Ensemble Method for Estimating Individualized Treatment Effects *

Kevin Wu Han
Stanford University

Han Wu
Stanford University

Draft version March 2022

Abstract

In many medical and business applications, researchers are interested in estimating individualized treatment effects using data from a randomized experiment. For example in medical applications, doctors learn the treatment effects from clinical trials and in technology companies, researchers learn them from A/B testing experiments. Although dozens of machine learning models have been proposed for this task, it is challenging to determine which model will be best for the problem at hand because ground-truth treatment effects are unobservable. In contrast to several recent papers proposing methods to *select* one of these competing models, we propose an algorithm for *aggregating* the estimates from a diverse library of models. We compare ensembling to model selection on 43 benchmark datasets, and find that ensembling wins almost every time. Theoretically, we prove that our ensemble model is (asymptotically) at least as accurate as the best model under consideration, even if the number of candidate models is allowed to grow with the sample size.

1   Introduction

The conditional average treatment effect (CATE) function $\tau(\cdot)$ is an object of central importance in precision medicine and targeted marketing [Gutierrez and Gérardy, 2017b, Kent et al., 2018]. For any feature profile x, $\tau(x)$ gives the average effect of a treatment (relative to control) among individuals with features x. Estimates of $\tau(\cdot)$ can be used to target a drug or advertisement to the subpopulation where it will have the largest impact.

In recent years, causal inference researchers have proposed dozens of methods that adapt machine learning algorithms so that they can be used to estimate CATE functions from experimental data. For example, random forests, gradient boosting, and deep neural networks have all been repurposed for this task [Wager and Athey, 2018, Powers et al., 2017, Johansson et al., 2016, Shalit et al., 2017]. Going further, several papers have proposed general-purpose "meta-learning frameworks" that allow analysts to adapt *any* supervised learning algorithm to estimate CATE functions [Künzel et al., 2019, Nie and Wager, 2017]. Today, fitting half a dozen different CATE models is only a few hours of work for a skilled data analyst.

That being said, deciding *which* of these competing models to use on the problem at hand remains challenging. Evidence from large-scale simulation studies and causal inference

---

competitions show that no single method performs the best in every problem [Künzel et al., 2019, Schuler et al., 2018, Dorie et al., 2019]. Unfortunately, standard model selection techniques from supervised learning cannot be applied to solve this problem, since individualized treatment effects are unobservable.

To address this problem, several authors have proposed "causal" variants of cross-validation which can be used for model selection [Alaa and Van Der Schaar, 2019, Schuler et al., 2017]. In contrast, we propose a method that bypasses model selection entirely and instead computes an ensemble of candidate models. For bounded outcomes, we show that no validation-set-based model selection rule can (asymptotically) outperform our ensemble model. This result holds even if the number of candidate models grows faster than the sample size, and we do not make any assumptions about the quality or correlation structure of the candidate models. We assess the empirical performance of ensembling by applying our method to 43 benchmark semi-synthetic datasets, and find that ensembling almost always outperforms model selection. Our practical recommendation for data analysts estimating CATE models is to fit a wide variety of diverse models, and then use ensembling rather than model selection to obtain a CATE model for deployment.

2 Formal Problem Statement

2.1 The potential outcomes framework

We adopt potential outcomes model of causality introduced by Neyman and Rubin [Rubin, 2005]. In this model, we posit that, in addition to an individual's observed features $X_i \in \mathcal{X}$, he or she is endowed with two unobserved *potential outcomes* $Y_{1i}, Y_{0i} \in \mathbb{R}$. Here, Y_{1i} is the outcome we would observe if individual i were assigned to the "treatment" condition, and Y_{0i} is the outcome we would observe if individual i were assigned to the "control" condition. The *treatment effect* for individual i is defined as the contrast between these two potential outcomes, $\tau_i = Y_{1i} - Y_{0i}$. Since only one of the two potential outcomes can be observed for any individual, τ_i is unobservable.

2.2 Assumptions on the randomized experiment

In this paper, we focus on randomized experiments. We do so because of two reasons. First, randomized experiments are ubiquitous in applications. Thousands of A/B tests are being performed in technology companies every day and randomized controlled trials are the gold standard in medical applications. Learning a better model for treatment effects leads to better personalization. Second, estimation from observational data typically relies on much stronger, unverifiable assumptions. It is our goal in this paper to minimize assumptions and keep things simple.

We assume that the analyst has access to n samples from a randomized experiment that was performed in the following fashion:

1. Participants in the experiment were sampled i.i.d. from the population of interest, i.e. $(X_i, Y_{1i}, Y_{0i}) \sim \mathbb{P}$ independently for each $i \in [n]$.

2. Independently of $\{(X_i, Y_{1i}, Y_{0i})\}_{i=1}^n$, treatment assignments $(Z_1, \cdots, Z_n) \in \{0,1\}^n$ were sampled from either a Bernoulli design or a Completely Randomized design. In a Bernoulli design, $\mathbb{P}(Z_i = 1) = p$ independently across $i \in [n]$. In a completely randomized design, the vector $(Z_1, \cdots, Z_n)$ is sampled uniformly from the set of binary

vectors with sum pn. In our notation, $Z_i = 1$ indicates that participant i was assigned to the "treatment" condition, and $Z_i = 0$ indicates that participant i was assigned to the "control" condition.

3. The observed outcome is $Y_i = Z_i Y_{1i} + (1 - Z_i) Y_{0i}$. The final dataset available to the analyst is $\{(X_i, Y_i, Z_i)\}_{i=1}^n$.

2.3 The estimation problem

We study the problem of estimating the conditional average treatment effect function $\tau(\cdot)$, which can be defined as:

$$\tau(x) := \mathbb{E}[Y_{1i} - Y_{0i} | X_i = x] = \mathbb{E}[\tau_i | X_i = x] \tag{1}$$

Under the experimental assumptions stated above, the function $\tau(\cdot)$ is identified (up to $\mathbb{P}$-almost sure equivalence).

We assume that the analyst has a library of K candidate algorithms for estimating $\tau(\cdot)$, where each algorithm $\mathcal{A}_k$ maps a dataset $\{(X_i, Y_i, Z_i)\}$ (of any size, not necessarily n) to an estimated CATE function $\hat{\tau}_k : \mathcal{X} \to \mathbb{R}$. The estimation problem is to use these algorithms $\{\mathcal{A}_k\}_{k=1}^K$ to construct a single function $\hat{\tau}_s(\cdot)$ that is close to the true CATE function $\tau(\cdot)$. We will measure the accuracy of an estimator $\hat{\tau}(\cdot)$ using its $\mathcal{L}^2(\mathbb{P}_X)$ distance to $\tau(\cdot)$, $||\hat{\tau} - \tau||_2 = \sqrt{\mathbb{E}[(\hat{\tau}(X_{\text{new}}) - \tau(X_{\text{new}}))^2]}$. This is the most commonly used accuracy measure in the literature on heterogeneous treatment effect estimation, and is called the *precision in estimating heterogeneous effects* by Hill [Hill, 2011].

3 Our Proposal

3.1 The causal stacking algorithm

Our proposal, which we call *causal stacking*, can be simply stated. First, split the data into a training and an averaging set. Then, use the CATE algorithms $\mathcal{A}_k$ to learn estimated CATE models $\hat{\tau}_k(\cdot)$ on the training set. Finally, use the averaging set to find a weight vector $\hat{w}$ such that $\tau_s = \hat{w}^\top \hat{\tau}_{1:K}$ minimizes an estimate of $||\hat{\tau}_s - \tau||_2^2$. A more complete description of the method is given in Algorithm 1[1].

3.2 Practical recommendations

We have a few practical recommendations regarding the use of Algorithm 1. We will further elaborate in section 6.

- **Number of Candidate Models**. Causal stacking gives more gains when we have larger number of candidate models. For best performance, we recommend trying a wide variety of machine-learning algorithms and meta-learning frameworks.

- **Prediction models**. Although the theoretical guarantees on the causal stacking algorithm do not rely on the correctness or consistency of the models $\hat{\mu}_1$ and $\hat{\mu}_0$, we have found that Algorithm 1 performs better when these models have low out-of-sample MSE. There are two standard approaches to fitting regression models in causal

---

[1]The non-negativity constraint and ℓ_1 constraint are for interpretability and help in theoretical analysis. We consider extensions in empirical study section

---

Algorithm 1 Causal stacking

1: **Input:** Data $\{(X_i, Y_i, Z_i)\}_{i=1}^n$, CATE algorithms $\{\mathcal{A}_k\}_{k=1}^K$.

2: Partition the data into a training set $\mathcal{S}_{\text{train}}$ containing $100(1 - \alpha)\%$ of the data and an averaging set $\mathcal{S}_{\text{avg}}$ containing $100\alpha\%$ of the data. If the experimental design was Bernoulli, the partition can be random. If the experimental design was Completely Randomized, the fraction of treated units in the averaging set should be p.

3: **for** $t \in \{0, 1\}$ **do**

4:     Using the data in $\mathcal{S}_{\text{train}}$, fit a regression model $\hat{\mu}_t$ that predicts Y_{ti} using X_i, i.e. $\hat{\mu}_t$ is an estimate of $\mathbb{E}[Y_i | X_i, W_i = t]$

5: **end for**

6: **for all** $k \in [K]$ **do**

7:     Set $\hat{\tau}_k \leftarrow \mathcal{A}_k(\mathcal{S}_{\text{train}})$.

8: **end for**

9: Solve the optimization problem

$$\hat{w} = \underset{w \succeq 0, ||w||_1 = 1}{\operatorname{argmin}} \frac{1}{|\mathcal{S}_{\text{avg}}|} \sum_{i \in \mathcal{S}_{\text{avg}}} (\hat{\tau}_i - w^\top \hat{\tau}_{1:K}(X_i))^2 \tag{2}$$

where $\hat{\tau}_i$ is defined by:

$$\begin{aligned} \hat{\tau}_i = {}& [\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)] \\ & + \frac{[Y_i - \hat{\mu}_1(X_i)]Z_i}{p} - \frac{[Y_i - \hat{\mu}_0(X_i)](1 - Z_i)}{1 - p} \end{aligned} \tag{3}$$

10: **Output:** $\hat{\tau}_s(\cdot) = \hat{w}^\top \hat{\tau}_{1:K}(\cdot)$.

---

    inference: (i) use only data from treated units to fit $\hat{\mu}_1$ and only data from control units to fit $\hat{\mu}_0$; (ii) fit a single model $\hat{\mu}(Z_i, X_i)$ using all the data in $\mathcal{S}_{\text{train}}$, then set $\hat{\mu}_1(X_i) = \hat{\mu}(1, X_i)$ and $\hat{\mu}_0(X_i) = \hat{\mu}(0, X_i)$. We have found that these approaches work well on different datasets, so it is worthwhile to try both and choose the approach with smaller error in cross-validation.

- **Refitting on the entire dataset.** After computing the weight vector $\hat{w}$, one might be tempted to refit the algorithms $\{\mathcal{A}_k\}$ on the *entire* dataset of n samples to construct estimated CATE models $\{\hat{\tau}_k^*\}$ which are then averaged using the first-stage weights. The resulting CATE model $\hat{\tau}_s^*(\cdot) = \hat{w}^\top \hat{\tau}_{1:K}^*$ would then be deployed. Although this approach sounds like it is more sample-efficient, $\hat{\tau}_s^*$ is not guaranteed to perform better than $\hat{\tau}_s$. CATE models based on trees or neural networks can be quite sensitive to the input dataset and the parameter initialization, so $\mathcal{A}_k(\mathcal{S}_{\text{train}})$ may bear little resemblance to $\mathcal{A}_k(\mathcal{S}_{\text{train}} \cup \mathcal{S}_{\text{avg}})$.

4 Related Work

In supervised learning problems, it is well-known that model averaging could outperform model selection [Wolpert, 1992]. For example, all of the top submissions in the "Netflix Prize" competition used extensive ensembling [Bell and Koren, 2007], and model averaging

is now considered essential to achieve competitive performance in Kaggle tournaments.

Recently, model averaging methods have also made their way into causal inference algorithms, although usually only to optimize a regression model, a propensity score model, or a forecast. For example, in the panel data context, Athey et al. [2019] propose using model ensembling to improve predictions of potential outcomes. In the causal inference competition of Dorie et al. [2019], several top-performing submissions used some form of ensembling to optimize potential outcome models, notably the Super Learner + tMLE proposal based on van der Laan et al. [2007] and van der Laan [2010].

The only prior work which we are aware of that proposes stacking for CATE function estimation is Nie & Wager [Nie and Wager, 2017], where one of the state of art CATE models R-learner is introduced. Their proposal for stacking is based on the decomposition which inspires the R-learner and is quite similar to ours[2], but they use the loss function (4) in place of the loss function (2) and optimize over all nonnegative weight vectors, not just those constrained to sum to one. We call this method R-Stacking.

$$R(w) = \sum_{i \in \mathcal{S}_{\mathrm{avg}}} [Y_i - \hat{\mu}(X_i) - (Z_i - p)(w^\top \hat{\tau}_{1:K}(X_i))]^2 \tag{4}$$

In (4), $\hat{\mu}$ is any regression function that predicts Y_i using X_i. In balanced experiments ($p = 0.5$), noting that $(Z_i - p)^2$ is indeed a constant, it is straightforward to show that $R(\cdot)$ is equivalent (up to a constant scaling) to our stacking loss function (2) with $\hat{\mu}_1 \equiv \hat{\mu}_0 \equiv \hat{\mu}$. We give a thorough comparison of causal stacking and R-Stacking in Section 6.

The stacking loss function (2) in Algorithm 1 is simply a feasible approximation of the "idealized" validation-set estimate of $||\hat{\tau}_s - \tau||_2^2$, which would be

$$\frac{1}{|\mathcal{S}_{\mathrm{avg}}|} \sum_{i \in \mathcal{S}_{\mathrm{avg}}} (\tau_i - \hat{\tau}_s(X_i))^2 \tag{5}$$

Since τ_i is unobservable, we plug in an estimate $\hat{\tau}_i$. The estimate $\hat{\tau}_i$ that we use was proposed in Aronow and Middleton [2013] and shown to be unbiased for $\tau(X_i)$ under any choice of the regression functions $\hat{\mu}_0$ and $\hat{\mu}_1$. This estimator is closely related to the AIPW/doubly-robust estimator from the literature on observational studies. The idea of using a plug-in estimator to approximate the idealized validation set procedure has appeared before in the literature on CATE model selection. For example, Saito and Yasui [2019] and Gutierrez and Gérardy [2017a] explore this approach with specific choices of $\hat{\mu}_0$ and $\hat{\mu}_1$. Rolling and Yang [2014] study (5) using a matching-based estimator of τ_i. Alaa and Van Der Schaar [2019] propose a perturbation of the plug-in estimate based on influence functions from semiparametric efficiency theory. On the theoretical side, the analysis in van der Laan and Dudoit [2003] shows that – under certain technical conditions – minimizing the plug-in validation estimate is a consistent model selection rule.

5 Theoretical Analysis

In this section, we will state a result saying that – for bounded outcomes – replacing the averaging step of Algorithm 1 by *any* model selection rule cannot (asymptotically) lead to an improvement, even if the number of candidate models K grows sub-exponentially with the

---

[2]We are describing a slightly simplified version of the proposal of Nie & Wager which uses only a single split of the data.

sample size n. The proof is simple, and is based on classical results on the assumption-free risk-consistency of ℓ_1-constrained functional aggregation [Juditsky and Nemirovski, 2000, Greenshtein and Ritov, 2004].

Proposition 1. *(Finite-sample bound)*
Assume that $Y_{0i}, Y_{1i}, \hat{\tau}_k, \hat{\mu}_1$, and $\hat{\mu}_0$ are all uniformly bounded by a constant L. Then, with probability at least $1 - \delta$, the output $\hat{\tau}_s$ of the causal stacking algorithm satisfies:

$$||\hat{\tau}_s - \tau||_2^2 \leq \min_{k \in [K]} ||\hat{\tau}_k - \tau||_2^2 + 12L^2 \sqrt{\frac{\log([K+1]^2/\delta)}{\alpha n}}$$

Corollary 2. *(Asymptotic implication)*
Assume that $Y_{0i}, Y_{1i}, \hat{\tau}_k, \hat{\mu}_1$ and $\hat{\mu}_0$ are uniformly bounded. Assume that the number of candidate models K_n grows sub-exponentially with n. Then if $\hat{\tau}_{k^*} \in \{\hat{\tau}_1, \cdots, \hat{\tau}_K\}$ is the model selected by any model selection rule, we have the asymptotic result:

$$\max\{0, ||\hat{\tau}_s - \tau||_2^2 - ||\hat{\tau}_{k^*} - \tau||_2^2\} \xrightarrow{p} 0 \tag{6}$$

In other words, no model selection algorithm based on a validation set can asymptotically outperform ensembling.

The proofs may be found in the supplementary material. We will make a few remarks about both the assumptions and the conclusions of Proposition 1 and Corollary 2.

- The condition that the regression models $\hat{\mu}_1, \hat{\mu}_0$ and candidate CATE models $\hat{\tau}_k$ are bounded imposes no restriction beyond boundedness of the potential outcomes, since one can always truncate these models to produce a new regression/CATE models which satisfy the boundedness assumption and are at least as accurate as the originals.

- The result (6) is most interesting when $\min_k ||\hat{\tau}_k - \tau||_2^2$ does *not* tend to zero, i.e. the true CATE function is so complex that no candidate model is consistent. We doubt that any CATE algorithm can achieve vanishingly small error in any real problem with more than a handful of covariates. Standard results from statistical minimax theory show that no CATE algorithm $\mathcal{A}_k$ can achieve (worst-case) error better than $\mathcal{O}(n^{-1/2})$ without extremely strong continuity assumptions on $\tau(\cdot)$ unless $\dim(\mathcal{X})$ is very small[3]. This means that, in high-dimensional problems with only moderate smoothness, using $\hat{\tau}_s$ instead of the best candidate model does not degrade the worst-case rate of convergence.

- Ideally, we would also like to claim that causal stacking is competitive with the best algorithm fit on the *entire* training dataset, whereas Proposition 1 only compares it to the best model fit on a $\mathcal{S}_{\text{train}}$. However, without imposing further assumptions on the algorithms $\mathcal{A}_1, \cdots, \mathcal{A}_K$, it is not possible to prove any such result. The output of $\mathcal{A}_k$ on a dataset of n samples need not be related in any way to the output of $\mathcal{A}_k$ on a dataset with αn samples.

---

[3]Estimating a CATE function is at least as hard as nonparametric regression, since a double (X_i, Y_i) can always be augmented into a triple $(X_i, 0, Y_{1i})$. The problem of recovering the regression function from the original data is the same as the problem of recovering the CATE function from the augmented data. Standard results on minimax estimation over nonparametric function classes can then be applied to derive the stated lower bound, c.f. Tsybakov [2008].

6 Empirical Study

We performed an extensive simulation study to assess the performance of causal stacking relative to two other approaches:

1. *Oracle model selection.* The oracle model selection procedure selects among candidate models by choosing the model that minimizes the *idealized* validation-set risk estimate (7).

$$\frac{1}{|\mathcal{S}_{\mathrm{avg}}|} \sum_{i \in \mathcal{S}_{\mathrm{avg}}} (\tau(X_i) - \hat{\tau}_k(X_i))^2 \tag{7}$$

   In other words, this procedure has oracle knowledge of the true CATE values on the validation set in the model selection step.

2. *R-Stacking.* The R-Stacking procedure proposed by Nie & Wager [Nie and Wager, 2017] optimizes the loss (4) over the class of nonnegative weight vectors. We use XGBoost model to fit the regression function $\hat{\mu}$ in (4).

To facilitate a fair comparison with R-Stacking, we also used XGBoost models to fit the regressions $\hat{\mu}_1$ and $\hat{\mu}_0$ required to implement Algorithm 1.

6.1 Datasets

We applied our method to all 43 "high-heterogeneity" datasets from the 2016 Atlantic Causal Inference Competition [Dorie, 2016]. These semi-synthetic datasets use covariates drawn from a subset of the Collaborative Perinatal Project, but potential outcomes Y_{1i}, Y_{0i} and CATE functions $\tau(\cdot)$ were generated synthetically. More detail on the construction of these datasets can be found in Dorie et al. [2019]. We filtered out a small number of unmanageable categorical features, leaving 52 covariates remaining.

For each dataset, we used 2000 observations for training CATE models and 1000 observations for averaging. The remaining 1802 observations were reserved as a test set for evaluating the averaged models. In the training and validation data, the assignments Z_i were generated i.i.d. from a Bernoulli(p) distribution, and we considered $p \in \{0.1, 0.3, 0.5\}$. Fifty replications were performed per dataset, with both the train/averaging/split and treatment assignments being rerandomized across replications. Our final quality measure is the average test MSE across the 50 replications.

6.2 Candidate models

In our simulations, we considered nine candidate CATE algorithms. These algorithms span a variety of meta-learning "frameworks," covering T-, S-, X-, and R-learning [Künzel et al., 2019, Nie and Wager, 2017]. They also employ a variety of machine learning algorithms.

1. *SVM S-learner.* This method uses the training data to fit a single support vector regression $\hat{\mu}$ making predictions $\hat{Y}_i = \hat{\mu}(Z_i, X_i)$. We used a Gaussian kernel in the support vector regression. The fitted CATE function is defined by $\hat{\tau}(x) = \hat{\mu}(1, x) - \hat{\mu}(0, x)$.

2. *XGBoost T-learner.* This method fits two gradient boosting models $\hat{\mu}_1$ and $\hat{\mu}_0$ using data from treated and control units respectively. The fitted CATE function is $\hat{\tau}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x)$. We used 100 boosting iterations, and all other parameters were set at their default values.

3. *Random Forest T-learner.* This is the same as the XGBoost T-learner, but the regression functions are fit using random forests. We used the default parameter settings in the `ranger` implementation of random forests.

4. *Regression tree S-learner.* This is the same as the SVM S-learner, but $\hat{\mu}$ is fit using CART.

5. *Lasso T-learner.* This is the same as the XGBoost T-learner, but the regression functions are fit using Lasso. The ℓ_1-penalty parameters were chosen using 5-fold cross-validation.

6. *Random Forest X-learner.* This method fits two regression models $\hat{\mu}_1$ and $\hat{\mu}_0$ in the same way as the Random Forest T-learner. After fitting these models, individualized treatment effects $\hat{\tau}_i$ are estimated for each unit by using the regression models to impute the unobserved potential outcomes. Two new random forest models $\hat{\gamma}_1$ and $\hat{\gamma}_0$ are fit to predict $\hat{\tau}_i$ using X_i based on data from treated and control units, respectively. The final CATE model is $\hat{\tau}(x) = (1 - p)\hat{\gamma}_1(x) + p\hat{\gamma}_0(x)$.

7. *XGBoost R-learner.* In this method, we first fit an XGBoost model $\hat{\mu}$ predicting Y_i using X_i. Then, we fit a separate XGBoost model $\hat{\tau}$ to minimize the weighted least-squares objective (8).

$$\sum_{i \in \mathcal{S}_{\text{train}}} (Z_i - p)^2 \left(\frac{Y_i - \hat{\mu}(X_i)}{Z_i - p} - \hat{\tau}(X_i) \right)^2 \tag{8}$$

We used 100 boosting iterations for each model, and all other parameters were set at their default values.

8. *Lasso R-learner.* This is the same as the XGBoost R-learner, except the models $\hat{\mu}$ and $\hat{\tau}$ were fit using Lasso with 5-fold cross-validation to select the tuning parameter.

9. *Constant.* In this model, we set $\hat{\tau}$ to be the constant function that always predicts the difference in mean outcomes between treated and control units in the training dataset.

6.3 Results

6.3.1 Causal stacking vs. oracle model selection

Our simulations suggested three main takeaways in the comparison between causal stacking and oracle model selection:

1. Causal stacking outperforms overall.

In every setting of the treatment assignment probability ($p = 0.1, 0.3, 0.5$), causal stacking outperformed oracle model selection on the vast majority of datasets in terms of average mean-squared error on the test set. Table 1 summarizes these results.
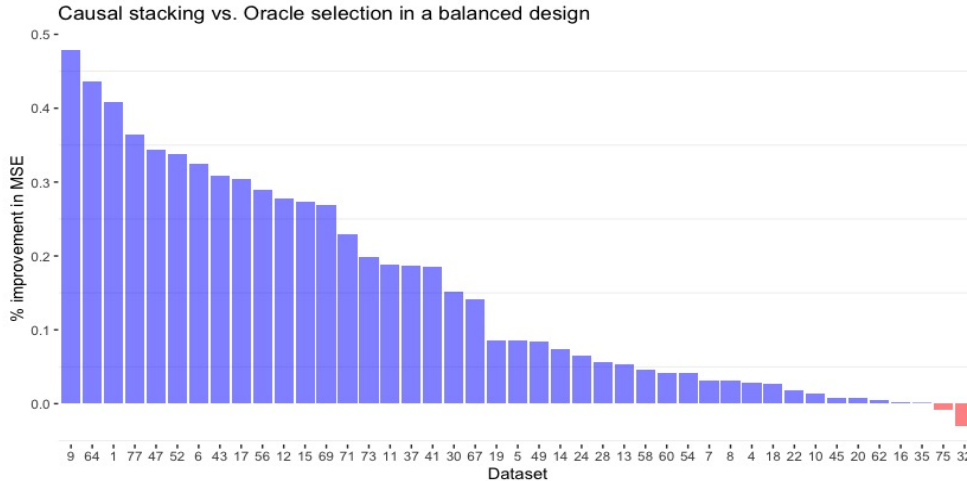
Figure 1: Causal stacking vs. oracle model selection in the balanced design ($p = 0.5$). The % improvement in test-set MSE is defined as $1 - \mathsf{MSE}_{\text{stacking}}/\mathsf{MSE}_{\text{oracle}}$. In datasets 32 and 75, the decline in MSE is less than 3%.

Table 1: Causal stacking vs. oracle: % of datasets won

$\mathbb{P}(Z_i = 1)$	CAUSAL STACKING	ORACLE
0.1	81%	19%
0.3	91%	9%
0.5	95%	5%

 2. Causal stacking shines in balanced experiments.
Figure 1 shows the percent improvement of causal stacking relative to oracle model selection on all 43 datasets under the balanced design ($p = 0.5$). Causal stacking improves the MSE by as much as 48%, and never degrades MSE by more than 3%.

 3. No harm from superstars.
In our experiments, we found that a single CATE algorithm (Algorithm 2: XGBoost T-learner) far outperformed its competition and was chosen by the oracle more often than all other algorithms combined. For example, the model selection oracle picked the XGBoost T-learner on over 95% of replications on dataset 5. Although this situation is often considered to be unfavorable for ensemble methods, we did not find this to be the case. Figure 2 plots the average weights chosen by causal stacking on dataset 5 for each of the nine candidate CATE algorithms. Although the dominant model gets the bulk of the weight, contributions from algorithms $3 - 6$ helped causal stacking outperform the model selection oracle on this dataset. We also performed all of our simulations without the dominant XGBoost T-learner model, and found results essentially identical to those from Table 1.

6.3.2 Causal stacking vs. R-Stacking

Whether or not causal stacking performs better or R-stacking performs better is more context dependent. Causal stacking outperformed R-stacking in many simulations, and R-stacking outperformed causal stacking in many others. We have a few general pieces of
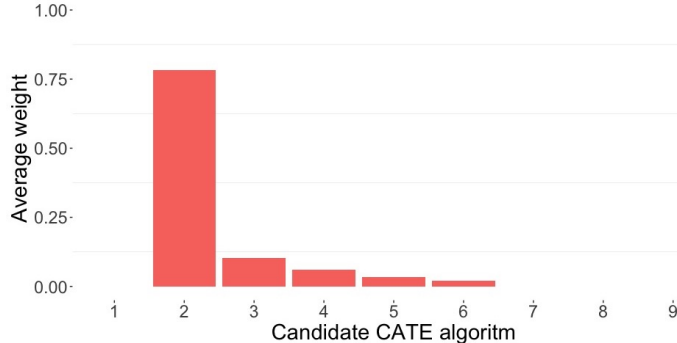
Figure 2: The average causal stacking weights on dataset 5. Models 1, 7, 8, and 9 received essentially zero weight, due to the sparsity properties of ℓ_1-constraints.

advice on how to choose between these two methods:

Table 2: Causal stacking vs. R-Stacking: % of datasets won

$\mathbb{P}(Z_i = 1)$	CAUSAL STACKING	R-STACKING
0.10	35%	65%
0.30	60%	40%
0.50	60%	40%

1. Balanced designs favor causal stacking.
Table 2 shows the relative win rates of causal stacking vs. R-stacking across the three experimental designs. Causal stacking does better in balanced experiments whereas R-stacking does better in imbalanced experiments. This agrees with our intuition, since in balanced experiments, R-stacking is (essentially) a special case of causal stacking with regression models $\hat{\mu}_1$ and $\hat{\mu}_0$ constrained to be the same[4]. When there are enough samples in both the treatment and control group to estimate different response functions, the added flexibility can reduce bias. In imbalanced experiments, one of the two regression functions is fitted using a relatively small sample, and the added variance may harm performance. In addition, the weights $1/p$ are larger if p is small, so the loss function has higher variance, which degrades the performance of causal stacking.

2. With worse-performing candidate models, R-stacking is preferred.
In our experiments without the dominant XGBoost T-learner model, the R-stacking procedure performed better than causal stacking (in terms of percent of datasets won) even in balanced experiments. We speculate that this may be due to the capacity control induced by the additional constraint in causal stacking that weights sum to one. With a weaker library of models, greater flexibility may be needed to fit the underlying CATE function. Figure 3 shows the average weights of both causal stacking and R-stacking on dataset 5 without Algorithm 2; a noticeable feature is that the R-stacking procedure uses weights with sum far exceeding one. This extra capacity also allows the R-stacking procedure to employ five candidate models, whereas causal stacking only uses three. In fact, as our next set of experiments show, removing the constraint generally improves the performance in

---

[4]The R-stacking procedure uses conic combinations instead of convex combinations.
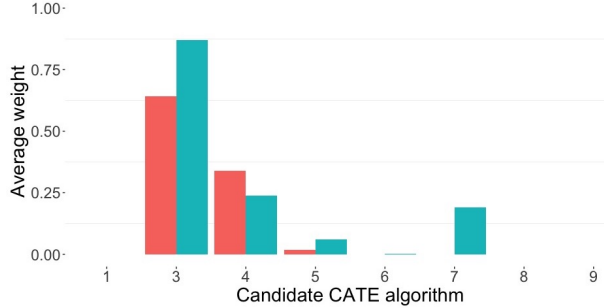
more balanced designs.

Figure 3: Causal stacking vs. R-stacking weights (red) on dataset 5, without algorithm 2. The R-stacking model takes advantage of its extra capacity and chooses weights with total sum substantially exceeding one.

6.3.3   Causal stacking vs. Other variants

One may consider some other variants of Algorithm 1. An immediate idea is to remove either one of the constraints or all constraints in (2). In our experiments, we found that removing the ℓ_1 constraint actually improves the performance in more balanced designs. Table 3 summarizes the results. We see that in more balanced experiments removing the constraint has a significant improvement. Also, removing both the ℓ_1 constraint and the nonnegative constraint is unstable in our experiments due to co-linearity caused by similarities between some of the models. However, we do see improvements in datasets that do not have the numerical issue.

Table 3: Causal stacking vs. Causal Stacking (no ℓ_1 constraint): % of datasets won

$\mathbb{P}(Z_i = 1)$	CAUSAL STACKING	NO ℓ_1 CONSTRAINT
0.10	58%	42%
0.30	37%	63%
0.50	26%	74%

Table 4: Causal Stacking (no ℓ_1 constraint) vs. R-Stacking: % of datasets won

$\mathbb{P}(Z_i = 1)$	NO ℓ_1 CONSTRAINT	R-STACKING
0.10	35%	65%
0.30	77%	23%
0.50	84%	16%

Table 4 shows the comparison between causal stacking without ℓ_1 constraint with R-stacking. We see that removing ℓ_1 constraint makes causal stacking beat R-stacking more often in more balanced designs and still suffer from added variance of fitting two regression functions as explained in the comparison between causal stacking and R-stacking.

6.3.4 Elaboration on recommendations

Finally, we elaborate on recommendations we gave previously. First, we consider dropping those models that are seldom selected in causal stacking, i.e. we only consider models that are selected quite often. In our experiments, XGBoost T-learner, Random Forest T-learner and Regression tree S-learner are three dominant models. We performed causal stacking with only these three models and compared it with the causal stacking with all nine models. Table 5 shows the results. It is clear from the table that though the other six models were selected only a few times, they do play a role and make the performance better in every setting of the treatment assignment probability we consider.

Table 5: Causal Stacking vs. Causal Stacking (only model 2, 3 and 4): % of datasets won

$\mathbb{P}(Z_i = 1)$	ALL MODELS	ONLY MODEL 2, 3, 4
0.10	72%	28%
0.30	65%	35%
0.50	63%	37%

Second, the quality of two regression models used in (3) matters. We compared causal stacking that uses XGBoost to predict $\{\mu_t\}_{t=0,1}$ with causal stacking that uses linear model to predict $\{\mu_t\}_{t=0,1}$. In our experiments, we found that the latter one significantly degrades the performance. Table 6 summarizes the results.

Table 6: Causal Stacking (XGBoost $\hat{\mu}_t$) vs. Causal Stacking (linear $\hat{\mu}_t$): % of datasets won

$\mathbb{P}(Z_i = 1)$	XGBOOST $\hat{\mu}_t$	LINEAR $\hat{\mu}_t$
0.10	81%	19%
0.30	93%	7%
0.50	95%	5%

7 Discussion

In Bell et al. [2007], the authors of the winning solution to the Netflix prize wrote:

> *Predictive accuracy is substantially improved when blending multiple predictors. Our experience is that most efforts should be concentrated in deriving substantially different approaches, rather than refining a single technique.*

In our paper, we argued that the same wisdom can be applied to the problem of estimating CATE functions. Specifically, we have the following practical recommendations for data analysts working on this problem:

- When estimating CATE functions, it is better to fit a diverse array of candidate models and then use model averaging (with either causal stacking or R-stacking) to *combine* these models rather than to use a validation set to *select* a model.

- On the 43 datasets we studied, X-learner and T-learner were top performers. We recommend that data analysts include these algorithms in their list of candidates.

For future work, we would like to better understand which situations lead to causal stacking performing better than R-learner stacking and vice versa. Although we provided some speculation as to why the R-learner stacking procedure outperformed causal stacking when the dominant XGBoost T-learner model was removed from consideration, we do not yet have any rigorous understanding of this phenomenon. Other factors to explore include the role of sample size and the number of candidate models K; when the sample size is small or the library of candidate models K is large, we suspect that the extra capacity control in causal stacking will help to control overfitting the validation set. Another interesting direction for future work is to explore more sophisticated ensembling techniques, including causal variants of boosting and bagging.

References

Ahmed Alaa and Mihaela Van Der Schaar. Validating causal inference models via influence functions. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 191–201, Long Beach, California, USA, 09–15 Jun 2019. PMLR.

Peter M. Aronow and Joel A. Middleton. A Class of Unbiased Estimators of the Average Treatment Effect in Randomized Experiments. *Journal of Causal Inference*, 1(1):135–154, June 2013.

Susan Athey, Mohsen Bayati, Guido Imbens, and Zhaonan Qu. Ensemble methods for causal effects in panel data settings, 2019.

R. M. Bell and Y. Koren. Scalable collaborative filtering with jointly derived neighborhood interpolation weights. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pages 43–52, Oct 2007. doi: 10.1109/ICDM.2007.90.

Robert M. Bell, Yehuda Koren, and Chris Volinsky. The bellkor solution to the netflix prize. 2007.

Vincent Dorie. aciccomp2016. https://github.com/vdorie/aciccomp/tree/master/2016, 2016.

Vincent Dorie, Jennifer Hill, Uri Shalit, Marc Scott, and Dan Cervone. Automated versus do-it-yourself methods for causal inference: Lessons learned from a data analysis competition. *Statist. Sci.*, 34(1):43–68, 02 2019. doi: 10.1214/18-STS667.

Eitan Greenshtein and Ya'Acov Ritov. Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. *Bernoulli*, 10(6):971–988, 12 2004. doi: 10.3150/bj/1106314846.

Pierre Gutierrez and Jean-Yves Gérardy. Causal inference and uplift modelling: A review of the literature. In Claire Hardgrove, Louis Dorard, Keiran Thompson, and Florian Douet-teau, editors, *Proceedings of The 3rd International Conference on Predictive Applications and APIs*, volume 67 of *Proceedings of Machine Learning Research*, pages 1–13, Microsoft NERD, Boston, USA, 11–12 Oct 2017a. PMLR.

Pierre Gutierrez and Jean-Yves Gérardy. Causal inference and uplift modelling: A review of the literature. In Claire Hardgrove, Louis Dorard, Keiran Thompson, and Florian Douetteau, editors, *Proceedings of The 3rd International Conference on Predictive Applications and APIs*, volume 67 of *Proceedings of Machine Learning Research*, pages 1–13, Microsoft NERD, Boston, USA, 11–12 Oct 2017b. PMLR.

Jennifer L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011. doi: 10.1198/jcgs.2010.08162.

Fredrik D. Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pages 3020–3029. JMLR.org, 2016.

Anatoli Juditsky and Arkadii Nemirovski. Functional aggregation for nonparametric regression. *Ann. Statist.*, 28(3):681–712, 05 2000. doi: 10.1214/aos/1015951994.

David M Kent, Ewout Steyerberg, and David van Klaveren. Personalized evidence based medicine: predictive approaches to heterogeneous treatment effects. *BMJ*, 363, 2018. ISSN 0959-8138. doi: 10.1136/bmj.k4245.

Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, Feb 2019. ISSN 1091-6490. doi: 10.1073/pnas.1804597116.

Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects, 2017.

Scott Powers, Junyang Qian, Kenneth Jung, Alejandro Schuler, Nigam Haresh Shah, Trevor J. Hastie, and Robert Tibshirani. Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in medicine*, 37 11:1767–1787, 2017.

Craig A. Rolling and Yuhong Yang. Model selection for estimating treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(4):749–769, 2014.

Donald B. Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005. ISSN 01621459.

Yuta Saito and Shota Yasui. Counterfactual cross-validation: Effective causal model selection from observational data, 2019.

Alejandro Schuler, Ken Jung, Robert Tibshirani, Trevor Hastie, and Nigam Shah. Synthvalidation: Selecting the best causal inference method for a given dataset, 2017.

Alejandro Schuler, Michael Baiocchi, Robert Tibshirani, and Nigam Shah. A comparison of methods for model selection when estimating individual treatment effects, 2018.

Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 3076–3085. JMLR.org, 2017.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008. ISBN 0387790519.

Mark van der Laan. Targeted maximum likelihood based causal inference: Part i. *The International Journal of Biostatistics*, 2010.

Mark van der Laan and Sandrine Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples, 2003.

Mark van der Laan, Eric Polley, and Alan Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 2007.

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, 2018. doi: 10.1080/01621459.2017.1319839.

David H. Wolpert. Stacked generalization. *Neural Networks*, 5(2):241 – 259, 1992. ISSN 0893-6080.

A   Proofs

A.1   Proof of Proposition 1

Proof. We will show the stronger result that the inequality holds even conditional on on the training set $\mathcal{S}_{\text{train}}$. Therefore, we may consider $\hat{\tau}_1, \cdots, \hat{\tau}_K, \hat{\mu}_1$ and $\hat{\mu}_0$ to be fixed functions, and all expectations in this proof will only be taken with respect to the randomness in $\mathcal{S}_{\text{avg}}$. Define the function $\gamma : \mathbb{R}^K \to \mathbb{R}^{K+1}$ by $\gamma(w) = (1, -w)$. The stacking loss function $\hat{\ell}(w)$ can be written as:

$$\hat{\ell}(w) = \frac{1}{|\mathcal{S}_{\text{avg}}|} \sum_{i \in \mathcal{S}_{\text{avg}}} (\hat{\tau}_i - w^\top \hat{\tau}_{1:K}(X_i))^2 = \gamma(w)^\top \underbrace{\left[\frac{1}{|\mathcal{S}_{\text{avg}}|} \sum_{i \in \mathcal{S}_{\text{avg}}} \hat{\tau}_{0:K}(X_i)\hat{\tau}_{0:K}(X_i)^\top \right]}_{:=\hat{\Sigma}} \gamma(w)$$

where $\hat{\tau}_{0:K}(X_i) = (\hat{\tau}_i, \hat{\tau}_1(X_i), \cdots, \hat{\tau}_K(X_i))$. Let $\Sigma := \mathbb{E}[\hat{\Sigma}]$, and notice that the entries of $\hat{\Sigma}$ are all sample averages of terms bounded in absolute value by L^2. Therefore, Hoeffding's inequality and a union bound imply that:

$$\mathbb{P}\left(\max_{1 \leq i,j \leq K+1} |\hat{\Sigma}_{ij} - \Sigma_{ij}| > \sqrt{2}L^2 \sqrt{\frac{\log([K+1]^2/\delta)}{|\mathcal{S}_{\text{avg}}|}} \right) \leq \delta \tag{9}$$

On the complement of the event (9), the stacking loss $\hat{\ell}(w)$ is close to its expectation $\ell(w) := \mathbb{E}[\hat{\ell}(w)]$ simultaneously over the set $\Delta_K := \{w \in \mathbb{R}^K : w \succeq 0, ||w||_1 = 1\}$.

$$
\begin{aligned}
\sup_{w \in \Delta_K} |\hat{\ell}(w) - \ell(w)| &= \sup_{w \in \Delta_K} |\gamma(w)^\top (\hat{\mathbf{\Sigma}} - \mathbf{\Sigma})\gamma(w)| \\
&\leq \left(\sup_{w \in \Delta_K} ||\gamma(w)||_1^2 \right) \left(\max_{1 \leq i,j \leq K+1} |\hat{\mathbf{\Sigma}} - \mathbf{\Sigma}| \right) \\
&\leq 4\sqrt{2}L^2 \sqrt{\frac{\log([K+1]^2/\delta)}{|\mathcal{S}_{\text{avg}}|}}
\end{aligned}
$$

This implies that $\ell(\hat{w})$ is not much larger than $\ell(w_{\text{opt}})$, where $w_{\text{opt}} := \underset{w \in \Delta_K}{\operatorname{argmin}} \, \ell(w)$.

$$
\begin{aligned}
\ell(\hat{w}) &\leq \hat{\ell}(\hat{w}) + |\ell(\hat{w}) - \hat{\ell}(\hat{w})| \\
&\leq \hat{\ell}(w_{\text{opt}}) + |\ell(\hat{w}) - \hat{\ell}(\hat{w})| \\
&\leq \ell(w_{\text{opt}}) + |\ell(w_{\text{opt}}) - \hat{\ell}(w_{\text{opt}})| + |\ell(\hat{w}) - \hat{\ell}(\hat{w})| \\
&\leq \ell(w_{\text{opt}}) + 8\sqrt{2}L^2 \sqrt{\frac{\log([K+1]^2/\delta)}{|\mathcal{S}_{\text{avg}}|}} \tag{10}
\end{aligned}
$$

To finish the proof, we will use the fact that $\ell(w) = ||w^\top \hat{\tau}_{1:K} - \tau||_2^2 + \text{Var}(\hat{\tau}_i)$ for any w. We can see this by first decomposing $\ell(w)$ into three terms:

$$
\begin{aligned}
\ell(w) &= \mathbb{E}\left[\frac{1}{|\mathcal{S}_{\text{avg}}|} \sum_{i \in \mathcal{S}_{\text{avg}}} (\hat{\tau}_i - w^\top \hat{\tau}_{1:K}(X_i))^2 \right] \\
&= \mathbb{E}[\mathbb{E}\left[(\hat{\tau}_i - w^\top \hat{\tau}_{1:K}(X_i))^2 | X_i \right]] \\
&= \mathbb{E}[\mathbb{E}[(\hat{\tau}_i - \tau(X_i))^2 + (\tau(X_i) - w^\top \hat{\tau}_{1:K}(X_i))^2 - 2(\hat{\tau}_i - \tau(X_i))(\tau(X_i) - w^\top \hat{\tau}_{1:K}(X_i))|X_i]] \\
&= \underbrace{\mathbb{E}[\mathbb{E}[(\hat{\tau}_i - \tau(X_i))^2|X_i]]}_{(i)} + \underbrace{||\tau - w^\top \hat{\tau}_{1:K}||_2^2}_{(ii)} - \underbrace{2\mathbb{E}[(\tau(X_i) - w^\top \hat{\tau}_{1:K}(X_i))\mathbb{E}[(\hat{\tau}_i - \tau(X_i))|X_i]]}_{(iii)}
\end{aligned}
$$

Conditional on X_i, $\hat{\tau}_i$ is unbiased for $\tau(X_i)$:

$$
\begin{aligned}
\mathbb{E}[\hat{\tau}_i|X_i] &= [\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)] + \mathbb{E}\left[\frac{(Y_i - \hat{\mu}_1(X_i))Z_i}{p}\bigg|X_i\right] - \mathbb{E}\left[\frac{(Y_i - \hat{\mu}_0)(1 - Z_i)}{1-p}\bigg|X_i\right] \\
&= [\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)] + \mathbb{E}[Y_{1i} - \hat{\mu}_1(X_i)|X_i]\mathbb{E}\left[\frac{Z_i}{p}\right] - \mathbb{E}\mathbb{E}[Y_{0i} - \hat{\mu}_0(X_i)|X_i]\mathbb{E}\left[\frac{1-Z_i}{1-p}\bigg|X_i\right] \\
&= \mathbb{E}[Y_{1i} - Y_{0i}|X_i] \\
&= \tau(X_i)
\end{aligned}
$$

Therefore, the term (iii) vanishes and the term (i) simplifies to $\text{Var}(\hat{\tau}_i)$ by the law of total variance. This establishes the identity $\ell(w) = ||w^\top \hat{\tau}_{1:K} - \tau||_2^2 + \text{Var}(\hat{\tau}_i)$, which can be applied on both sides of (10) to obtain the final result. $\qquad\square$