

Predicting Counterfactuals from Large Historical Data and Small Randomized Trials

Nir Rosenfeld
Hebrew University of
Jerusalem and
Microsoft Research
nir.rosenfeld@mail.huji.ac.il

Yishay Mansour
Tel Aviv University and
Microsoft Research
mansour@microsoft.com

Elad Yom-Tov
Microsoft Research
eladyt@microsoft.com

ABSTRACT

When a new treatment is considered for use, whether a pharmaceutical drug or a search engine ranking algorithm, a typical question that arises is, will its performance exceed that of the current treatment? The conventional way to answer this counterfactual question is to estimate the effect of the new treatment in comparison to that of the conventional treatment by running a controlled, randomized experiment. While this approach theoretically ensures an unbiased estimator, it suffers from several drawbacks, including the difficulty in finding representative experimental populations as well as the cost of running randomized trials. Moreover, such trials neglect the huge quantities of available control-condition data, which in principle can be utilized for the harder task of *predicting* individualized effects.

In this paper we propose a discriminative framework for predicting the outcomes of a new treatment from a large dataset of the control condition and data from a small (and possibly unrepresentative) randomized trial comparing new and old treatments. Our learning objective, which requires minimal assumptions on the treatments, models the relation between the outcomes of the different conditions. This allows us to not only estimate mean effects but also to generate individual predictions for examples outside the small randomized sample.

We demonstrate the utility of our approach through experiments in three areas: search engine operation, treatments to diabetes patients, and market value estimation of houses. Our results demonstrate that our approach can reduce the number and size of the currently performed randomized controlled experiments, thus saving significant time, money and effort on the part of practitioners.

1. INTRODUCTION

Novel medical treatments, new government policies, and innovative website designs are all examples of changes to an existing method of interaction with people that need to be evaluated for their effectiveness before they can be put into

use. The gold standard for testing such interventions are randomized controlled trials (RCTs) [6]. RCTs are widely used in medicine: Approximately 200,000 RCTs were conducted in the 1990's alone [6]. Internet website operators were early adopters of RCTs [9]. Most large Internet companies are known to run thousands of RCTs every year [8].

RCTs work by randomly assigning every subject to either a control group or a treatment group. The average measurement of the result variable for each group is then an unbiased estimator of its corresponding population mean. Given these, unbiased estimators of other desirable quantities such as the mean treatment effect can be easily constructed.

This approach, while appealing, has several drawbacks. First, for the estimators to be truly unbiased, subjects must be sampled i.i.d. from the general population of interest. Not only is this unrealistic and seldom the case, but often times the sample represents a very specific sub-population, which can lead to extremely biased estimates. This is especially evident in clinical trials, where subjects (who typically volunteer to take part in an experiment) are often those suffering from severe symptoms, which no other treatment helped, or are simply more prone to volunteer.

Second, as controlled trials are expensive and time consuming, sample sizes tend to be small. This greatly limits the amount of information available to researchers and practitioners for drawing conclusions, generating predictions, and deciding on policies. The small samples are typically sufficient for constructing estimators with reasonably low variance, but are seldom enough for generating predictors for individual outcomes. For instance, in search engine A/B tests, the decision of whether to use an alternative results ranker (or even whether to continue running the experiment) is often based on the average measures of click-through rate (CTR) or similar measures, and not on predictions regarding specific queries. Larger samples should potentially allow for the application of high-end learning algorithms.

Third, the price paid for guaranteeing that the estimators are unbiased is that only data from the controlled trial can be used. This completely discards the huge quantities of data that are often times available for the control condition, which in most cases is just the current policy. For instance, consider the case of predicting whether administering a new drug would prove better than the current standard for a given patient. A predictor trained only on the results of a small-scale clinical trial should prove to be inferior to one which also takes into account all the past medical records corresponding to the currently applied drug. Using only the

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License.
WWW 2017 Companion, April 3–7, 2017, Perth, Australia.
ACM 978-1-4503-4914-7/17/04.
<http://dx.doi.org/10.1145/3041021.3054190>



trial results seems wasteful in terms of data, and this is in general suboptimal for prediction.

Given the above, the question we pose here is the following: how can we design a learning algorithm for generating predictors in counterfactual settings, which takes as input both a small randomized trial dataset, and a large labeled dataset of the control population? Answering the above question is the motivation behind this paper.

It is important to note that although the setting we discuss is of a counterfactual nature, our focus here is on *prediction*. In pursuing the goal of generating high-accuracy predictors, we knowingly forfeit the ability to explain the underlying causal mechanism. The latter objective has been the focus of an abundant body of works, most based on the framework of *causal inference* [13]. We argue here that there is an inherent tradeoff between interpretability and predictive performance, and that when the goal is to optimize accuracy, a direct approach is preferred. Our work follows the more recent line of work where a discriminative loss-centric approach is applied in counterfactual settings [17, 18, 7].

The paper is organized in the following manner. We begin by covering related material in Sec. 2. We present notations and our problem statement in Secs. 3 and 4, respectively. Sec. 5 contains a detailed description of the core of our approach, followed by Sec. 6 in which several extensions are presented. Sec. 7 contains several experiments on real data. We conclude with a discussion in Sec. 8.

2. RELATED MATERIAL

Our setting draws relations to several lines of work. The fundamental property of the prediction task we consider is that it is counterfactual in nature. Causal inference [13] is a standard framework for estimating the causal relation between variables, in a way which can then be used to answer counterfactual questions. In order to achieve this, methods for causal inference are usually based on simple, interpretable models from which actionable conclusions can be drawn [4]. Our approach is different in that it focuses on prediction by introducing an ad-hoc loss function for parametrized predictors. Classic causal-inference models on the other hand do not always allow for arbitrary features, nor is it always straightforward to learn or to generate predictions from a given model.

Alternatively, counterfactual questions can be answered if data can be collected under a random policy [11, 4, 10]. In practice, true randomization is hard to obtain for business reasons (in Internet settings) or for ethical reasons (in medical and social domains). Due of this, it has been proposed to treat data collected under different settings as randomized, and use it to answer counterfactual questions [12].

More recently, notions from causal inference have been incorporated into discriminative learning methods, with the declared goal of minimizing loss. In analogy to Empirical Risk Minimization, the principle of Counterfactual Risk Minimization is proposed in [17, 18]. The suggested method offers a discriminative learning objective based on inverse propensity scores, where the variance is controlled by a regularization term or by self-normalization. In contrast to our setup, this method requires that in addition to examples x and labels y , each sample must also include its logged propensity score. Other works use doubly-robust methods which are based on propensity scores as well [1]. Some parametric

non-linear methods for estimating treatment effect are based on Bayesian regression trees [5] and random forests [21].

A parallel discriminative approach to counterfactual prediction is based on the notion of domain adaptation [2]. Following the work of [15], the authors of [7] observe that generalizing from the observed factual distribution to the unobserved counterfactual distribution is a special case of covariance shift, and in general of domain adaptation. Therefore, the non-convex representation learning method in [7] incorporates a discrepancy-based regularization term which encourages a label-invariant representation. In contrast, our method regularizes the relation between the control and treatment variables themselves, conditioned on a the original shared representation. Moreover, while the method of [7] requires large amounts of data for both labels, our method is tailored for a setting where the treatment variable is rare.

3. NOTATIONS

Our setup is similar to a standard supervised learning setup where we are given a sample set of examples x and labels y , but with some additions. We assume examples are from a general domain \mathcal{X} , and denote by $\mathcal{X}' \subseteq \mathcal{X}$ the sub-domain of examples that take part in the controlled trials. Our setup includes two label domains, denoted by \mathcal{Y}_C for the control variable and by \mathcal{Y}_T for the treatment variable. Throughout the paper we use the terms label, variable, and experimental outcome interchangeably.

Instantiations of examples are denoted by $x \in \mathcal{X}$, and of labels are denoted by $y_C \in \mathcal{Y}_C$ and $y_T \in \mathcal{Y}_T$. We assume there exists a single governing joint distribution $D_{\mathcal{X}, \mathcal{Y}_C, \mathcal{Y}_T}$ for tuples (x, y_C, y_T) , though we have no direct access to it, nor do we observe such tuples. Rather, for a given example $x \sim D_{\mathcal{X}}$ drawn from the marginal distribution, we observe either the control variable $y_C \sim D_{\mathcal{Y}_C | \mathcal{X}=x}$ or the treatment variable $y_T \sim D_{\mathcal{Y}_T | \mathcal{X}=x}$, drawn from their respective conditional distributions.

In the setting we consider, we are given as input three sample sets of example-label pairs:

1. $S_C = \left\{ (x^{(i)}, y_C^{(i)}) \right\}_{i=1}^{M_C}$ sampled i.i.d. from $D_{\mathcal{X}, \mathcal{Y}_C}$
2. $S'_C = \left\{ (x^{(i)}, y_C^{(i)}) \right\}_{i=1}^{M'_C}$ sampled i.i.d. from $D_{\mathcal{X}', \mathcal{Y}_C}$
3. $S'_T = \left\{ (x^{(i)}, y_T^{(i)}) \right\}_{i=1}^{M'_T}$ sampled i.i.d. from $D_{\mathcal{X}', \mathcal{Y}_T}$

The first set S_C is a large sample of the general population with labels for the control variable, representing past data accumulated by running the default policy. The sets S'_C and S'_T are smaller and represent the results of the controlled trial for the control and treatment groups.¹ We therefore assume that $M'_C, M'_T \ll M_C$ and that $M'_C \approx M'_T$. More concretely, we assume that while M_C is sufficiently large to learn a reasonably accurate predictor for the control variable, M'_T is insufficiently small for adequate learning of the treatment variable. Note that we are not guaranteed to have any x for which we observe both y_C and y_T . This is a fundamental problem in counterfactual settings, and makes estimating the individual treatment effect $y_{\Delta} = y_T - y_C$ especially challenging [23].

¹Note that the i.i.d. assumption mimics the procedure of random subject assignment found in RCTs.




Figure 1: A graphical depiction of the problem setting. Our goal is to predict treatment outcomes for a general-population test sample S_T using: (1) a small randomized controlled trial from a (possibly biased) sub-population X' , where either of the two possible treatments are randomly given, creating two subsets, (S'_C, S'_T) , and (2) a large historical dataset of control outcomes S_C . Our method links the two datasets using a minimal set of modeling assumptions (3).

4. PROBLEM STATEMENT

Recall that our goal is to construct a framework for learning predictors by leveraging both the small randomized trial data (S'_C, S'_T) and the large control-labeled dataset S_C . The main task we consider is predicting the treatment variable y_T for new, unobserved examples from a test set $S_T \sim D_{X, Y_T}$. In other words, we'd like our predictor to generalize well to \mathcal{Y}_T on the general population. The challenge here is that our data contains only a small number of treatment labels. The solution we present in Sec. 5 utilizes all the available data by modeling the relation between the control and treatment variable. An illustration of our setup is given in Fig. 1.

A related task that is of high interest is to predict the individualized treatment effect $y_\Delta = y_T - y_C$ [14, 20]. Accurate predictions of y_Δ can in principle aid decision makers in deciding what treatment to apply. Such predictions can also be used to estimate the mean treatment effect $\mathbb{E}[y_\Delta]$, and by so offer an alternative to conventional estimators used in randomized trials. As we show in Sec. 5, the individualized treatment effect y_Δ plays a central role in our learning objective for all tasks we consider.

5. METHOD

At the core of our method lie only two simple modeling assumptions: that predictions for both conditions are of the form $\hat{y} = \langle w, x \rangle$,² and that the models w_C, w_T for the control and treatment conditions, respectively, should be similar under some notion.

For ease of exposition, consider first a regression task where $x \in \mathbb{R}^d$, $y_C, y_T \in \mathbb{R}$, and our goal is to minimize the squared loss of a linear predictor for the treatment variable, namely $\hat{y}_T = \langle w_T, x \rangle$. In Sec. 6 we show that our method applies to both regression and classification, to a wide array of loss functions, and to some non-linear predictors as well.

²In general we assume that predictors are linear in *some* feature representation $\varphi(x)$, as we describe in Sec. 6.3.

Since our task is to predict the treatment outcome y_T of a given sample x , a reasonable place to start would be in considering a learning objective over the sample set S'_T , as it is the only one for which we have treatment labels. Applying the squared loss and adding ℓ_2 regularization gives us:

$$\min_{w_T \in \mathbb{R}^d} \frac{1}{M'_T} \sum_{i \in S'_T} \left(\langle w_T, x^{(i)} \rangle - y_T^{(i)} \right)^2 + \lambda \|w_T\|_2^2 \quad (1)$$

As in any discriminative objective, the number of samples greatly affects the quality of generalization of the learned predictor. Unfortunately, for the above objective and under our assumptions, S'_T will not prove to be sufficiently large for training a high-accuracy predictor for the treatment variable. Put simply, our data does not include enough labeled instances from \mathcal{Y}_T .

Our approach remedies this deficiency by artificially augmenting S'_T with samples that serve as a proxy for treatment labels. As a first step, we will add to the objective in Eq. (1) the samples from S_C , our largest available dataset:

$$\min_{w_T \in \mathbb{R}^d} \frac{\gamma}{M'_T} \sum_{i \in S'_T} \left(\langle w_T, x^{(i)} \rangle - y_T^{(i)} \right)^2 + \frac{(1-\gamma)}{M_C} \sum_{i \in S_C} \left(\langle w_T, x^{(i)} \rangle - y_C^{(i)} \right)^2 + \lambda \|w_T\|_2^2 \quad (2)$$

where $\gamma \in [0, 1]$ controls the relative weight of each dataset in the training objective. For ease of notation, we overload S_C to include all of the available control condition examples, namely assume $S'_C \subset S_C$.

At a first glance using control outcomes y_C when trying to predict the treatment outcome y_T may seem peculiar. Nonetheless, work in multi-task learning has shown that training a single predictor over several labels is beneficial in practice when the conditional distribution of different labels is similar [3]. However, even if the control and treatment distributions do share similarities, our focus here is rather on their differences. We therefore do not suffice with Eq. (2), and in place of the control labels $y_C^{(i)}$ we insert proxy treatment labels $\tilde{y}_T^{(i)}$ which we define next.

Denote by $\Delta = y_C - y_T$ the negative of the individual treatment effect y_Δ , namely the difference between the control and treatment variables. We have already set \hat{y}_T to be a linear function of x with weights w_T ; extending this to \hat{y}_C with weights w_C gives us a linear estimator for Δ :

$$\hat{\Delta} = \hat{y}_C - \hat{y}_T = \langle w_C, x \rangle - \langle w_T, x \rangle = \langle w_\Delta, x \rangle \quad (3)$$

where $w_\Delta = w_C - w_T$. Replacing \hat{y}_C with the true labels y_C from S_C and rearranging gives us our proxy:

$$\tilde{y}_T = y_C - \hat{\Delta} = y_C - \langle w_\Delta, x \rangle \quad (4)$$

Note that this derivation is possible due to our view of the tuple (x, y_T, y_C) as jointly distributed. This is in contrast to the more conventional approach where the distribution is modeled using tuples of the form (x, ν, y_ν) , where $\nu \in \{C, T\}$ is the experimental condition and y_ν is the outcome under that condition [7]. Our formulation induces a joint distribution over pairs $(x, \Delta) \sim D_{X, \Delta}$, which we utilize.

Plugging back into Eq. (2) and further regularizing gives:

$$\begin{aligned} \min_{w_T, w_\Delta} \frac{\gamma}{M'_T} \sum_{i \in S'_T} \left(\langle w_T, x^{(i)} \rangle - y_T^{(i)} \right)^2 + \\ \frac{(1-\gamma)}{M_C} \sum_{i \in S_C} \left(\langle w_T, x^{(i)} \rangle - \tilde{y}_T^{(i)} \right)^2 + \\ \lambda \|w_T\|_2^2 + \eta R(w_\Delta) \end{aligned} \quad (5)$$

where R is a regularization function, and $\eta \in \mathbb{R}$ will be described shortly. Note that \tilde{y}_T is in fact a function of w_Δ .

To gain insight into the above construction, we next analyze the learning objective under an alternative formulation. Notice that by Eqs. (3) and (4), the second loss term and the additional regularization term in Eq. (5) are equivalently:³

$$\sum_{i \in S_C} \left(\langle w_C, x^{(i)} \rangle - y_C^{(i)} \right)^2, \quad \eta R(w_T - w_C)$$

Under this representation, the choice of R and η respectively determine the nature and magnitude of similarity between w_T and w_C . For instance, setting $R = \|\cdot\|_2^2$ will encourage w_T and w_C to be close under a Euclidian metric, while setting $R = \|\cdot\|_1$ will induce sparsity on w_Δ , meaning that w_T and w_C will be different only on a small subset of entries.

This gives an intuitive interpretation of our assumption on the similarity of w_C and w_T via w_Δ ; we assume that w_C models the baseline effect, while w_Δ models the deviation of the treatment effect as expressed by w_T . This aligns well with the problem setup. Since S_C is large, it should allow for a good fit to the baseline effect of the control condition. Given this, the fewer samples in S'_T should now suffice to fit the deviated treatment effect. This is especially true for high-dimensional data, where learning requires a large number of samples. We will return to this in Sec. 6.3;

The value of η sets the de-facto linkage strength of the two loss terms in Eq. (5). Setting $\eta = 0$ will allow w_C to be arbitrarily far away from w_T , which will lead to a disjoint objective - minimizing w_T over S'_T and w_C over S_C independently. On the other hand, setting $\eta = \infty$ will constrain $w_T = w_C$ and hence revert the objective back to Eq. (2).

While η controls the relation between w_T and w_C , γ signifies the importance of each sample set for training w_T to generalize well to the treatment variable. While S'_T contains actual treatment labels but is small, S_C is sufficiently large but contains only control labels (used as proxies for the treatment variables). The purpose of γ is therefore to allow us to balance these complementary properties. Setting $\gamma = 1$ will revert the objective back to Eq. (1), while setting $\gamma = 0$ will result in a training objective based only on S_C . In effect, the above notions model our belief in how, and how well, \tilde{y}_T serves as a proxy for y_T .

6. EXTENSIONS

In the above section, we presented our method for a regression task under a squared loss function and an ℓ_2 regularization term. Note however that our only modeling assumption was that both \hat{y}_T and \hat{y}_C (and accordingly $\hat{\Delta}$) admitted to a linear form under some joint feature representation. This simple assumption allows us to apply our method to more

³This is similar to the regularization term of the Fused Lasso approach [19] used for time-series prediction.

general learning settings and to non-linear predictors. For some cases, it also provides a closed-form solution.

6.1 Classification

In our analysis thus far we have assumed a specific ridge regression setting, with a squared loss and labels in \mathbb{R} . Our results however are not restricted to regression, and in fact apply to general loss functions $L(\langle w, x \rangle, y)$ over linear predictors, and to general regularization terms $Q(w)$ of the predictor's parameters. The general form of Eq. (5) becomes:

$$\begin{aligned} \min_{w_T, w_\Delta} \frac{\gamma}{M'_T} \sum_{i \in S'_T} L(\langle w_T, x^{(i)} \rangle, y_T^{(i)}) + \\ \frac{(1-\gamma)}{M_C} \sum_{i \in S_C} L(\langle w_T - w_\Delta, x^{(i)} \rangle, y_C^{(i)}) + \\ \lambda Q(w_T) + \eta R(w_\Delta) \end{aligned} \quad (6)$$

and applies to classification as well as to regression. For losses used in margin-based binary classification methods, y_Δ signifies the difference in distances to the margin, rather than the difference in the actual outcome. For an appropriate definition of $y_\Delta = y_T - y_C$, Eq. (6) can also be applied in principle to multi-class and multi-label classification and to structured prediction. However, note that in such classification settings, the interpretation of y_Δ as the individual treatment effect no longer holds. For other tasks the role of the regularization term R may also change.

6.2 Closed form solution

When applying our method to ridge regression (as in Sec. 5), setting $R(\cdot) = \|\cdot\|_2^2$ allows for a closed form solution of the objective in Eq. (5). This is accomplished by transforming the objective into a canonical ridge regression form:

$$\min_w \|w^\top X - Y\|_2^2 + \alpha \|w\|_2^2 \quad (7)$$

for which the solution is:

$$\hat{w} = (X^\top X + \alpha I)^{-1} X^\top Y \quad (8)$$

We now show how to construct the data matrix X , label vector Y , and regularization constant α , so that the minimizer of Eq. (5) can be extracted from \hat{w} .

Since the objective in Eq. (5) includes the minimization over both w_T and w_Δ , we first set w to be their concatenation, namely $w = (w_T, w_\Delta) \in \mathbb{R}^{2d}$. Under this expanded representation, we next set:

$$\begin{aligned} i \in S'_T: \quad X_i. &= c_1 \cdot (x^{(i)}, \mathbf{0}) & Y_i &= c_1 y_T^{(i)} \\ i \in S_C: \quad X_i. &= c_2 \cdot (x^{(i)}, -c_3 x^{(i)}), & Y_i &= c_2 y_C^{(i)} \end{aligned} \quad (9)$$

where $\mathbf{0}$ is a vector of zeros of size d , and the constants are:

$$c_1 = \sqrt{\gamma/M'_T}, \quad c_2 = \sqrt{(1-\gamma)/M_C}, \quad c_3 = \sqrt{\lambda/\eta}$$

Finally, letting $\alpha = \lambda$ and plugging into Eq. (8) gives the solution for w_T and w_Δ of our original objective in Eq. (5).

6.3 Non-linear predictors

While linear predictors are easy to work with and often work well in practice, they lack the expressive power that non-linear predictors offer. As our method is not constrained to a specific representation, a straightforward way for incorporating non-linearity is via kernels, as we describe next.

Task	Measure	S'_T	S_C	$S_{T' \cup C}$	Δ	S_T
Stay length	Mean r^2	0.153	0.068	0.167	0.219	0.323
	% bench.	47%	21%	52%	68%	100%
Above median	Accuracy	0.711	0.709	0.711	0.725	0.749
	% bench.	95%	95%	95%	97%	100%

Table 1: Results of the prediction and classification tasks on the diabetes treatment dataset. The proposed method (Δ) reaches the highest accuracy, compared to methods which use subsets of the available data.

The construction in Eq. (9) shows how the regularization of w_Δ can be achieved by a simple expansion of the feature representation. In cases where $R = Q$ decomposes, a similar procedure can be applied to a more general settings. When $\gamma = 1/2$, setting the expanded features $\phi(x) = (x, \mathbf{0})$ for $x \in S'_T$ and $\phi(x) = (x, -cx)$ for $x \in S_C$ with $c = \sqrt{\lambda/\eta}$ allows for R and Q to share a single constant λ , and due to decomposability define a single regularization function over the new expanded model $\tilde{w} \in \mathbb{R}^{2d}$.

Since the above holds for any feature representation $\varphi(x)$, kernel-supporting methods can be readily applied. For a linear kernel $K(x, x') = \langle x, x' \rangle$, the expanded kernel \bar{K} is:

$$\begin{aligned} \bar{K}(x, x') &= \langle \phi(x), \phi(x') \rangle = g(x, x') \cdot K(x, x'), \\ g(x, x') &= \begin{cases} c^2 & x, x' \in S_C \\ 1 & \text{o.w.} \end{cases} \end{aligned} \quad (10)$$

For a general feature representation $\varphi(x)$, the closure of kernels under addition gives us:

$$\bar{K}(x, x') = g(x, x') \cdot \langle \varphi(x), \varphi(x') \rangle \quad (11)$$

Hence, our method can be applied to wide class of regularized kernel methods. General values of γ can be incorporated into losses which support differential sample weights, such as kernel ridge regression, SVMs, SVRs, and others.

Due to their high dimensionality, kernel methods, as well as many other non-linear predictors, require a considerable number of samples to learn properly. This makes using kernels only on the small S'_T unrealistic, while applying them to S_C does not optimize for the treatment condition. As mentioned in Sec. 5, our method utilizes the large number of samples in S_C to learn the baseline effect of the control condition. This makes learning with kernels feasible, while still taking advantage of the treatment samples in S'_T .

Finally, we note that since many non-linear deep architectures include a linear output layer, our method can potentially be applied to such. In a similar fashion to the construction in Sec. 6.2, such an architecture should include two linear output layers - one for w_T and one for w_Δ - and corresponding regularization terms. We leave the exploration of such an approach for future work.

7. EXPERIMENTS

In this section we evaluate the performance of our method on three counterfactual prediction tasks: A simulated medical clinical trial, a web search engine experiment, and a social choice question. Since our learning goals include predictions regarding the treatment variable y_T , our data must

Task	Measure	S'_T	S_C	$S_{T' \cup C}$	Δ	S_T
Value	Mean r^2	0.564	0.651	0.660	0.688	0.716
	% bench.	79%	91%	92%	96%	100%
Top decile	Accuracy	0.849	0.831	0.853	0.861	0.875
	% bench.	97%	95%	97%	98%	100%

Table 2: Results of the prediction and classification tasks on the housing dataset. The proposed method (Δ) reaches the highest accuracy and 96% or more of the benchmark (S_T), which uses data from the entire treatment dataset.

contain a large pool of ground-truth labels for this class. This is a necessary condition for ensuring a sound evaluation procedure. Unfortunately, for the same reasons that motivate our work, most datasets do not include many labeled treatment instances, as they are typically hard, expensive, and time consuming to acquire.

To this end, we focus on three datasets. The first dataset contains information on the clinical status of approximately 100,000 diabetes patients. Our task is to predict the length of hospitalization for each patient, given their treatments so far. The second dataset comprises of a large collection of around 20,000 houses along with their attributes. Our task is to estimate the market price of a house given its attributes. While the datasets themselves were not collected by a randomized trial procedure, we partition the records into control and treatment conditions in a way which emulates a realistic controlled trial scenario. This allows us to validate our predictions on the treatment variable.

The third dataset is from the domain of search engine operation. Search engines regularly modify and improve their ranking algorithm and other parameters such as the user interface. In many cases decisions are based on the results of a large number of A/B tests, where the current ranking algorithm is compared to a new alternative. In this setting, early and accurate predictions of query-centric measures like the *click-through rate* (CTR) and its derivatives are of great importance. Thus, we focus on this prediction, which can also assist in early termination of treatments which are predicted to be as good as (or worse than) the current treatment.

At its core, our method provides a way to model the linkage between a small randomized trial and a large historical dataset. Our goal in this section is therefore to evaluate the added benefit of using our model when such data is available. In Sec. 6, we describe why and how our method can be applied to a large set of loss functions and predictors. It therefore makes sense to compare the performance of our method to methods which use the same loss and predictor class, while considering different unlinked combinations of the available datasets. Specifically, we compare our method (denoted as Δ in the tables) to training only on S'_T , only on S_C , and on the union of both sets $S_{T' \cup C} = S'_T \cup S_C$. As for other linear methods described in Sec. 2, [17, 18] assume that the data includes loss terms and propensity scores, while the linear method in [7] does not outperform standard ridge regression.

We evaluated performance on two tasks: predicting **individual treatment outcomes** (y_T), and predicting the **individual treatment effect** (y_Δ). As mentioned, all methods were evaluated on a held-out test set which included treatment labels, where results were averaged over 10 ran-

Task	Measure	25:75 split					75:25 split				
		S'_T	S_C	$S_{T' \cup C}$	Δ	S_T	S'_T	S_C	$S_{T' \cup C}$	Δ	S_T
Individual treatment outcome y_T	Mean r^2	-0.04	0.19	0.21	0.26	0.30	0.26	0.22	0.24	0.33	0.36
	% bench.	-	65%	72%	89%	100%	72%	63%	68%	94%	100%
Individual treatment effect y_Δ	Mean r^2	-0.08	0.18	0.20	0.24	0.26	0.22	0.22	0.23	0.31	0.32
	% bench.	-	70%	76%	90%	100%	69%	67%	72%	95%	100%
Average effect $\mathbb{E}[y_\Delta]$	Abs. diff.	0.23	0.27	0.24	0.14	0.06	0.05	0.24	0.23	0.07	0.05

Table 3: Search engine ranking results for predicting individualized treatment outcomes and effects (higher is better) and the average treatment effect for an A/B test (lower is better) using different sample training sets. The proposed method (Δ) links both samples by enforcing similarity. Results are averaged over all A/B tests. S_T is used as a high-end benchmark.

dom splits of the data. All of the results presented in Tables 1, 2, and 3 are significant with a p -value of well below 0.0001 under the Friedman test.⁴ For all tasks we used ridge regression as a learning objective, and applied ℓ_2 regularization for our method. Meta-parameters were chosen on a small held-out validation set. Our high-end benchmark for performance is based on learning over a large treatment-labeled training set S_T . As this type of data is typically hard to obtain, this benchmark serves as an empirical upper bound on the overall achievable accuracy.

7.1 Hospitalization of diabetes patients

The Diabetes dataset⁵ contains data from 10 years (1999-2008) of clinical care at 130 US hospitals and integrated delivery networks [16]. The 'new' treatment which we attempt to estimate is whether prescription of diabetes medications prior to hospitalization would have changed the hospitalization length. We focused on patients for which the reason of admission was unknown, as these represent the difficult cases, of which there were 4,785 patients in the data. We simulated a clinical trial by randomly selecting 25% of the population into \mathcal{X}' , some of whom were prescribed diabetes medications, and some who were not, and used a 25:75 split.

We evaluate performance on two tasks: predicting the hospitalization length, and predicting whether the length would be above the median. The results of these experiments are shown in Table 1. As the results show, our method provides better predictions of the treatment effect and outcome, compared to methods which are based on subsets of available data. Moreover, this prediction is close in its quality to that achieved by a learner which uses the actual treatment information.

7.2 House pricing Dataset

The House Sales in King County dataset⁶ contains records of 21,613 houses sold in King County, USA, a region which includes Seattle. Along with the market price of each house, the data includes 19 numerical and categorical attributes for each house including the number and types of rooms, size, number of floors, and geographic location. Of special interest is an attribute which determines whether the house was renovated or not. By considering this as a treatment

indicator variable and partitioning accordingly, we can simulate the following counterfactual question: Does renovating increase a house's value, and if so, by how much?

As houses were not randomly assigned to each condition, the data does not represent a true randomized controlled trial. This of course raises questions as to whether predictions can be used to answer the above question. Nonetheless, our method still applies here, as we do not assume a random assignment, but rather use it as motivation.

We evaluated performance on two tasks: a regression task in which we predict the value of a house, and a classification task in which we predict whether the value is in the top decile. In both tasks we assign all houses which did not undergo renovation to the control condition, and all those which did to the treatment condition. These amounted to 20,699 and 914 houses, respectively. We used 75:25 train-test splits, and set the experimental sub-population \mathcal{X}' to be all houses in a random subset of zip codes, representing about 25% of all zip codes. This is similar to a setting where only certain areas residential areas participate in a survey.

Results for all tasks are presented in Table 2. We report the mean R^2 for the task of predicting a house's value, and mean accuracy for predicting the attribution to the top decile. Results show that the highest accuracy in both prediction tasks is obtained by the proposed method. Moreover, these accuracies are within a few percentage points from the accuracy obtained when using data from benchmark dataset. Thus, the proposed method can replace the use of a large RCT, which would be expensive and difficult to execute, in this kind of setting.

7.3 Search Engine Dataset

We collected queries submitted to the Bing search engine on June 1st 2016 which were randomly assigned to internal A/B tests, and whose frequency was at least 1,000. As examples x we took all of the queries that appeared in the control condition, and focused on tests which included all of these queries. Our dataset contains 277 comparable treatment conditions and one control condition, each with 1,572 distinct query examples, for a total of 437,016 instances.

Features included both categories of the queries and features of the query words. Categorization was determined using a proprietary classifier [24] developed by the Microsoft Bing team to assign each query into a set of 63 categories, including, for example, commerce, tourism, video games, weather-related, and adult-themed queries. The classifier is used by Bing to determine whether to display special re-

⁴Save for the classification task in Table 1 where $p = 0.009$.

⁵<https://archive.ics.uci.edu/ml/datasets/Diabetes+130-US+hospitals+for+years+1999-2008>

⁶<https://www.kaggle.com/harlfoxem/housesalesprediction>

sults such as instant answers. Queries can be classified into multiple categories (e.g., purchase of flight tickets would be classified into both tourism and commerce). Word-based features included some basic attributes such as the number of characters, number of tokens, minimal and maximal token size, and a numeric token indicator. In addition, a bag-of-words representation of the tokens was computed, and applied using the feature hashing trick [22].

To generate labels, for each query and for each experimental condition, we computed CTR estimates by pooling all relevant query instances. Since the distribution of CTR is highly skewed, and since some queries have CTR=0, we set label values to $y = \log(\text{CTR} + \epsilon)$ for $\epsilon = 10^{-10}$. The above process ensured that for a given x our data included both y_C and y_T , from which y_Δ was computed.

A distinct characteristic of this dataset is that for most search phrases, the data includes both y_C (the CTR under the default ranker) and y_T (the CTR under the new ranker). This is because, for common search phrases, responses will be recorded under both control and treatment conditions. The above allows us to directly evaluate the individual treatment effect y_Δ . Moreover, since the data includes a large collection of alternative rankers (but only one default ranker), we can compare the effect of many treatments to the same control condition.

Running search engine A/B tests is an expensive procedure. Tests are therefore often limited in time and resources, and are allocated only a small fraction of the overall traffic. This causes CTR estimates to be unrepresentative, as high-frequency queries will be assigned to an A/B test more often than low-frequency queries, which may not appear in trials at all. Therefore, our proposed algorithm can potentially shorten A/B test by reaching a conclusion as to the benefit of a new ranking algorithm using only the popular queries (which are easy to collect), but inferring the benefit for rare queries as well, solely based on a complete historical record of the control condition and a small random trial over a non-representative population (e.g., popular queries).

Since our estimation procedure requires a full set of ground-truth treatment labels, we can use only query instances which participated in trials. This requires us to mimic the above setting using A/B test data alone, by constraining S'_T to contain only queries whose frequency is in the top τ^{th} -quantile. For instance, by setting $\tau = 0.75$, we guarantee that S'_T contains only queries with frequency in the top quartile. To keep S'_T small, we further discard a random 25% of the qualified examples.

For comparison of the proposed method we use the results of learning on S_T with $\tau = 0.75$. This means learning on most of the available data, necessitating a long A/B test to collect the rarer queries. We refer to this learning task as the benchmark. Results for all prediction tasks appear in Table 3. For individualized treatment outcomes and effects we report the mean R^2 and its fraction of the benchmark. For average treatment effect we report the absolute difference from the true effect, which we estimate on a large distinct dataset. As can be seen, our method significantly outperforms learning using the available subsets of the data by a significant margin. The accuracy of the proposed algorithm is not far from that of the benchmark, reaching approximately 90% or more of the potential accuracy for both tasks.

To explore the effect of trial duration (represented here by thresholding frequency), we repeated the above procedure




Figure 2: Accuracy of the regression algorithm as a function of the fraction of queries used for training, by datasets used for training. As the fraction of queries increases (τ decreases), S'_T includes more queries. This mimics a setting where the length of an A/B test determines the frequency threshold of observed items. The proposed method (Δ) quickly reaches performance close to that of the benchmark, which uses the treatment information for all queries.

for various values of τ . To accentuate results, we focus on the the top 10% of trials for which the difference between conditions was a-priori most significant, and employed a 50:50 train-test ratio. Results are presented in Fig. 2. Our method enjoys a fast growth rate in accuracy, and should potentially allow for shorter trial lengths, or for early stopping of trials, when the new treatment is deemed to be inferior to existing treatments.

8. DISCUSSION

Randomized controlled trials (RCTs) are the gold standard for testing new treatments and interventions. RCTs are widely used by Internet websites, by medical authorities, and, increasingly, by governments. However, RCTs are difficult and expensive to run. Nowadays, historical data is available in many settings where RCTs are considered. However, as these data were collected using an existing policy, utilizing these data has proven difficult.

In this paper we proposed a new algorithm for using historical data in conjunction with the results of small RCTs, to counterfactually infer the outcomes of large RCTs. Our method can additionally be used as an early stopping criterion for RCTs, when the method predicts that the benefit of a new treatment will not be larger than those of the existing treatment. Thus, our method can provide benefit to existing RCTs.

The proposed method is based on two assumptions: The first is that the outcome of each treatment can be predicted using a linear predictor. The second is that the difference between the predictor of the current treatment and the predictor of the proposed treatment is not large. In Sec. 6.3 we showed extensions to the method which overcome the first assumption. We hypothesize that predictions can be improved by using robust regression, or through inclusion of a confidence measure for each point in our data. Such extensions are left for future work.

9. REFERENCES

- [1] H. Bang and J. M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.
- [2] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- [3] S. Bickel, J. Bogojeska, T. Lengauer, and T. Scheffer. Multi-task learning for HIV therapy screening. In *Proceedings of the 25th international conference on Machine learning*, pages 56–63. ACM, 2008.
- [4] L. Bottou, J. Peters, J. Q. Candela, D. X. Charles, M. Chickering, E. Portugaly, D. Ray, P. Y. Simard, and E. Snelson. Counterfactual reasoning and learning systems: the example of computational advertising. *Journal of Machine Learning Research*, 14(1):3207–3260, 2013.
- [5] H. A. Chipman, E. I. George, and R. E. McCulloch. BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, pages 266–298, 2010.
- [6] L. Haynes, O. Service, B. Goldacre, and D. Torgerson. Test, learn, adapt: Developing public policy with randomised controlled trials. Technical report, UK, 2012.
- [7] F. D. Johansson, U. Shalit, and D. Sontag. Learning representations for counterfactual inference. In *Proceedings of the 33rd International Conference on Machine Learning, ICML*, pages 3020–3029, 2016.
- [8] R. Kohavi, A. Deng, B. Frasca, R. Longbotham, T. Walker, and Y. Xu. Trustworthy online controlled experiments: Five puzzling outcomes explained. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 786–794. ACM, 2012.
- [9] R. Kohavi, R. Longbotham, D. Sommerfield, and R. M. Henne. Controlled experiments on the web: survey and practical guide. *Data mining and knowledge discovery*, 18(1):140–181, 2009.
- [10] L. Li, S. Chen, J. Kleban, and A. Gupta. Counterfactual estimation and optimization of click metrics in search engines: A case study. In *Proceedings of the 24th International Conference on World Wide Web*, pages 929–934, 2015.
- [11] L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670. ACM, 2010.
- [12] L. Li, J. Y. Kim, and I. Zitouni. Toward predicting the outcome of an A/B experiment for search relevance. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pages 37–46. ACM, 2015.
- [13] J. Pearl. *Causality*. Cambridge university press, 2009.
- [14] D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- [15] B. Schölkopf, D. Janzing, J. Peters, E. Sgouritsa, K. Zhang, and J. M. Mooij. On causal and anticausal learning. In *Proceedings of the 29th International Conference on Machine Learning, ICML*, 2012.
- [16] B. Strack, J. P. DeShazo, C. Gennings, J. L. Olmo, S. Ventura, K. J. Cios, and J. N. Clore. Impact of HbA1c measurement on hospital readmission rates: analysis of 70,000 clinical database patient records. *BioMed research international*, 2014.
- [17] A. Swaminathan and T. Joachims. Counterfactual risk minimization. In *Proceedings of the 24th International Conference on World Wide Web*, pages 939–941. ACM, 2015.
- [18] A. Swaminathan and T. Joachims. The self-normalized estimator for counterfactual learning. In *Advances in Neural Information Processing Systems*, pages 3231–3239, 2015.
- [19] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
- [20] M. J. van der Laan and M. L. Petersen. Causal effect models for realistic individualized treatment and intention to treat rules. *The International Journal of Biostatistics*, 3(1), 2007.
- [21] S. Wager and S. Athey. Estimation and inference of heterogeneous treatment effects using random forests. *arXiv preprint arXiv:1510.04342*, 2015.
- [22] K. Q. Weinberger, A. Dasgupta, J. Langford, J. Attenberg, and A. J. Smola. Feature hashing for large scale multitask learning. In *Proceedings of the 26th International Conference on Machine Learning (ICML-09)*, page 140, 2009.
- [23] J. Weiss, F. Kuusisto, K. Boyd, J. Liu, and D. Page. Machine learning for treatment assignment: Improving individualized risk attribution. In *AMIA Annual Symposium Proceedings*, volume 2015, page 1306. American Medical Informatics Association, 2015.
- [24] E. Yom-Tov, R. W. White, and E. Horvitz. Seeking insights about cycling mood disorders via anonymized search logs. *Journal of medical Internet research*, 16(2):e65, 2014.