, Sharon Lise Normand, and Corwin M. Zigler. Approaches to treatment effect heterogeneity in the presence of confounding. *Statistics in Medicine*, 38(15):2797–2815, 7 2019. ISSN 10970258. doi: 10.1002/sim.8143.
- [106] Alexander R. Luedtke and Mark J. Van Der Laan. Evaluating the impact of treating the optimal subgroup. *Statistical Methods in Medical Research*, 26(4):1630–1640, 8 2017. ISSN 14770334. doi: 10.1177/0962280217708664.
- [107] Yaobin Ling, Pulakesh Upadhyaya, Luyao Chen, Xiaoqian Jiang, and Yejin Kim. Emulate randomized clinical trials using heterogeneous treatment effect estimation for personalized treatments: methodology review and benchmark. *Journal of Biomedical Informatics*, 137: 104256, 2023.
- [108] David M Kent, Jessica K Paulus, David Van Klaveren, Ralph D’Agostino, Steve Goodman, Rodney Hayward, John PA Ioannidis, Bray Patrick-Lake, Sally Morton, Michael Pencina, et al. The predictive approaches to treatment effect heterogeneity (path) statement. *Annals of internal medicine*, 172(1):35–45, 2020.
- [109] Amit Sharma and Emre Kiciman. Dowhy: An end-to-end library for causal inference. *arXiv preprint arXiv:2011.04216*, 2020. doi: 10.48550/arxiv.2011.04216.
- [110] Soorajnath Boominathan, Michael Oberst, Helen Zhou, Sanjat Kanjilal, and David Sontag. Treatment Policy Learning in Multiobjective Settings with Fully Observed Outcomes. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1937–1947, 2020. ISBN 9781450379984. doi: 10.1145/3394486.3403245.
- [111] Andreas D Meid, Lucas Wirbka, ARMIN Study Group, Andreas Groll, and Walter E Haefeli. Can machine learning from real-world data support drug treatment decisions? a prediction modeling case for direct oral anticoagulants. *Medical Decision Making*, 42(5):587–598, 2022.
- [112] Maya L Petersen and Mark J van der Laan. Causal models and learning from data: integrating causal modeling and statistical estimation. *Epidemiology (Cambridge, Mass.)*, 25(3):418, 2014.
- [113] David W. Bates, Andrew Auerbach, Peter Schulam, Adam Wright, and Suchi Saria. Reporting and Implementing Interventions Involving Machine Learning and Artificial Intelligence. *Annals of internal medicine*, 172(11):S137–S144, 6 2020. ISSN 15393704. doi: 10.7326/M19-0872.
- [114] Pranav Rajpurkar, Emma Chen, Oishi Banerjee, and Eric J Topol. AI in health and medicine. *Nature medicine*, 28(1):31–38, 2022.
- [115] Steve Harris, Tim Bonnici, Thomas Keen, Watjana Lilaonitkul, Mark J White, and Nel Swanepoel. Clinical deployment environments: Five pillars of translational machine learning for health. *Frontiers in Digital Health*, 4:939292, 2022.

- [116] Angela Zhang, Lei Xing, James Zou, and Joseph C Wu. Shifting machine learning for healthcare from development to deployment and from models to data. *Nature Biomedical Engineering*, 6(12):1330–1345, 2022.
- [117] Will Douglas Heaven. Hundreds of AI tools have been built to catch Covid. none of them helped. *MIT Technology Review*. Retrieved October, 6:2021, 2021.
- [118] JM Smit, JH Krijthe, WMR Kant, JA Labrecque, M Komorowski, DAMPJ Gommers, J van Bommel, MJT Reinders, and ME van Genderen. Causal inference using observational intensive care unit data: a scoping review and recommendations for future practice. *npj Digital Medicine*, 6(1):221, 2023.
- [119] Aparna Balagopalan, Ioana Baldini, Leo Anthony Celi, Judy Gichoya, Liam G McCoy, Tristan Naumann, Uri Shalit, Mihaela van der Schaar, and Kiri L Wagstaff. Machine learning for healthcare that matters: Reorienting from technical novelty to equitable impact. *PLOS Digital Health*, 3(4):e0000474, 2024.
- [120] Melissa D. McCradden, Elizabeth A. Stephenson, and James A. Anderson. Clinical research underlies ethical integration of healthcare artificial intelligence. *Nature Medicine* 2020 26:9, 26(9):1325–1326, 9 2020. ISSN 1546-170X. doi: 10.1038/s41591-020-1035-9.
- [121] Batia Mishan Wiesenfeld, Yin Aphinyanaphongs, and Oded Nov. AI model transferability in healthcare: a sociotechnical perspective. *Nature Machine Intelligence*, 4(10):807–809, 2022.
- [122] Shalmali Joshi, Iñigo Urteaga, Wouter AC van Amsterdam, George Hripcsak, Pierre Elias, Benjamin Recht, Noémie Elhadad, James Fackler, Mark P Sendak, Jenna Wiens, et al. AI as an intervention: improving clinical outcomes relies on a causal approach to AI development and validation. *Journal of the American Medical Informatics Association*, page ocae301, 2025.
- [123] Joachim A Behar, Jeremy Levy, and Leo Anthony Celi. Generalization in medical ai: a perspective on developing scalable models. *arXiv preprint arXiv:2311.05418*, 2023.
- [124] Donald B Rubin. Causal Inference Through Potential Outcomes and Principal Stratification: Application to Studies with "Censoring" Due to Death 1. *Statistical Science*, 21(3):299–309, 2006. doi: 10.1214/088342306000000114.
- [125] Stuart J Pocock, Cono A Ariti, Timothy J Collier, and Duolao Wang. The win ratio: a new approach to the analysis of composite endpoints in clinical trials based on clinical priorities. *European heart journal*, 33(2):176–182, 2012.
- [126] Mathieu Even and Julie Josse. Rethinking the win ratio: A causal framework for hierarchical outcome analysis. *arXiv preprint arXiv:2501.16933*, 2025.
- [127] Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 65(2):331–355, 2003.
- [128] Omer Gottesman, Fredrik Johansson, Matthieu Komorowski, Aldo Faisal, David Sontag, Finale Doshi-Velez, and Leo Anthony Celi. Guidelines for reinforcement learning in healthcare. *Nature medicine*, 25(1):16–18, 2019.

- [129] David Madras, Toni Pitassi, and Richard Zemel. Predict responsibly: improving fairness and accuracy by learning to defer. *Advances in Neural Information Processing Systems*, 31, 2018.
- [130] Hussein Mozannar and David Sontag. Consistent estimators for learning to defer to an expert. In *International Conference on Machine Learning*, pages 7076–7087. PMLR, 2020.
- [131] Marah Ghoummaid and Uri Shalit. When to act and when to ask: Policy learning with deferral under hidden confounding. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [132] Ruijiang Gao and Mingzhang Yin. Confounding-robust policy improvement with human-AI teams. *arXiv preprint arXiv:2310.08824*, 2023.
- [133] Gagan Bansal, Besmira Nushi, Ece Kamar, Eric Horvitz, and Daniel S Weld. Is the most accurate AI the best teammate? optimizing AI for teamwork. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(13):11405–11414, 2021.
- [134] Maia Jacobs, Melanie F Pradier, Thomas H McCoy Jr, Roy H Perlis, Finale Doshi-Velez, and Krzysztof Z Gajos. How machine-learning recommendations influence clinician treatment selections: the example of antidepressant selection. *Translational psychiatry*, 11(1):108, 2021.
- [135] Julien Meyer, April Khademi, Bernard Têtu, Wencui Han, Pria Nippak, and David Remisch. Impact of artificial intelligence on pathologists’ decisions: an experiment. *Journal of the American Medical Informatics Association*, 29(10):1688–1695, 2022.
- [136] Samuel G Finlayson, Adarsh Subbaswamy, Karandeep Singh, John Bowers, Annabel Kupke, Jonathan Zittrain, Isaac S Kohane, and Suchi Saria. The clinician and dataset shift in artificial intelligence. *New England Journal of Medicine*, 385(3):283–286, 2021.
- [137] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünnner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, pages 7599–7609. PMLR, 2020.
- [138] Eric J Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal learning. *arXiv preprint arXiv:2009.10982*, 2020.
- [139] Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. A selective review of negative control methods in epidemiology. *Current epidemiology reports*, 7(4):190–202, 2020.
- [140] Erik Sverdrup and Yifan Cui. Proximal causal learning of conditional average treatment effects. In *International Conference on Machine Learning*, pages 33285–33298. PMLR, 2023.
- [141] Mihai Gheorghiade and Peter S. Pang. Acute Heart Failure Syndromes. *Journal of the American College of Cardiology*, 53(7):557–573, 2 2009. ISSN 07351097. doi: 10.1016/j.jacc.2008.10.041.
- [142] John G Laffey and Brian P Kavanagh. Negative trials in critical care: why most research is probably wrong. *The Lancet. Respiratory medicine*, 6(9):659–660, 9 2018. ISSN 2213-2619. doi: 10.1016/S2213-2600(18)30279-0.

- [143] John J.V. McMurray, Milton Packer, Akshay S. Desai, Jianjian Gong, Martin P. Lefkowitz, Adel R. Rizkala, Jean L. Rouleau, Victor C. Shi, Scott D. Solomon, Karl Swedberg, and Michael R. Zile. Angiotensin–Neprilysin Inhibition versus Enalapril in Heart Failure. *New England Journal of Medicine*, 371(11):993–1004, 9 2014. ISSN 0028-4793. doi: 10.1056/NEJMoa1409077.
- [144] Barry M. Massie, Christopher M. O’Connor, Marco Metra, Piotr Ponikowski, John R. Teerlink, Gad Cotter, Beth Davison Weatherley, John G.F. Cleland, Michael M. Givertz, Adriaan Voors, Paul DeLucca, George A. Mansoor, Christina M. Salerno, Daniel M. Bloomfield, and Howard C. Dittrich. Rolofylline, an Adenosine A<sub>1</sub>-Receptor Antagonist, in Acute Heart Failure. *New England Journal of Medicine*, 363(15):1419–1428, 10 2010. ISSN 0028-4793. doi: 10.1056/NEJMoa0912613.
- [145] C.M. O’Connor, R.C. Starling, A.F. Hernandez, P.W. Armstrong, K. Dickstein, V. Hasselblad, G.M. Heizer, M. Komajda, B.M. Massie, J.J.V. McMurray, M.S. Nieminen, C.J. Reist, J.L. Rouleau, K. Swedberg, K.F. Adams, S.D. Anker, D. Atar, A. Battler, R. Botero, N.R. Bohidar, J. Butler, N. Clausell, R. Corbalán, M.R. Costanzo, U. Dahlstrom, L.I. Deckelbaum, R. Diaz, M.E. Dunlap, J.A. Ezekowitz, D. Feldman, G.M. Felker, G.C. Fonarow, D. Genovesi, S.S. Gottlieb, J.A. Hill, J.E. Hollander, J.G. Howlett, M.P. Hudson, R.D. Kociol, H. Krum, A. Laucevicius, W.C. Levy, G.F. Méndez, M. Metra, S. Mittal, B.-H. Oh, N.L. Pereira, P. Ponikowski, W.H.W. Tang, S. Tanomsup, J.R. Teerlink, F. Triposkiadis, R.W. Troughton, A.A. Voors, D.J. Whellan, F. Zannad, and R.M. Califf. Effect of Nesiritide in Patients with Acute Decompensated Heart Failure. *New England Journal of Medicine*, 365(1):32–43, 7 2011. ISSN 0028-4793. doi: 10.1056/NEJMoa1100171.
- [146] Alexander Goldberg, Elena Kogan, Haim Hammerman, Walter Markiewicz, and Doron Aronson. The impact of transient and persistent acute kidney injury on long-term outcomes after acute myocardial infarction. *Kidney International*, 76(8):900–906, 10 2009. ISSN 0085-2538. doi: 10.1038/KI.2009.295.
- [147] Javed Butler, Daniel E Forman, William T Abraham, Stephen S Gottlieb, Evan Loh, Barry M Massie, Christopher M O’Connor, Michael W Rich, Lynne Warner Stevenson, Yongfei Wang, James B Young, and Harlan M Krumholz. Relationship between heart failure treatment and development of worsening renal function among hospitalized patients. *American Heart Journal*, 147(2):331–338, 2 2004. ISSN 00028703. doi: 10.1016/j.ahj.2003.08.012.
- [148] Jesse A. Kane, Joseph K. Kim, Syed Abbas Haidry, Louis Saliccioli, and Jason Lazar. Discontinuation/dose reduction of angiotensin-converting enzyme inhibitors/angiotensin receptor blockers during acute decompensated heart failure in African-American patients with reduced left-ventricular ejection fraction. *Cardiology (Switzerland)*, 137(2):121–125, 2017. ISSN 14219751. doi: 10.1159/000457946.
- [149] Michael Böhm, Andreas Link, Danlin Cai, Markku S. Nieminen, Gerasimos S. Filippatos, Reda Salem, Alain Cohen Solal, Bidan Huang, Robert J. Padley, Matti Kivikko, and Alexandre Mebazaa. Beneficial association of  $\beta$ -blocker therapy on recovery from severe acute heart failure treatment: Data from the Survival of Patients with Acute Heart Failure in Need of Intravenous Inotropic Support trial. *Critical Care Medicine*, 39(5):940–944, 5 2011. ISSN 15300293. doi: 10.1097/CCM.0b013e31820a91ed.

## Appendix

### A Questions for clinical-algorithmic team collaboration

We propose that the system be defined by discussion of the following points with the clinical team.

1. *What is the treatment decision for which the clinical team would like assistance from an algorithmic model?*
2. (a) *At what point in the clinical workflow would the clinical team want the recommendation?*  
(b) *Does the clinical recommendation time point correspond to the time where the decision is currently made?*

### B Which covariates should be used

Creating and validating a causal graph that captures all of the hundreds of potential variables involved in a clinical problem is often challenging and time-consuming. We thus propose an alternative that could be useful for many cases of interest.

To simplify the task, we start by attempting to identify the potential confounders. Towards this goal, we ask experts to answer the following questions:

- a. Which factors plausibly affect the treatment decisions as they occurred in the historical data?
- b. Of the above factors, which also plausibly affect the outcome?

The answer to the second question can also come from the scientific literature at large. Factors that plausibly affect both the outcome and the treatment decisions as they occurred in the cases present in our data are potential confounders. Next, we need to make sure these potential confounders are indeed represented in the data, or, if not available directly, at least have a good proxy. E.g., one can use hematocrit and hemoglobin biomarkers as proxies for congestion in AHF patients [70]. If we find an important confounder which is *not* recorded in the data and has no good proxy, then we must either find a way to obtain the required confounder, or pursue a completely different analysis.

Assuming we are satisfied that the bulk of confounding factors is represented in the data, we should consider which covariates we should *discard* due to the possibility of inducing bias. We recommend the following procedure for each covariate in the dataset:

- a. Is it affected by the treatment decision, or does it occur after the treatment decision? If it is, it should be discarded from future analysis.

If not:

- b. Does it plausibly affect the treatment decision as reflected in the available historical data (or is a proxy for such a variable)?
- c. Does it plausibly affect the outcome (or is a proxy for such a variable)?

Having discarded the covariates for which the answer to **a** is “yes”, we are left with the following: Covariates for which the answer to **b** AND **c** is “yes”; these are potential confounders and must be included in the analysis. Covariates for which the answer to **b** is “no” and **c** is “yes”; these are (potentially) what are known as effect modifiers, and should also be included in the analysis as their inclusion can reduce variance [73]. Covariates for which the answer to **b** is “yes” and to **c** is “no”; these are potential instruments, and should be discarded from analysis as their inclusion might *increase* variances [71, 73]. Finally, covariates which do not plausibly affect outcome or treatment, and are not proxies for such variables, should usually be discarded from the analysis as they are unrelated to the problem at hand.

While the above is a simplification of the more rigorous process of identifying a backdoor blocking set and other causal identification schemes [13], we find that in many cases this is a more realistic endeavor than constructing a large, accurate, causal graph.

We also note that in principle, proxy variables should be used differently from variables which are direct causes of treatment or outcome. However, in practice, we believe this distinction is often far from clear, and there are not many methods for rigorously dealing with proxy variables when the proxy mechanism and graph are unknown [138–140]. This is an area where we believe further research is called for.

## C Simulation details

Given real patient data  $\{(x_i, t_i)\}_{i=1}^n$  of  $n$  patients with  $d$  covariates, we wish to simulate two potential outcomes  $Y^0, Y^1$  for each patient; we outline the process in Algorithm SA1. Our goal is to create a learnable yet somewhat realistic simulated CATE function. Towards that end we simulate the two expected potential outcomes  $\mathbb{E}[Y^1 | x]$  and  $\mathbb{E}[Y^0 | x]$  as linear functions of the covariates:  $w_1^\top x$ , and  $w_0^\top x$ , respectively. This leads to a linear CATE function of the form  $\tau(x) = (w_1 - w_0)^\top x = \Delta^\top x$  where  $\Delta = w_1 - w_0$ . We generate the vectors  $w_0, w_1$  so that their difference  $\Delta$  will have a meaningful connection to the underlying data, as we now explain.

We start by fitting a logistic propensity model  $\hat{e}(x) = p(T = 1 | x)$  using the real treatment assignments, obtaining a coefficient vector which we denote  $\beta_{prop}$ . We consider this vector as an embedding of the clinical knowledge driving treatment assignment. We then further generate a random vector with the same norm as  $\beta_{prop}$ , and combine the two using a weighting parameter  $\lambda \in [0, 1]$  to generate  $\Delta$ . The CATE function is thus a linear function which is a weighted combination of the propensity score coefficient vector and a random vector. The weighting factor  $\lambda$  represents the degree to which the clinicians’ understanding of the important factors regarding the outcome is correct: For  $\lambda = 1$ , the simulated  $Y^1$  for patients who receive treatment  $T = 1$  will tend to have a higher potential outcome, meaning clinicians made the correct treatment decision (assuming here that higher outcomes are better).

We draw the potential outcome values themselves from a Gaussian distribution with means  $w_1^\top x, w_0^\top x$  and variances  $\varepsilon_1, \varepsilon_0$ . We set the variances to be 1.2 SD times the variance of  $w_t^\top x$  across patients for each outcome arm  $t$ . Finally, we re-scale the resulting average CATE to match a pre-defined value that represents a clinically reasonable value.

The full simulation procedure and settings are presented in Algorithm SA1.

---

**Algorithm SA1:** Simulation Parameters
 

---

**Require:**  $\lambda \in [0, 1]$ , desired effect  $C \in \mathbb{R}$

**Require:**  $\{(x_i, t_i)\}_{i=1}^n$  (baseline data)

- 1:  $\beta_{\text{prop}} \leftarrow$  coefficients from logistic regression for  $p(t_i = 1|x_i)$
- 2:  $\beta_{\text{rand}} \sim \mathcal{N}\left(0, \left(\frac{1}{\sqrt{d}}\right)^2\right)$
- 3: Normalize:
- 4:  $\beta_{\text{prop}} \leftarrow \frac{\beta_{\text{prop}}}{\|\beta_{\text{prop}}\|_2}$
- 5:  $\beta_{\text{rand}} \leftarrow \frac{\beta_{\text{rand}}}{\|\beta_{\text{rand}}\|_2}$
- 6:  $\Delta \leftarrow \sqrt{\lambda} \beta_{\text{prop}} + \sqrt{1 - \lambda} \beta_{\text{rand}}$
- 7:  $a \leftarrow \frac{1}{n} \sum_{i=1}^n |x_i \cdot \Delta|$
- 8:  $\Delta \leftarrow \Delta \cdot \frac{C}{a}$
- 9:  $w_0 \sim \mathcal{N}\left(0, \left(\frac{1}{\sqrt{d}}\right)^2\right)$
- 10:  $\sigma_0 \leftarrow \text{SD}(w_0^\top X)$
- 11:  $w_1 \leftarrow \Delta + w_0$
- 12:  $\sigma_1 \leftarrow \text{SD}(w_1^\top X)$
- 13: Initialize empty vectors  $Y_0$  and  $Y_1$
- 14: **for**  $i = 1$  to  $n$  **do**
- 15:  $\varepsilon_{0,i} \sim \mathcal{N}\left(0, (1.2 \sigma_0)^2\right)$
- 16:  $Y_{0,i} \leftarrow w_0^\top X_i + \varepsilon_{0,i}$
- 17:  $\varepsilon_{1,i} \sim \mathcal{N}\left(0, (1.2 \sigma_1)^2\right)$
- 18:  $Y_{1,i} \leftarrow w_1^\top X_i + \varepsilon_{1,i}$
- 19: **end for**
- 20: **return** Policy  $\pi^*(x)$  defined as

$$\pi^*(x) = \begin{cases} 1, & \text{if } \Delta^\top x < 0, \\ 0, & \text{otherwise.} \end{cases}$$

- 21: **return** Vectors  $Y_0$  and  $Y_1$
-



## D Policy value Formulation

As detailed in Section 2.5 and Section 3.11, the policy value  $V(\pi)$  is defined in Eq. (5). In this work we suggest using a flavor of methods suggested by [55, 92], where each sample  $i$  contributes weight  $w_i$ :  $w_i = \frac{\mathbb{I}_{t_i=\pi(x_i)}}{p^*(t_i=\pi(x_i)|x_i)}$ , and we provide the DR and IPW estimands:

$$\begin{aligned}\tilde{V}_{\text{DR}}(\pi) &= \frac{1}{\sum_{i=1}^n w_i} \sum_{i=1}^n w_i \cdot (y - \hat{y}^{\pi(x_i)}) + \frac{1}{n} \sum_{i=1}^n \hat{y}^{\pi(x_i)}, \\ \tilde{V}_{\text{IPW}}(\pi) &= \frac{1}{\sum_{i=1}^n w_i} \sum_{i=1}^n w_i \cdot y,\end{aligned}$$

where  $p^*(t_i = \pi(x_i)|x_i)$  is the propensity score model trained on all the data, and  $\hat{y}^{\pi(x_i)}$  is the estimated potential outcome under the policy  $\pi(x_i)$ , given by plug-in estimator. The plug-in estimator can be any outcome model trained in the outcome estimation stage (Section 3.6).

## E Rank graph Details

As described in Section 3.11, we suggest using a rank graph comparison. To perform such analysis, we suggest the following steps. Given policies, the threshold for treatment assignment for a given policy is set to be the  $q^{\text{th}}$  percentile of the CATE value, and the policy value is calculated based on this percentile. I.e., instead of policy  $\pi_{\mathcal{A}}(x_i) = \text{sign}(\widehat{\tau}_{\mathcal{A}}(x_i))$ , we define a set of new policies, where each policy is defined according to the  $q^{\text{th}}$  percentile of the CATE value of policy  $\mathcal{A}$ :  $\pi_{\mathcal{A}}^{q\text{-rank}}(x_i) = \mathbb{I}(\widehat{\tau}_{\mathcal{A}}(x_i) > Q(\widehat{\tau}_{\mathcal{A}}, q))$ , where  $\pi_{\mathcal{A}}^{q\text{-rank}}$  is defined for each  $q$ , we call this the  $q$ -rank policy of  $\mathcal{A}$ . Policy values for each rank- $q$  policy are calculated for  $q$  values ranging from 0 to 1 (with step size  $\delta$ ), and these values are compared to the random policy and dichotomies policy. The x-axis of the graph is set to be the proportion of patients assigned  $T = 1$  for each policy and quantile value.

## F AKI and ADHF

**Acute decompensated heart failure (ADHF).** We focus on ADHF as an important use case for learning causal decision support models from healthcare data. ADHF is the most common cause of hospital admission in persons aged 65 years, accounting for over 1,000,000 hospitalizations each year in the U.S. alone [141]. The prognosis of patients admitted with ADHF is dismal [141]. Seemingly a single homogeneous clinical phenotype presenting with shortness of breath and fluid overload, these patients are a very heterogeneous group in terms of the underlying pathophysiology. This leads to a major challenge in tailoring treatment, as a single uniform approach may prove ineffective or harmful in subgroups of these patients [142]. Indeed, that may explain the inconclusiveness of virtually all recent randomized controlled trials (RCTs) in ADHF patients and the lack of evidence-based guidelines for their management [143–145].

**Acute Kidney Injury (AKI)** is a common and potentially dangerous complication for patients suffering from ADHF. AKI is characterized by a sudden episode of kidney failure or kidney damage that happens within a few hours or a few days. AKI, usually defined as an increase in the presence of the biomarker creatinine on the order of  $> 0.3 \frac{\text{mg}}{\text{dL}}$  to  $> 0.5 \frac{\text{mg}}{\text{dL}}$  from baseline. It causes a build-up

of waste products in the blood and compromises the kidneys' ability to maintain the right balance of fluid in the body. AKI occurs in 20-30% of patients with ADHF, and is associated with greater length of hospital stay, increased chance of hospital readmission, and death [146, 147].

A physician treating an ADHF patient with AKI faces a therapeutic dilemma. On the one hand, ADHF entails fluid accumulation (congestion) requiring diuretic therapy and excessive neurohormonal activation, since sustaining proper kidney function during therapeutic interventions is vital to the alleviation of congestion. On the other hand, these very treatments aimed at alleviating congestion might be a cause or an aggravating factor in AKI, which can harm the patient's health. This requires a delicate balancing act on the side of the treating physician, a difficulty which is made worse by the fact that there are no guidelines for patients who develop AKI in the setting of ADHF.

Currently, the prevailing treatments for ADHF patients with AKI are cessation of diuretics and fluids infusion [8]. In addition, the physician may elect to reduce or discontinue neurohormonal inhibitors (e.g. beta blockers) and to initiate inotrope therapy. The latter decision, although common among physicians treating ADHF, has no proven benefit in the short term and is undoubtedly harmful in the long term [148, 149]. Therefore, when AKI occurs in an ADHF patient, the treating physician is faced with multiple options to respond by modifying therapy, with no guidance as to how these decisions affect renal and overall clinical outcomes. Indeed, in a recent pilot study of 277 ADHF patients at the Rambam hospital, there was remarkable variability in the observed management of patients with AKI [8].

## F.1 Data description

	Missing		Overall	Test	Train	Validation
<b>n</b>			2157	530	1305	322
firstadm, n (%)	0	0	942 (43.7)	234 (44.2)	565 (43.3)	143 (44.4)
	1		1215 (56.3)	296 (55.8)	740 (56.7)	179 (55.6)
admyear, n (%)	1.0	0	198 (9.2)	44 (8.3)	123 (9.4)	31 (9.6)
	2.0		205 (9.5)	51 (9.6)	124 (9.5)	30 (9.3)
	3.0		256 (11.9)	63 (11.9)	147 (11.3)	46 (14.3)
	4.0		292 (13.5)	78 (14.7)	182 (13.9)	32 (9.9)
	5.0		322 (14.9)	80 (15.1)	198 (15.2)	44 (13.7)
	6.0		312 (14.5)	70 (13.2)	187 (14.3)	55 (17.1)
	7.0		297 (13.8)	71 (13.4)	181 (13.9)	45 (14.0)
	8.0		273 (12.7)	73 (13.8)	162 (12.4)	38 (11.8)
	0.0		2 (0.1)		1 (0.1)	1 (0.3)
age, mean (SD)		0	76.9 (11.5)	76.2 (12.1)	77.5 (11.1)	75.6 (12.2)
gender, n (%)	0	0	1140 (52.9)	263 (49.6)	715 (54.8)	162 (50.3)
	1		1017 (47.1)	267 (50.4)	590 (45.2)	160 (49.7)
Weight, mean (SD)		0	80.9 (14.5)	81.3 (14.5)	80.5 (14.8)	82.0 (13.1)
Hight, mean (SD)		0	164.8 (6.3)	165.1 (7.4)	164.6 (6.0)	165.0 (5.0)
bmi, mean (SD)		0	30.3 (17.6)	31.3 (33.3)	29.9 (7.7)	30.2 (4.4)
htn, n (%)	0	0	610 (28.3)	149 (28.1)	365 (28.0)	96 (29.8)
	1		1547 (71.7)	381 (71.9)	940 (72.0)	226 (70.2)
firsttemp, mean (SD)		0	37.1 (10.3)	37.4 (14.5)	36.8 (2.9)	37.8 (18.4)
dm, n (%)	0	0	963 (44.6)	235 (44.3)	570 (43.7)	158 (49.1)
	1		1194 (55.4)	295 (55.7)	735 (56.3)	164 (50.9)
smk, n (%)	0	0	1637 (75.9)	405 (76.4)	994 (76.2)	238 (73.9)
	1		285 (13.2)	63 (11.9)	171 (13.1)	51 (15.8)
	2		235 (10.9)	62 (11.7)	140 (10.7)	33 (10.2)
ihd, n (%)	0	0	1605 (74.4)	405 (76.4)	955 (73.2)	245 (76.1)
	1		552 (25.6)	125 (23.6)	350 (26.8)	77 (23.9)
vhd, n (%)	0	0	1859 (86.2)	458 (86.4)	1125 (86.2)	276 (85.7)
	1		298 (13.8)	72 (13.6)	180 (13.8)	46 (14.3)
af, n (%)	0	0	1204 (55.8)	308 (58.1)	719 (55.1)	177 (55.0)
	1		953 (44.2)	222 (41.9)	586 (44.9)	145 (45.0)
Hyperlipidemia, n (%)	0	0	882 (40.9)	211 (39.8)	542 (41.5)	129 (40.1)
	1		1275 (59.1)	319 (60.2)	763 (58.5)	193 (59.9)
copd, n (%)	0	0	1841 (85.4)	453 (85.5)	1109 (85.0)	279 (86.6)
	1		316 (14.6)	77 (14.5)	196 (15.0)	43 (13.4)
crf, n (%)	0	0	1171 (54.3)	294 (55.5)	697 (53.4)	180 (55.9)
	1		986 (45.7)	236 (44.5)	608 (46.6)	142 (44.1)
bb, n (%)	0	0	547 (25.4)	131 (24.7)	319 (24.4)	97 (30.1)
	1		1610 (74.6)	399 (75.3)	986 (75.6)	225 (69.9)
acei, n (%)	0	0	733 (34.0)	179 (33.8)	445 (34.1)	109 (33.9)
	1		1424 (66.0)	351 (66.2)	860 (65.9)	213 (66.1)
arf, n (%)	0	0	1298 (60.2)	312 (58.9)	796 (61.0)	190 (59.0)
	1		859 (39.8)	218 (41.1)	509 (39.0)	132 (41.0)
antiplt, n (%)	0	0	1315 (61.0)	317 (59.8)	801 (61.4)	197 (61.2)
	1		842 (39.0)	213 (40.2)	504 (38.6)	125 (38.8)
anticoagulants, n (%)	0	0	1547 (71.7)	394 (74.3)	932 (71.4)	221 (68.6)
	1		610 (28.3)	136 (25.7)	373 (28.6)	101 (31.4)
furosemide, n (%)	0	0	623 (28.9)	159 (30.0)	371 (28.4)	93 (28.9)
	1		1534 (71.1)	371 (70.0)	934 (71.6)	229 (71.1)
zaroxolin, n (%)	0	0	2098 (97.3)	511 (96.4)	1278 (97.9)	309 (96.0)
	1		59 (2.7)	19 (3.6)	27 (2.1)	13 (4.0)
thiamin, n (%)	0	0	2104 (97.5)	518 (97.7)	1272 (97.5)	314 (97.5)
	1		53 (2.5)	12 (2.3)	33 (2.5)	8 (2.5)
insulin, n (%)	0	0	1672 (77.5)	410 (77.4)	1012 (77.5)	250 (77.6)
	1		485 (22.5)	120 (22.6)	293 (22.5)	72 (22.4)
tn, mean (SD)		0	0.3 (1.9)	0.3 (1.2)	0.4 (2.1)	0.4 (2.1)
hb, mean (SD)		0	11.3 (1.9)	11.4 (2.0)	11.3 (1.9)	11.4 (2.1)
first_mcv, mean (SD)		0	86.6 (7.0)	86.7 (7.4)	86.6 (6.8)	86.5 (6.9)
first_rdw, mean (SD)		0	15.9 (2.0)	15.9 (2.0)	15.9 (2.0)	15.8 (1.9)
wbc, mean (SD)		0	10.9 (9.2)	10.8 (11.6)	11.0 (8.8)	10.6 (5.9)
alb, mean (SD)		0	3.1 (0.5)	3.1 (0.5)	3.1 (0.5)	3.2 (0.5)
ast, mean (SD)		0	55.2 (359.1)	46.4 (164.4)	57.3 (428.6)	61.0 (274.0)

	Missing	Overall	Test	Train	Validation
alt, mean (SD)	0	62.9 (250.8)	71.1 (296.6)	60.0 (241.6)	61.2 (201.1)
ggt, mean (SD)	0	98.5 (141.6)	98.1 (124.8)	100.7 (157.0)	90.3 (93.7)
alkp, mean (SD)	0	113.5 (72.0)	110.7 (58.8)	115.6 (79.7)	109.8 (56.9)
first_p, mean (SD)	0	4.4 (1.2)	4.4 (1.2)	4.4 (1.2)	4.4 (1.3)
crp, mean (SD)	0	46.8 (28.9)	47.5 (32.0)	46.4 (27.5)	47.1 (28.7)
first_cl, mean (SD)	0	102.5 (5.2)	102.3 (5.4)	102.6 (5.2)	102.8 (5.0)
glu, mean (SD)	0	174.3 (83.0)	174.0 (83.7)	174.4 (81.8)	174.5 (86.6)
dbp_last, mean (SD)	0	64.4 (14.6)	64.7 (14.7)	64.2 (14.5)	64.6 (15.1)
dbp_first, mean (SD)	0	76.3 (16.9)	76.1 (16.3)	76.5 (17.3)	75.9 (16.3)
dbp_mean, mean (SD)	0	68.8 (10.9)	68.9 (10.8)	68.9 (10.9)	68.5 (11.5)
dbp_median, mean (SD)	0	68.2 (11.6)	68.3 (11.1)	68.2 (11.6)	68.0 (12.2)
dbp_quantile_90, mean (SD)	0	80.8 (13.2)	80.6 (12.6)	81.1 (13.4)	79.9 (13.6)
dbp_quantile_10, mean (SD)	0	57.3 (11.0)	57.7 (11.1)	57.2 (10.7)	57.2 (11.6)
dbp_std, mean (SD)	0	11.6 (4.9)	11.3 (4.6)	11.9 (5.1)	11.1 (4.3)
sbp_last, mean (SD)	0	123.0 (23.7)	123.8 (24.7)	123.0 (23.6)	122.1 (22.5)
sbp_first, mean (SD)	0	145.6 (32.4)	145.3 (31.2)	146.0 (33.4)	144.6 (30.1)
sbp_mean, mean (SD)	0	131.7 (21.2)	132.3 (21.8)	131.9 (21.0)	130.0 (21.1)
sbp_median, mean (SD)	0	130.7 (22.0)	131.1 (22.5)	131.0 (21.8)	129.0 (21.9)
sbp_quantile_90, mean (SD)	0	151.0 (26.4)	150.9 (26.2)	151.7 (26.6)	147.8 (25.8)
sbp_quantile_10, mean (SD)	0	113.2 (19.6)	114.3 (20.3)	113.0 (19.4)	112.3 (19.6)
sbp_std, mean (SD)	0	18.8 (9.3)	18.1 (7.8)	19.3 (9.1)	18.2 (12.1)
pp_last, mean (SD)	0	58.6 (19.8)	59.0 (20.4)	58.8 (19.8)	57.4 (18.6)
pp_first, mean (SD)	0	69.3 (24.6)	69.1 (23.8)	69.6 (25.3)	68.7 (23.1)
pp_mean, mean (SD)	0	62.9 (17.2)	63.4 (17.5)	63.1 (17.1)	61.5 (16.9)
pp_median, mean (SD)	0	62.3 (18.0)	62.6 (18.4)	62.5 (17.9)	60.7 (17.4)
pp_quantile_90, mean (SD)	0	77.2 (20.7)	77.3 (20.9)	77.6 (20.8)	75.3 (20.3)
pp_quantile_10, mean (SD)	0	49.2 (15.8)	50.0 (16.1)	49.2 (15.8)	48.0 (15.7)
pp_std, mean (SD)	0	14.1 (7.5)	13.6 (5.9)	14.2 (6.8)	14.1 (11.4)
MAP_last, mean (SD)	0	83.9 (15.6)	84.4 (15.9)	83.8 (15.4)	83.8 (15.7)
MAP_first, mean (SD)	0	99.4 (20.1)	99.2 (19.4)	99.6 (20.7)	98.8 (19.0)
MAP_mean, mean (SD)	0	89.8 (12.8)	90.0 (13.0)	89.9 (12.7)	89.0 (13.2)
MAP_median, mean (SD)	0	89.2 (13.4)	89.5 (13.3)	89.2 (13.3)	88.5 (14.0)
MAP_quantile_90, mean (SD)	0	103.0 (15.9)	103.0 (15.5)	103.4 (16.1)	101.4 (15.8)
MAP_quantile_10, mean (SD)	0	77.1 (12.3)	77.6 (12.7)	76.9 (12.1)	76.7 (12.6)
MAP_std, mean (SD)	0	12.8 (5.6)	12.4 (5.1)	13.2 (5.9)	12.3 (5.5)
hr_last, mean (SD)	0	79.3 (18.6)	79.7 (18.6)	78.7 (17.9)	80.9 (21.0)
hr_first, mean (SD)	0	85.8 (20.7)	85.8 (20.7)	85.4 (20.4)	87.5 (21.7)
hr_mean, mean (SD)	0	80.9 (15.0)	81.1 (14.7)	80.7 (14.9)	81.5 (15.9)
hr_median, mean (SD)	0	80.1 (15.9)	80.4 (15.7)	79.8 (15.8)	81.0 (16.8)
hr_quantile_90, mean (SD)	0	92.2 (19.0)	92.6 (18.9)	91.9 (18.5)	93.3 (20.8)
hr_quantile_10, mean (SD)	0	70.3 (13.8)	70.4 (13.3)	70.3 (13.8)	70.5 (14.5)
hr_std, mean (SD)	0	11.1 (6.6)	11.3 (6.6)	11.1 (6.3)	11.2 (7.5)
bnp_last, mean (SD)	0	1224.2 (882.1)	1251.6 (930.5)	1219.5 (868.7)	1198.4 (855.5)
bnp_first, mean (SD)	0	1224.7 (878.8)	1244.7 (921.9)	1222.7 (866.0)	1199.6 (859.5)
bnp_mean, mean (SD)	0	1222.5 (879.3)	1246.3 (925.4)	1219.1 (865.7)	1197.1 (857.3)
bnp_median, mean (SD)	0	1222.5 (879.3)	1246.3 (925.4)	1219.1 (865.7)	1197.1 (857.3)
bnp_std, mean (SD)	0	246.9 (59.2)	246.5 (50.6)	247.4 (67.9)	245.5 (23.9)
bnp_quantile_90, mean (SD)	0	1231.8 (881.9)	1255.2 (928.6)	1228.8 (868.8)	1205.6 (857.1)
bnp_quantile_10, mean (SD)	0	1216.4 (878.0)	1240.6 (923.0)	1212.7 (864.5)	1191.8 (857.5)
bun_last, mean (SD)	0	41.0 (23.7)	40.8 (23.7)	40.7 (23.3)	42.2 (25.5)
bun_first, mean (SD)	0	38.7 (23.4)	38.2 (23.2)	38.8 (23.2)	39.2 (24.7)
bun_mean, mean (SD)	0	39.6 (23.1)	39.2 (23.0)	39.5 (22.8)	40.5 (24.7)
bun_median, mean (SD)	0	39.4 (23.1)	39.0 (23.1)	39.4 (22.8)	40.3 (24.7)
bun_std, mean (SD)	0	4.0 (3.1)	4.1 (3.1)	4.0 (3.2)	4.0 (2.8)
bun_quantile_90, mean (SD)	0	41.8 (23.8)	41.4 (23.6)	41.6 (23.5)	42.7 (25.3)
bun_quantile_10, mean (SD)	0	37.5 (22.9)	37.0 (22.8)	37.5 (22.5)	38.3 (24.5)
crea_last, mean (SD)	0	1.8 (1.0)	1.8 (1.0)	1.8 (1.0)	1.8 (1.2)
crea_first, mean (SD)	0	1.8 (1.0)	1.8 (1.1)	1.7 (1.0)	1.8 (1.1)
crea_mean, mean (SD)	0	1.8 (1.0)	1.8 (1.0)	1.8 (1.0)	1.8 (1.1)
crea_median, mean (SD)	0	1.8 (1.0)	1.8 (1.0)	1.8 (1.0)	1.8 (1.1)
crea_std, mean (SD)	0	0.1 (0.1)	0.1 (0.1)	0.1 (0.1)	0.1 (0.1)
crea_quantile_90, mean (SD)	0	1.8 (1.0)	1.8 (1.1)	1.8 (1.0)	1.9 (1.1)
crea_quantile_10, mean (SD)	0	1.7 (1.0)	1.7 (1.0)	1.7 (1.0)	1.7 (1.1)
hct_last, mean (SD)	0	34.2 (5.7)	34.4 (6.1)	34.0 (5.5)	34.5 (6.2)
hct_first, mean (SD)	0	34.4 (5.7)	34.6 (6.0)	34.3 (5.5)	34.7 (6.3)
hct_mean, mean (SD)	0	34.2 (5.5)	34.5 (5.8)	34.1 (5.2)	34.5 (5.9)
hct_median, mean (SD)	0	34.2 (5.5)	34.4 (5.8)	34.1 (5.3)	34.4 (6.0)
hct_std, mean (SD)	0	1.9 (1.3)	1.8 (1.2)	1.9 (1.2)	1.9 (1.4)

		Missing	Overall	Test	Train	Validation
hct_quantile_90, mean (SD)		0	35.3 (5.5)	35.5 (5.8)	35.1 (5.2)	35.6 (5.9)
hct_quantile_10, mean (SD)		0	33.2 (5.8)	33.5 (6.0)	33.0 (5.5)	33.3 (6.3)
k_last, mean (SD)		0	4.2 (0.6)	4.2 (0.6)	4.2 (0.7)	4.1 (0.6)
k_first, mean (SD)		0	4.4 (0.7)	4.4 (0.7)	4.4 (0.6)	4.3 (0.6)
k_mean, mean (SD)		0	4.3 (0.6)	4.3 (0.6)	4.3 (0.6)	4.2 (0.6)
k_median, mean (SD)		0	4.2 (0.6)	4.3 (0.6)	4.2 (0.6)	4.2 (0.6)
k_std, mean (SD)		0	0.3 (0.2)	0.3 (0.2)	0.3 (0.2)	0.3 (0.2)
k_quantile_90, mean (SD)		0	4.5 (0.6)	4.5 (0.7)	4.5 (0.6)	4.5 (0.6)
k_quantile_10, mean (SD)		0	4.0 (0.6)	4.0 (0.6)	4.0 (0.6)	4.0 (0.6)
na_last, mean (SD)		0	138.0 (4.8)	137.7 (4.7)	138.0 (4.8)	138.3 (5.0)
na_first, mean (SD)		0	136.9 (4.9)	136.8 (4.9)	136.9 (4.9)	137.0 (4.6)
na_mean, mean (SD)		0	137.5 (4.3)	137.4 (4.2)	137.6 (4.3)	137.7 (4.3)
na_median, mean (SD)		0	137.6 (4.3)	137.4 (4.3)	137.6 (4.4)	137.7 (4.3)
na_std, mean (SD)		0	2.2 (1.3)	2.2 (1.3)	2.2 (1.3)	2.2 (1.4)
na_quantile_90, mean (SD)		0	139.3 (4.4)	139.1 (4.3)	139.3 (4.4)	139.5 (4.6)
na_quantile_10, mean (SD)		0	135.8 (4.6)	135.6 (4.7)	135.8 (4.6)	136.0 (4.6)
EGFR_last, mean (SD)		0	40.7 (21.4)	40.6 (20.9)	40.3 (21.0)	42.3 (23.4)
EGFR_first, mean (SD)		0	42.5 (22.4)	42.8 (22.2)	42.0 (21.9)	44.2 (24.2)
EGFR_mean, mean (SD)		0	42.0 (21.8)	42.3 (21.7)	41.5 (21.3)	43.7 (23.6)
EGFR_median, mean (SD)		0	41.9 (21.8)	42.3 (21.8)	41.4 (21.3)	43.5 (23.6)
EGFR_quantile_90, mean (SD)		0	44.1 (23.2)	44.5 (23.5)	43.5 (22.6)	45.9 (25.0)
EGFR_quantile_10, mean (SD)		0	40.0 (20.7)	40.1 (20.5)	39.5 (20.4)	41.5 (22.4)
EGFR_std, mean (SD)		0	3.8 (3.3)	4.0 (4.0)	3.8 (3.0)	3.9 (3.4)
crea_d_point, mean (SD)		0	2.3 (1.1)	2.4 (1.2)	2.3 (1.1)	2.4 (1.2)
crea_diff, mean (SD)		0	0.6 (0.4)	0.6 (0.4)	0.6 (0.4)	0.6 (0.4)
time_to_treat, mean (SD)		0	7.7 (9.7)	8.1 (9.9)	7.4 (9.7)	8.1 (9.6)
1_ind, n (%)	1.0	0	2156 (100.0)	530 (100.0)	1304 (99.9)	322 (100.0)
	0.0		1 (0.0)		1 (0.1)	
2_ind, n (%)	0.0	0	449 (20.8)	100 (18.9)	281 (21.5)	68 (21.1)
	1.0		1708 (79.2)	430 (81.1)	1024 (78.5)	254 (78.9)
3_ind, n (%)	0.0	0	1198 (55.5)	290 (54.7)	729 (55.9)	179 (55.6)
	1.0		959 (44.5)	240 (45.3)	576 (44.1)	143 (44.4)
4_ind, n (%)	0.0	0	1782 (82.6)	440 (83.0)	1071 (82.1)	271 (84.2)
	1.0		375 (17.4)	90 (17.0)	234 (17.9)	51 (15.8)
5_ind, n (%)	0.0	0	1736 (80.5)	419 (79.1)	1069 (81.9)	248 (77.0)
	1.0		421 (19.5)	111 (20.9)	236 (18.1)	74 (23.0)
6_ind, n (%)	0.0	0	1996 (92.5)	490 (92.5)	1211 (92.8)	295 (91.6)
	1.0		161 (7.5)	40 (7.5)	94 (7.2)	27 (8.4)
7_ind, n (%)	0.0	0	2143 (99.4)	528 (99.6)	1298 (99.5)	317 (98.4)
	1.0		14 (0.6)	2 (0.4)	7 (0.5)	5 (1.6)
8_ind, n (%)	0.0	0	2148 (99.6)	529 (99.8)	1299 (99.5)	320 (99.4)
	1.0		9 (0.4)	1 (0.2)	6 (0.5)	2 (0.6)
1_ind_hosp, n (%)	0.0	0	524 (24.3)	124 (23.4)	320 (24.5)	80 (24.8)
	1.0		1633 (75.7)	406 (76.6)	985 (75.5)	242 (75.2)
2_ind_hosp, n (%)	0.0	0	524 (24.3)	125 (23.6)	304 (23.3)	95 (29.5)
	1.0		1633 (75.7)	405 (76.4)	1001 (76.7)	227 (70.5)
3_ind_hosp, n (%)	0.0	0	1226 (56.8)	296 (55.8)	745 (57.1)	185 (57.5)
	1.0		931 (43.2)	234 (44.2)	560 (42.9)	137 (42.5)
4_ind_hosp, n (%)	0.0	0	1674 (77.6)	409 (77.2)	1008 (77.2)	257 (79.8)
	1.0		483 (22.4)	121 (22.8)	297 (22.8)	65 (20.2)
5_ind_hosp, n (%)	0.0	0	1781 (82.6)	441 (83.2)	1083 (83.0)	257 (79.8)
	1.0		376 (17.4)	89 (16.8)	222 (17.0)	65 (20.2)
6_ind_hosp, n (%)	0.0	0	2153 (99.8)	528 (99.6)	1303 (99.8)	322 (100.0)
	1.0		4 (0.2)	2 (0.4)	2 (0.2)	
7_ind_hosp, n (%)	0.0	0	2122 (98.4)	526 (99.2)	1287 (98.6)	309 (96.0)
	1.0		35 (1.6)	4 (0.8)	18 (1.4)	13 (4.0)
8_ind_hosp, n (%)	0.0	0	2140 (99.2)	527 (99.4)	1294 (99.2)	319 (99.1)
	1.0		17 (0.8)	3 (0.6)	11 (0.8)	3 (0.9)
returning_patient, n (%)	0.0	0	1215 (56.3)	296 (55.8)	740 (56.7)	179 (55.6)
	1.0		756 (35.0)	192 (36.2)	453 (34.7)	111 (34.5)
	2.0		157 (7.3)	35 (6.6)	97 (7.4)	25 (7.8)
relative_date_first_tn, mean (SD)	3.0		29 (1.3)	7 (1.3)	15 (1.1)	7 (2.2)
	0	0	0.6 (2.2)	0.6 (2.3)	0.5 (2.0)	0.7 (3.0)
	1.0	0	90 (4.2)	22 (4.2)	49 (3.8)	19 (5.9)
relative_date_first_hb, n (%)	0		2044 (94.8)	501 (94.5)	1242 (95.2)	301 (93.5)
	2.0		16 (0.7)	6 (1.1)	9 (0.7)	1 (0.3)
	3.0		4 (0.2)	1 (0.2)	3 (0.2)	
	12.0		1 (0.0)		1 (0.1)	

		Missing	Overall	Test	Train	Validation
		8.0	1 (0.0)		1 (0.1)	
		6.0	1 (0.0)			1 (0.3)
		0.0	2041 (94.6)	499 (94.2)	1241 (95.1)	301 (93.5)
		1.0	90 (4.2)	22 (4.2)	49 (3.8)	19 (5.9)
		2.0	19 (0.9)	8 (1.5)	10 (0.8)	1 (0.3)
		3.0	4 (0.2)	1 (0.2)	3 (0.2)	
		12.0	1 (0.0)		1 (0.1)	
		8.0	1 (0.0)		1 (0.1)	
		6.0	1 (0.0)			1 (0.3)
relative_date_first_rdw, mean (SD)		0	1.5 (1.8)	1.5 (1.7)	1.6 (1.8)	1.5 (1.6)
		0.0	2041 (94.6)	499 (94.2)	1241 (95.1)	301 (93.5)
		1.0	90 (4.2)	22 (4.2)	49 (3.8)	19 (5.9)
		2.0	19 (0.9)	8 (1.5)	10 (0.8)	1 (0.3)
		3.0	4 (0.2)	1 (0.2)	3 (0.2)	
		12.0	1 (0.0)		1 (0.1)	
		8.0	1 (0.0)		1 (0.1)	
		6.0	1 (0.0)			1 (0.3)
relative_date_first_wbc, n (%)						
		0	2.0 (1.6)	1.9 (1.3)	2.1 (1.7)	2.0 (1.4)
		0	1.6 (1.6)	1.7 (1.7)	1.7 (1.6)	1.5 (1.4)
		0	2.1 (1.8)	2.1 (1.8)	2.2 (2.0)	2.0 (1.4)
		0	2.2 (1.9)	2.2 (2.0)	2.2 (2.0)	2.0 (1.4)
		0	2.1 (1.8)	2.1 (1.8)	2.2 (2.0)	2.1 (1.4)
		0	2.1 (1.7)	2.1 (1.6)	2.1 (1.7)	2.2 (1.7)
		0	3.5 (4.9)	3.7 (6.8)	3.4 (3.9)	3.6 (5.0)
		0	0.6 (1.1)	0.6 (1.2)	0.6 (1.1)	0.6 (1.1)
relative_date_first_alb, mean (SD)		0	2.0 (1.6)	1.9 (1.3)	2.1 (1.7)	2.0 (1.4)
relative_date_first_ast, mean (SD)		0	1.6 (1.6)	1.7 (1.7)	1.7 (1.6)	1.5 (1.4)
relative_date_first_ggt, mean (SD)		0	2.1 (1.8)	2.1 (1.8)	2.2 (2.0)	2.0 (1.4)
relative_date_first_alkp, mean (SD)		0	2.1 (1.8)	2.1 (1.8)	2.2 (2.0)	2.1 (1.4)
relative_date_first_p, mean (SD)		0	2.1 (1.7)	2.1 (1.6)	2.1 (1.7)	2.2 (1.7)
relative_date_first_crp, mean (SD)		0	3.5 (4.9)	3.7 (6.8)	3.4 (3.9)	3.6 (5.0)
relative_date_first_cl, mean (SD)		0	0.6 (1.1)	0.6 (1.2)	0.6 (1.1)	0.6 (1.1)
		0.0	2052 (95.1)	503 (94.9)	1246 (95.5)	303 (94.1)
		1.0	86 (4.0)	23 (4.3)	46 (3.5)	17 (5.3)
		2.0	12 (0.6)	3 (0.6)	8 (0.6)	1 (0.3)
		3.0	4 (0.2)	1 (0.2)	3 (0.2)	
		6.0	2 (0.1)		1 (0.1)	1 (0.3)
		8.0	1 (0.0)		1 (0.1)	
d_point_relative_date, mean (SD)		0	3.9 (4.3)	4.0 (4.7)	3.8 (3.8)	4.2 (5.1)
Weight_missing, n (%)		0	1041 (48.3)	270 (50.9)	616 (47.2)	155 (48.1)
		1	1116 (51.7)	260 (49.1)	689 (52.8)	167 (51.9)
Hight_missing, n (%)		0	497 (23.0)	122 (23.0)	293 (22.5)	82 (25.5)
		1	1660 (77.0)	408 (77.0)	1012 (77.5)	240 (74.5)
bmi_missing, n (%)		0	496 (23.0)	121 (22.8)	293 (22.5)	82 (25.5)
		1	1661 (77.0)	409 (77.2)	1012 (77.5)	240 (74.5)
firsttemp_missing, n (%)		0	2153 (99.8)	529 (99.8)	1302 (99.8)	322 (100.0)
		1	4 (0.2)	1 (0.2)	3 (0.2)	
tn_missing, n (%)		0	1862 (86.3)	454 (85.7)	1131 (86.7)	277 (86.0)
		1	295 (13.7)	76 (14.3)	174 (13.3)	45 (14.0)
hb_missing, n (%)		0	2156 (100.0)	530 (100.0)	1304 (99.9)	322 (100.0)
		1	1 (0.0)		1 (0.1)	
first_mcv_missing, n (%)		0	2156 (100.0)	530 (100.0)	1304 (99.9)	322 (100.0)
		1	1 (0.0)		1 (0.1)	
first_rdw_missing, n (%)		0	2080 (96.4)	516 (97.4)	1254 (96.1)	310 (96.3)
		1	77 (3.6)	14 (2.6)	51 (3.9)	12 (3.7)
wbc_missing, n (%)		0	2156 (100.0)	530 (100.0)	1304 (99.9)	322 (100.0)
		1	1 (0.0)		1 (0.1)	
alb_missing, n (%)		0	2084 (96.6)	510 (96.2)	1263 (96.8)	311 (96.6)
		1	73 (3.4)	20 (3.8)	42 (3.2)	11 (3.4)
ast_missing, n (%)		0	2073 (96.1)	508 (95.8)	1254 (96.1)	311 (96.6)
		1	84 (3.9)	22 (4.2)	51 (3.9)	11 (3.4)
alt_missing, n (%)		0	2048 (94.9)	507 (95.7)	1236 (94.7)	305 (94.7)
		1	109 (5.1)	23 (4.3)	69 (5.3)	17 (5.3)
ggt_missing, n (%)		0	2040 (94.6)	503 (94.9)	1232 (94.4)	305 (94.7)
		1	117 (5.4)	27 (5.1)	73 (5.6)	17 (5.3)
alkp_missing, n (%)		0	2048 (94.9)	506 (95.5)	1237 (94.8)	305 (94.7)
		1	109 (5.1)	24 (4.5)	68 (5.2)	17 (5.3)
first_p_missing, n (%)		0	2076 (96.2)	517 (97.5)	1253 (96.0)	306 (95.0)
		1	81 (3.8)	13 (2.5)	52 (4.0)	16 (5.0)
crp_missing, n (%)		0	327 (15.2)	79 (14.9)	203 (15.6)	45 (14.0)
		1	1830 (84.8)	451 (85.1)	1102 (84.4)	277 (86.0)
first_cl_missing, n (%)		0	2049 (95.0)	502 (94.7)	1237 (94.8)	310 (96.3)
		1	108 (5.0)	28 (5.3)	68 (5.2)	12 (3.7)
glu_missing, n (%)		0	2157 (100.0)	530 (100.0)	1305 (100.0)	322 (100.0)
dbp_ts_missing, n (%)		0	2145 (99.4)	525 (99.1)	1298 (99.5)	322 (100.0)

		Missing	Overall	Test	Train	Validation
	1		12 (0.6)	5 (0.9)	7 (0.5)	
sbp_ts_missing, n (%)	0	0	2145 (99.4)	525 (99.1)	1298 (99.5)	322 (100.0)
	1		12 (0.6)	5 (0.9)	7 (0.5)	
pp_ts_missing, n (%)	0	0	2145 (99.4)	525 (99.1)	1298 (99.5)	322 (100.0)
	1		12 (0.6)	5 (0.9)	7 (0.5)	
MAP_ts_missing, n (%)	0	0	2145 (99.4)	525 (99.1)	1298 (99.5)	322 (100.0)
	1		12 (0.6)	5 (0.9)	7 (0.5)	
hr_ts_missing, n (%)	0	0	2146 (99.5)	526 (99.2)	1298 (99.5)	322 (100.0)
	1		11 (0.5)	4 (0.8)	7 (0.5)	
bnp_ts_missing, n (%)	0	0	1214 (56.3)	307 (57.9)	727 (55.7)	180 (55.9)
	1		943 (43.7)	223 (42.1)	578 (44.3)	142 (44.1)
bun_ts_missing, n (%)	0	0	2157 (100.0)	530 (100.0)	1305 (100.0)	322 (100.0)
	1		2153 (99.8)	529 (99.8)	1302 (99.8)	322 (100.0)
hct_ts_missing, n (%)	0	0	4 (0.2)	1 (0.2)	3 (0.2)	
k_ts_missing, n (%)	0	0	2139 (99.2)	524 (98.9)	1296 (99.3)	319 (99.1)
	1		18 (0.8)	6 (1.1)	9 (0.7)	3 (0.9)
na_ts_missing, n (%)	0	0	2156 (100.0)	530 (100.0)	1304 (99.9)	322 (100.0)
	1		1 (0.0)		1 (0.1)	
relative_date_first_tn_missing, n (%)	0	0	1873 (86.8)	457 (86.2)	1137 (87.1)	279 (86.6)
	1		284 (13.2)	73 (13.8)	168 (12.9)	43 (13.4)
relative_date_first_hb_missing, n (%)	0	0	2156 (100.0)	530 (100.0)	1304 (99.9)	322 (100.0)
	1		1 (0.0)		1 (0.1)	
relative_date_first_mcv_missing, n (%)	0	0	2156 (100.0)	530 (100.0)	1304 (99.9)	322 (100.0)
	1		1 (0.0)		1 (0.1)	
relative_date_first_rdw_missing, n (%)	0	0	2080 (96.4)	516 (97.4)	1254 (96.1)	310 (96.3)
	1		77 (3.6)	14 (2.6)	51 (3.9)	12 (3.7)
relative_date_first_wbc_missing, n (%)	0	0	2156 (100.0)	530 (100.0)	1304 (99.9)	322 (100.0)
	1		1 (0.0)		1 (0.1)	
relative_date_first_alb_missing, n (%)	0	0	2084 (96.6)	510 (96.2)	1263 (96.8)	311 (96.6)
	1		73 (3.4)	20 (3.8)	42 (3.2)	11 (3.4)
relative_date_first_ast_missing, n (%)	0	0	2074 (96.2)	508 (95.8)	1254 (96.1)	312 (96.9)
	1		83 (3.8)	22 (4.2)	51 (3.9)	10 (3.1)
relative_date_first_alt_missing, n (%)	0	0	2048 (94.9)	507 (95.7)	1236 (94.7)	305 (94.7)
	1		109 (5.1)	23 (4.3)	69 (5.3)	17 (5.3)
relative_date_first_ggt_missing, n (%)	0	0	2040 (94.6)	503 (94.9)	1232 (94.4)	305 (94.7)
	1		117 (5.4)	27 (5.1)	73 (5.6)	17 (5.3)
relative_date_first_alkp_missing, n (%)	0	0	2048 (94.9)	506 (95.5)	1237 (94.8)	305 (94.7)
	1		109 (5.1)	24 (4.5)	68 (5.2)	17 (5.3)
relative_date_first_p_missing, n (%)	0	0	2076 (96.2)	517 (97.5)	1253 (96.0)	306 (95.0)
	1		81 (3.8)	13 (2.5)	52 (4.0)	16 (5.0)
relative_date_first_crp_missing, n (%)	0	0	327 (15.2)	79 (14.9)	203 (15.6)	45 (14.0)
	1		1830 (84.8)	451 (85.1)	1102 (84.4)	277 (86.0)
relative_date_first_cl_missing, n (%)	0	0	2150 (99.7)	528 (99.6)	1300 (99.6)	322 (100.0)
	1		7 (0.3)	2 (0.4)	5 (0.4)	
outcome, mean (SD)		0	0.2 (1.5)	0.2 (1.6)	0.2 (1.5)	0.2 (1.7)
crea_outcome, mean (SD)		0	2.2 (1.3)	2.2 (1.4)	2.2 (1.2)	2.2 (1.4)
c1, n (%)	0.0	0	1065 (49.4)	252 (47.5)	651 (49.9)	162 (50.3)
	1.0		1092 (50.6)	278 (52.5)	654 (50.1)	160 (49.7)

**Table S.1:** Covariants characteristics, used for policy construction. See Table S.2 for descriptions of feature name abbreviations.


Abbreviation	Description
firstadm	Indicator of first admission (0/1)
admyear	Year of admission
age	Age in years (mean and standard deviation)
gender	Gender (0: Male, 1: Female)
Weight, Hight	Patient weight (kg) and height (cm)
bmi	Body Mass Index
htn	Hypertension indicator
firsttemp	First recorded temperature
dm	Diabetes mellitus indicator
smk	Smoking status (0: non-smoker, 1: former, 2: current)

Abbreviation	Description
ihd	Ischemic heart disease
vhd	Valvular heart disease
af	Atrial fibrillation
Hyperlipidemia	High lipid levels
copd	Chronic obstructive pulmonary disease
crf	Chronic renal failure
bb	Beta Blockers
acei	ACE Inhibitors
arf	Angiotensin Receptor Blockers
antiplt	Antiplatelet agents
anticoagulants	Anticoagulant use
furosemide	Use of furosemide
zaroxolin	Zaroxolin administration
thiamin	Thiamin (Vitamin B1) administration
insulin	Insulin use
tn, hb	Laboratory tests (e.g., total nitrogen, hemoglobin)
first_mcv, rdw	Mean corpuscular volume and red cell distribution width
wbc	White blood cell count
alb	Albumin
ast, alt, ggt, alkp	Liver function tests (AST, ALT, GGT, ALKP)
first_p	First phosphorus level
crp	C-reactive protein
first_cl	First chloride level
glu	Glucose
dbp, sbp, pp	Blood pressure and pulse pressure (various measures)
MAP	Mean arterial pressure
hr	Heart rate
bnp	Brain natriuretic peptide
bun	Blood urea nitrogen
crea	Creatinine
hct	Hematocrit
k, na	Potassium, Sodium
EGFR	Estimated glomerular filtration rate
crea_d_point/diff	Measures related to creatinine changes
time.to.treat	Time (in days) to treatment
<b>Drug Indicators:</b>	
1_ind	Diuretics (administered during hospitalization)
2_ind	Beta Blockers
3_ind	ACE Inhibitors
4_ind	Angiotensin Receptor Blockers
5_ind	MRA (Mineralocorticoid Receptor Antagonists)
6_ind	Fluids
7_ind	Diuretics and ARB
8_ind	Diuretics and ACE
N_ind_hosp	Indicates whether the patient had the drug N at the time of admission
returning_patient	Indicator for returning patients, derived from the admission serial number. Assigns: 0 for (0,1] (first admission), 1 for (1,4] (1-4 previous admission), 2 for (4,8] (4-8 previous admissions), and 3 for (8,∞) (8+ previous admissions).
relative_date_first_*	Relative time (in days) from admission to the first measurement of various tests
Weight_missing, Hight_missing, bmi_missing, etc.	Indicators for missing data in the respective variables
crea_outcome	Last creatinine value recorded within 7 days from the decision point
outcome (RTB)	Indicator of return to baseline creatinine level
c1	Treatment adjustment for diuretics (increase or decrease)

**Table S.2:** Dictionary of Abbreviations for Data Description.




## F.2 Case study: propensity score model




**Figure S.1:** The features with the highest SHAP values in the XGBoost propensity score model.

	Brier	AUROC	Accuracy	Precision	Recall	Fscore
Train	0.153	0.878	0.798	0.8	0.799	0.798
Validation	0.220	0.698	0.671	0.674	0.670	0.669

**Table S.3:** Goodness-of-fit metrics of the XGBoost model used for propensity score estimation. The results on the train and validation set.



(a) Propensity distribution before trimming




(b) Propensity distribution after trimming

**Figure S.2**


### F.3 Simulation results for case study

Following Section 4.4.3, Fig. S.3 and Fig. S.4 show a rank-graph and a box-plot graph of the policy values of various policies on the simulated data. For both policy value estimations, we use plain Logistic Regression (with L2 regularization) to estimate the “true” propensity score on all the data ( $p^*(t = \pi(x)|x)$ ). In DR evaluation the XGB T-learner predication was used as the plug-in estimator  $\hat{g}^{\pi(x)}$ .

We see that we have a clear separation between the policy value of the “Doctors” treatment assignment, the policy values of two policies that performed poorly (Causal Forest and propensity-score based policy), and the other learned policies, which all performed well (note that lower is better in this case).



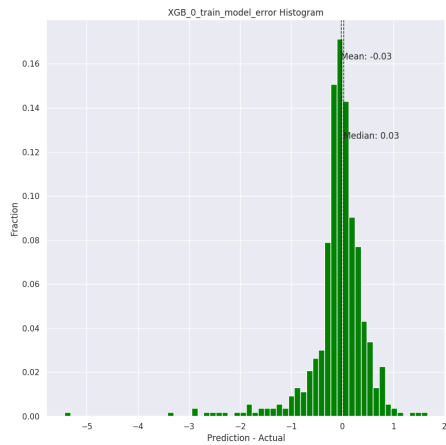
**Figure S.3:** Policy value rank graph for simulation data. Lower values are better.




**Figure S.4:** Box-plots of simulation policy values, per policy. Lower values are better.

## F.4 Outcome models Results

In Fig. S.5 we present the error distribution of the XGBoost model on the train set. The results indicate that the model mostly predicts well.




(a) Error distributions of XGBoost model that predicts creatinine in the “Decrease” population.



(b) Error distributions of XGBoost model that predicts creatinine in the “Increase” population.

Figure S.5

In Fig. S.6 we present the most important features, according to SHAP, for predicting creatinine in the underline XGBoost outcome models.



(a) Important features, according to SHAP values of XGB model that predicts creatinine in the “Decrease” population.

(b) Important features, according to SHAP values of XGB model that predicts creatinine in the “Increase” population.

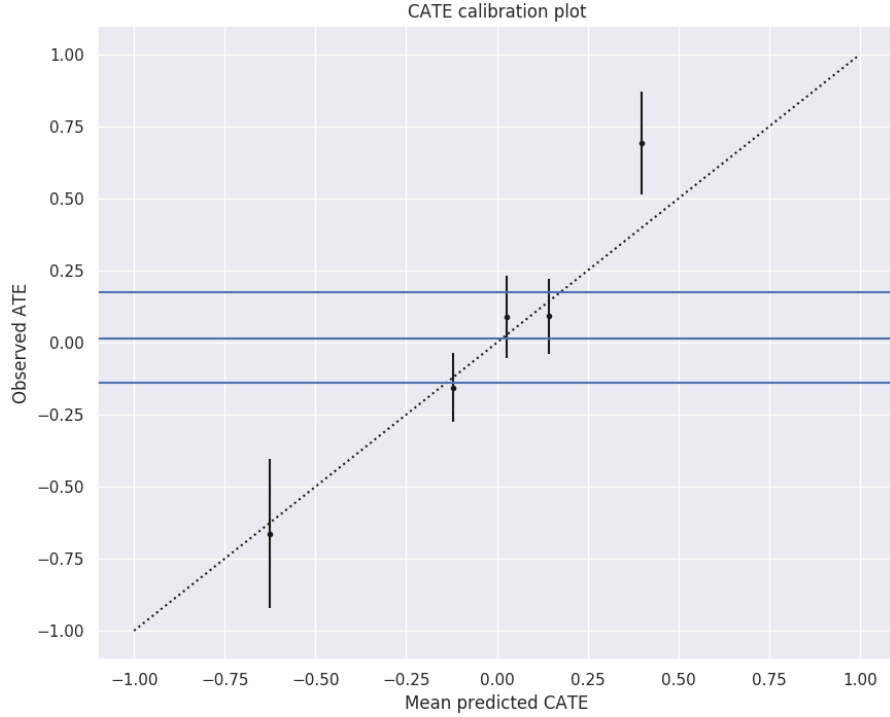
Figure S.6

	BART	XGBoost	Ridge	Lasso	Average	Causal Forest	Majority	Consensus	Propensity	DragoNet
<b>BART</b>	1.000	0.583	0.375	0.246	0.836	0.413	0.300	0.515	0.124	0.344
<b>XGBoost</b>	0.583	1.000	0.221	0.370	0.753	0.510	0.284	0.542	0.042	0.305
<b>Ridge</b>	0.375	0.221	1.000	0.104	0.686	0.165	0.179	0.353	-0.004	0.259
<b>Lasso</b>	0.246	0.370	0.104	1.000	0.470	0.623	0.216	0.513	0.060	0.250
<b>Average</b>	0.836	0.753	0.686	0.470	1.000	0.557	0.346	0.657	0.079	0.412
<b>Causal Forest</b>	0.413	0.510	0.165	0.623	0.557	1.000	0.279	0.520	0.117	0.388
<b>Majority</b>	0.300	0.284	0.179	0.216	0.346	0.279	1.000	0.263	0.474	0.181
<b>Consensus</b>	0.515	0.542	0.353	0.513	0.657	0.520	0.263	1.000	0.067	0.317
<b>Propensity</b>	0.124	0.042	-0.004	0.060	0.079	0.117	0.474	0.067	1.000	0.020
<b>DragoNet</b>	0.344	0.305	0.259	0.250	0.412	0.388	0.181	0.317	0.020	1.000

Table S.4: Pearson correlation between the CATE estimates on the train set

	BART	XGBoost	Ridge	Lasso	Average	Causal Forest	Majority	Consensus	Propensity	DragoNet
<b>BART</b>	1.000	0.368	0.274	0.170	0.632	0.239	0.251	0.456	0.081	0.284
<b>XGBoost</b>	0.368	1.000	0.169	0.270	0.534	0.307	0.250	0.503	0.022	0.252
<b>Ridge</b>	0.274	0.169	1.000	0.079	0.494	0.134	0.183	0.349	-0.003	0.249
<b>Lasso</b>	0.170	0.270	0.079	1.000	0.346	0.373	0.181	0.456	0.020	0.176
<b>Average</b>	0.632	0.534	0.494	0.346	1.000	0.359	0.286	0.613	0.047	0.353
<b>Causal Forest</b>	0.239	0.307	0.134	0.373	0.359	1.000	0.248	0.464	0.071	0.305
<b>Majority</b>	0.251	0.250	0.183	0.181	0.286	0.248	1.000	0.263	0.386	0.181
<b>Consensus</b>	0.456	0.503	0.349	0.456	0.613	0.464	0.263	1.000	0.052	0.317
<b>Propensity</b>	0.081	0.022	-0.003	0.020	0.047	0.071	0.386	0.052	1.000	0.014
<b>DragoNet</b>	0.284	0.252	0.249	0.176	0.353	0.305	0.181	0.317	0.014	1.000

Table S.5: Kendall correlation between the CATE estimates on the train set



**Figure S.7:** CATE calibration graph of XGBoost T-learner model.


## F.5 Deferral Insights

To further understand the population that got deferral, in Section 4.4.6, a policy summarization is required. Towards this goal, we train a  $L_1$  regularized logistic regression model, on the label whether the patient was part of the deferral set. In Fig. S.8 we present the SHAP summary plot [79], with the 20 most important values. Then, in Table S.6 we present the group statistics difference, based on the features in the summary plot. The analysis suggests that deferred patients have, among other things, lower values in the 90th quantile of creatinine ('crea\_quantile\_90'), lower 10th quantile of diastolic blood pressure ('dbp\_quantile\_10'), and a lower proportion of first-time admissions (51.7 vs. 58.4). Conversely, the standard deviation of estimated glomerular filtration rate ('EGFR\_std') is higher in the deferred group, suggesting that greater variability in kidney function contribute to increased uncertainty in treatment recommendations.

## F.6 Policy value results

As noted in Section 4.1, the aim of this test case is constructing individual-level treatment policy. This recommendation is based on the CATE value, to determine which treatment would better affect the patient state, which in this case means having a higher predicted RTB ratio. We examined 12 such policies – 5 model-based (XGBoost, Causal Forest, Ridge, Lasso, BART), 3 ensemble policies (Average, Majority, Consensus based on the 5 aforementioned models), 2 based on fixed rules (either Increase or Decrease dosage to all), a policy based on the propensity score, and current treatment policy (i.e. the policy actually observed in the data). As detailed in Section 3, we estimated the policy value for each rule in various methods on the held out data.

For each policy we performed a 10K bootstrap where for each round we randomly select patients (with replacements) and evaluate the policy value with both the IPW and DR estimators (see Section 3.10 and Appendix D). We use logistic regression with  $L_2$  regularization to estimate the propensity score on all the data. In DR evaluation the XGBoost T-learner predictions were used as the plug-in estimator  $\hat{g}^{\pi(x)}$ . We present the results in Fig. S.9 ("Inclusive" deferral set) and Fig. S.10 ("Conservative" deferral set).



**Figure S.8:** SHAP summary plot of the most important features of a logistic regression model trained on the deferral assignment.

		Missing	Overall	Grouped by deferred	
				Rec	Def
n			1139	822	317
na_last, mean (SD)	0		137.8 (4.8)	137.8 (4.9)	137.8 (4.6)
crea_quantile_90, mean (SD)	0		1.8 (1.0)	1.9 (1.1)	1.7 (0.7)
dbp_quantile_10, mean (SD)	0		57.2 (10.9)	57.4 (11.1)	56.8 (10.6)
hct_quantile_90, mean (SD)	0		35.0 (5.3)	35.0 (5.5)	35.0 (5.0)
bun_median, mean (SD)	0		40.0 (23.1)	40.6 (23.7)	38.6 (21.3)
hr_quantile_90, mean (SD)	0		91.6 (18.5)	92.0 (18.6)	90.3 (18.4)
bun_first, mean (SD)	0		39.5 (23.5)	40.2 (24.1)	37.8 (21.8)
MAP_quantile_10, mean (SD)	0		77.0 (12.3)	77.1 (12.5)	76.6 (11.8)
EGFR_last, mean (SD)	0		40.0 (21.1)	39.5 (21.8)	41.3 (19.0)
firstadm, n (%)	0	0	495 (43.5)	342 (41.6)	153 (48.3)
	1		644 (56.5)	480 (58.4)	164 (51.7)
na_quantile_90, mean (SD)	0		139.2 (4.6)	139.1 (4.7)	139.3 (4.3)
hct_last, mean (SD)	0		33.9 (5.6)	33.9 (5.7)	34.1 (5.3)
EGFR_std, mean (SD)	0		3.8 (3.0)	3.6 (3.0)	4.1 (3.0)
first_p, mean (SD)	0		4.4 (1.2)	4.5 (1.2)	4.2 (1.1)
hr_std, mean (SD)	0		10.9 (6.1)	10.8 (5.9)	11.1 (6.7)
k_last, mean (SD)	0		4.2 (0.7)	4.2 (0.7)	4.2 (0.7)
hb, mean (SD)	0		11.2 (1.9)	11.2 (1.9)	11.3 (1.8)
tn_missing, n (%)	0	0	981 (86.1)	695 (84.5)	286 (90.2)
	1		158 (13.9)	127 (15.5)	31 (9.8)
bun_quantile_10, mean (SD)	0		38.1 (22.8)	38.6 (23.4)	36.9 (21.3)
k_first, mean (SD)	0		4.4 (0.6)	4.4 (0.7)	4.4 (0.6)

**Table S.6:** The group differences between the deferred (“Def”) group and the vs the patients that would get a recommendation (“Rec”) and the union of both (“Overall”), on the train set. The chosen features are based on the largest SHAP values of a logistic regression model.




	$n$	Mean	Std	Min	25%	50%	75%	Max
<b>Current</b>	10000	22.00%	6.94%	-5.87%	17.27%	22.07%	26.68%	49.24%
<b>Random</b>	10000	23.72%	11.05%	-31.00%	16.51%	24.02%	31.32%	60.83%
<b>Keep/Increase</b>	10000	13.35%	11.91%	-36.48%	5.54%	13.56%	21.43%	59.53%
<b>Decrease</b>	10000	37.22%	9.32%	-5.19%	31.01%	37.61%	43.53%	68.16%
<b>XGBoost</b>	10000	40.53%	10.37%	-2.50%	33.67%	40.70%	47.60%	76.01%
<b>Ridge</b>	10000	45.45%	9.19%	10.11%	39.42%	45.61%	51.72%	77.73%
<b>DragonNet</b>	10000	36.59%	9.86%	-10.71%	29.89%	36.74%	43.37%	69.59%
<b>Lasso</b>	10000	30.88%	10.02%	-9.47%	24.39%	31.08%	37.70%	64.49%
<b>Causal Forset</b>	10000	21.63%	11.53%	-26.73%	14.10%	22.10%	29.63%	65.46%
<b>BART</b>	10000	15.21%	12.34%	-38.57%	6.87%	15.61%	23.83%	63.08%
<b>Propensity</b>	10000	26.20%	8.43%	-10.10%	20.62%	26.27%	31.86%	56.83%
<b>Average</b>	10000	35.40%	10.90%	-10.44%	28.15%	35.64%	43.00%	76.30%
<b>Majority</b>	10000	27.38%	8.14%	-7.80%	22.08%	27.54%	32.87%	56.31%
<b>Consensus</b>	10000	34.38%	9.58%	-7.81%	28.08%	34.56%	40.91%	70.33%


**Table S.7:** Doubly robust (DR) estimates of policy values over 10K bootstraps rounds on held-out data, showing all policies: *Keep/Increase*: all patients given “increase”, *Current*: current treatment by doctors, *Random*: randomly assigning treatment at the same proportion as current treatment, *Decrease*: all patient given “decrease”, *XGBoost*: T-learner of XGBoost models, *Ridge*: T-learner of Ridge models, *DrangoNet*: A policy based on DragoNet estimator, *Causal Forest*: A policy based on Causal forest estimator, *BART*: T-learner of BART model, *Lasso*: T-learner of Lasso models, *Propensity*: a policy based on patient’s propensity score. Ensemble models – *Average*: a policy based on the average of the above models, *Majority*: a policy based on majority vote of the above models, and *Consensus*: a policy based on consensus between the above models

	$n$	Mean	Std	Min	25%	50%	75%	Max
<b>Current</b>	10000	22.00%	6.94%	-5.87%	17.27%	22.07%	26.68%	49.24%
<b>Random</b>	10000	25.41%	10.95%	-28.42%	18.29%	25.70%	32.88%	64.96%
<b>Keep/Increase</b>	10000	13.32%	11.87%	-35.76%	5.66%	13.51%	21.30%	62.24%
<b>Decrease</b>	10000	41.15%	9.09%	1.29%	35.18%	41.33%	47.46%	76.50%
<b>XGBoost</b>	10000	40.60%	9.96%	0.55%	34.05%	40.77%	47.29%	75.83%
<b>Ridge</b>	10000	43.47%	9.19%	5.99%	37.48%	43.60%	49.71%	78.35%
<b>DragonNet</b>	10000	37.97%	9.88%	-8.29%	31.44%	38.01%	44.67%	71.74%
<b>Lasso</b>	10000	31.26%	9.91%	-8.81%	24.75%	31.42%	37.89%	65.33%
<b>Causal Forset</b>	10000	26.12%	11.63%	-26.12%	18.70%	26.79%	34.04%	69.50%
<b>BART</b>	10000	15.91%	11.60%	-37.21%	8.17%	16.26%	24.05%	55.68%
<b>Propensity</b>	10000	31.81%	8.77%	-3.41%	25.87%	31.88%	37.71%	65.44%
<b>Average</b>	10000	35.45%	10.75%	-10.72%	28.35%	35.55%	42.85%	74.17%
<b>Majority</b>	10000	27.72%	8.06%	-9.30%	22.36%	27.94%	33.16%	55.45%
<b>Consensus</b>	10000	33.92%	9.32%	-7.32%	27.74%	34.00%	40.33%	68.25%


**Table S.8:** Inverse Propensity Weighting (IPW) estimates of policy values over 10K bootstraps rounds on held-out data, showing all policies: *Keep/Increase*: all patients given “increase”, *Current*: current treatment by doctors, *Random*: randomly assigning treatment at the same proportion as current treatment, *Decrease*: all patient given “decrease”, *XGBoost*: T-learner of XGBoost models, *Ridge*: T-learner of Ridge models, *DrangoNet*: A policy based on DragoNet estimator, *Causal Forest*: A policy based on Causal forest estimator, *BART*: T-learner of BART model, *Lasso*: T-learner of Lasso models, *Propensity*: a policy based on patient’s propensity score. Ensemble models – *Average*: a policy based on the average of the above models, *Majority*: a policy based on majority vote of the above models, and *Consensus*: a policy based on consensus between the above models



**Figure S.9:** Policy value box-plot, result of running 10K bootstraps evaluation on held-out data, showing all policies: *Keep/Increase*: all patients given “increase”, *Current*: current treatment by doctors, *Random*: randomly assigning treatment at the same proportion as current treatment, *Decrease*: all patient given “decrease”, *XGboost*: T-learner of XGBoost models, *Ridge*: T-learner of Ridge models, *DrangoNet*: A policy based on DrangoNet estimator, *Causal Forest*: A policy based on Causal forest estimator, *BART*: T-learner of BART model, *Lasso*: T-learner of Lasso models, *Propensity*: a policy based on patient’s propensity score. Ensemble models – *Average*: a policy based on the average of the above models, *Majority*: a policy based on majority vote of the above models, and *Consensus*: a policy based on consensus between the above models



**Figure S.10:** Policy value box-plot, result of running 10K bootstraps evaluation on held-out data, showing all policies, under conservative deferral set: *Keep/Increase*: all patients given “increase”, *Current*: current treatment by doctors, *Random*: randomly assigning treatment at the same proportion as current treatment, *Decrease*: all patient given “decrease”, *XGboost*: T-learner of XGBoost models, *Ridge*: T-learner of Ridge models, *DrangoNet*: A policy based on DragoNet estimator, *Causal Forest*: A policy based on Causal forest estimator, *BART*: T-learner of BART model, *Propensity*: a policy based on patient’s propensity score, *Average*: a policy based on the average of the above models.



**Figure S.11:** “Outcome Tree” under conservative deferral set: A graph representing the mean RTB value for each policy group, using XGBoost T-learner as the policy. Results on the “Conservative” set. The gray boxes represent the number of patients in each subgroup, and the green boxes represent their mean RTB value and SEM in parentheses.