
Classifying Treatment Responders Under Causal Effect Monotonicity

Nathan Kallus¹

Abstract

In the context of individual-level causal inference, we study the problem of predicting whether someone will respond or not to a treatment based on their features and past examples of features, treatment indicator (e.g., drug/no drug), and a binary outcome (e.g., recovery from disease). As a classification task, the problem is made difficult by not knowing the example outcomes under the opposite treatment indicators. We assume the effect is monotonic, as in advertising’s effect on a purchase or bail-setting’s effect on reappearance in court: either it would have happened regardless of treatment, not happened regardless, or happened only depending on exposure to treatment. Predicting whether the latter is latently the case is our focus. While previous work focuses on conditional average treatment effect estimation, formulating the problem as a classification task rather than an estimation task allows us to develop new tools more suited to this problem. By leveraging monotonicity, we develop new discriminative and generative algorithms for the responder-classification problem. We explore and discuss connections to corrupted data and policy learning. We provide an empirical study with both synthetic and real datasets to compare these specialized algorithms to standard benchmarks.

1. Introduction

In many domains where personalization is of interest, such as healthcare and marketing, a central problem is individual-level causal inference on treatment effects, which are the differences in outcome if a treatment is applied and if not applied. The aim is to learn a function that, given a rich set of features describing an individual, predicts the causal effect of an intervention on the individual, such as the ef-

fect of a pharmaceutical drug on their mortality or of an advertisement on whether they purchase the product. Compared to aggregate average causal effects on a whole population, such fine-grained predictions can better describe how an intervention would affect a specific individual and help determine whether it should be applied in their case. Learning such a function from either experimental or observational data has been the subject of much recent research. (Künzel et al., 2017; Shalit et al., 2017; Wager & Athey, 2017, among others; see Section 1.4). The key difficulty in this task arises due to what is often termed the *Fundamental Problem of Causal Inference*: that for any individual in the data, the data only contains the outcome given that the treatment was in fact applied or not, and it does not contain the counterfactual outcome under the opposite scenario. This difficulty arises in both experimental and observational data, although the former has the benefit of a priori eliminating any potential additional biases due to treatment selection via randomization.

In many applications, the outcomes of interest are binary. In medicine, we are often interested in mortality, recovery, readmission, or disease remission. In advertising, we are interested in whether or not a user purchases, visits, clicks, etc. The same holds in many other applications in domains ranging from criminal justice to education policymaking. In deciding whether to release a defendant on their own recognizance (and not require bail), a judge is interested in whether or not a given defendant will fail to reappear in court. In designing job training programs, a key outcome of interest is whether the recipient secures employment afterward or not. Moreover, in many applications and for many outcomes of interest, the treatment’s effect can only be monotonic: while it may or may not have any effect, exposing someone to an advertisement does not make them *less* likely to visit; while anticoagulants may have an unknown effect on an individual stroke patient’s mortality, it can only make the occurrence of a haemorrhage *more* likely; and requiring to post bail does not make a defendant *less* likely to reappear.

When outcomes are binary and effects are monotonic, the individual-level causal inference question boils down to just whether a given individual will respond to the treatment or not. In all of the above settings, either the outcome event of interest would have happened anyway, not happened

¹School of Operations Research and Information Engineering and Cornell Tech, Cornell University. Correspondence to: Nathan Kallus <kallus@cornell.edu>.

anyway, or happened if and only if treatment was applied. We call instances that fall in the latter category responders and those in the former two category non-responders. An instance can be finer than an individual because it refers to a particular realization, where an individual could have a probability of being in any one of these categories.

In this paper, we study the problem of learning to classify responders in settings with binary outcomes and monotonic effects. This is a unique classification problem as it suffers from the problem that the labels are not observed in the training data: due to the fundamental problem of causal inference, we do not know the counterfactual outcome bit and whether it is the same as the observed outcome or flipped. By leveraging monotonicity we develop a new discriminative approach based on minimizing a surrogate loss for the responder-classification task. Using a hinge loss and kernelizing the decision function, this gives rise to an algorithm we term RespSVM. We discuss the approach from a corrupted-label perspective as well as what happens if monotonicity fails. Based on the corrupted-label perspective, we further develop a new generative approach that gives rise to a new cross-entropy loss that we use in an algorithm we term RespNet. We then explore how these algorithms compare to standard benchmarks from individual-level effect estimation. Our empirical study includes both synthetic and real-data examples and shows that, when outcomes are binary and classifying response is of interest, specialized algorithms such the ones we develop can provide better performance.

1.1. Problem setup

We consider a population of instances where each instance is associated with the following random variables:

- $X \in \mathbb{R}^p$, features to be used to predict outcomes and treatment response (also known as baseline covariates);
- $Y(+1) \in \{-1, +1\}$, outcome if treatment is applied;
- $Y(-1) \in \{-1, +1\}$, outcome if treatment is *not* applied.

Note that we can also conceive of $Y(+1), Y(-1)$ as the potential outcomes of any two alternative interventions, $+1$ and -1 . Here we identify intervention $+1$ with applying the treatment and -1 with *not* applying treatment only for the sake of exposition. Note also that if we would rather consider an instance as having some *probability* of having some particular outcomes rather than having certain binary outcomes, we can just simply augment the population appropriately with each binary-outcome scenario.

The causal effect of treatment in each instance is defined as the difference in outcome if treatment is applied or not

$$\text{Causal effect: } Y(+1) - Y(-1)$$

Our standing assumption, as motivated in the introduction, is that the causal effect is nonnegative:

Assumption 1 (Monotonicity). $Y(+1) \geq Y(-1)$.

Note that if our assumed monotonicity went the other way (treatment can only decrease outcome or keep it the same), we can simply negate the outcome (*i.e.*, swap the physical meanings of having a $+1$ or -1 outcome) in order to conform to Assumption 1.

Under Assumption 1, we can exhaustively classify units into three categories:

- Responders: $Y(+1) = +1, Y(-1) = -1$
- Type-1 non-responders: $Y(+1) = Y(-1) = -1$
- Type-2 non-responders: $Y(+1) = Y(-1) = +1$

Assumption 1 simply eliminates the fourth possibility of having $Y(+1) = -1, Y(-1) = +1$. Notice that all responders have a causal effect of 2 and all non-responders have a causal effect of 0.

1.2. The data and the classification task

Letting

$$R = \begin{cases} +1 & Y(+1) > Y(-1) \text{ (responder)} \\ -1 & \text{Otherwise (non-responder)} \end{cases}$$

and

$$\rho(X) = \mathbb{P}(R = +1 \mid X),$$

we consider the classification task of predicting whether a unit is a responder or not. That is, the binary classification task with features X and binary label R .

The training data we have for this classification task does *not* consist of example pairs of features and labels, however. Instead, the training data consists of n observations of units that were either exposed or not to the treatment and the outcome corresponding to this exposure. Specifically, our observations are of the random variables

- $X \in \mathbb{R}^p$, features as before;
- $T \in \{-1, +1\}$, an indicator of whether the unit was exposed ($+1$) or not (-1) to treatment; and
- $Y = Y(T) \in \{-1, +1\}$, the corresponding outcome.

And, our training data consist of n observations, $X_1, T_1, Y_1, \dots, X_n, T_n, Y_n$, of the variables X, T, Y .

We focus on the case where this data came from an experiment. We therefore assume that treatment selection is unconfounded in that

$$Y(+1), Y(-1) \perp\!\!\!\perp T \mid X,$$

as would be the case under randomization. We further define the randomization probability

$$e(X) = \mathbb{P}(T = 1 \mid X) \quad \text{and} \quad Q = \frac{1}{2} + (e(X) - \frac{1}{2}) T.$$

For experimental data, $e(X)$ is known by design and is often constant, usually equal to $1/2$. Observational data are characterized by the setting where unconfoundedness

is an assumption rather than a design choice and $e(X)$ is unknown. One reduction for using any of the approaches we discuss on observational data is to assume unconfoundedness holds and estimate $e(X)$ from the data and impute its value. There may also be other reductions, for example leveraging orthogonalized (doubly robust) estimation (Chernozhukov et al., 2016), but for the sake of clarity we focus on known $e(X)$ and Q . Note also that Assumption 1 is necessary for the identifiability of ρ (Imbens & Angrist, 1994; Manski, 1997; Tian & Pearl, 2000).

Given the above data, we are interested in learning a classifier $f : \mathbb{R}^p \rightarrow \{-1, +1\}$ to predict R from X . To assess the quality of classifiers, we focus on the (weighted) misclassification loss. For $\theta \in [0, 1]$, define

$$\begin{aligned} L_\theta(f) &= \theta \mathbb{P}(\text{false positive}) + (1 - \theta) \mathbb{P}(\text{false negative}) \\ &= \theta \mathbb{P}(f(X) = +1, R = -1) \\ &\quad + (1 - \theta) \mathbb{P}(f(X) = -1, R = +1). \end{aligned}$$

We will usually focus on the case $\theta = 1/2$, for which $L_{1/2}(f) = (1 - \text{Accuracy}(f))/2$. Note that, given the true conditional probability $\rho(X)$, the minimizer of $L_\theta(f)$ over all functions f , also known as the Bayes-optimal classifier, is $f_\theta^*(X) = \text{sign}(\rho(X) - \theta)$. This gives the standard reduction of the classification problem to estimating and thresholding conditional probabilities of labels. However, estimating these conditional probabilities may not be necessary for successful classification and may not be the best approach.

1.3. Relationship to CATE

The conditional average treatment effect (CATE) is the conditional expectation of the causal effect given features:

$$\tau(X) = \mathbb{E}[Y(+1) - Y(-1) \mid X] = \mathbb{E}[YT/Q \mid X], \quad (1)$$

where the latter equality arises immediately from unconfoundedness (Athey & Imbens, 2016). As a conditional expectation, CATE can be understood as the *best* predictor of the causal effect in terms of squared error over all functions of X . CATE can of course be defined as in Eq. (1) even if outcomes are not binary. As reviewed in the next section, learning CATE from observations of X, T, Y has been the subject of much recent research.

When outcomes are binary and effects monotonic, we have the following relationship:

Lemma 1. *Under Assumption 1, $\rho(X) = \tau(X)/2$.*

Proof. Since the causal effect is 2 for responders and 0 for non-responders, we have $Y(+1) - Y(-1) = 2\mathbb{I}[R = +1]$ and hence $2\rho(X) = \mathbb{E}[2\mathbb{I}[R = +1] \mid X] = \tau(X)$. \square

This allows for a naïve reduction from learning $f_\theta^*(X)$ to learning $\tau(X)$ using a plug-in approach: given an estimate

$\hat{\tau}$ of τ , return $\hat{f}(X) = \text{sign}(\hat{\tau}(X) - 2\theta)$. This, however, does not directly optimize the classification loss and may fail in producing asymptotically optimal classifiers if $\hat{\tau}$ is not consistent for τ . In our empirical results (Section 4), we will use this reduction to benchmark our algorithms against a variety of existing CATE-learning algorithms.

1.4. Related literature

Many recent advances have been made for the important problem of estimating CATE from X, T, Y data. One basic approach to estimating CATE is to estimate $\mathbb{E}[Y \mid X, T = +1]$ using some regression method on the treated data, similarly estimate $\mathbb{E}[Y \mid X, T = -1]$ on the untreated data, and return the difference, which is sometimes known as ‘‘T-learner’’ (Künzel et al., 2017). More sophisticated methods attempt to learn the difference directly. Athey & Imbens (2016); Wager & Athey (2017) study adapting tree- and forest-based methods to this problem. Johansson et al. (2018); Shalit et al. (2017) develop a neural network architecture for learning CATE with a shared representation as well as generalization bounds that motivate new regularizers. Künzel et al. (2017); Nie & Wager (2017) develop meta-learners that combine base learners for the outcome regressions and treatment model to learn CATE. Assumption 1 also implies the shape constraint $\tau(X) \geq 0$, which can be used as a constraint to improve CATE estimation (Aronow & Carnegie, 2013; Huang et al., 2012). All of the above methods estimate τ . These estimates can then be used to classify responders based on the above plug-in approach. However, this does not directly address the misclassification loss.

Another strand of literature has focused on the problem of policy evaluation and learning from X, T, Y data (Bottou et al., 2013; Dudík et al., 2011; Kallus, 2018; Kallus & Zhou, 2018a;b; Strehl et al., 2010; Swaminathan & Joachims, 2015). In policy evaluation the target is to estimate the average outcome that would be induced in the population if a certain policy were implemented, that is a mapping from covariates to treatment assignment. In policy learning the target is to find a policy with large average outcome. In Section 2.3, we explain that our discriminative classifiers essentially arise from formulating the classification problem as a policy learning problem.

Monotonicity is also a common assumption that arises in instrumental variable (IV) analysis with binary instruments and treatments (Angrist & Pischke, 2008). In such models, we assume that there is an instrument (*e.g.*, encouragement to enroll in a program) that only affects the outcome (*e.g.*, some measurement after program) via its effect on treatment (*e.g.*, enrollment). The instrument is often assumed to be monotonic in its effect on treatment take-up and instrument responders are known as compliers (here we instead use

“responder” because we focus on effect on outcomes). If the instrument is valid and its effect monotonic, then the local average treatment effect (LATE) of the treatment on compliers is identifiable and can be estimated using the Wald estimator: the ratio of the the instrument’s effect on the outcome and on the treatment (Angrist et al., 1996). Monotonicity has been shown to be critical to identifiability in the IV setting in the presence of heterogeneous effects (Imbens & Angrist, 1994). We could conceivably do the same after conditioning everywhere on covariates X to obtain a conditional LATE (CLATE) (see e.g. Aronow & Carnegie, 2013; Athey et al., 2019). The ratio is then between the CATE of the instrument on the outcome and the CATE of the instrument on on the treatment, and the latter (but not the former) indeed assumes monotonic effect. But for use in this conditional Wald estimator, we would actually be interested in estimating CATE itself rather than learning a classifier. Kennedy et al. (2018), however, use a classifier given by thresholding such a CATE estimate in order to focus on subgroups where compliance is high for better interpretability. The discriminative classifiers developed herein can instead be used in their method.

2. A Discriminative Approach using Surrogate Losses

We next proceed to develop a new discriminative approach to classifying treatment responders from X, T, Y data. The approach is based on leveraging monotonicity to reformulate the weighted misclassification loss $L_\theta(f)$ in terms of an expectation over observable quantities, interpreting this expectation as an average loss, and minimizing an empirical average of surrogate losses.

We begin by reformulating the weighted misclassification loss under causal effect monotonicity.

Lemma 2. *Under Assumption 1,*

$$L_\theta(f) = \frac{1}{4} \underbrace{\mathbb{E}[f(X)(2\theta - YT/Q)]}_{L'_\theta(f)} + \frac{1}{4} \mathbb{E}[2\theta + (1 - 2\theta)YT/Q].$$

Proof. We have that

$$\begin{aligned} & \mathbb{P}(\text{false negative}) \\ &= \mathbb{E}[\mathbb{I}[f(X) = -1] \mathbb{I}[R = +1]] \\ &= \frac{1}{2} \mathbb{E}[(1 - f(X)) \rho(X)] \quad (\text{iter. expectations}) \\ &= \frac{1}{4} \mathbb{E}[(1 - f(X)) \mathbb{E}[YT/Q \mid X]] \quad (\text{Lemma 1}) \\ &= \frac{1}{4} \mathbb{E}[(1 - f(X)) (YT/Q)] \quad (\text{iter. expectations}). \end{aligned}$$

A symmetric argument similarly shows $\mathbb{P}(\text{false positive}) = \frac{1}{4} \mathbb{E}[(1 + f(X)) (2 - YT/Q)]$. Combining the two, respectively weighted by $1 - \theta$ and θ , and collecting terms yields the above result. \square

Lemma 2 decomposes the misclassification loss $L_\theta(f)$ into two parts: a part that depends on f ($L'_\theta(f)$) and a part that is independent of f . It therefore shows that optimizing $L_\theta(f)$ is the same as optimizing $L'_\theta(f)$.

Notice, moreover, that we can rewrite

$$L'_\theta(f) = \mathbb{E}[W\ell(f(X), Z)], \quad (2)$$

where $Z = \text{sign}(YT/Q - 2\theta)$, $W = |YT/Q - 2\theta|$, and

$$\ell(\hat{z}, z) = \begin{cases} +1 & \hat{z} \neq z \\ -1 & \hat{z} = z \end{cases}.$$

This shows that $L'_\theta(f)$ has the form of a weighted misclassification loss for the problem of trying to predict Z from X , where each instance is weighted by W .

2.1. A corrupted label perspective

We now give an interpretation of this reformulation from the perspective of a classification problem with corrupted label data. For the sake of exposition, suppose $Q = 1/2$, that is, the data came from a Bernoulli trial with equal treatment probabilities. Then we have that

$$Z = YT, \quad W = 2(1 - Z\theta).$$

That is, Z is a binary label indicating whether $Y = T$ or $Y \neq T$, and examples where $Y \neq T$ get $\frac{1+\theta}{1-\theta} \in [1, \infty]$ times the weight that examples where $Y = T$ get. For example, if $\theta = 1/2$, then this weight ratio is 3 to 1.

To understand this disparity, we will consider Z as a surrogate label for the responder status R . To see that Z can serve as a surrogate label for R note that by Lemma 2, f_θ^* minimizes $L'_\theta(f)$ and hence can also be understood as a classifier for Z . Next, note that an example with responder status $R = +1$ will by definition have $Y = Y(T) = T$ and therefore $Z = +1$. On the other hand, an example with responder status $R = -1$ will either have $Z = -1$ if by chance the T coin flip (recall $Q = 1/2$) ends up opposite to the unit’s non-responder type and otherwise $Z = +1$, so Z will be ± 1 equiprobably. Therefore, Z can be seen as a corrupted form of R , which aligns with R whenever R is positive and gets scrambled whenever R is negative. As such, negative examples with $Z = -1$ are seen as more definitive and therefore carry more weight.

In Section 3, we also use this corrupted label perspective to develop a generative approach and a new cross-entropy loss.

2.2. Weighted surrogate loss minimization

We now present our first proposal for a responder-classification algorithm. The reformulation in Eq. (2) suggests the following discriminative responder-classification

algorithm based on re-weighting surrogate-loss-based classification algorithms. Let $\mathcal{H} \subseteq [\mathbb{R}^p \rightarrow \mathbb{R}]$ be a function class representing score functions, let ℓ' be a surrogate classification loss, and let $\Omega : \mathcal{H} \rightarrow \mathbb{R}_+$ be a potential regularizer. Then return the classifier

$$\hat{f}(x) = \text{sign}(\hat{h}(x)) \quad (3)$$

where \hat{h} is the solution

$$\hat{h} \in \underset{h \in \mathcal{H}}{\text{argmin}} \frac{1}{n} \sum_{i=1}^n W_i \ell'(h(X_i), Z_i) + \Omega(h). \quad (4)$$

For example, if $\ell'(\hat{r}, z) = \max(0, 1 - z\hat{r})$, \mathcal{H} is a reproducing kernel Hilbert space, and $\Omega(h) = \lambda \|h\|^2$ is the squared norm in that space, then we get a sample-weighted support vector machine (Scholkopf & Smola, 2001). We call the corresponding responder-classification algorithm **RespSVM**. As another example, if $\ell'(\hat{r}, z) = \log(1 + e^{\hat{r}}) - (1 + z)\hat{r}/2$ and \mathcal{H} is all neural networks of a given architecture then we get a sample-weighted neural network ($\Omega(h)$ may be nothing or it may be the sum of squared weights for weight decay). We call the corresponding responder-classification algorithm **RespNet-disc**, or **RespLR-disc** in case of a linear architecture with no hidden layers.

2.3. What happens if monotonicity fails? A policy learning perspective

While it is self-evident when one is in a setting where outcome data are binary, monotonicity is always an assumption. Moreover since we do not see counterfactual outcomes, it may not have observable ramifications and may not be testable. This raises an important question: what happens if the monotonicity assumption fails? Can we still meaningfully interpret the classifier \hat{f} that we learn in Eq. (3)?

The next result shows how we can give an interpretation based on policy learning.

Lemma 3. *Minimizing $L'_\theta(f)$ is the same as maximizing $U_\theta(f) = \mathbb{E}[Y(f(X))] - 2\theta\mathbb{P}(f(X) = 1)$.*

Proof. This follows from the facts that $\mathbb{E}[Y(f(X))] = \mathbb{E}[\frac{1+f(X)}{2}Y(+1) - \frac{1-f(X)}{2}Y(-1)] = \frac{1}{2}\mathbb{E}[Y(+1) + Y(-1)] + \mathbb{E}[f(X)YT/Q]$ (invoking Eq. (1)) and $\mathbb{P}(f(X) = 1) = \mathbb{E}[\frac{1+f(X)}{2}] = \frac{1}{2} + \frac{1}{2}\mathbb{E}[f(X)]$. \square

We can interpret Lemma 3 as follows. Suppose outcomes are rewards, where positive outcomes are preferred to negative outcomes. Suppose the function f is a policy mapping features to a decision to apply treatment (+1) or not (-1). And, suppose the cost of applying treatment is 2θ . Then $U_\theta(f)$ is the total average rewards minus costs incurred by following the policy f . Then, regardless of monotonicity

holding or not, by minimizing $L'_\theta(f)$ (or an empirical surrogate version thereof) we are seeking a policy that achieves a good rewards-costs trade off.

The policy learning perspective provides a useful frame even when monotonicity holds. Notice that if monotonicity were true, then $\tau(X) \geq 0$ and, in this reward interpretation of outcomes, every unit can only benefit from treatment. Correspondingly, if treatment had no cost ($\theta = 0$) and monotonicity held, we would always set $f(X) = +1$. Indeed, this would minimize false negatives. However, if we were also concerned with false positives, we would not always predict positive. Indeed, even if everyone could only stand to benefit from treatment, if there was a cost to treatment and there was an uncertainty as to whether treatment would actually make a difference in a certain context, then perhaps the treatment should not be applied.

This perspective is closely related to the reduction by Beygelzimer & Langford (2009); Zhao et al. (2012) of maximizing $\mathbb{E}[Y(f(X))]$ to cost-sensitive classification, which reduces to weighted misclassification in the binary treatment case. However, since effects are monotonic, it is clear that $f(x) = +1$ constant maximizes the above. To introduce the cost to treatment in this framework one would shift all treated outcomes by θ . Doing this, however, produces a *different* set of weights that depend not just on the label value of Z but rather depend simultaneously on Y and T .

3. A Generative Approach

We next present a new generative approach to classifying treatment responders from X, T, Y data. In this approach, we will actually estimate the conditional probability ρ directly using maximum likelihood. The approach is closely related to the corrupted label perspective we presented in Section 2.1. Without loss of generality, we can consider the data as being generated by first drawing (X, R) and then corrupting the R label to produce (X, Z) . We can then use maximum likelihood to fit a generative model to the probability of observing the label $Z = YT$, which we can phrase in terms of ρ .

For this section, we assume that treatment assignment is equiprobable so that $Q = 1/2$. Alternatively, this can be achieved by weighting each unit by $1/Q$ to create a pseudo-population where this is the case. Under this assumption we can relate

Lemma 4. *Suppose $Q = 1/2$ and Assumption 1 holds. Then*

$$\mathbb{P}(Z = z | X) = \frac{1+z\rho(X)}{2}. \quad (5)$$

Proof. Notice that since $R = 1 \implies Z = 1$, we have $\mathbb{P}(Z = z | X, R = 1) = \frac{1+z}{2}$. Let $A = \frac{Y(+1)+Y(-1)}{2}$, which is -1 for type-1 non-responders, 0 for responders,

and +1 type-2 non-responders. By unconfoundedness and $Q = 1/2$, we have $\mathbb{P}(T = t | X, A) = 1/2$. Therefore,

$$\begin{aligned} & \mathbb{P}(Z = z | X, R = -1) \\ &= \mathbb{P}(A = -1 | X, A \neq 0) \mathbb{P}(T = -z | X, A = -1) \\ & \quad + \mathbb{P}(A = +1 | X, A \neq 0) \mathbb{P}(T = z | X, A = +1) = \frac{1}{2}. \end{aligned}$$

Hence, marginalizing over R ,

$$\mathbb{P}(Z = z | X) = \rho(X) \frac{(1+z)}{2} + (1 - \rho(X)) \frac{1}{2},$$

which gives the result. \square

Lemma 4 shows how the corrupted label Z is related to R in their conditional probabilities. Note that, if Q were not equal to $\frac{1}{2}$ then $\mathbb{P}(A = +1 | X, A \neq 0)$ would not cancel out in the above and would remain as a nuisance parameter in the below estimation approach, that is, we would not be able to avoid also estimating the probabilities of being each type of non-responder. Having $Q = \frac{1}{2}$ (e.g., by creating an appropriate weighted pseudo-population if it is not already the case) enables us to ignore this nuisance and focus just on $\rho(X)$.

3.1. MLE of ρ under monotonicity

Based on Lemma 4, we can formulate the negative log likelihood of observing the labels Z_1, \dots, Z_n given the covariate design X_1, \dots, X_n as a function of ρ as a parameter:

$$\text{NLL}_n(\rho) = - \sum_{i=1}^n \log \left(\frac{1+Z_i \rho(X_i)}{2} \right). \quad (6)$$

Then, given a class of probability models $\mathcal{R} \subseteq [\mathbb{R}^p \rightarrow [0, 1]]$ and potentially a regularizer $\Omega : \mathcal{R} \rightarrow \mathbb{R}_+$, we can estimate ρ by optimizing a regularized maximum likelihood

$$\hat{\rho} \in \operatorname{argmin}_{\rho \in \mathcal{R}} \text{NLL}_n(\rho) + \Omega(\rho). \quad (7)$$

Specifically, we focus on using this for neural networks, where \mathcal{R} is neural networks of a given architecture (with a sigmoid activation at its output). We call the corresponding responder-classification algorithm **RespNet-gen**, or **RespLR-gen** in case of no hidden layers

3.2. Comparison to weighted cross-entropy loss

In Section 2.2, one proposal was to reweight and minimize the cross entropy loss. While the cross entropy loss serves both as a surrogate for misclassification *and* as the negative likelihood objective for classification (when probabilities are set to the expit of the score), it is not the case for our problem even after reweighing. If $Q = \frac{1}{2}$, the weighted cross entropy (WCE) loss applied to (the logit of) ρ at a particular observation (X, Z) as proposed in Section 2.2 would be (after scaling by $1/3$)

$$-1 \times \frac{1-Z}{2} \log(1 - \rho(X)) - \frac{1}{3} \times \frac{1+Z}{2} \log(\rho(X)). \quad (8)$$

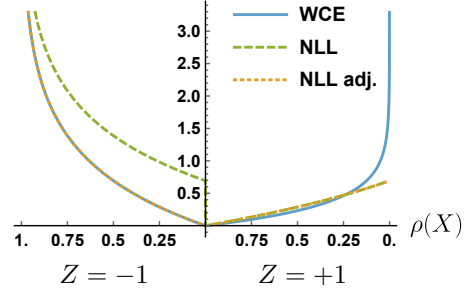


Figure 1. Comparison of the weighted cross-entropy loss (WCE; Eq. (8)), the negative log likelihood (NLL; negative log of Eq. (5)), and the latter after subtracting $\log(2)$ from the $Z = -1$ branch. Note the tick marks on the horizontal axis.

This is distinct from the negative log of the likelihood, Eq. (5). The difference between these is shown in Fig. 1, where we have stitched together the two cases $Z = \pm 1$. The most noticeable feature is that the WCE touches 0 in both cases whereas NLL does not reach 0 when $Z = -1$. Indeed, even if the label R is perfectly predictable from X , the label Z is not. And, when $Z = -1$ we know that $R = -1$ necessarily, in which case observing $Z = \pm 1$ was actually equiprobable and therefore the probability of $Z = -1$ can never be 1 (in fact, it is bounded by $1/2$) and hence NLL, its negative log, does not touch 0. But, fixing the label data, this amounts to constant shift in the loss function relative to ρ . Once we remove this shift (i.e., subtract $\sum_{i=1}^n \frac{1-Z_i}{2} \log 2$ from $\text{NLL}_n(\rho)$ in Eq. (6)), we obtain the curve denoted by NLL adj. in Fig. 1. This matches WCE exactly in the $Z = -1$ case but is much flatter in the $Z = +1$ case and does not approach infinity as $\rho(X) \rightarrow 0$. This permits the misclassification of $Z = +1$ labels, which indeed could have arisen from either $R = \pm 1$ and therefore should not necessarily rule out $\rho(X) = 0$.

Note, nonetheless, that taking the conditional expectation of Eq. (8) given X , differentiating by $\rho(X)$, and solving, gives back Eq. (5) again. Same holds for the NLL loss. This shows that both approaches would be Fisher consistent. In practice, as explored in Section 4, we find that the generative approach (RespNet-Gen) and its bounded loss in the $Z = +1$ case outperforms WCE (RespNet-Disc), although other surrogate losses such as hinge perform well.

4. Empirical Studies

4.1. Synthetic datasets

We first explore responder-classification on two synthetic datasets where we can more clearly illustrate and explain the behavior of different algorithms. We consider two scenarios for various covariate dimensions d . In *both* scenarios we let $X \sim \mathcal{N}(0, I_d)$ be drawn a standard d -dimensional normal and $T = \pm 1$ be drawn by an even coin flip. For each scenario we define $\rho(X)$ and $\alpha(X)$ below. To generate a

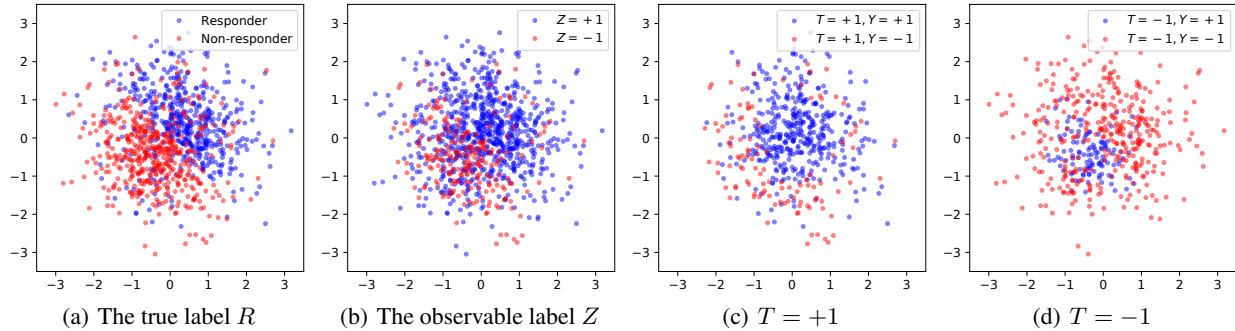


Figure 2. Linear scenario

data point given X and T , we draw $R = \pm 1$ as Bernoulli per $\rho(X)$ and $A = \pm 1$ as Bernoulli per $\alpha(X)$. We then let $Y = \mathbb{I}[R = +1]T + \mathbb{I}[R = -1]A$.

We describe the two scenarios below:

- Linear scenario:** $\rho(X) = 0.15 + 0.7\mathbb{I}[X_1 > 0]$, where X_1 is the first coordinate of X ; and $\alpha(X) = 1 - F_{\text{Beta}(4,4)}(F_{\chi_d^2}(\|X\|_2^2))$, where $F_{\text{Beta}(4,4)}$, $F_{\chi_d^2}$ are the CDFs of a Beta and Chi-squared random variables, respectively. The parameters are chosen so that there are equal numbers of responder and non-responders and of type-1 and type-2 non-responders.
- Spherical scenario:** $\rho(X) = F_{\text{Beta}(4,4)}(F_{\chi_d^2}(\|X\|_2^2))$; and $\alpha(X) = 0.15 + 0.7\mathbb{I}[\bigotimes_{j=1}^{d/2}(X_{2j-1} + X_{2j}) > 0]$, where \bigotimes denotes the exclusive or (XOR) operation.

An example draw with $d = 2$ and $n = 1000$ for the linear scenario is plotted in Fig. 2. Panel (a) shows the true responder label, to which we do not have access at training. Panel (b) shows the $Z = YT$ label, which we can observe. These figures illustrate why in order to solve the responder classification problem in panels (a) we should up-weight the $Z = -1$ examples. Panels (c) and (d) show the data in the treated and untreated groups, respectively. These show why it can often be harder to fit $\mathbb{P}(Y | X, T = \pm 1)$ separately and take difference. In particular, this approach requires we actually estimate this probability rather than just lead some classifier in each of the treated and untreated sample.

We next consider the performance of various approaches on these datasets. We focus on accuracy at predicting responder status, and hence set $\theta = 1/2$. As benchmarks, we consider classifying responders by thresholding an estimate of CATE, as described in Section 1.3. We consider three prominent approaches to estimating CATE: using the differences of random forest regressions (RF; sklearn defaults as planned for v0.22, which increases default number of trees), using the causal forest method (CF; Wager & Athey, 2017, using R package `grf` and defaults), and using a TARNet (Shalit et al., 2017) with one shared hidden layer with $2d$ neurons and a hidden layer of d neurons for each potential outcome

with ELU activations in the interior and sigmoid activations at the outputs. We compare these to the following variants of our methods: RespSVM with linear kernel and 5-fold cross validation (CV) to choose regularization (with L'_θ for scoring); RespSVM with RBF kernel and 5-fold CV to choose regularization and length-scale; and unregularized generative and discriminative RespNets with either no hidden layers (RespLR) or two hidden layers with $2d$ and d neurons each and ELU interior activations (RespNet). RespNets and TARNets are implemented using Keras and TensorFlow and trained with Adam for 100 epochs.

In Figs. 3 and 4, we plot average, 10th, and 90th percentile accuracies of these methods in predicting $R^* = f_{\frac{1}{2}}^*(X)$ as we vary n , d , and the scenario, each over 100 replications. In each scenario, we see a growing divergence between methods as we increase the dimension. In the linear case, the best methods overall are RespSVM (either kernel) and RespLR with the biggest improvements seen when d is large and/or n is small. In the spherical scenario, all linear models naturally fail and the best method overall is RespSVM with RBF kernel followed by RespNet-gen. Overall, the performance of RespNet-disc is similar to TARNet, which is improved upon by using the generative loss instead.

The results showcase that directly targeting the responder classification problem can be beneficial when it is of interest. Note that the results do not imply that these existing algorithms are not good CATE learners; only that they can be improved upon in the special though common setting where outcomes are binary and effects are monotonic.

4.2. Predicting response in decision to have third kid

We next study the application of our methods to the data derived from 1980 census. Following Angrist & Evans (1996), we construct a dataset of married couples with at least two children. We consider the treatment variable to be whether the biological sex of the two children at birth is the same and the outcome variable to be whether the couple has a third child or not. Thus, we are concerned with

Classifying Treatment Responders Under Causal Effect Monotonicity

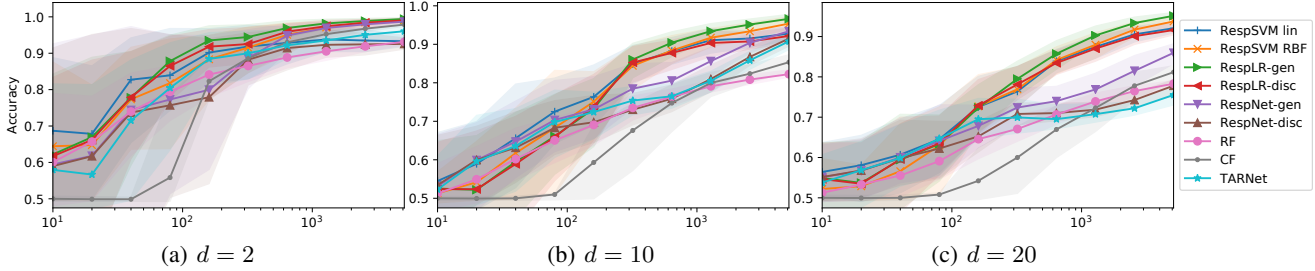


Figure 3. Accuracy results in the linear scenario as n varies




Figure 4. Accuracy results in the spherical scenario as n varies

Table 1. Results for child rearing example

Method	L_θ (in 0.01)	% 1st	% 2nd	% 3rd
RespSVM lin	49 ± 2.7	100%		
RespLR-gen	57 ± 2.4		100%	
RespLR-disc	58 ± 2.3			2%
LR	58 ± 2.3			92%
RF	58 ± 2.3			6%

predicting whether the couple will respond or not to this treatment, and treatment is assigned at random equiprobably. (Angrist & Evans, 1996 were originally interested in the effect of childbearing on women’s participation in the labor force using the above as an instrument. Here we are only concerned in the choice of having a third kid as an outcome.) As features we consider the ethnicity of the mother and of the father, their income and employment status, their ages at marriage, their ages at census, their ages at having their first kid and at having their second, their year of marriage, and the education level of the mother.

To compare the different methods, we repeatedly draw two sets of 100,000 units (observations of X, T, Y). On one, we train each of our methods (after normalizing each column) and on the other we evaluate the false positives, false negatives, and weighted loss, using the result of Lemma 2 to estimate these. We set $\theta = 0.0612$ to be equal to $\mathbb{E}[Z]$ so that always predicting positive or negative has the same weighted misclassification loss (0.0574). This allows us to focus on non-trivial improvements in balanced classification performance. We focus on the linear and forest-based methods as in the last section and add the difference of logistic

regressions (LR, which is equivalent to TARNet with no hidden layers). (We were not able to run CF on this size dataset, but this was likely due to limitations of the `ipy2` package.) In Table 1 we tabulate the average and standard deviation of L_θ over 50 replications and how often each of the methods produce the best, second best, or third best result. We find, as before, that the best performing methods are RespSVM and RespNet-gen (here with no hidden layers).

Finally, as an example of inference using these approaches, we consider the distribution of coefficients in the ResLR-gen model and construct 95% Studentized bootstrap confidence intervals (Efron & Tibshirani, 1994). We find that the only variables without a statistically significant influence on response at 0.05 significance are: father being white vs other, mother being black vs other, the age of the father at marriage, and the education of the mother being strictly more than high school vs no high school.

5. Conclusions

Predicting individual-level causal effects is an important problem. In this paper we specifically studied the arguably common setting where outcomes are binary and effect is monotonic, in which case this problem reduces to determining whether someone will respond to treatment. We formulated this as a classification problem, rather than a CATE estimation problem, and used this, together with monotonicity, to develop new methods for predicting individual-level causal effects. In their common but specialized setting they outperformed standard benchmarks.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 1846210.

References

- Angrist, J. D. and Evans, W. N. Children and their parents' labor supply: Evidence from exogenous variation in family size. Technical report, National bureau of economic research, 1996.
- Angrist, J. D. and Pischke, J.-S. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press, 2008.
- Angrist, J. D., Imbens, G. W., and Rubin, D. B. Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434): 444–455, 1996.
- Aronow, P. M. and Carnegie, A. Beyond late: Estimation of the average treatment effect with an instrumental variable. *Political Analysis*, 21(4):492–506, 2013.
- Athey, S. and Imbens, G. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016.
- Athey, S., Tibshirani, J., Wager, S., et al. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.
- Beygelzimer, A. and Langford, J. The offset tree for learning with partial labels. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 129–138. ACM, 2009.
- Bottou, L., Peters, J., Candela, J. Q., Charles, D. X., Chikering, M., Portugaly, E., Ray, D., Simard, P. Y., and Snelson, E. Counterfactual reasoning and learning systems: the example of computational advertising. *Journal of Machine Learning Research*, 14(1):3207–3260, 2013.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., et al. Double machine learning for treatment and causal parameters. *arXiv preprint arXiv:1608.00060*, 2016.
- Dudík, M., Langford, J., and Li, L. Doubly robust policy evaluation and learning. *arXiv preprint arXiv:1103.4601*, 2011.
- Efron, B. and Tibshirani, R. J. *An introduction to the bootstrap*. CRC press, 1994.
- Huang, Y., Gilbert, P. B., and Janes, H. Assessing treatment-selection markers using a potential outcomes framework. *Biometrics*, 68(3):687–696, 2012.
- Imbens, G. W. and Angrist, J. D. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475, 1994.
- Johansson, F. D., Kallus, N., Shalit, U., and Sontag, D. Learning weighted representations for generalization across designs. *arXiv preprint arXiv:1802.08598*, 2018.
- Kallus, N. Balanced policy evaluation and learning. In *Advances in Neural Information Processing Systems*, pp. 8909–8920, 2018.
- Kallus, N. and Zhou, A. Confounding-robust policy improvement. In *Advances in Neural Information Processing Systems*, pp. 9269–9279, 2018a.
- Kallus, N. and Zhou, A. Policy evaluation and optimization with continuous treatments. *arXiv preprint arXiv:1802.06037*, 2018b.
- Kennedy, E. H., Balakrishnan, S., and G'Sell, M. Sharp instruments for classifying compliers and generalizing causal effects. *arXiv preprint arXiv:1801.03635*, 2018.
- Künzel, S., Sekhon, J., Bickel, P., and Yu, B. Meta-learners for estimating heterogeneous treatment effects using machine learning. *arXiv preprint arXiv:1706.03461*, 2017.
- Manski, C. F. Monotone treatment response. *Econometrica: Journal of the Econometric Society*, pp. 1311–1334, 1997.
- Nie, X. and Wager, S. Learning objectives for treatment effect estimation. *arXiv preprint arXiv:1712.04912*, 2017.
- Scholkopf, B. and Smola, A. J. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- Shalit, U., Johansson, F. D., and Sontag, D. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pp. 3076–3085, 2017.
- Strehl, A., Langford, J., Li, L., and Kakade, S. M. Learning from logged implicit exploration data. In *Advances in Neural Information Processing Systems*, pp. 2217–2225, 2010.
- Swaminathan, A. and Joachims, T. Counterfactual risk minimization: Learning from logged bandit feedback. In *ICML*, pp. 814–823, 2015.
- Tian, J. and Pearl, J. Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1-4):287–313, 2000.
- Wager, S. and Athey, S. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 2017.

Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.