

# TRANSFER LEARNING FOR ESTIMATING CAUSAL EFFECTS USING NEURAL NETWORKS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

We develop new algorithms for estimating heterogeneous treatment effects, combining recent developments in transfer learning for neural networks with insights from the causal inference literature. By taking advantage of transfer learning, we are able to efficiently use different data sources that are related to the same underlying causal mechanisms. We compare our algorithms with those in the extant literature using extensive simulation studies based on large-scale voter persuasion experiments and the MNIST database. Our methods can perform an order of magnitude better than existing benchmarks while using a fraction of the data.

## 1 INTRODUCTION

The rise of massive datasets that provide fine-grained information about human beings and their behavior provides unprecedented opportunities for evaluating the effectiveness of treatments. Researchers want to exploit these large and heterogeneous datasets, and they often seek to estimate how well a given treatment works for individuals conditioning on their observed covariates. This problem is important in medicine (where it is sometimes called personalized medicine) (Henderson et al., 2016; Powers et al., 2018), digital experiments (Taddy et al., 2016), economics (Athey and Imbens, 2016), political science (Green and Kern, 2012), statistics (Tian et al., 2014), and many other fields. Although a large number of articles are being written on this topic, many outstanding questions remain. We present the first paper that applies transfer learning to this problem.

In the simplest case, treatment effects are estimated by splitting a training set into a treatment and a control group. The treatment group receives the treatment, while the control group does not. The outcomes in those groups are then used to construct an estimator for the Conditional Average Treatment Effect (CATE), which is defined as the expected outcome under treatment minus the expected outcome under control given a particular feature vector (Athey and Imbens, 2015). This is a challenging task because, for every unit, we either observe its outcome under treatment or control, but never both. Assumptions, such as the random assignment of treatment and additional regularity conditions, are needed to make progress. Even with these assumptions, the resulting estimates are often noisy and unstable because the CATE is a vector parameter. Recent research has shown that it is important to use estimators which consider both treatment groups simultaneously (Künzel et al., 2017; Wager and Athey, 2017; Nie and Wager, 2017; Hill, 2011). Unfortunately, these recent advances are often still insufficient to train robust CATE estimators because of the large sample sizes required when the number of covariates is not small.

In this paper, we show how these difficulties in estimating the CATE can sometimes be overcome through the use of transfer learning. In particular, we provide several strategies for utilizing ancillary datasets that are related to the causal mechanism under investigation. Examples of such datasets include observations from: experiments in different locations on different populations, different treatment arms, different outcomes, and non-experimental observational studies. We show that, by transferring information from these ancillary datasets, CATE estimators can converge to better solutions with fewer samples. This is particularly important for CATE estimation, as the cost of collecting additional data is quite high and often requires real-world data collection.

**Our contributions are as follows:**

1. **We introduce the new problem of transfer learning for estimating heterogeneous treatment effects.**

2. **MLRW Transfer for CATE Estimation** adapts the idea of meta-learning regression weights (MLRW) to CATE estimation. By using a learned initialization, regression problems can be optimized much more quickly than with random initializations. Though a variety of MLRW algorithms exist, it is not immediately obvious how one should use these methods for CATE estimation. The principal difficulty is that CATE estimation requires the simultaneous estimation of outcomes under both treatment and control, but we only observe one of the outcomes for any individual unit. Most MLRW transfer methods optimize on a per-task basis to estimate a single quantity. We show that one can overcome this problem with clever use of the Reptile algorithm (Nichol et al., 2018).
3. **We provide several additional methods for transfer learning for CATE estimation:** warm start, frozen-features, multi-head, and joint training.
4. **We apply our methods to difficult data problems and show that they perform better than existing benchmarks.** We reanalyze a set of large field experiments that evaluate the effect of a mailer on voter turnout in the 2014 U.S. midterm elections (Gerber et al., 2017). This includes 17 experiments with 1.96 million individuals in total. We also simulate several randomized controlled trials using image data of handwritten digits found in the MNIST database (LeCun, 1998). We show that our methods, **MLRW** in particular, obtain better than state-of-the-art performance in estimating CATE, and that they require far fewer observations than extant methods.
5. **We provide open source code for our algorithms.**<sup>1</sup>

## 2 CATE ESTIMATION

We begin by formally introducing the CATE estimation problem. Following the potential outcomes framework (Rubin, 1974), assume there exists a single experiment wherein we observe  $N$  i.i.d. distributed units from some super population,  $(Y_i(0), Y_i(1), X_i, W_i) \sim \mathcal{P}$ .  $Y_i(0) \in \mathbb{R}$  denotes the potential outcome of unit  $i$  if it is in the control group,  $Y_i(1) \in \mathbb{R}$  is the potential outcome of  $i$  if it is in the treatment group,  $X_i \in \mathbb{R}^d$  is a  $d$ -dimensional feature vector, and  $W_i \in \{0, 1\}$  is the treatment assignment. For each unit in the treatment group ( $W_i = 1$ ), we only observe the outcome under treatment,  $Y_i(1)$ . For each unit under control ( $W_i = 0$ ), we only observe the outcome under control. Crucially, there cannot exist overlap between the set of units for which  $W_i = 1$  and the set for which  $W_i = 0$ . It is impossible to observe both potential outcomes for any unit. This is commonly referred to as the fundamental problem of causal inference.

However, not all hope is lost. We can still estimate the Conditional Average Treatment Effect (CATE) of the treatment. Let  $x$  be an individual feature vector. Then the CATE of  $x$ , denoted  $\tau(x)$ , is defined by

$$\tau(x) = \mathbb{E}[Y(1) - Y(0)|X = x].$$

Estimating  $\tau$  is impossible without making further assumptions on the distribution of  $(Y_i(0), Y_i(1), X_i, W_i)$ . In particular, we need to place two assumptions on our data.

**Assumption 1 (Strong Ignorability, Rosenbaum and Rubin (1983))**

$$(Y_i(1), Y_i(0)) \perp W|X.$$

**Assumption 2 (Overlap)** Define the propensity score of  $x$  as,

$$e(x) := \mathbb{P}(W = 1|X = x).$$

Then there exists constant  $0 < e_{min}, e_{max} < 1$  such that for all  $x \in \text{Support}(X)$ ,

$$0 < e_{min} < e(x) < e_{max} < 1.$$

In words,  $e(x)$  is bounded away from 0 and 1.

---

<sup>1</sup>The software will be released once anonymity is no longer needed. We can also provide an anonymized copy to reviewers upon request.

Assumption 1 ensures that there is no unobserved confounder, a random variable which influences both the probability of treatment and the potential outcomes, which would make the CATE unidentifiable. The assumption is particularly strong and difficult to check in applications. Meanwhile, Assumption 2 rectifies the situation wherein a certain part of the population is always treated or always in the control group. If, for example, all women were in the control group, one cannot identify the treatment effect for women. Though both assumptions are strong, they are nevertheless satisfied by design in randomized controlled trials. While the estimators we discuss would be sensible in observational studies when the assumptions are satisfied, we warn practitioners to be cautious in such studies, especially when the number of covariates is large (D’Amour et al., 2017).

Given these two assumptions, there exist many valid CATE estimators. The crux of these methods is to estimate two quantities: the control response function,

$$\mu_0(x) = \mathbb{E}[Y(0)|X = x],$$

and the treatment response function,

$$\mu_1(x) = \mathbb{E}[Y(1)|X = x].$$

If we denote our learned estimates as  $\hat{\mu}_0(x)$  and  $\hat{\mu}_1(x)$ , then we can form the CATE estimate as the difference between the two

$$\hat{\tau}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x).$$

The astute reader may be wondering why we don’t simply estimate  $\mu_0$  and  $\mu_1$  with our favorite function approximation algorithm at this point and then all go home. After all, we have access to the ground truths  $\mu_0$  and  $\mu_1$  and the corresponding inputs  $x$ . In fact, it is commonplace to do exactly that. When people directly estimate  $\mu_0$  and  $\mu_1$  with their favorite model, we call the procedure a T-learner (Künzel et al., 2017). Common choices of models include linear models and random forests, though neural networks have recently been considered (Nie and Wager, 2017).

A practitioner of deep learning might find the T-learner quite trivial. After all, it amounts to using neural networks to fit two quantities,  $\mu_0$  and  $\mu_1$ . However, it is important to note that the T-learner is a baseline method. We use it in this paper only to ease exposition, especially as it relates to transfer learning. The T-learner has many drawbacks (Athey and Imbens, 2015). It is almost always an inefficient estimator. For example, it will often perform poorly when one can borrow information across the treatment conditions. For these reasons, more sophisticated learners such as the S, X, T, R, and Y learners are almost always used instead of the T-learner (Hill, 2011; Athey and Imbens, 2016; Nie and Wager, 2017; Künzel et al., 2017; Stadie et al., 2018). Although much of our exposition will focus on transfer learning in the context of the T-learner, in practice we extend the discussed methods to these other more advanced learners, as shown in the Evaluation section. Descriptions of these more advanced estimators are given in the appendix.

### 3 TRANSFER LEARNING

In this section, we consider a scenario wherein one has access to many related causal inference experiments. The goal is to use the results from some old experiments to obtain faster training with less data on other new experiments. Since direct transfer between different populations is wrought with difficulty, we will instead achieve transfer by using previous experiments to help find an initialization for new experiments which leads to faster optimization.<sup>2</sup>

We consider two kinds of algorithms. First, there are transfer algorithms that sit on top of existing CATE estimators. These transfer algorithms take a CATE estimation strategy, such as the S-learner, and provide a recipe for transforming it into a transfer learning CATE estimator. The second class of algorithms does not sit on top of existing CATE estimation strategies. Instead, they are built from the ground up to take advantage of transfer learning. These algorithms are joint training and MLRW.

Across all experiments, the input space  $X$  is the same. Let  $i$  index an experiment. Each experiment has its own distinct outcome when treatment is received,  $\mu_{1,i}(x)$ , and when no treatment is received,  $\mu_{0,i}(x)$ . Together, these quantities define the CATE  $\tau_{\cdot,i}(x) = \mu_{1,i}(x) - \mu_{0,i}(x)$ . In standard CATE estimation, we define a strategy that takes  $x$  as input and outputs predictions  $\hat{\mu}_{0,i}(x)$  and  $\hat{\mu}_{1,i}(x)$ .

<sup>2</sup>By faster optimization, we mean that starting from the learned initialization will allow the problem to be solved in fewer optimization epochs and with less data than starting with random weights.

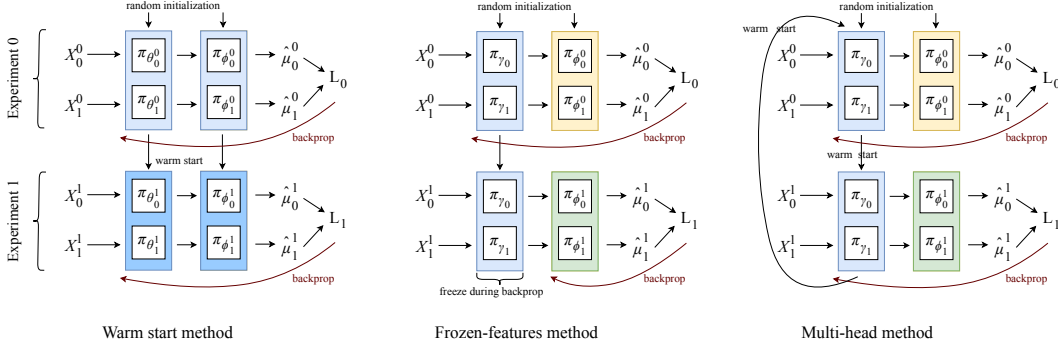


Figure 1: Warm start, frozen-features, and multi-head methods for CATE transfer learning. For these figures, we use the T-learner as the base learner for simplicity. All three methods attempt to reuse neural network features from previous experiments.

In transfer learning, the hope is that we can transfer knowledge between experiments. The model parameters that allowed us predict  $\mu_{0,i}(x)$ ,  $\mu_{1,i}(x)$  and  $\tau_{\cdot,i}(x)$  from experiment  $i$  should help us predict  $\mu_{1,j}(x)$ ,  $\mu_{0,j}(x)$ , and  $\tau_{\cdot,j}(x)$  from experiment  $j$ .

Let  $\pi_\theta$  be a generic expression for a neural network parameterized by  $\theta$ . Parameters will have two subscripts. The index on the left indicates if their neural network predicts treatment or control (0 for control and 1 for treatment). The index on the right is for the experiment. For example,  $\theta_{0,2}$  parametrizes  $\pi_{\theta_{0,2}}(x)$  to predict  $\mu_{0,2}(x)$ , the outcome under control for Experiment 2. All of the transfer algorithms described here are presented in full detail as pseudo-code in the appendix.

### 3.0.1 ALGORITHMS THAT EXTEND EXISTING CATE ESTIMATORS TO TRANSFER LEARNING

These algorithms extend existing CATE estimation techniques to the transfer learning setting. The following exposition is largely motivated by transfer learning with the T-Learner as a base CATE estimator. This is only for ease of exposition. The discussed procedures can extend to other, more complicated, CATE estimators such as the R, X, Y, and S learners.

**Warm start:** Experiment 0 predicts  $\pi_{\theta_{0,0}}(x) = \hat{\mu}_{0,0}(x)$  and  $\pi_{\theta_{1,0}}(x) = \hat{\mu}_{1,0}(x)$  to form the CATE estimator  $\hat{\tau}_{\cdot,0} = \hat{\mu}_{1,0}(x) - \hat{\mu}_{0,0}(x)$ . Suppose  $\theta_{0,0}, \theta_{1,0}$  are fully trained and produce a good CATE estimate. For experiment 1, the input space  $X$  is identical to the input space for experiment 0, but the outcomes  $\mu_{0,1}(x)$  and  $\mu_{1,1}(x)$  are different. However, we suspect the underlying data representations learned by  $\pi_{\theta_{0,0}}$  and  $\pi_{\theta_{1,0}}$  are still useful. Hence, rather than randomly initializing  $\theta_{0,1}$  and  $\theta_{1,1}$  for experiment 1, we set  $\theta_{0,1} = \theta_{0,0}$  and  $\theta_{1,1} = \theta_{1,0}$ . We then train  $\pi_{\theta_{0,1}}(x) = \hat{\mu}_{0,1}(x)$  and  $\pi_{\theta_{1,1}}(x) = \hat{\mu}_{1,1}(x)$ . See Figure 1 and Algorithm 8 in the appendix.

**Frozen-features:** Begin by training  $\pi_{\theta_{0,0}}$  and  $\pi_{\theta_{1,0}}$  to produce good CATE estimates for experiment 0. Assuming  $\theta_{0,0}$  and  $\theta_{1,0}$  have more than  $k$  layers, let  $\gamma_0$  be the parameters corresponding to the first  $k$  layers of  $\theta_{0,0}$ . Define  $\gamma_1$  analogously. Since we think the features encoded by  $\pi_{\gamma_i}(X)$  would make a more informative input than the raw features  $X$ , we want to use those features as a transformed input space for  $\pi_{\theta_{0,1}}$  and  $\pi_{\theta_{1,1}}$ . To wit, set  $z_0 = \pi_{\gamma_0}(x)$  and  $z_1 = \pi_{\gamma_1}(x)$ . Then form the estimates  $\pi_{\theta_{0,1}}(z_0) = \hat{\mu}_{0,1}(x)$  and  $\pi_{\theta_{1,1}}(z_1) = \hat{\mu}_{1,1}(x)$ . During training of experiment 1, we only backpropagate through  $\theta_{0,1}$  and  $\theta_{1,1}$  and not through  $\gamma_0$  and  $\gamma_1$ . See Figure 1 and Algorithm 9 in the appendix.

**Multi-head:** In this setup, all experiments share base layers that are followed by experiment-specific layers. The intuition is that the base layers should learn general features, and the experiment-specific layers should transform those features into estimates of  $\mu_{j,i}(x)$ . More concretely, let  $\gamma_0$  and  $\gamma_1$  be shared base layers for estimating  $\mu_{0,\cdot}(x)$  and  $\mu_{1,\cdot}(x)$  respectively. Set  $z_0 = \pi_{\gamma_0}(x_0)$  and  $z_1 = \pi_{\gamma_1}(x_1)$ . The base layers are followed by experiment-specific layers  $\phi_{0,i}$  and  $\phi_{1,i}$ . Let  $\theta_{j,i} = [\gamma_j, \phi_{j,i}]$ . Then  $\pi_{\theta_{j,i}}(x) = \pi_{\phi_{j,i}}(\pi_{\gamma_j}(x)) = \pi_{\phi_{j,i}}(z_j) = \hat{\mu}_{j,i}(x)$ . Training alternates between experiments: each  $\theta_{0,i}$  and  $\theta_{1,i}$  is trained for some small number of iterations, and then the experiment and head being trained are switched. Every head is usually trained several times. See Figure 1 Algorithm 10 in the appendix.

**SF Reptile transfer for CATE estimators:** Pick your favorite CATE estimator. The goal is to learn an initialization for that CATE estimator’s weights that leads to fast convergence on new experiments. More concretely, starting from good initializers  $\theta_0$  and  $\theta_1$ , one can train neural networks  $\pi_{\theta_0}$  and  $\pi_{\theta_1}$  to estimate  $\mu_{0,i}(x)$  and  $\mu_{1,i}(x)$  much faster and with less data than starting from random initializations. To learn these good initializations, we use a transfer learning technique called Reptile. The idea is to perform experiment-specific inner updates  $U(\theta)$  and then aggregate them into outer updates of the form  $\theta_{\text{new}} = \epsilon \cdot U(\theta) + (1 - \epsilon) \cdot \theta$ . In this paper, we consider a slight variation of Reptile. In standard Reptile,  $\epsilon$  is either a scalar or correlated to per-parameter weights furnished via SGD. For our problem, we would like to encourage our network layers to learn at different rates. The hope is that the lower layers can learn more general, slowly-changing features like in the frozen features method, and the higher layers can learn comparatively faster features that more quickly adapt to new tasks after ingesting the stable lower-level features. To accomplish this, we take the path of least resistance and make  $\epsilon$  a vector which assigns a different learning rate to each neural network layer. Because our intuition involves slow and fast weights, we will refer to this modification in this paper as SF Reptile: Slow Fast Reptile. Though this change is seemingly small, we found it boosted performance on our problems. See Algorithm 11.

### 3.0.2 TRANSFER LEARNING ALGORITHMS THAT DO NOT EXTEND CATE ESTIMATION STRATEGIES

**Joint training:** All predictions share base layers  $\theta$ . From these base layers, there are two heads per experiment  $i$ : one to predict  $\mu_{0,i}(x)$  and one to predict  $\mu_{1,i}(x)$ . Every head and the base features are trained simultaneously by optimizing with respect to the loss function  $\mathcal{L} = \sum_i \|(\hat{\mu}_{0,i}(x) - \mu_{0,i}(x))\| + \|(\hat{\mu}_{1,i}(x) - \mu_{1,i}(x))\|$  and minimizing over all weights. This will encourage the base layers to learn generally applicable features and the heads to learn features specific to predicting a single  $\mu_{j,i}(x)$ . See Algorithm 6.

**MLRW transfer:** In this method, there exists one single set of weights  $\theta$ . There are no experiment-specific weights. Furthermore, we do not use separate networks to estimate  $\mu_0$  and  $\mu_1$ . Instead,  $\pi_\theta$  is trained to estimate one  $\mu_{i,j}(x)$  at a time. We train  $\theta$  with SF Reptile so that in the future  $\pi_\theta$  requires minimal samples to fit  $\mu_{i,j}(x)$  from any experiment. To actually form the CATE estimate, we use a small number of training samples to fit  $\pi_\theta$  to  $\mu_{0,i}(x)$  and then a small number of training samples to fit  $\pi_\theta$  to  $\mu_{1,i}(x)$ . We call  $\theta$  **meta-learned regression weights (MLRW)** because they are meta-learned over many experiments to quickly regress onto any  $\mu_{i,j}(x)$ . The full MLRW algorithm is presented as Algorithm 5.

## 4 EVALUATION

We evaluate our transfer learning estimators on both real and simulated data. In our data example, we consider the important problem of voter encouragement. Analyzing a large data set of 1.96 million potential voters, we show how transfer learning across elections and geographic regions can dramatically improve our CATE estimators. **To the best of our knowledge, this is the first successful demonstration of transfer learning for CATE estimation.** The simulated data has been intentionally chosen to be different in character from our real-world example. In particular, the simulated input space is images and the estimated outcome variable is continuous.

### 4.1 GOTV EXPERIMENT

To evaluate transfer learning for CATE estimation on real data, we reanalyze a set of large field experiments with more than 1.96 million potential voters (Gerber et al., 2017). The authors conducted 17 experiments to evaluate the effect of a mailer on voter turnout in the 2014 U.S. Midterm Elections. The mailer informs the targeted individual whether or not they voted in the past four major elections (2006, 2008, 2010, and 2012), and it compares their voting behavior with that of the people in the same state. The mailer finishes with a reminder that their voting behavior will be monitored. The idea is that social pressure—i.e., the social norm of voting—will encourage people to vote. The likelihood of voting increases by about 2.2% (s.e.=0.001) when given the mailer.

Each of the experiments targets a different state. This results in different populations, different ballots, and different electoral environments. In addition to this, the treatment is slightly different in each

experiment, as the median voting behavior in each state is different. However, there are still many similarities across the experiments, so there should be gains from transferring information.

In this example, the input  $X$  is a voter’s demographic data including age, past voting turnout in 2006, 2008, 2009, 2010, 2011, 2012, and 2013, marital status, race, and gender. The treatment response function  $\hat{\mu}_1(x)$  estimates the voting propensity for a potential voter who receives a mailer encouraging them to vote. The control response function  $\hat{\mu}_0$  estimates the voting propensity if that voter did not receive a mailer. The CATE  $\tau$  is thus the change in the probability of voting when a unit receives a mailer. The complete dataset has this data over 17 different states. Treating each state as a separate experiment, we can perform transfer learning across them.

| $x$             | outcome                        | $\mu_0$  | $\mu_1$  | $\tau$  |
|-----------------|--------------------------------|--|--|---|
| A voter profile | The voter’s propensity to vote | The voter’s propensity to vote when they do not receive a mailer | The voter’s propensity to vote when they do receive a mailer | Change in the voter’s propensity to vote after receiving a mailer |

Being able to estimate the treatment effect of sending a mailer is an important problem in elections. We may wish to only treat people whose likelihood of voting would significantly increase when receiving the mailer, to justify the cost for these mailers. Furthermore, we wish to avoid sending mailers to voters who will respond negatively to them. This negative response has been previously observed and is therefore feasible and a relevant problem—e.g., some recipients call their Secretary of State’s office or local election registrar to complain (Mann, 2010; Michelson, 2016).

Evaluating a CATE estimator on real data is difficult. The primary difficulty is that we do not get to observe the true CATE for any unit, due to the fundamental problem of causal inference. By definition, only one of the two outcomes is observed for any unit. One could use the original features and simulate the outcome features, but this would require us to create a response model. Instead, we estimate the "truth" on the real data using linear models (version 1) or random forests (version 2). We then construct the data based on these estimates. For a detailed description, see Appendix A.2. We then ask the question: How do the various methods perform when they have less data than the entire sample?

## RESULTS

We evaluate all the algorithms discussed in section 3 on the GOTV dataset. For the algorithms in section 3.0.1 that require a base CATE estimator, we use the Y learner because we found it delivered the best performance.<sup>3</sup> For baselines, we compare against the non-transfer Y-learner and the S learner with random forests.<sup>4</sup> In previous work, state of the art results on this problem have been achieved with both non-transfer tree-based estimators such as S-RF (Künzel et al., 2017; Green and Kern, 2012) and neural-network-based learners such as the R and Y-learners (Nie and Wager, 2017; Stadie et al., 2018).

**The best estimator is MLRW.** This algorithm consistently converges to a very good solution with very few observations. Looking at Tables 1, 2, and 3, we observe that MLRW is the best performing transfer learner for GOTV version 1 in 8 out of 17 trials. In GOTV version 2, it is the best in 11 out of 17 trials. In Figure 2, its average performance is dominant over all other algorithms. We hypothesize that this method does best because it does not try to artificially bottleneck the flow of information between outcomes and experiments. MLRW also seems more resilient to data-poisoning when it encounters outlier data, though we did not concretely test against this. We also observe that multi-head, frozen-features, and SF all generally improve upon non-transfer baselines. The faster learning rate of these algorithms indicates that positive transfer between experiments is occurring. Warm start, however, does not work well and often even leads to worse results than the baseline estimators. This is consistent with prior findings on warm start (Finn et al., 2017).

<sup>3</sup>See Tables 1, 2, and 3 and Figures 5, 6, and 7 for full results with the X, Y, S, and T-learners as base CATE estimators.

<sup>4</sup>We use the S learner because the Y learner is not compatible with random forests.




Figure 2: Social Pressure and Voter Turnout, Version 1 (real data as a linear model) and Version 2 (real data as random forest). Our transfer learning algorithms far exceed the previous state of the art non-transfer baselines, which are represented here as S-RF and Baseline (Y-NN).

#### 4.2 MNIST EXAMPLE

In the previous experiment, we observed that the MLRW estimator performed most favorably, and transfer learning significantly improved upon the baseline. To confirm that this conclusion is not specific to voter persuasion studies, we intentionally consider a very different type of data. Recently, (Nie and Wager, 2017) introduced a simulation study wherein MNIST digits are rotated by some number of degrees  $\alpha$ ; with  $\alpha$  furnished via a single data generating process that depends on the value of the depicted digit. They then attempt to do CATE estimation to measure the heterogeneous treatment effect of a digit’s label.

Motivated by this example, we develop a data generating process using MNIST digits wherein transfer learning for CATE estimation is applicable. In our example, the input  $X$  is an MNIST image. We have  $k$  data-generating processes which return different outcomes for each input when given either treatment or control. Thus, under some fixed data-generating process,  $\mu_0$  represents the outcome when the input image  $X$  is given the control,  $\mu_1$  represents the outcome when  $X$  is given the treatment, and  $\tau$  is the difference in outcomes given the placement of  $X$  in the treatment or control group. Each data-generating process has different response functions ( $\mu_0$  and  $\mu_1$ ) and thus different CATEs ( $\tau$ ), but each of these functions only depends on the label presented in the image  $X$ . We thus hope that transfer learning could expedite the process of learning features which are indicative of the label. See Appendix A for full details of the data generation process. In Figure 3, we confirm that a transfer learning strategy outperforms its non-transfer learning counterpart, even on image data. We also see that MLRW performs well, though in this case multi-head is competitive. We also see that several of the transfer methods are worse than non-transfer baselines.




Figure 3: MNIST task. The baseline is the S-learner. All transfer CATE estimators for this task are built on top the S-learner, rather than the Y-learner, because we found it delivered better performance for this problem.

## 5 RELATED WORKS AND DISCUSSION

### 5.0.1 RELATED WORKS

In this paper, we proposed the problem of transfer learning for CATE estimation. One immediate question the reader may be left with is why we chose the transfer learning techniques we did. We only considered two common types of transfer: (1) Basic fine tuning and weights sharing techniques common in the computer vision literature (Welinder et al., 2010; Saenko and Darrell, 2010; Bourdev et al., 2011; Donahue et al., 2014; Koch, 2015), (2) Techniques for learning an initialization that can be quickly optimized (Finn et al., 2017; Ravi and Larochelle, 2017; Nichol et al., 2018). However, many further techniques exist. Yet, transfer learning is an extensively studied and perennial problem (Schmidhuber, 1992; Bengio et al., 1992; Thrun, 1996; Thrun and Pratt, 1998; Taylor and Stone, 2009; Silver et al., 2013). In (Vinyals et al., 2016), the authors attempt to combine feature embeddings that can be utilized with non-parametric methods for transfer. (Snell et al., 2017) is an extension of this work that modifies the procedure for sampling examples from the support set during training. (Andrychowicz et al., 2016) and related techniques try to meta-learn an optimizer that can more quickly solve new tasks. (Rusu et al., 2016) attempts to overcome forgetting during transfer by systematically introducing new network layers with lateral connections to old frozen layers. (Munkhdalai and Yu, 2017) uses networks with memory to adapt to new tasks. We invite the reader to review (Finn et al., 2017) for an excellent overview of the current transfer learning landscape. Though the majority of the discussed techniques could be extended to CATE estimation, our implementations of (Rusu et al., 2016; Andrychowicz et al., 2016) proved difficult to tune and consequently learned very little. Furthermore, we were not able to successfully adapt (Snell et al., 2017) to the problem of regression. We decided to instead focus our attention on algorithms for obtaining good initializations, which were easy to adapt to our problem and quickly delivered good results without extensive tuning.

On the topic of using neural networks to improve causal inference algorithms, a flurry of relevant work exists (Ramachandra, 2018; Magliacane et al., 2017; Johansson et al., 2016; Louizos et al., 2017; Alaa et al., 2017; Shalit et al., 2017; Nie and Wager, 2017). We found that these papers either did not allow us to better estimate the CATE, or else provided worse performance than the baseline methods we did consider in this paper. Extending transfer to other causal inference algorithms is an ongoing and interesting area of research.

### 5.0.2 CLOSING REMARKS

We are left with several open questions. Can transfer learning still be applied to CATE estimation when the experiment input spaces differ? How should one properly deal with missing and incomplete data? Do there exist better methods for interpretability, highlighting which features are most important for transfer and why? Can these techniques be extended to causal models outside of CATE estimation? How can one properly encode causal relationships into a neural network? Answering these questions would have a positive impact on fields such as causal inference, deep learning, and reinforcement learning.



## REFERENCES

- Alaa, A., Weisz, M., and M., V. D. S. (2017). Deep counterfactual networks with propensity-dropout. *1706.05966*.
- Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M. W., Pfau, D., Schaul, T., and de Freitas, N. (2016). Learning to learn by gradient descent by gradient descent. *Neural Information Processing Systems (NIPS)*.
- Athey, S. and Imbens, G. W. (2015). Machine learning methods for estimating heterogeneous causal effects. *stat*, 1050(5).
- Athey, S. and Imbens, G. W. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences of the United States of America*, 113(27):7353–60.
- Bengio, S., Bengio, Y., Cloutier, J., and Gecsei, J. (1992). On the optimization of a synaptic learning rule. *Biological Neural Networks*.
- Bourdev, L., Maji, S., and Malik, J. (2011). Describing people: A poselet-based approach to attribute classification. *ICCV*.
- D’Amour, A., Ding, P., Feller, A., Lei, L., and Sekhon, J. (2017). Overlap in observational studies with high-dimensional covariates. *arXiv preprint arXiv:1711.02582*.
- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., and Darrell, T. (2014). Decaf: A deep convolutional activation feature for generic visual recognition. *International Conference on Machine Learning (ICML)*.
- Finn, C., Abbeel, P., and Levine, S. (2017). Model-agnostic metalearning for fast adaptation of deep networks. *ICML*.
- Gerber, A. S., Huber, G. A., Fang, A. H., and Gooch, A. (2017). The generalizability of social pressure effects on turnout across high-salience electoral contexts: Field experimental evidence from 1.96 million citizens in 17 states. *American Politics Research*, 45(4):533–559.
- Green, D. P. and Kern, H. L. (2012). Modeling heterogeneous treatment effects in survey experiments with bayesian additive regression trees. *Public opinion quarterly*, 76(3):491–511.
- Henderson, N. C., Louis, T. A., Wang, C., and Varadhan, R. (2016). Bayesian analysis of heterogeneous treatment effects for patient-centered outcomes research. *Health Services and Outcomes Research Methodology*, 16(4):213–233.
- Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240.
- Johansson, F. D., Shalit, U., and Sontag, D. (2016). Learning representations for counterfactual inference. *ICML*.
- Koch, G. (2015). Siamese neural networks for one-shot image recognition. *ICML Deep Learning Workshop*.
- Künzel, S., Sekhon, J., Bickel, P., and Yu, B. (2017). Meta-learners for estimating heterogeneous treatment effects using machine learning. *arXiv preprint arXiv:1706.03461*.
- LeCun, Y. (1998). The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>.
- Louizos, C., Shalit, U., Mooij, J., Sontag, D., Zemel, R., and Welling, M. (2017). Causal effect inference with deep latent-variable models. *NIPS*.
- Magliacane, S., Van Ommen, T., Claassen, T., Bongers, S., Versteeg, P., and Mooij, J. (2017). Domain adaptation by using causal inference to predict invariant conditional distributions. *1707.06422*.
- Mann, C. B. (2010). Is there backlash to social pressure? a large-scale field experiment on voter mobilization. *Political Behavior*, 32(3):387–407.

- Michelson, M. R. (2016). The risk of over-reliance on the institutional review board: An approved project is not always an ethical project. *PS: Political Science & Politics*, 49(02):299–303.
- Munkhdalai, T. and Yu, H. (2017). Meta networks. *ICML*.
- Nichol, A., Achiam, J., and Schulman, J. (2018). On first-order meta-learning algorithms. *CoRR*, abs/1803.02999.
- Nie, X. and Wager, S. (2017). Learning objectives for treatment effect estimation. *arXiv preprint arXiv:1712.04912*.
- Powers, S., Qian, J., Jung, K., Schuler, A., Shah, N. H., Hastie, T., and Tibshirani, R. (2018). Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in medicine*.
- Ramachandra, V. (2018). Deep learning for causal inference. *1803.00149*.
- Ravi, S. and Larochelle, H. (2017). Optimization as a model for few-shot learning. *International Conference on Learning Representations (ICLR)*.
- Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., and Hadsell, R. (2016). Progressive neural networks. *CoRR*, vol. abs/1606.04671.
- Saenko, K., K. B. F. M. and Darrell, T. (2010). Adapting visual category models to new domains. *ECCV*.
- Schmidhuber, J. (1992). Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*.
- Shalit, U., Johansson, F., and Sontag, D. (2017). Estimating individual treatment effect: generalization bounds and algorithms. *ICML*.
- Silver, Yand, and Li (2013). Lifelong machine learning systems: Beyond learning algorithms. *DAAAI Spring Symposium-Technical Report, 2013*.
- Snell, J., Swersky, K., and Zemel, R. (2017). Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*.
- Stadie, B. C., Künzel, S. R., Vemuri, N., Ramakrishnan, V., Sekhon, J. S., and Abbeel, P. (2018). Estimating heterogeneous treatment effects with the y-learner. *arXiv*.
- Taddy, M., Gardner, M., Chen, L., and Draper, D. (2016). A nonparametric bayesian analysis of heterogeneous treatment effects in digital experimentation. *Journal of Business & Economic Statistics*, 34(4):661–672.
- Taylor and Stone (2009). Transfer learning for reinforcement learning domains: A survey. *DAAAI Spring Symposium-Technical Report, 2013*.
- Thrun (1996). Is learning the n-th thing any easier than learning the first? *NIPS*.
- Thrun, S. and Pratt, L. (1998). Learning to learn. *Springer Science and Business Media*.
- Tian, L., Alizadeh, A. A., Gentles, A. J., and Tibshirani, R. (2014). A simple method for estimating interactions between a treatment and a large number of covariates. *Journal of the American Statistical Association*, 109(508):1517–1532.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., and et al (2016). Matching networks for one shot learning. *Neural Information Processing Systems (NIPS)*.
- Wager, S. and Athey, S. (2017). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*.
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., and Perona, P. (2010). Caltech-ucsd birds 200. technical report cns-tr-2010-001. *California Institute of Technology*.

## A APPENDIX: SIMULATION STUDIES AND APPLICATION

### A.1 MNIST SIMULATION

For our MNIST simulation study (Section 4.2), we used the MNIST database (LeCun, 1998) which contains labeled handwritten images. We follow here the notation of Nie and Wager (2017), who introduce a very similar simulation study which is not trying to evaluate transfer learning for CATE estimation, but instead emulates a RCT with the goal to evaluate different CATE estimators.

The MNIST data set contains labeled image data  $(X_i, C_i)$ , where  $X_i$  denotes the raw image of  $i$  and  $C_i \in \{0, \dots, 9\}$  denotes its label. We create  $k$  Data Generating Processes (DGPs),  $D_1, \dots, D_k$ , each of which specifies a distribution of  $(Y_i(0), Y_i(1), W_i, X_i)$  and represents different CATE estimation problems.

In this simulation, we let  $W_i = 0$  if the image  $X_i$  is placed in the control, and  $W_i = 1$  if the image  $X_i$  is placed in the treatment.  $Y_i(W_i)$  quantifies the the outcome of  $X_i$  under  $W_i$ .

To generate a DGP  $D_j$ , we first sample weights in the following way,

$$\begin{aligned} m^j(0), m^j(1), \dots, m^j(9) &\stackrel{iid}{\sim} \text{Unif}(-3, 3), \\ t^j(0), t^j(1), \dots, t^j(9) &\stackrel{iid}{\sim} \text{Unif}(-1, 1), \\ p^j(0), p^j(1), \dots, p^j(9) &\stackrel{iid}{\sim} \text{Unif}(0.3, 0.7), \end{aligned}$$

and we define the response functions and the propensity score as

$$\begin{aligned} \mu_0^j(C_i) &= m^j(C_i) + 3C_i, \\ \mu_1^j(C_i) &= \mu_0^j(C_i) + t^j(C_i), \\ e^j(C_i) &= p^j(C_i). \end{aligned}$$

To generate  $(Y_i(0), Y_i(1), W_i, X_i)$  from  $D_j$ , we first sample a  $(X_i, C_i)$  from the MNIST data set, and we then generate  $Y_i(0), Y_i(1)$ , and  $W_i$  in the following way:

$$\begin{aligned} \varepsilon_i &\stackrel{iid}{\sim} \mathcal{N}(0, 1) \\ Y_i(0) &= \mu_0(C_i) + \varepsilon_i \\ Y_i(1) &= \mu_1(C_i) + \varepsilon_i \\ W_i &\sim \text{Bern}(e(C_i)). \end{aligned}$$

During training,  $X_i, W_i$ , and  $Y_i$  are made available to the convolutional neural network, which then predicts  $\hat{\tau}$  given a test image  $X_i$  and a treatment  $W_i$ .  $\tau$  is the difference in the outcome given the difference in treatment and control.

Having access to multiple DGPs can be interpreted as having access to prior experiments done on a similar population of images, allowing us to explore the effects of different transfer learning methods when predicting the effect of a treatment in a new image.

### A.2 GOTV DATA EXAMPLE AND SIMULATION

In this section, we describe how the simulations for the GOTV example in the main paper were done and we discuss the results of a much bigger simulation study with 51 experiments which is summarized in Tables 1, 2, and 3.

#### A.2.1 DATA GENERATING PROCESSES FOR OUR REAL WORLD EXAMPLE

For our data example, we took one of the experiments conducted by Gerber et al. (2017). The study took place in 2014 in Alaska and 252,576 potential voters were randomly assigned in a control and a treatment group. Subjects in the treatment group were sent a mailer as described in the main text and their voting turnout was recorded.

To evaluate the performance of different CATE estimators we need to know the true CATEs, which are unknown due to the fundamental problem of causal inference. To still be able to evaluate CATE

estimators researchers usually estimate the potential outcomes using some machine learning method and then generate the data from this estimate. This is to some extent also a simulation, but unlike classical simulation studies it is not up to the researcher to determine the data generating distribution. The only choice of the researcher lies in the type of estimator she uses to estimate the response functions. To avoid being misled by artifacts created by a particular method, we used a linear model in **Real World Data Set 1** and random forests estimator in **Real World Data Set 2**.

Specifically, we generate for each experiment a true CATE and we simulate new observed outcomes based on the real data in four steps.

1. We first use the estimator of choice (e.g., a random forests estimator) and train it on the treated units and on the control units separately to get estimates for the response functions,  $\mu_0$  and  $\mu_1$ .
2. Next, we sample  $N$  units from the underlying experiment to get the features and the treatment assignment of our samples  $(X_i, W_i)_{i=1}^N$ .
3. We then generate the true underlying CATE for each unit using  $\tau_i = \tau(X_i) = \mu_1(X_i) - \mu_0(X_i)$ .
4. Finally we generate the observed outcome by sampling a Bernoulli distributed variable around mean  $\mu_i$ .

$$Y_i^{obs} \sim \text{Bern}(\mu_i), \quad \mu_i = \begin{cases} \mu_0(X_i) & \text{if } W = 0, \\ \mu_1(X_i) & \text{if } W = 1. \end{cases}$$

After this procedure, we have 17 data sets corresponding to the 17 experiments for which we know the true CATE function, which we can now use to evaluate CATE estimators and CATE transfer learners.

#### A.2.2 DATA GENERATING PROCESSES FOR OUR SIMULATION STUDY

Simulations motivated by real-world experiments are important to assess whether our methods work well for voter persuasion data sets, but it is important to also consider other settings to evaluate the generalizability of our conclusions.

To do this, we first specify the control response function,  $\mu_0(x) = \mathbb{E}[Y(0)|X = x] \in [0, 1]$ , and the treatment response function,  $\mu_1(x) = \mathbb{E}[Y(1)|X = x] \in [0, 1]$ .

We then use each of the 17 experiments to generate a simulated experiment in the following way:

1. We sample  $N$  units from the underlying experiment to get the features and the treatment assignment of our samples  $(X_i, W_i)_{i=1}^N$ .
2. We then generate the true underlying CATE for each unit using  $\tau_i = \tau(X_i) = \mu_1(X_i) - \mu_0(X_i)$ .
3. Finally we generate the observed outcome by sampling a Bernoulli distributed variable around mean  $\mu_i$ .

$$Y_i^{obs} \sim \text{Bern}(\mu_i), \quad \mu_i = \begin{cases} \mu_0(X_i) & \text{if } W = 0, \\ \mu_1(X_i) & \text{if } W = 1. \end{cases}$$

The experiments range in size from 5,000 units to 400,000 units per experiment and the covariate vector is 11 dimensional and the same as in the main part of the paper. We will present here three different setups.

**Simulation LM** (Table 1): We choose here  $N$  to be all units in the corresponding experiment. Sample  $\beta^0 = (\beta_1^0, \dots, \beta_d^0) \stackrel{iid}{\sim} \mathcal{N}(0, 1)$  and  $\beta^1 = (\beta_1^1, \dots, \beta_d^1) \stackrel{iid}{\sim} \mathcal{N}(0, 1)$  and define,

$$\begin{aligned} \mu_0(x) &= \text{logistic}(x\beta^0), \\ \mu_1(x) &= \text{logistic}(x\beta^1). \end{aligned}$$

**Simulation RF** (Table 2): We choose here  $N$  to be all units in the corresponding experiment.

1. Train a random forests estimator on the real data set and define  $\mu_0$  to be the resulting estimator,
2. Sample a covariate  $f$  (e.g., age),
3. Sample a random value in the support of  $f$  (e.g., 38),
4. Sample a shift  $s \sim \mathcal{N}(0, 4)$ .

Now define the potential outcomes as follows:

$$\begin{aligned}\mu_0(x) &= \text{trained Random Forests algorithm} \\ \mu_1(x) &= \text{logistic}(\text{logit}(\mu_0(x) + s * 1_{f \geq v}))\end{aligned}$$

**Simulation RFt** (Table 3): This experiment is the same as Simulation RF, but use only one percent of the data,  $N = \frac{\#units}{100}$ .

## B PSEUDO CODE FOR CATE ESTIMATORS

In this section, we will present pseudo code for the CATE estimators in this paper. We present code for the meta learning algorithms in Section C. We denote by  $Y^0$  and  $Y^1$  the observed outcomes for the control and the treated group. For example,  $Y_i^1$  is the observed outcome of the  $i$ th unit in the treated group.  $X^0$  and  $X^1$  are the features of the control and treated units, and hence,  $X_i^1$  corresponds to the feature vector of the  $i$ th unit in the treated group.  $M_k(Y \sim X)$  is the notation for a regression estimator, which estimates  $x \mapsto \mathbb{E}[Y|X = x]$ . It can be any regression/machine learning estimator, but in this paper we only choose it to be a neural network or random forest.

These algorithms first appeared in (Künzel et al., 2017; Stadie et al., 2018). We reproduce them here for completeness.

---

### Algorithm 1 T-learner

---

- 1: **procedure** T-LEARNER( $X, Y^{obs}, W$ )
  - 2:      $\hat{\mu}_0 = M_0(Y^0 \sim X^0)$
  - 3:      $\hat{\mu}_1 = M_1(Y^1 \sim X^1)$
  - 4:      $\hat{\tau}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x)$
  - 5: **end procedure**
- 

$M_0$  and  $M_1$  are here some, possibly different machine learning/regression algorithms.

---

### Algorithm 2 S-learner

---

- 1: **procedure** S-LEARNER( $X, Y^{obs}, W$ )
  - 2:      $\hat{\mu} = M(Y^{obs} \sim (X, W))$
  - 3:      $\hat{\tau}(x) = \hat{\mu}(x, 1) - \hat{\mu}(x, 0)$
  - 4: **end procedure**
- 

$M(Y^{obs} \sim (X, W))$  is the notation for estimating  $(x, w) \mapsto \mathbb{E}[Y|X = x, W = w]$  while treating  $W$  as a 0,1-valued feature.

---

### Algorithm 3 X-learner

---

- 1: **procedure** X-LEARNER( $X, Y^{obs}, W, g$ )
  - 2:      $\hat{\mu}_0 = M_1(Y^0 \sim X^0)$  ▷ Estimate response function
  - 3:      $\hat{\mu}_1 = M_2(Y^1 \sim X^1)$
  - 4:      $\tilde{D}_i^1 = Y_i^1 - \hat{\mu}_0(X_i^1)$  ▷ Compute imputed treatment effects
  - 5:      $\tilde{D}_i^0 = \hat{\mu}_1(X_i^0) - Y_i^0$
  - 6:      $\hat{\tau}_1 = M_3(\tilde{D}^1 \sim X^1)$  ▷ Estimate CATE for treated and control
  - 7:      $\hat{\tau}_0 = M_4(\tilde{D}^0 \sim X^0)$
  - 8:      $\hat{\tau}(x) = g(x)\hat{\tau}_0(x) + (1 - g(x))\hat{\tau}_1(x)$  ▷ Average the estimates
  - 9: **end procedure**
- 

$g(x) \in [0, 1]$  is a weighing function which is chosen to minimize the variance of  $\hat{\tau}(x)$ . It is sometimes possible to estimate  $\text{Cov}(\tau_0(x), \tau_1(x))$ , and compute the best  $g$  based on this estimate. However, we have made good experiences by choosing  $g$  to be an estimate of the propensity score, but also choosing it to be constant and equal to the ratio of treated units usually leads to a good estimator of the CATE.

---

**Algorithm 4** Y-Learner Pseudo Code

---

```
1: if  $W_i == 0$  then  
2:   Update the network  $\pi_{\theta_0}$  to predict  $Y_i^{obs}$   
3:   Update the network  $\pi_{\theta_1}$  to predict  $Y_i^{obs} + \pi_{\tau}(X_i)$   
4:   Update the network  $\pi_{\tau}$  to predict  $\pi_{\theta_1}(X_i) - Y_i^{obs}$   
5: end if  
6: if  $W_i == 1$  then  
7:   Update the network  $\pi_{\theta_0}$  to predict  $Y_i^{obs} - \pi_{\tau}(X_i)$   
8:   Update the network  $\pi_{\theta_1}$  to predict  $Y_i^{obs}$   
9:   Update the network  $\pi_{\tau}$  to predict  $Y_i^{obs} - \pi_{\theta_0}(X_i)$   
10: end if
```

---

This process describes training the Y-Learner for one step given a data point  $(Y_i^{obs}, X_i, W_i)$

## C EXPLICIT TRANSFER LEARNING ALGORITHMS FOR CATE ESTIMATION

## C.1 MLRW TRANSFER FOR CATE ESTIMATION

**Algorithm 5** MLRW Transfer for Cate Estimation.

---

```

1: Let  $\mu_0^{(i)}$  and  $\mu_1^{(i)}$  be the outcome under treatment and control for experiment  $i$ .
2: Let numexps be the number of experiments.
3: Let  $\pi_\theta$  be an  $N$  layer neural network parameterized by  $\theta = [\theta_0, \dots, \theta_N]$ .
4: Let  $\epsilon = [\epsilon_0, \dots, \epsilon_N]$  be a vector, where  $N$  is the number of layers in  $\pi_\theta$ .
5: Let outeriters be the total number of training iterations.
6: Let inneriters be the number of inner loop training iterations.
7: for oiter < outeriters do
8:   for i < numexps do
9:     Sample  $X_0$  and  $X_1$ : control and treatment units from experiment  $i$ 
10:    for  $j = [0, 1]$  do ▷ j iterating over treatment and control
11:      Let  $U_0(\theta) = \theta$ 
12:      for k < inneriters do
13:         $\mathcal{L} = \|\pi_{U_k(\theta)}(X_j) - \mu_j(X_j)\|$ 
14:        Compute  $\nabla_\theta \mathcal{L}$ .
15:        Use ADAM with  $\nabla_\theta \mathcal{L}$  to obtain  $U_{k+1}(\theta)$ .
16:         $U_k(\theta) = U_{k+1}(\theta)$ 
17:      end for
18:      for  $p < N$  do
19:         $\theta_p = \epsilon_p \cdot U_k(\theta_p) + (1 - \epsilon_p) \cdot \theta_p$ .
20:      end for
21:    end for
22:  end for
23: end for
24: To Evaluate CATE estimate, do
25:  $C = []$ 
26: for i < numexps do
27:   Sample  $X_0$  and  $X_1$ : control and treatment units from experiment  $i$ 
28:   Sample  $X$ : test units from experiment  $i$ .
29:   for  $j = [0, 1]$  do ▷ j iterating over treatment and control
30:     for  $k < inneriters$  do
31:        $\mathcal{L} = \|\pi_{U_k(\theta)}(X_j) - \mu_j(X_j)\|$ 
32:       Compute  $\nabla_\theta \mathcal{L}$ .
33:       Use ADAM with  $\nabla_\theta \mathcal{L}$  to obtain  $U_{k+1}(\theta)$ .
34:        $U_k(\theta) = U_{k+1}(\theta)$ 
35:     end for
36:      $\hat{\mu}_j = \pi_{U_k(\theta)}(X)$ 
37:   end for
38:    $\hat{\tau}_i = \hat{\mu}_0 - \hat{\mu}_1$ 
39:   C.append( $\hat{\tau}_i$ )
40: end for
41: return  $C$ 

```

---





## C.3 T-LEARNER TRANSFER CATE ESTIMATORS

Here, we present full pseudo code for the algorithms from Section 3 using the T-learner as a base learner. All of these algorithms can be extended to other learners including  $S$ ,  $R$ ,  $X$ , and  $Y$ . See the released code for implementations.

**Algorithm 7** Vanilla T-learner (also referred to as Baseline T-learner)

---

```

1: Let  $\mu_0$  and  $\mu_1$  be the outcome under treatment and control.
2: Let  $X$  be the experimental data. Let  $X_t$  be the test data.
3: Let  $\pi_{\theta_0}$  and  $\pi_{\theta_1}$  be a neural networks parameterized by  $\theta_0$  and  $\theta_1$ .
4: Let  $\theta = \theta_0 \cup \theta_1$ .
5: Let numiters be the total number of training iterations.
6: Let batchsize be the number of units sampled. We use 64.
7: for  $i < \text{numiters}$  do
8:   Sample  $X_0$  and  $X_1$ : control and treatment units. Sample batchsize units.
9:    $\mathcal{L}_0 = \|\pi_{\theta_0}(X_0) - \mu_0(X_0)\|$ 
10:   $\mathcal{L}_1 = \|\pi_{\theta_1}(X_1) - \mu_1(X_1)\|$ 
11:   $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1$ 
12:  Compute  $\nabla_{\theta}\mathcal{L} = \frac{\partial\mathcal{L}}{\partial\theta}$ .
13:  Apply ADAM with gradients given by  $\nabla_{\theta}\mathcal{L}$ .
14: end for
15:  $\hat{\mu}_0 = \pi_{\theta_0}(X_t)$ 
16:  $\hat{\mu}_1 = \pi_{\theta_1}(X_t)$ 
17: return CATE estimate  $\hat{\tau} = \hat{\mu}_1 - \hat{\mu}_0$ 

```

---

**Algorithm 8** Warm Start T-learner

---

```

1: Let  $\mu_0^i$  and  $\mu_1^i$  be the outcome under treatment and control for experiment  $i$ .
2: Let  $X^i$  be the data for experiment  $i$ . Let  $X_t^i$  be the test data for experiment  $i$ .
3: Let  $\pi_{\theta_0}$  and  $\pi_{\theta_1}$  be a neural networks parameterized by  $\theta_0$  and  $\theta_1$ .
4: Let  $\theta = \theta_0 \cup \theta_1$ .
5: Let numiters be the total number of training iterations.
6: Let batchsize be the number of units sampled. We use 64.
7: for  $i < \text{numiters}$  do
8:   Sample  $X_0^0$  and  $X_1^0$ : control and treatment units for experiment 0. Sample batchsize units.
9:    $\mathcal{L}_0 = \|\pi_{\theta_0}(X_0^0) - \mu_0(X_0^0)\|$ 
10:   $\mathcal{L}_1 = \|\pi_{\theta_1}(X_1^0) - \mu_1(X_1^0)\|$ 
11:   $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1$ 
12:  Compute  $\nabla_{\theta}\mathcal{L} = \frac{\partial\mathcal{L}}{\partial\theta}$ .
13:  Apply ADAM with gradients given by  $\nabla_{\theta}\mathcal{L}$ .
14: end for
15: for  $i < \text{numiters}$  do
16:   Sample  $X_0^1$  and  $X_1^1$ : control and treatment units for experiment 1. Sample batchsize units.
17:    $\mathcal{L}_0 = \|\pi_{\theta_0}(X_0^1) - \mu_0(X_0^1)\|$ 
18:    $\mathcal{L}_1 = \|\pi_{\theta_1}(X_1^1) - \mu_1(X_1^1)\|$ 
19:    $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1$ 
20:   Compute  $\nabla_{\theta}\mathcal{L} = \frac{\partial\mathcal{L}}{\partial\theta}$ .
21:   Apply ADAM with gradients given by  $\nabla_{\theta}\mathcal{L}$ .
22: end for
23:  $\hat{\mu}_0 = \pi_{\theta_0}(X_t^1)$ 
24:  $\hat{\mu}_1 = \pi_{\theta_1}(X_t^1)$ 
25: return CATE estimate  $\hat{\tau} = \hat{\mu}_1 - \hat{\mu}_0$ 

```

---

**Algorithm 9** Frozen Features T-learner

- 
- 1: Let  $\mu_0^i$  and  $\mu_1^i$  be the outcome under treatment and control for experiment  $i$ .
  - 2: Let  $X^i$  be the data for experiment  $i$ . Let  $X_t^i$  be the test data for experiment  $i$ .
  - 3: Let  $\pi_\rho$  be a generic expression for a neural network parameterized by  $\rho$ .
  - 4: Let  $\theta_0^0, \theta_0^1, \theta_1^0, \theta_1^1$  be neural network parameters. The subscript indicates the outcome that  $\theta$  is associated with predicting (0 for control and 1 for treatment) and the superscript indexes the experiment.
  - 5: Let  $\gamma_0$  be the first  $k$  layers of  $\pi_{\theta_0^0}$ . Define  $\gamma_1$  analogously.
  - 6: Let  $\theta^i = \theta_0^i \cup \theta_1^i$ .
  - 7: Let numiters be the total number of training iterations.
  - 8: Let batchsize be the number of units sampled. We use 64.
  - 9: **for**  $i < \text{numiters}$  **do**
  - 10:     Sample  $X_0^0$  and  $X_1^0$ : control and treatment units for experiment 0. Sample batchsize units.
  - 11:      $\mathcal{L}_0 = \|\pi_{\theta_0^0}(X_0^0) - \mu_0(X_0^0)\|$
  - 12:      $\mathcal{L}_1 = \|\pi_{\theta_1^0}(X_1^0) - \mu_1(X_1^0)\|$
  - 13:      $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1$
  - 14:     Compute  $\nabla_\theta \mathcal{L} = \frac{\partial \mathcal{L}}{\partial \theta}$ .
  - 15:     Apply ADAM with gradients given by  $\nabla_{\theta^0} \mathcal{L}$ .
  - 16: **end for**
  - 17: **for**  $i < \text{numiters}$  **do**
  - 18:     Sample  $X_0^1$  and  $X_1^1$ : control and treatment units for experiment 1. Sample batchsize units.
  - 19:     Compute  $Z_0^1 = \pi_\gamma(X_0^1)$  and  $Z_1^1 = \pi_\gamma(X_1^1)$
  - 20:      $\mathcal{L}_0 = \|\pi_{\theta_0^1}(Z_0^1) - \mu_0(Z_0^1)\|$
  - 21:      $\mathcal{L}_1 = \|\pi_{\theta_1^1}(Z_1^1) - \mu_1(Z_1^1)\|$
  - 22:      $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1$
  - 23:     Compute  $\nabla_{\theta^1} \mathcal{L} = \frac{\partial \mathcal{L}}{\partial \theta^1}$ . Do not compute gradients with respect to  $\theta^0$  parameters.
  - 24:     Apply ADAM with gradients given by  $\nabla_{\theta^1} \mathcal{L}$ .
  - 25: **end for**
  - 26: Compute  $Z_t^1 = \pi_\gamma(X_t^1)$ .
  - 27:  $\hat{\mu}_0 = \pi_{\theta_0^1}(Z_t^1)$
  - 28:  $\hat{\mu}_1 = \pi_{\theta_1^1}(Z_t^1)$
  - 29: **return** CATE estimate  $\hat{\tau} = \hat{\mu}_1 - \hat{\mu}_0$
-

**Algorithm 10** Multi-Head T-learner

- 
- 1: Let  $\mu_0^i$  and  $\mu_1^i$  be the outcome under treatment and control for experiment  $i$ .
  - 2: Let  $X^i$  be the data for experiment  $i$ . Let  $X_t^i$  be the test data for experiment  $i$ .
  - 3: Let  $\pi_\rho$  be a generic expression for a neural network parameterized by  $\rho$ .
  - 4: Let  $\theta_0$  be base neural network layers shared by all experiments for predicting outcomes under control.
  - 5: Let  $\theta_1$  be base neural network layers shared by all experiments for predicting outcomes under treatment.
  - 6: Let  $\phi_0^{(i)}$  be neural network layers receiving  $\pi_{\theta_0}(x_0^i)$  as input and predicting  $\mu_0^{(i)}(x_0^i)$  in experiment  $i$ .
  - 7: Let  $\phi_1^{(i)}$  be neural network layers receiving  $\pi_{\theta_1}(x_1^i)$  as input and predicting  $\mu_1^{(i)}(x_1^i)$  in experiment  $i$ .
  - 8: Let  $\omega_0^{(i)} = [\theta, \phi_0^{(i)}]$  be all trainable parameters used to predict  $\mu_0^i$ .
  - 9: Let  $\omega_1^{(i)} = [\theta, \phi_1^{(i)}]$  be all trainable parameters used to predict  $\mu_1^i$ .
  - 10: Let  $\Omega^i = \omega_0^{(i)} \cup \omega_1^{(i)}$ .
  - 11: Let numiters be the total number of training iterations.
  - 12: Let batchsize be the number of units sampled. We use 64.
  - 13: Let numexps be the number of experiments.
  - 14: **for**  $i < \text{numiters}$  **do**
  - 15:     **for**  $j < \text{numexps}$  **do**
  - 16:         Sample  $X_0^j$  and  $X_1^j$ : control and treatment units for experiment  $j$ . Sample batchsize units.
  - 17:         Compute  $Z_0^j = \pi_{\theta_0}(X_0^j)$  and  $Z_1^j = \pi_{\theta_1}(X_1^j)$
  - 18:         Compute  $\hat{\mu}_0^j = \pi_{\phi_0^j}(z_0^j)$  and  $\hat{\mu}_1^j = \pi_{\phi_1^j}(z_1^j)$
  - 19:          $\mathcal{L}_0 = \|\hat{\mu}_0^j - \mu_0^j(X_0^j)\|$
  - 20:          $\mathcal{L}_1 = \|\hat{\mu}_1^j - \mu_1^j(X_1^j)\|$
  - 21:          $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1$
  - 22:         Compute  $\nabla_{\Omega^i} \mathcal{L} = \frac{\partial \mathcal{L}}{\partial \Omega^i}$ .
  - 23:         Apply ADAM with gradients given by  $\nabla_{\theta} \mathcal{L}$ .
  - 24:     **end for**
  - 25: **end for**
  - 26: Let  $C = []$
  - 27: **for**  $j < \text{numexps}$  **do**
  - 28:     Compute  $Z_0^j = \pi_{\theta_0}(X_t^j)$  and  $Z_1^j = \pi_{\theta_1}(X_t^j)$
  - 29:     Compute  $\hat{\mu}_0^j = \pi_{\phi_0^j}(z_0^j)$  and  $\hat{\mu}_1^j = \pi_{\phi_1^j}(z_1^j)$
  - 30:     Estimate CATE  $\hat{\tau} = \hat{\mu}_1^j - \hat{\mu}_0^j$ .
  - 31:      $C.append(\hat{\tau})$
  - 32: **end for**
  - 33: **return**  $C$
-

**Algorithm 11** SF Reptile T-learner

---

```

1: Let  $\mu_0^i$  and  $\mu_1^i$  be the outcome under treatment and control for experiment  $i$ .
2: Let  $X^i$  be the data for experiment  $i$ . Let  $X_t^i$  be the test data for experiment  $i$ .
3: Let  $\pi_{\theta_0}$  and  $\pi_{\theta_1}$  be a neural networks parameterized by  $\theta_0$  and  $\theta_1$ .
4: Let  $\theta = \theta_0 \cup \theta_1$ .
5: Let  $\epsilon = [\epsilon_0, \dots, \epsilon_N]$  be a vector, where  $N$  is the number of layers in  $\pi_{\theta_i}$ .
6: Let numouteriters be the total number of outer training iterations.
7: Let numinneriters be the total number of inner training iterations.
8: Let numexps be the number of experiments.
9: Let batchsize be the number of units sampled. We use 64.
10: for iouter < numouteriters do
11:   for i < numexps do
12:      $U_0(\theta_0) = \theta_0$ 
13:      $U_0(\theta_1) = \theta_1$ .
14:     for k < numinneriters do
15:       Sample  $X_0^i$  and  $X_1^i$ : control and treatment units. Sample batchsize units.
16:        $\mathcal{L}_0 = \|\pi_{U_k(\theta_0)}(X_0^i) - \mu_0(X_0^i)\|$ 
17:        $\mathcal{L}_1 = \|\pi_{U_k(\theta_1)}(X_1^i) - \mu_1(X_1^i)\|$ 
18:        $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1$ 
19:       Compute  $\nabla_{\theta}\mathcal{L} = \frac{\partial\mathcal{L}}{\partial\theta}$ .
20:       Use ADAM with gradients given by  $\nabla_{\theta}\mathcal{L}$  to obtain  $U_{k+1}(\theta_0)$  and  $U_{k+1}(\theta_1)$ .
21:       Set  $U_k(\theta_0) = U_{k+1}(\theta_0)$  and  $U_k(\theta_1) = U_{k+1}(\theta_1)$ 
22:     end for
23:     for  $p < N$  do
24:        $\theta_p = \epsilon_p \cdot U_k(\theta_p) + (1 - \epsilon_p) \cdot \theta_p$ .
25:     end for
26:   end for
27: end for
28: To Evaluate CATE estimate, do
29:  $C = []$ .
30: for i < numexps do
31:    $U_0(\theta_0) = \theta_0$ 
32:    $U_0(\theta_1) = \theta_1$ .
33:   for k < numinneriters do
34:     Sample  $X_0^i$  and  $X_1^i$ : control and treatment units. Sample batchsize units.
35:      $\mathcal{L}_0 = \|\pi_{U_k(\theta_0)}(X_0^i) - \mu_0(X_0^i)\|$ 
36:      $\mathcal{L}_1 = \|\pi_{U_k(\theta_1)}(X_1^i) - \mu_1(X_1^i)\|$ 
37:      $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1$ 
38:     Compute  $\nabla_{\theta}\mathcal{L} = \frac{\partial\mathcal{L}}{\partial\theta}$ .
39:     Use ADAM with gradients given by  $\nabla_{\theta}\mathcal{L}$  to obtain  $U_{k+1}(\theta_0)$  and  $U_{k+1}(\theta_1)$ .
40:     Set  $U_k(\theta_0) = U_{k+1}(\theta_0)$  and  $U_k(\theta_1) = U_{k+1}(\theta_1)$ 
41:   end for
42:    $\hat{\mu}_0^i = \pi_{U_k(\theta_0)}(X_0^i)$ 
43:    $\hat{\mu}_1^i = \pi_{U_k(\theta_1)}(X_1^i)$ 
44:    $\hat{\tau}^i = \hat{\mu}_1^i - \hat{\mu}_0^i$ 
45:    $C.append(\hat{\tau}^i)$ .
46: end for
47: return  $C$ .

```

---

## D FULL RESULTS

Below, we include the full results for the GOTV and MNIST experiments. In particular, we use show results for transfer CATE learners with S, T, X, and Y base learners. We also provide full tables of more comprehensive results for all methods and all train-test splits.




Figure 5: Social Pressure and Voter Turnout, Version 1. Our results far exceed the previous state of the art, which are represented here as S-RF, T-RF, and the baseline method for S-NN and T-NN. Our new methods are Y-NN and the transfer learning methods: warm, frozen, multi-head, joint, SF Reptile, and MLRW.

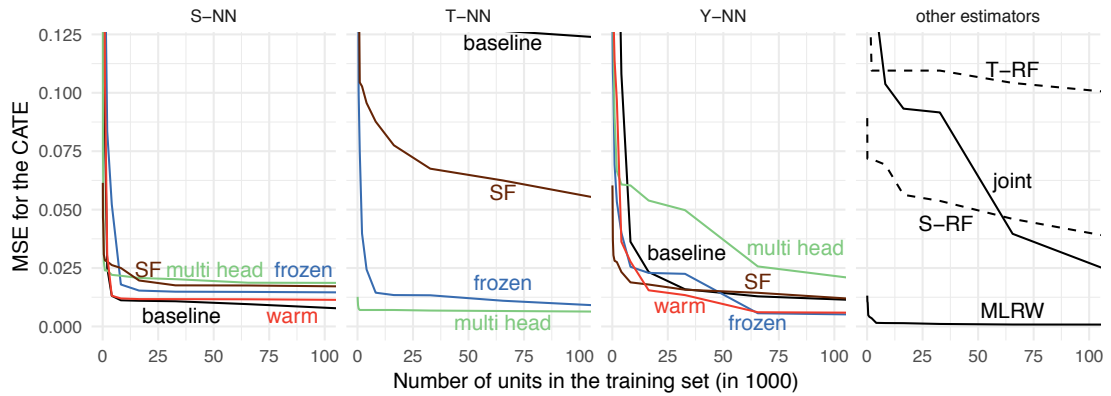


Figure 6: Social Pressure and Voter Turnout, Version 2. Our results exceed the previous state of the art results, which are represented here as S-RF, T-RF, and the baseline method for S-NN and T-NN. Our new methods are Y-NN and the transfer learning methods: warm, frozen, multi-head, joint, SF Reptile, and MLRW.




Figure 7: MNIST task



| Method      | LM-1         | LM-2        | LM-3      | LM-4         | LM-5         | LM-6         | LM-7         | LM-8         | LM-9         | LM-10        | LM-11     | LM-12        | LM-13        | LM-14        |
|-------------|--------------|-------------|-----------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|-----------|--------------|--------------|--------------|
| <b>t-lm</b> | 15.95        | 7.82        | 20.14     | 6.62         | 46.15        | 17.24        | 9.88         | 10.65        | 44.77        | 7.63         | 8.43      | 9            | 11.21        | 20.9         |
| <b>s-rf</b> | 19.13        | 13.18       | 7.62      | 13.27        | 15.66        | 15.58        | 11.44        | 16.3         | 11.34        | 16.53        | 12.57     | 14.11        | 13.49        | 18.56        |
| <b>t-rf</b> | 20.04        | 13.56       | 7.79      | 13.62        | 16           | 15.99        | 11.76        | 16.65        | 11.54        | 17.11        | 13.66     | 14.38        | 13.67        | 18.96        |
| <b>R-NN</b> | 18.95        | 5.19        | 9.33      | 28.98        | 3.04         | 10.03        | 4.88         | 16.34        | 7.91         | 14.56        | 19.23     | 3.5          | 10.4         | 15.53        |
| <b>S-NN</b> | 6.74         | 6.17        | 2.75      | 5.76         | 5.68         | 6.27         | 4.23         | 7.17         | 4.15         | 6.77         | 4.88      | 5.9          | 6.63         | 9.87         |
| frozen      | 3.3          | 3.05        | 2.34      | 3.83         | 3.6          | 4.89         | 4.56         | 8.47         | 5.87         | 5.22         | 5.43      | 6.77         | 5.15         | 5.58         |
| multi head  | 4.85         | 38.65       | 7.54      | 55.92        | 11.36        | 3.98         | 50.76        | 7.74         | 5.72         | 7.73         | 8.64      | 31.06        | 30.81        | 18.08        |
| SF          | 6.84         | 5.29        | 3.77      | 6.86         | 5.94         | 7.19         | 4.6          | 8.08         | 5.61         | 7.12         | 4.34      | 6.05         | 8.59         | 10.75        |
| baseline    | 7.29         | 6.44        | 4.08      | 7.25         | 6.74         | 7.17         | 6.03         | 8.8          | 5.63         | 7.6          | 5.82      | 6.53         | 8.55         | 12.02        |
| warm        | <b>T-NN</b>  | 6.53        | 6.73      | 4.49         | 6.35         | 7.23         | 6.61         | 5.99         | 7.59         | 6.79         | 5.99      | 6.58         | 7.38         | 9.39         |
| frozen      | 2.73         | 2.34        | 1.34      | 2.11         | 2.49         | 2.3          | 1.97         | 2.73         | 1.94         | 2.36         | 1.95      | 2.32         | 2.65         | 3.64         |
| multi head  | SF           | 20.72       | 27.71     | 8.54         | 20.18        | 23.06        | 15.35        | 14.27        | 13.05        | 27.94        | 40.03     | 16.3         | 20.81        | 6.78         |
| SF          | baseline     | 22.76       | 5.34      | 5.98         | 5.84         | 5.16         | 10.37        | 10.9         | 7.26         | 10.18        | 6.26      | 5.69         | 6.71         | 10.31        |
| baseline    | warm         | 23.46       | 7.2       | 5.75         | 6.41         | 5.21         | 11.56        | 12.21        | 8.77         | 6.81         | 8.93      | 6.15         | 7.25         | 18.98        |
| warm        | <b>X-NN</b>  | 6.63        | 14.04     | 10.73        | 19.57        | 17.94        | 14           | 13.67        | 18.83        | 14.36        | 12.25     | 16.61        | 39.96        | 32.81        |
| frozen      | 1.19         | 11.38       | 84.13     | 174.18       | 1.87         | 19.55        | 62.12        | 22.7         | 9.67         | <b>0.85*</b> | 3.34      | 5.03         | <b>0.94*</b> | 111.42       |
| multi head  | SF           | 18.83       | 10.72     | 10.3         | 10.61        | 10.11        | 12.5         | 21.37        | 11.12        | 16.33        | 9.81      | 13.8         | 9.85         | 10.15        |
| SF          | baseline     | 19.52       | 8.25      | 4.68         | 5.06         | 6.6          | 11.5         | 10.78        | 12           | 6.25         | 13.11     | 7.27         | 9.2          | 18.45        |
| baseline    | warm         | 20.06       | 8.54      | 5.57         | 6.37         | 10.77        | 9.34         | 11.36        | 8.03         | 8.66         | 10.96     | 7.52         | 11.57        | 16.7         |
| warm        | <b>Y-NN</b>  | 1.54        | <b>1*</b> | 2.1          | 3.3          | <b>1.11*</b> | 46.07        | 27.63        | 7.48         | 7.21         | 1.02      | <b>1.15*</b> | 0.97         | 43.82        |
| frozen      | 0.92         | 2.21        | 1.26      | 15.86        | 1.19         | 20.47        | 1.76         | 2.75         | 3.96         | 9.4          | 19.11     | 1.26         | 35.43        | 9.29         |
| multi head  | SF           | 5.68        | 12        | 16.57        | 38           | 5.54         | 3.49         | 8.2          | 2.61         | 7.21         | 8.48      | 12.98        | 9.58         | 5.08         |
| SF          | baseline     | <b>0.9*</b> | 1.31      | 5.24         | 31.43        | 6.8          | <b>1.23*</b> | 7.5          | <b>1.07*</b> | 1.35         | 8.19      | 1.2          | 4.11         | 7.72         |
| baseline    | warm         | 48.91       | 1.26      | 5.79         | 3.61         | 29.43        | 2.71         | 3.94         | 21.76        | 13.25        | 15.78     | 17.45        | 1.12         | 29.01        |
| warm        | <b>joint</b> | 13.75       | 5.81      | 3.46         | 4.12         | 3.46         | 12.22        | 9.58         | 14.96        | 4.75         | 8.55      | 13.13        | 9.91         | 11.68        |
| joint       | <b>MLRW</b>  | 1           | 1.41      | <b>1.05*</b> | <b>1.94*</b> | 1.97         | 1.26         | <b>1.03*</b> | <b>0.9*</b>  | 1.85         | <b>1*</b> | 1.15         | 1.57         | <b>2.75*</b> |

Table 1: MSE in percent for different CATE estimators.

| Method       | RF-1         | RF-2         | RF-3         | RF-4        | RF-5         | RF-6         | RF-7        | RF-8         | RF-9         | RF-10        | RF-11        | RF-12        | RF-13 | RF-14        |
|--------------|--------------|--------------|--------------|-------------|--------------|--------------|-------------|--------------|--------------|--------------|--------------|--------------|-------|--------------|
| <b>t-lm</b>  | 17.95        | 26.73        | 2.48         |             | 4.84         | 1.76         | 6.2         | 7.55         | 24.58        | 3.5          | 2.35         | 2.1          | 3.39  | 7.41         |
| <b>s-rf</b>  | 7.18         | 7            | 4.18         | 6.86        | 8.43         | 9.15         | 6.03        | 9.28         | 5.95         | 8.49         | 6.7          | 8.1          | 6.46  | 8.64         |
| <b>t-rf</b>  | 10.95        | 7.76         | 4.69         | 7.79        | 9.03         | 10.07        | 6.67        | 10.08        | 6.3          | 10.35        | 8.19         | 9.28         | 6.73  | 10.04        |
| <b>R-NN</b>  | 11.41        | 1.18         | 5.7          | 1.26        | 0.7          | 3.79         | 9.11        | 8.01         | 7.17         | 5.14         | 4.15         | 0.52         | 1.26  | 6.88         |
| <b>S-NN</b>  | 1.56         | 0.66         | 0.7          | 0.77        | 0.87         | 0.83         | 0.91        | 0.89         | 0.59         | 0.87         | 0.6          | 0.92         | 0.91  | 1.23         |
| frozen       | 0.69         | 1.23         | 0.4          | 1.59        | 1.58         | 0.65         | 0.81        | 0.57         | <b>0.28*</b> | 1.98         | 1.08         | 1.24         | 2.27  | 3.66         |
| multi head   | 1.36         | 9.64         | 4.26         | 4.53        | 6            | 6.74         | 11.68       | 2.8          | 12.01        | 41.6         | 45.72        | 4.59         | 1.98  | 6.75         |
| SF           | 0.91         | 1.28         | 0.84         | 0.93        | 1.78         | 2.05         | 1.16        | 2.04         | 1.61         | 1.06         | 1.29         | 1.49         | 1.99  | 2.69         |
| baseline     | 1.08         | 1.3          | 0.94         | 1.16        | 1.85         | 2.22         | 1.4         | 2.15         | 1.65         | 1.12         | 1.37         | 1.42         | 1.83  | 2.52         |
| warm         | 1.55         | 1.02         | 0.62         | 0.95        | 1.11         | 0.99         | 0.89        | 1.18         | 0.86         | 1.03         | 0.89         | 1.01         | 1.14  | 1.49         |
| <b>T-NN</b>  | 0.66         | 0.58         | 0.35         | 0.53        | 0.63         | 0.53         | 0.47        | 0.63         | 0.49         | 0.57         | 0.51         | 0.55         | 0.64  | 0.91         |
| frozen       | 11.66        | 27.5         | 6.89         | 22.59       | 20.99        | 12.04        | 19.77       | 7.85         | 5.39         | 17.31        | 17.15        | 9.1          | 30.42 | 3.55         |
| multi head   | 7.83         | 4.25         | 1.44         | 1.47        | 1.35         | 1.33         | 2.05        | 8.55         | 1.79         | 3.21         | 2.29         | 3.62         | 2.03  | 8.7          |
| SF           | 12.74        | 3.32         | 1.33         | 1.58        | 1.06         | 1.6          | 2.19        | 9.28         | 1.52         | 2.77         | 3.73         | 4.99         | 1.64  | 8.74         |
| baseline     | 2.53         | 22.89        | 55.57        | 44.11       | 4.18         | 5.72         | 33.18       | 2.62         | 7.59         | 4.96         | 4.41         | 2.01         | 3.01  | 1.23         |
| warm         | 3.45         | 2.53         | 47.09        | 39.7        | 2            | 27.62        | 10.39       | 0.78         | 11.8         | 10.85        | 0.93         | 9.72         | 11.74 | 30.62        |
| <b>X-NN</b>  | 4.34         | 1.85         | 5.47         | 6.82        | 4.21         | 11.89        | 7.91        | 5.02         | 4.67         | 6.82         | 16.39        | 6.99         | 5.22  | 2.11         |
| frozen       | 4.09         | 2.05         | 1.67         | 0.73        | 0.58         | 1.16         | 4.97        | 9.23         | 4.05         | 1.49         | 5.17         | 2.15         | 4.07  | 8.05         |
| multi head   | 2.51         | 1.73         | 1.72         | 1.48        | 0.97         | 1.04         | 7.81        | 4.35         | 4.59         | 1.34         | 3.86         | 1.37         | 2.08  | 4.83         |
| SF           | 0.42         | <b>0.32*</b> | 10.65        | 0.72        | 36           | 1.62         | 0.82        | 0.87         | 15.19        | 1.61         | 0.77         | 46.3         | 28.37 | 1.98         |
| baseline     | 29.41        | 9.71         | 2.74         | 12.27       | 48.87        | 16.59        | 50.21       | 73.74        | 1.65         | 1.76         | 3.52         | 24.26        | 0.76  | 216.69       |
| warm         | 5.45         | 2.51         | 1.53         | 4.76        | 9.61         | 20.81        | 3.68        | 4.84         | 10.8         | 2.87         | 2.78         | 1.36         | 9.07  | 4.08         |
| <b>Y-NN</b>  | 0.71         | 1.06         | 1.54         | 0.57        | 2.44         | 0.47         | 0.73        | 1.38         | 1.25         | 1.77         | 0.71         | 0.29         | 0.63  | 1.98         |
| frozen       | 0.9          | 27.01        | 0.52         | 22.7        | 1.37         | 24.94        | <b>0.3*</b> | 2.08         | 12.36        | 3.39         | 9.58         | 2.74         | 2.47  | 11.48        |
| multi head   | 33.47        | 1.3          | 0.61         | 0.61        | 5.15         | 0.47         | 2.68        | 2.67         | 0.37         | <b>0.37*</b> | 7.76         | 28.14        | 7.48  | 10.04        |
| SF           | <b>0.37*</b> | 1.12         | <b>0.18*</b> | <b>0.3*</b> | <b>0.46*</b> | <b>0.21*</b> | 0.42        | <b>0.45*</b> | 0.37         | 1.18         | <b>0.35*</b> | <b>0.23*</b> | 0.26  | <b>0.16*</b> |
| baseline     | 0.9          | 27.01        | 0.52         | 22.7        | 1.37         | 24.94        | <b>0.3*</b> | 2.08         | 12.36        | 3.39         | 9.58         | 2.74         | 2.47  | 11.48        |
| warm         | 33.47        | 1.3          | 0.61         | 0.61        | 5.15         | 0.47         | 2.68        | 2.67         | 0.37         | <b>0.37*</b> | 7.76         | 28.14        | 7.48  | 10.04        |
| <b>joint</b> | <b>0.37*</b> | 1.12         | <b>0.18*</b> | <b>0.3*</b> | <b>0.46*</b> | <b>0.21*</b> | 0.42        | <b>0.45*</b> | 0.37         | 1.18         | <b>0.35*</b> | <b>0.23*</b> | 0.26  | <b>0.16*</b> |
| <b>MLRW</b>  | <b>0.37*</b> | 1.12         | <b>0.18*</b> | <b>0.3*</b> | <b>0.46*</b> | <b>0.21*</b> | 0.42        | <b>0.45*</b> | 0.37         | 1.18         | <b>0.35*</b> | <b>0.23*</b> | 0.26  | <b>0.16*</b> |

Table 2: MSE in percent for different CATE estimators.

| Method       | RFt-1        | RFt-2        | RFt-3        | RFt-4        | RFt-5        | RFt-6        | RFt-7        | RFt-8        | RFt-9        | RFt-10      | RFt-11       | RFt-12       | RFt-13 | RFt-14       |
|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|-------------|--------------|--------------|--------|--------------|
| <b>t-lm</b>  | 22.65        | 2.74         | 2.3          | 5.07         | 4.78         | 8.3          | 9.77         | 38.73        | 23.94        | 8.94        | 83.37        | 4.26         | 4.31   | 31.76        |
| <b>s-rf</b>  | 3.54         | 7.04         | 4.39         | 6.4          | 8.09         | 6.41         | 11.44        | 6.33         | 6.05         | 6.52        | 5.09         | 9.29         | 11.64  | 4.28         |
| <b>t-rf</b>  | 13.3         | 13.55        | 8.41         | 11.9         | 14.11        | 15.48        | 14.42        | 12.54        | 10.19        | 18.99       | 12.48        | 14.64        | 13.61  | 12.02        |
| <b>R-NN</b>  | 62.08        | 17.68        | 5.22         | 19.91        | 14.43        | 20.53        | 18.51        | 49.63        | 8.99         | 86.89       | 54.17        | 23.86        | 3.02   | 18.98        |
| <b>S-NN</b>  | 2.49         | 51.39        | 37.47        | 49.5         | 43.17        | 29.19        | 21.81        | 27.25        | 61.34        | 27.02       | 17.25        | 64.22        | 24.94  | 11.24        |
| frozen       | 1.4          | 0.57         | 0.73         | 1.48         | 1.54         | 1.27         | 1.1          | 1.63         | 0.88         | 1.07        | 1.53         | 0.79         | 0.74   | 1.59         |
| multi head   | 0.84         | 68.31        | 4.12         | 10.77        | 17.28        | 20.06        | 10.22        | 1.08         | 8.95         | 9.11        | 5.97         | 19.67        | 11.64  | 8.68         |
| SF           | 14.66        | 2.08         | 0.95         | 1.38         | 2.24         | 7.3          | 3.17         | 1.08         | 0.88         | 5.29        | 1.44         | 2.37         | 2.85   | 1.78         |
| baseline     | 17.64        | 2.76         | 1.38         | 2.31         | 1.99         | 8.64         | 4.32         | 1.37         | 0.85         | 5.5         | 3.76         | 5.44         | 3.13   | 3.49         |
| warm         | 2.57         | 1.69         | 1.12         | 1.55         | 1.81         | 1.6          | 1.52         | 1.86         | 1.43         | 1.63        | 1.47         | 1.68         | 1.85   | 2.27         |
| <b>T-NN</b>  | 0.69         | 0.55         | 0.34         | 0.49         | 0.63         | 0.64         | 0.51         | 0.64         | 0.5          | 0.53        | <b>0.46*</b> | <b>0.55*</b> | 0.66   | 0.8          |
| frozen       | 18.04        | 9.54         | 7.33         | 6.61         | 22.11        | 9.21         | 20.2         | 10.65        | 15.06        | 13.87       | 27.64        | 17.96        | 14.08  | 4.49         |
| multi head   | 43.37        | 14.55        | 9.48         | 19.35        | 15.68        | 20.91        | 16           | 41.72        | 6.99         | 30.02       | 45.59        | 16.03        | 4.79   | 18.63        |
| SF           | 69.41        | 19.81        | 12.99        | 24.12        | 20.28        | 37.07        | 19.64        | 39.37        | 9.36         | 40.65       | 37.59        | 18.62        | 6.87   | 32.52        |
| baseline     | 12.41        | 221.49       | 215.01       | 131.46       | 94.89        | 328.53       | 199.48       | 41.14        | 398.33       | 83.75       | 132          | 434.25       | 125.88 | 168.92       |
| warm         | 12.24        | 0.79         | 3.22         | 12.46        | 2.87         | 12.11        | 21.38        | 1.86         | 1.48         | 18.46       | 1.82         | 99.38        | 14.67  | 28.43        |
| <b>X-NN</b>  | 3.98         | 11.03        | 4.33         | 5.31         | 5.42         | 4.67         | 14.03        | 5.99         | 7            | 7           | 6.42         | 7.21         | 4.89   | 2.34         |
| frozen       | 60.44        | 15.03        | 4.93         | 7.08         | 9.51         | 21.22        | 14.7         | 16.83        | 8            | 80.71       | 23.89        | 25.5         | 4.67   | 17.1         |
| multi head   | 30.46        | 13.02        | 5.11         | 9.95         | 9.21         | 13.12        | 9.95         | 12.8         | 4.48         | 50.45       | 19.27        | 12.92        | 8.1    | 15.73        |
| SF           | 1.72         | <b>0.33*</b> | 0.87         | 2.47         | 1.8          | 5.44         | 1.35         | 19.37        | 37.98        | 0.81        | 0.7          | 33.11        | 1.34   | 13.64        |
| baseline     | 5.95         | 11.12        | 0.41         | 2.94         | 2.49         | 24.73        | 3.44         | 58.63        | 12.99        | 0.41        | 9.55         | 0.6          | 3.93   | 0.49         |
| warm         | 3.98         | 14.56        | 12.85        | 4.3          | 3.66         | 9.78         | 11.13        | 3.03         | 5.78         | 3.37        | 11.87        | 11.33        | 4.51   | 2.64         |
| <b>Y-NN</b>  | 0.88         | 10.92        | 0.28         | 1.32         | 1.89         | 11.35        | 8.42         | 1.6          | 5.66         | 0.54        | 0.54         | 2.01         | 0.61   | 2.13         |
| frozen       | 8.03         | 0.48         | 0.36         | 0.71         | 2.53         | 1.21         | <b>0.31*</b> | 1.77         | <b>0.41*</b> | 0.74        | 23.48        | 1.4          | 1.15   | 17.72        |
| multi head   | 102.83       | 38.47        | 133.88       | 15.31        | 37.32        | 410.04       | 38.28        | 14.96        | 5.02         | 11.97       | 33.71        | 26.33        | 11.99  | 147.45       |
| baseline     | <b>0.31*</b> | 1.25         | <b>0.21*</b> | <b>0.27*</b> | <b>0.35*</b> | <b>0.29*</b> | 0.33         | <b>0.61*</b> | 0.61         | <b>0.4*</b> | 0.57         | 1.32         | 0.8    | <b>0.24*</b> |
| warm         |              |              |              |              |              |              |              |              |              |             |              |              |        |              |
| <b>joint</b> |              |              |              |              |              |              |              |              |              |             |              |              |        |              |
| <b>MLRW</b>  |              |              |              |              |              |              |              |              |              |             |              |              |        |              |

Table 3: MSE in percent for different CATE estimators.