

MALTS: Matching After Learning to Stretch

Harsh Parikh

*Department of Computer Science
Duke University
Durham, NC 27708-0129, USA.*

HARSH.PARIKH@DUKE.EDU

Cynthia Rudin

*Department of Computer Science
Duke University
Durham, NC 27708-0129, USA.*

CYNTHIA@CS.DUKE.EDU

Alexander Volfovsky

*Department of Statistical Science
Duke University
Durham, NC 27710, USA.*

ALEXANDER.VOLFOVSKY@DUKE.EDU

Editor: Russ Greiner

Abstract

We introduce a flexible framework that produces high-quality almost-exact matches for causal inference. Most prior work in matching uses ad-hoc distance metrics, often leading to poor quality matches, particularly when there are irrelevant covariates. In this work, we learn an interpretable distance metric for matching, which leads to substantially higher quality matches. The learned distance metric stretches the covariate space according to each covariate’s contribution to outcome prediction: this stretching means that mismatches on important covariates carry a larger penalty than mismatches on irrelevant covariates. Our ability to learn flexible distance metrics leads to matches that are interpretable and useful for the estimation of conditional average treatment effects.

Keywords: causal inference, matching, nearest neighbors, distance metric learning

1. Introduction

Matching methods are used throughout the social and health sciences to make causal conclusions where access to randomized trials is scarce but observational data are widely available. Matching methods construct sets of similar individuals, some of whom select into treatment and some of whom select into control, allowing for direct comparison of outcomes between the samples from these populations. These methods are particularly interpretable since they allow fine-grained troubleshooting of the data. For instance, examining a matched group of patients through chart review of their medical data and doctors’ notes may allow an analyst to determine whether the matched groups are indeed trustworthy, and if not, determine what other factors should be included in the analysis. Having high-quality matches also allows the user to estimate nonlinear treatment effects with lower bias than parametric approaches.

As a concrete example of the importance of match group quality, Table 1 presents a series of matched groups from the Lalonde dataset (LaLonde 1986, Dehejia and Wahba

1999). A simple visual inspection of the matched groups produced by standard-bearer methods like propensity score matching and prognostic score matching reveals that the units being considered similar by these methods are not similar on underlying covariates. On the other hand, the matches generated by our proposed method are qualitatively (and quantitatively) better. *The quality of the matches is our main consideration in this work.*

Table 1: Example control units in a matched group for a treated unit using (a) our approach (MALTS), (b) prognostic score (Hansen 2008), and (c) propensity score matching (Rosenbaum and Rubin 1983) for a query unit in the Lalonde dataset (top rows). Our method matched closely on covariates – age, education, whether the person had an academic degree, and income in 1975. In contrast, prognostic and propensity scores did not match closely on these factors.

Treatment		Covariates							Outcome	
<i>Unit ID</i>	Treated	Age	Education	Black	Hispanic	Married	No-Degree	Income-1975	Income-1978	
<i>Query: 1</i>	Yes	22		No	Yes	No	Yes	\$0	\$3596	
(a) Our Approach (MALTS)										
330	No	22	8	No	Yes	No	Yes	\$0	\$9921	
299	No	22	9	Yes	No	No	Yes	\$0	\$0	
416	No	22	9	Yes	No	No	Yes	\$0	\$12898	
(b) Prognostic Scores										
338	No	44	9	Yes	No	No	Yes	\$0	\$9722	
340	No	22	12	Yes	No	No	No	\$532	\$1333	
355	No	18	10	No	Yes	No	Yes	\$0	\$1859	
(c) Propensity Scores										
451	No	22	8	Yes	No	No	Yes	\$0	\$1391	
330	No	22	8	No	Yes	No	Yes	\$0	\$9921	
407	No	20	12	Yes	No	No	No	\$1371	\$20893	

Typically, matching methods place units that are close together into the same matched group, where closeness is measured in terms of a pre-defined distance (e.g., exact, coarsened exact, Euclidean, etc.), while maintaining balance constraints between treatment and control units. Despite its merits, this classical paradigm has flaws, namely that it relies heavily on a prespecified distance metric. The distance metric cannot be determined without an understanding of the importance of the variables; for instance, the quality of matches for any prespecified distance that weighs all covariates equally will degrade as the number of irrelevant covariates increases. This is true irrespective of the matching methodology employed. This issue has previously been referred to as the toenail problem (Wang et al. 2021, Dieng et al. 2019), where the inclusion of irrelevant covariates (like “toenail length”) with nonzero weights can worsen the metric for matching. A related concern is that the covariates may be scaled differently, where a given distance along one covariate has a different impact than the same distance along a different covariate; in this case, if the scaling or weights on the covariates are chosen poorly, the total distance metric can inadvertently be determined by less relevant covariates, again leading to lower quality matches.

Ideally, the distance metric would focus on important covariates that significantly contribute to the outcome, so that after matching, treatment effect estimates computed using the matched groups would be accurate. If the researcher knows how to choose the distance metric so that it yields accurate treatment effect estimates, it would solve the problem. However, there is no reason to believe that this is achievable in complex high-dimensional

data settings. Producing high dimensional functions to characterize data is a task at which humans are not naturally adept.

In this work, we propose a framework for matching where an interpretable distance measure between matched units is learned from a training set. As long as the distance metric generalizes from the training set to the full sample, we are able to compute high-quality matches and accurate estimates of conditional average treatment effects (CATEs) within the matched groups. One can use any form of distance metric to train, and in this work, we focus on exact matching for discrete variables and generalized Mahalanobis distances for continuous variables. By definition, the generalized Mahalanobis distance is determined by a matrix. If the matrix is diagonal, the distance calculation represents a stretch for each covariate. Irrelevant covariates will be compressed so that their values are always effectively zero. Highly relevant covariates will be stretched so that for two units to be considered a match, they must have very similar values for those covariates. In this way, diagonal matrices lead to very interpretable distance metrics. If the Mahalanobis distance matrix is not constrained to be diagonal, then it induces a stretch and rotation, leading to more flexible but less interpretable notions of distance.

The new framework is called Learning-to-Match, and the algorithm introduced in this work is called Matching After Learning to Stretch (MALTS). Figure 1 shows the main steps of MALTS, which are: divide the data into training and estimation sets, learn the distance metric on the training set, use the learned distance metric to perform nearest neighbor matching on estimation set, and use those matched groups to estimate conditional average treatment effects (CATEs). We tested MALTS against several other matching methods in simulation studies (Section 6), where ground truth CATEs are known. In these experiments, MALTS consistently achieves substantially better results than other matching methods including Genmatch, propensity score matching, and prognostic score matching for estimating CATEs. Even though our method is heavily constrained to produce interpretable matches, it performs at the same level as non-matching methods that are designed to fit extremely flexible but uninterpretable models directly to the response surface.

In Section 3, we introduce the learning-to-match framework and show that under a choice of smooth distance metric (Definition 1) we can estimate conditional average treatment effects accurately with high probability. Section 4 discusses MALTS’ optimization set up and training procedure that learns a smooth distance metric. In Section 5, we prove that the distance metric learned by MALTS is multi-robust (Definition 3) and generalizable (Definition 5). Thus, the distance metric estimated by MALTS’ framework facilitates the correct estimates of CATEs under SUTVA and positivity assumptions.

2. Related work

Since the 1970’s, the causal inference literature on matching methods has been concentrated on dimension reduction techniques (e.g., Rubin 1973a,b, 1976, Cochran and Rubin 1973). In this literature, the leading approach for dimension reduction uses the propensity score, which is the conditional probability of treatment given covariate information. Propensity score methods are designed for calculating average treatment effects (as opposed to conditional average treatment effects) and do not produce exact or almost-exact matches. When treatment is binary, they project data onto one dimension, and closeness of units in propen-

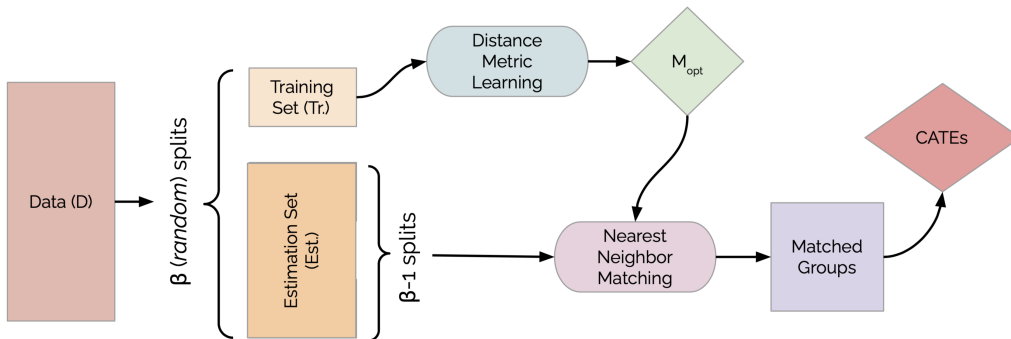


Figure 1: Schematic drawing of MALTS algorithm. The algorithm splits the data into random subsets and uses one of the subsets (training set) to learn a distance metric. It performs matching on the rest of the units (estimation set) using the learned distance metric to produce tightly matched groups and estimate conditional average treatment effects.

sity score does not imply their closeness in covariate space. As a result, the matches cannot directly be used for estimating heterogeneous treatment effects.

Other causal inference methods have been studied in the literature (Gu and Rosenbaum 1993, Imbens 2004), but almost all of them suffer from at least one of four possible problems: using a black box model that is uninterpretable (i.e., almost all machine learning methods), having a distance metric that is predefined (rather than learned), computational inefficiency, or not being applicable to CATE estimation (as we discussed with propensity scores). These issues cause the vast majority of matching methods to be ineffective in producing high quality interpretable CATE estimates. Regression methods can be used for CATE estimation, but only when the regression method is correctly specified – or in the case of doubly robust estimation (e.g., Farrell 2015), either the propensity model or the outcome model needs to be correctly specified. Machine learning approaches generalize regression approaches and can create models that are extremely flexible and predict outcomes accurately for both treatment and control groups (Hill 2011, Chernozhukov et al. 2018, Hahn et al. 2020). However, complicated regression methods lose the interpretability inherent to almost-exact matches and are difficult to troubleshoot and trust. In practice, MALTS performs similarly to (or better than) several machine learning methods in our experiments, despite being restricted to interpretable almost-exact matches with an interpretable distance metric.

A flexible setup for producing high-quality matches is provided by the optimal matching literature (Rosenbaum 2017). These are built on network flow algorithms and integer programming to produce matches that are constrained in user-defined ways (Zubizarreta 2012, Zubizarreta et al. 2014, Keele and Zubizarreta 2017, Resa and Zubizarreta 2016, Kallus 2017, Morucci et al. 2022). In all of these approaches, the user defines the distance metric (rather than learning it from data), potentially leading to poor quality matched groups. An alternative to optimal matching is coarsened exact matching (CEM, Iacus et al. 2012), an approach that requires users to specify explicit bins for all covariates on which to con-

struct matches. This requires users to know in advance that the outcomes are insensitive to movements within many high-dimensional bins, which is essentially equivalent to the user knowing the answer to the problem we investigate in this work. Large amounts of user choice to define these bins can also lead to unintentional user bias. By *learning* the stretching rather than asking the user to define it as in CEM, this bias is potentially reduced.

Zhao (2004) and Imbens (2004) discuss the choice of distance metric for matching. The approach by Zhao (2004) depends on the correlations between treatment choice, outcome and covariates. However, this approach assumes a model for the relationship between the outcome and covariates, or the treatment choice and covariates. Hence, under model misspecification, the estimator may not be consistent. MALTS learns a distance metric without any model assumptions.

The present work builds on work of Wang et al. (2021), Dieng et al. (2019) where a discrete distance metric is learned by considering the prediction quality of the covariate sets. That work does not pertain to continuous covariates, whereas ours does. There is substantial work on learning distance metrics (though not for causal inference, e.g., Goldberger et al. 2005, Weinberger et al. 2006, Weinberger and Saul 2009), where the goal is to learn a distance metric in latent space to separate different classes of data in supervised learning, often with a margin. This is different from our goal of matching for causal inference, but some of our proofs were inspired by this work in supervised learning.

A sister work, developed in parallel, is that of Morucci et al. (2020), which learns adaptively-sized hyperboxes as matched groups. MALTS was previously used on the ACIC 2018 Causal Inference Challenge Data (see Parikh et al. 2019). An extension of MALTS for multi-level treatments has been used to study the effect of seizures on the discharge status of critical ill patients (see Parikh et al. 2022).

3. Learning-to-Match Framework

Within this framework, we perform treatment effect estimation using following three stages: 1) learning a distance metric, 2) matching samples, and 3) estimating CATEs.

We denote the p dimensional covariate vector space as $\mathcal{X} \subset \mathbb{R}^p$ and the unidimensional outcome space by $\mathcal{Y} \subset \mathbb{R}$. Let \mathcal{T} be a finite label set of treatment indicators (in this paper we consider only the binary case). Let $\mathcal{Z} = \mathcal{X} \times \mathcal{Y} \times \mathcal{T}$ such that $z = (\mathbf{x}, y, t) \in \mathcal{Z}$ means that $\mathbf{x} \in \mathcal{X}$, $y \in \mathcal{Y}$ and $t \in \mathcal{T}$. Let μ be an unknown probability distribution over \mathcal{Z} such that $\forall z \in \mathcal{Z}$, $\mu(z) > 0$. We assume that \mathcal{X} is a compact convex space with respect to $\|\cdot\|_2$, thus there exists a constant \mathbf{C}_x such that $\|\mathbf{x}\|_2 \leq \mathbf{C}_x$. Also, $|y| \leq \mathbf{C}_y$. A distance metric is a symmetric, positive definite function with two arguments from \mathcal{X} such that $\mathbf{d} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$. A distance metric must obey the triangle inequality. Let \mathcal{S}_n denote a set of n observed units $\{s_1, \dots, s_n\}$ drawn i.i.d. from μ such that $\forall i$, $s_i \in \mathcal{Z}$. We parameterize \mathbf{d} with parameter $\mathcal{M}(\cdot)$, explicitly calling it $\mathbf{d}_{\mathcal{M}}$, and let $\mathcal{M}(\mathcal{S}_n)$ denote the parameter learned using MALTS methodology which is described in Section 4. For ease of notation, we will denote the observed sample of treated units as $\mathcal{S}_n^{(T)} := \{s_i^{(T)} = (\mathbf{x}_i, y_i, t_i) \mid s_i^{(T)} \in \mathcal{S}_n \text{ and } t_i = T\}$ and the observed sample of control units as $\mathcal{S}_n^{(C)} := \{s_i^{(C)} = (\mathbf{x}_i, y_i, t_i) \mid s_i^{(C)} \in \mathcal{S}_n \text{ and } t_i = C\}$.

We assume no unobserved confounders and standard ignorability assumptions, i.e., $\forall i$, $(Y_i^{(T)}, Y_i^{(C)}) \perp\!\!\!\perp T_i \mid (X_i = \mathbf{x}_i)$ (Rubin 2005) where $Y_i^{(T)}$ and $Y_i^{(C)}$ are potential outcomes for unit i under treatments (T) and (C) respectively, T_i is unit i 's treatment

choice and X_i corresponds to the vector of covariates for unit i . For each individual unit $s_i = (\mathbf{x}_i, y_i, t_i) \in \mathcal{Z}$ we define its conditional average treatment effect (or individualized treatment effect) as the difference of potential outcomes of unit i under the treatment and control, $\tau(\mathbf{x}_i) = \mathbb{E} \left[Y_i^{(T)} - Y_i^{(C)} | X_i = \mathbf{x}_i \right] = \mathbb{E} \left[Y_i^{(T)} | X_i = \mathbf{x}_i \right] - \mathbb{E} \left[Y_i^{(C)} | X_i = \mathbf{x}_i \right]$. We use the $\hat{Y}_{\mathbf{x}_i}^{(t)}$ to refer to the estimated conditional average potential outcome, $\mathbb{E} \left[Y_i^{(t)} | X_i = \mathbf{x}_i \right]$, for treatment $t \in \mathcal{T}$ and covariate level $\mathbf{x}_i \in \mathcal{X}$. $\hat{\tau}(\mathbf{x}_i)$ refers to the estimated conditional average treatment effect for covariate value \mathbf{x}_i .

Our goal is to minimize the expected loss between estimated treatment effects $\hat{\tau}(\mathbf{x})$ and true treatment effects $\tau(\mathbf{x})$ across target population $\mu(z)$ (this can either be a finite or super-population).

Let the population expected loss be:

$$\mathbb{E} [\ell(\hat{\tau}(\mathbf{x}), \tau(\mathbf{x}))] = \int \ell(\hat{\tau}(\mathbf{x}), \tau(\mathbf{x})) d\mu = \int \ell(\hat{Y}_{\mathbf{x}}^{(T)} - \hat{Y}_{\mathbf{x}}^{(C)}, \mathbb{E}[Y^{(T)} - Y^{(C)} | X = \mathbf{x}]) d\mu.$$

For a finite random i.i.d. sample $\{s_i = (\mathbf{x}_i, y_i, t_i)\}_{i=1}^n$ from the distribution μ , the finite sample version of the average loss can be written as

$$\frac{1}{n} \sum_{i=1}^n \ell \left(\hat{Y}_{\mathbf{x}_i}^{(T)} - \hat{Y}_{\mathbf{x}_i}^{(C)}, \mathbb{E} \left[Y_i^{(T)} | X_i = \mathbf{x}_i \right] - \mathbb{E} \left[Y_i^{(C)} | X_i = \mathbf{x}_i \right] \right).$$

However, we do not observed true values of $\mathbb{E} \left[Y_i^{(T)} | X_i = \mathbf{x}_i \right]$ and $\mathbb{E} \left[Y_i^{(C)} | X_i = \mathbf{x}_i \right]$.

Instead, we could estimate the upper bound of sample average loss as

$$\frac{1}{n} \sum_{i=1}^n t_i \ell \left(\hat{Y}_{\mathbf{x}_i}^{(T)}, y_i \right) + (1 - t_i) \ell \left(\hat{Y}_{\mathbf{x}_i}^{(C)}, y_i \right).$$

Here, we use can y_i for $t_i = 1$ as the unbiased estimate of $\mathbb{E} \left[Y_i^{(T)} | X_i = \mathbf{x}_i \right]$ and similar for $t_i = 0$.

For a unit s_i , we estimate the conditional average potential outcomes, $\hat{Y}_{\mathbf{x}_i}^{(T)}$ and $\hat{Y}_{\mathbf{x}_i}^{(C)}$, using the treated and control units' outcomes in the constructed *matched group* using the observed data. The *matched group* MG of unit s_i for treatment t' under the distance metric $\mathbf{d}_{\mathcal{M}}$ on covariate space is defined as a set of K nearest neighbors of s_i from set $\mathcal{S}_n^{(t')} = \{s_k | t_k = t', s_k \in \mathcal{S}_n\}$.

$$\text{MG}(s_i, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n^{(t')}, K) = KNN_{\mathcal{M}}^{\mathcal{S}_n}(s_i, t') := \left\{ s_k : \left[\sum_{s_l \in \mathcal{S}_n^{(t')}} \mathbb{1}(\mathbf{d}_{\mathcal{M}}(\mathbf{x}_l, \mathbf{x}_i) < \mathbf{d}_{\mathcal{M}}(\mathbf{x}_k, \mathbf{x}_i)) \right] < K \right\}. \quad (1)$$

We allow reuse of units in multiple matched groups. Thus for a chosen estimator ϕ ,

$$\hat{Y}_{\mathbf{x}_i}^{(t')} = \phi \left(\text{MG}(s_i, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n^{(t')}, K) \right) \quad (2)$$

where K is the size of the matched group $\text{MG}(s_i, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n^{(t')}, K)$. A simple example of ϕ is the mean estimator, i.e. $\phi \left(\text{MG}(s_i, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n^{(t')}, K) \right) = \frac{1}{K} \sum_{k \in \text{MG}(s_i, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n^{(t')}, K)} y_k$. However,

one can choose the estimator to be a weighted mean, linear regression or a non-parametric model like random-forest, within the matched group.

Our framework performs honest causal inference by learning a distance metric from a separate training set of data (not the estimation data considered in the averages above), and we denote this training set by \mathcal{S}_{tr} . To learn $\mathbf{d}_{\mathcal{M}}$, we minimize the following:

$$\mathcal{M}(\mathcal{S}_{tr}) \in \arg \min_{\mathcal{M}} \left[\sum_{s_i \in \mathcal{S}_{tr}^{(T)}} |y_i - \hat{Y}_{\mathbf{x}_i}^{(T)}| + \sum_{s_i \in \mathcal{S}_{tr}^{(C)}} |y_i - \hat{Y}_{\mathbf{x}_i}^{(C)}| \right],$$

where $\hat{Y}_{\mathbf{x}_i}^{(C)}$ and $\hat{Y}_{\mathbf{x}_i}^{(T)}$ are defined by Equations (1) and (2) including its dependence on the distance $\mathbf{d}_{\mathcal{M}}$, which is parameterized by \mathcal{M} , using the training data to create matched groups.

Once $\mathcal{M}(\mathcal{S}_{tr})$ is learned from the training set, it is used for matching (and estimation) on the estimation data.

3.1 Smooth Distance Metric and Treatment Effect Estimation

In this subsection, we discuss that if a distance metric is a smooth distance metric, then we can estimate the individualized treatment effect using a finite sample with high probability. First, let us define a smooth distance metric.

Definition 1 (Smooth Distance Metric) $\mathbf{d}_{\mathcal{M}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$ is a smooth distance metric if there exists a monotonically increasing bounded function $\delta_{\mathbf{d}_{\mathcal{M}}}(\cdot)$ with zero intercept, such that $\forall z_i, z_j \in \mathcal{Z}$ if $t_i = t_j$ and $\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j) \leq a$ then

$$|\mathbb{E}[Y_i | X_i = \mathbf{x}_i, T_i = t_i] - \mathbb{E}[Y_j | X_j = \mathbf{x}_j, T_j = t_j]| \leq \delta_{\mathbf{d}_{\mathcal{M}}}(a).$$

The concept of the smooth distance metric is analogous to commonly assumed Lipschitz continuity in the matching literature (Abadie and Imbens 2006). Note that because the range of Y is bounded, there always exists a choice of the function $\delta_{\mathcal{M}}(\cdot)$ such that a distance metric $d_{\mathcal{M}}$ is smooth. This choice of $\delta_{\mathcal{M}}(\cdot)$ controls the quality of inference from the matching as we see in Theorem 1 below.

Theorem 1 (Basic CATE Bound for Smooth Distance Metrics) Let $\{\mathcal{S}_n\}_{n=1}^{\infty}$ be a sequence of nested datasets, each of which includes n i.i.d. samples from $\mu(\mathcal{Z})$, $n = 1.. \infty$. Given a smooth distance metric $\mathbf{d}_{\mathcal{M}}$, covariate vector \mathbf{x} , and $\alpha > 0$, if there exists a small enough value of “ a ” and a large enough value of N such that $\mathcal{K}_n^{(t')}(\mathbf{x}) = \{z_k : \mathbf{d}_{\mathcal{M}}(\mathbf{X}_k, \mathbf{x}) < a, T_k = t', z_k \in \mathcal{S}_n\}$ is non-empty and $\alpha > 2\delta_{\mathbf{d}_{\mathcal{M}}}(a)$ for all $n \geq N$ and $t' \in \mathcal{T}$, then

$$P_{\{Y_i\}_{i=1}^n \sim \mu(\mathcal{Y}^n)} \left(|\hat{\tau}(\mathbf{x}) - \tau(\mathbf{x})| \geq \alpha \mid \{\mathbf{X}_i\}, \{T_i\} \right) \leq 4 \exp \left(\frac{-K_n(\mathbf{x}) \left(\frac{\alpha}{2} - \delta_{\mathbf{d}_{\mathcal{M}}}(a) \right)^2}{2\mathbf{C}_y} \right)$$

where $\hat{\tau}(\mathbf{x})$ is the estimated conditional average treatment effect using the matched sets $\mathcal{K}_n^{(1)}(\mathbf{x})$ and $\mathcal{K}_n^{(0)}(\mathbf{x})$, $\tau(\mathbf{x})$ is the true conditional average treatment effect, $K_n(\mathbf{x}) = \min_{t'} |\mathcal{K}_n^{(t')}(\mathbf{x})|$, and $\delta_{\mathbf{d}_{\mathcal{M}}}(a)$ is the bound from Definition 1 (definition of smooth distance metric).

Theorem 1 directly follows from Lemma 5 in the Appendix A which proves that for all $t' \in \mathcal{T}$ and $\mathbf{x} \in \mathcal{X}$, we can estimate average conditional potential outcomes, $\mathbb{E}[Y^{(t')}|X = \mathbf{x}]$, correctly with high probability using nearest neighbor matching under any smooth distance metric, and Lemma 6 in Appendix A which proves that estimating average conditional potential outcomes correctly with high probability leads to estimating CATEs, τ , correctly with high probability.

Our setup and Definition 1 are similar to one described by Kara et al. (2017). Our result in Lemma 5 proves the consistency for a uniform weighted nearest neighbor estimator where the weights are probability weights. The result is in congruence with consistency results by Stone (1977) and Jiang (2019); those works handled the special case where the weights are uniform probability weights instead of *any* probability weights.

Note that matching using any type of stretch norm that induces a smooth distance metric, including Mahalanobis distance (or its special case with an identity covariance matrix, the L_2 distance), to adjust for confounding produces consistent estimates of average treatment effects. Prognostic score (Hansen 2008) and other approaches that induce a smooth distance metric also produce consistent estimates of ATE.

4. Matching After Learning to Stretch (MALTS)

MALTS performs weighted nearest neighbors matching, where the weights for the nearest neighbors can be learned by minimizing the following objective. This objective is simply the loss of the in-sample nearest neighbor estimator:

$$\mathbf{W} \in \arg \min_{\mathbf{W}} \left[\sum_{i \in \mathcal{S}_{tr}^{(T)}} \left\| y_i - \sum_{s_l \in \mathcal{S}_{tr}^{(T)}, i \neq l} \widetilde{W}_{i,l} y_l \right\| \right] + \left[\sum_{i \in \mathcal{S}_{tr}^{(C)}} \left\| y_i - \sum_{l \in \mathcal{S}_{tr}^{(C)}, i \neq l} \widetilde{W}_{i,l} y_l \right\| \right] + \text{Reg}(\widetilde{W}), \quad (3)$$

where $\text{Reg}(\cdot)$ is a regularization function. We let $\widetilde{W}_{i,l}$ be a function of $\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_l)$. For example, the $\widetilde{W}_{i,l}$ can encode whether l belongs to i 's K -nearest neighbors. Alternatively, they can encode soft KNN weights where $\widetilde{W}_{i,l} \propto e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_l)}$. Thus, the intuition is to learn \mathbf{W} such that the in-sample nearest-neighbors estimator is as accurate as possible.

As a reminder of our notation, we consider distance metric $\mathbf{d}_{\mathcal{M}}$ parameterized by a set of parameters \mathcal{M} . We use Euclidean distances for continuous covariates, namely distances of the form $\|\mathcal{M}\mathbf{x}_a - \mathcal{M}\mathbf{x}_b\|_2$ where \mathcal{M} encodes the orientation of the data. In the past, \mathcal{M} has been hard-coded rather than learned; an example in the causal inference literature is the classical Mahalanobis distance (\mathcal{M} is fixed as the inverse covariance matrix for the observed covariates). This approach has been demonstrated to perform well in settings where all covariates are observed and the inferential target is the average treatment effect (Stuart 2010). We are interested instead in individualized treatment effects, and just as the choice of Euclidean norm in Mahalanobis distance matching depends on the estimand of interest, the stretch metric needs to be amended for this new estimand. We propose learning the parameters of a distance metric, \mathcal{M} , directly from the observed data rather than setting it beforehand. The parameters of distance metric \mathcal{M} can be learned such that \mathbf{W} minimizes the objective function on the training set.

In our framework, we can define ‘‘approximate closeness’’ differently for discrete covariates if desired. For continuous covariates, MALTS uses Euclidean distance, which is

also a reasonable metric to use for binary data (e.g., Mahalanobis-distance-matching papers recommend converting unordered categorical variables to binary indicators, see Stuart 2010); however, there are benefits to using other metrics, such as weighted Hamming distances, for comparison among sets of binary covariates. To accommodate a combination of Euclidean and Hamming distances, we parameterize our distance metric in terms of two components: one is a learned weighted Euclidean distance for continuous covariates while the other is a learned weighted Hamming distance for discrete covariates as in the FLAME and DAME algorithms (Wang et al. 2021, Dieng et al. 2019). These components are separately parameterized by matrices \mathcal{M}_c and \mathcal{M}_d respectively, $\mathcal{M} = [\mathcal{M}_c, \mathcal{M}_d]$ (here c indicates “continuous,” and d indicates “discrete”). Let $a = (a_c, a_d)$ and $b = (b_c, b_d)$ be the covariates for two individuals split into continuous and discrete pairs respectively.

Operationalizing Equation (3): To perform the step called “Distance Metric Learning” in Figure 1 we propose the following form for the distance metric:

$$\mathbf{d}_{\mathcal{M}}(a, b) = d_{\mathcal{M}_c}(a_c, b_c) + d_{\mathcal{M}_d}(a_d, b_d), \text{ where}$$

$$d_{\mathcal{M}_c}(a_c, b_c) = \|\mathcal{M}_c a_c - \mathcal{M}_c b_c\|_2, \quad d_{\mathcal{M}_d}(a_d, b_d) = \sum_{j=0}^{|a_d|} \mathcal{M}_d^{(j,j)} \mathbb{1}[a_d^{(j)} \neq b_d^{(j)}],$$

and $\mathbb{1}[A]$ is the indicator that event A occurred. We thus perform learned Hamming distance matching on the discrete covariates and learned-Mahalanobis-distance matching for continuous covariates.

MALTS performs an “honest” causal inference by splitting the observed sample dataset \mathcal{S}_n into a training set \mathcal{S}_{tr} (not for matching) and an estimation set \mathcal{S}_{est} (for matching). We learn $\mathcal{M}(\mathcal{S}_{tr})$ using the training set \mathcal{S}_{tr} such that in Equation (3), $\widetilde{W}_{i,l} = \frac{e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_l)}}{\sum_{s_k \in \mathcal{S}_{tr}^{(t)}} e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_k)}}$

and $Reg(\widetilde{W}) = \|\mathcal{M}\|_{\mathcal{F}}$ which defines **MALTS’ main implemented optimization problem**:

$$\mathcal{M}(\mathcal{S}_{tr}) \in \arg \min_{\mathcal{M}} \left(c \|\mathcal{M}\|_{\mathcal{F}} + \Delta_{\mathcal{S}_{tr}}^{(C)}(\mathcal{M}) + \Delta_{\mathcal{S}_{tr}}^{(T)}(\mathcal{M}) \right) \quad (4)$$

where $\|\cdot\|_{\mathcal{F}}$ is the Frobenius norm of the matrix, and:

$$\begin{aligned} \Delta_{\mathcal{S}_{tr}}^{(t)}(\mathcal{M}) &:= \frac{1}{|\mathcal{S}_{tr}^{(t)}|} \sum_{s_i \in \mathcal{S}_{tr}^{(t)}} \left| y_i - \sum_{s_l \in \mathcal{S}_{tr}^{(t)}} \frac{e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_l)}}{\sum_{s_k \in \mathcal{S}_{tr}^{(t)}} e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_k)}} y_l \right| \\ &= \frac{1}{|\mathcal{S}_{tr}^{(t)}|} \sum_{s_i \in \mathcal{S}_{tr}^{(t)}} \left| \sum_{s_l \in \mathcal{S}_{tr}^{(t)}} \frac{e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_l)}}{\sum_{s_k \in \mathcal{S}_{tr}^{(t)}} e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_k)}} (y_i - y_l) \right|. \end{aligned} \quad (5)$$

Matching and Estimation: To perform the step called “Nearest Neighbor Matching,” which produces “Matched Groups” that are used to estimate “CATEs” in Figure 1, we use the learned distance metric $\mathcal{M}(\mathcal{S}_{tr})$. To estimate conditional average treatment effects (CATEs) for each unit in the estimation set, we use its nearest neighbors from the same estimation set. Specifically, for any given unit s in the estimation set, we construct a K -nearest neighbor matched group $MG(s, \mathbf{d}_{\mathcal{M}(\mathcal{S}_{tr})}, \mathcal{S}_{est}, K)$ using control set $\mathcal{S}_{est}^{(C)}$ and treatment set

$\mathcal{S}_{est}^{(T)}$. For a choice of estimator ϕ , the estimated CATE for a treated unit $s = (\mathbf{x}_s, y_s, t_s = t')$ is calculated as follows:

$$\hat{\tau}(\mathbf{x}) = \phi\left(\text{MG}(s, \mathbf{d}_{\mathcal{M}(\mathcal{S}_{tr})}, \mathcal{S}_{est}^{(T)}, K)\right) - \phi\left(\text{MG}(s, \mathbf{d}_{\mathcal{M}(\mathcal{S}_{tr})}, \mathcal{S}_{est}^{(C)}, K)\right).$$

A simple example of ϕ is the empirical mean, i.e.,

$$\phi\left(\text{MG}(s, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n^{(t)}, K)\right) = \frac{1}{K} \sum_{k \in \text{MG}(s, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n^{(t)}, K)} y_k.$$

However, one can choose the estimator to be a weighted mean, linear regression or a non-parametric model like Random Forest. Particular choices of ϕ can also play a role in bias-adjustment to improve the matching estimator of the ATE as in Abadie and Imbens (2011).

For $\phi(\text{MG}(s, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n, K)) = \sum_{k \in \text{MG}(s, \mathbf{d}_{\mathcal{M}}, \mathcal{S}_n, K)} \widetilde{W}_k y_k$, if \widetilde{W}_k is chosen to be proportional to $e^{\mathbf{d}_{\mathcal{M}}(\mathbf{x}, \mathbf{x}_k)}$, then it leads to multi-robust (defined shortly) and generalizable CATE estimates via soft KNN (as shown in Theorem 2 and Theorem 4 below), while letting \widetilde{W}_k be proportional to $\mathbb{1}\left[s_k \in \text{KNN}_{\mathcal{M}(\mathcal{S}_{tr})}^{\mathcal{S}_{est}^{(C)}}\right]$ produces interpretable matched groups.

Hyperparameter choice: MALTS has four main hyperparameters: 1) K , which is the number of nearest neighbors used to estimate the counterfactual, which can be chosen by cross-validation. 2) n , the size of training set, i.e., the size of the split on the left of Figure 1. This can be chosen based on the amount of data relative to the number of features, though typically we choose it to be 10% of the data. 3) The maximum allowed diameter or caliper to prune bad matched groups. If the matches have a larger diameter, the matches are not tight and we may not be able to trust their estimates. The maximum diameter can be chosen by domain knowledge; the user defines how far apart points can be to make the matched group less interpretable. 4) The number of repeats refers to the number of times we shuffle the data and re-partition it for MALTS' training and estimation procedure. A larger number of repeats of the whole process helps with smoothing out the estimates over different train/test splits.

5. Robustness and Generalization of MALTS

In this section, we show that the MALTS framework correctly estimates the distance metric, facilitating correct estimates of CATEs under SUTVA and a positivity assumption. After basic definitions, and after showing that the learned distance metric and objective are bounded, we introduce and define the concepts of multi-robustness and generalizability of the learned distance metric. Multi-robustness implies that for any possible pair of points the empirical average loss is not far away from the population average loss. Theorem 2 proves that the distance metric learned by the MALTS algorithm is multi-robust. We use these results along with the error bound shown in Lemma 3, to show that MALTS' distance metric is generalizable, i.e., the population average loss and the empirical average loss on the observed data for the learned distance metric are close with high probability. Lastly, we show that MALTS' distance metric is asymptotically generalizable and that the empirical average loss approaches the population average loss as the size of the dataset goes to infinity.

Basic definitions of empirical loss and population loss. First, we define a pairwise loss for s_i and s_l so that it is only finite for treatment-treatment or control-control matched pairs,

$$\text{loss}[\mathcal{M}, s_i, s_l] := \begin{cases} e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_l)} |y_i - y_l| & \text{if } t_i = t_l \\ \infty & \text{otherwise.} \end{cases}$$

This loss is high for pairs of points that are close (i.e., with small $\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_l)$) when the outcomes y_i and y_l values are very different. Further, we define an empirical average pairwise loss over finite sample \mathcal{S}_n of size n as

$$L_{emp}(\mathcal{M}, \mathcal{S}_n) := \frac{1}{n^2} \sum_{(s_i, s_l) \in (\mathcal{S}_n \times \mathcal{S}_n)} \text{loss}[\mathcal{M}, s_i, s_l]$$

and define an average loss over population \mathcal{Z} as

$$L_{pop}(\mathcal{M}, \mathcal{Z}) := \mathbb{E}_{z_i, z_l \stackrel{i.i.d.}{\sim} \mu(\mathcal{Z})} [\text{loss}[\mathcal{M}, z_i, z_l]].$$

The search space over distance metrics is bounded. We show a basic result about the optimization-based approach we take to learn the distance metric. Specifically, we show that the learned distance metric will be in a bounded region of search space.

Now, because the learned $\mathcal{M}(\mathcal{S}_{tr})$ on the set \mathcal{S}_{tr} is the distance metric that minimizes the given objective function, we know that the following inequality is true, which states that the learned parameter has a lower training objective than that of the trivial parameter $\mathbf{0}$:

$$\left(c \|\mathcal{M}(\mathcal{S}_{tr})\|_{\mathcal{F}} + \Delta_{\mathcal{S}_{tr}}^{(C)}(\mathcal{M}(\mathcal{S}_{tr})) + \Delta_{\mathcal{S}_{tr}}^{(T)}(\mathcal{M}(\mathcal{S}_{tr})) \right) \leq \left(c \|\mathbf{0}\|_{\mathcal{F}} + \Delta_{\mathcal{S}_{tr}}^{(C)}(\mathbf{0}) + \Delta_{\mathcal{S}_{tr}}^{(T)}(\mathbf{0}) \right) =: g_0. \quad (6)$$

Denoting the right hand side of the inequality by g_0 we note that we can limit our search space over distance metrics \mathcal{M} that satisfy the following inequality:

$$\|\mathcal{M}\|_{\mathcal{F}} \leq \frac{g_0}{c}.$$

The objective function terms are bounded. The objective terms $\Delta_{\mathcal{S}_{tr}}^{(C)}$ and $\Delta_{\mathcal{S}_{tr}}^{(T)}$ (defined in Equation (5)) for learning the distance metric are also bounded, although it is not that easy to see this directly because their denominators are somewhat complicated, involving a sum over exponential terms. Here, we point out that because the learned distance metric is bounded, the objective's terms ($\Delta_{\mathcal{S}_{tr}}^{(C)}$ and $\Delta_{\mathcal{S}_{tr}}^{(T)}$) are also bounded. Specifically, their upper bound is proportional to the empirical average pairwise losses $L_{emp}(\mathcal{M}, \mathcal{S}_{tr}^{(C)})$ and $L_{emp}(\mathcal{M}, \mathcal{S}_{tr}^{(T)})$, defined above. Further, in Theorem 4, we show that for $t' \in \{T, C\}$ the empirical average loss $L_{emp}(\mathcal{M}, \mathcal{S}_{tr}^{(t')})$ is close to population average pairwise loss $L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')})$ with high probability. Following Equations (7) and (8) and Theorem 4, the objective terms $\Delta_{\mathcal{S}_{tr}}^{(C)}(\mathcal{M})$ and $\Delta_{\mathcal{S}_{tr}}^{(T)}(\mathcal{M})$ are upper-bounded by a term proportional to the population average pairwise loss with high probability.

$$\Delta_{\mathcal{S}_{tr}}^{(C)}(\mathcal{M}) \leq \frac{1}{|\mathcal{S}_{tr}^{(C)}|} \sum_{s_i \in \mathcal{S}_{tr}^{(C)}} \sum_{s_l \in \mathcal{S}_{tr}^{(C)}} \left| \frac{e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_l)}}{\sum_{s_k \in \mathcal{S}_{tr}^{(C)}} e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_k)}} (y_i - y_l) \right|$$

$$= \frac{1}{|\mathcal{S}_{tr}^{(C)}|} \sum_{s_i \in \mathcal{S}_{tr}^{(C)}} \frac{\sum_{s_l \in \mathcal{S}_{tr}^{(C)}} \text{loss}[\mathcal{M}, s_i, s_l]}{\sum_{s_k \in \mathcal{S}_{tr}^{(C)}} e^{-\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_k)}}.$$

We know that:

$$\forall i, k \mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_k) = [(\mathbf{x}_i - \mathbf{x}_k)' \mathcal{M}' \mathcal{M} (\mathbf{x}_i - \mathbf{x}_k)]^{1/2} \leq \|\mathbf{x}_i - \mathbf{x}_k\|_2 \|\mathcal{M}\|_{\mathcal{F}} \leq \frac{g_0 \mathbf{C}_x^2}{c}.$$

Together, the two previous lines imply:

$$\Delta_{\mathcal{S}_{tr}}^{(C)}(\mathcal{M}) \leq \frac{1}{\exp(-\frac{g_0 \mathbf{C}_x^2}{c}) |\mathcal{S}_{tr}^{(C)}|^2} \sum_{s_i \in \mathcal{S}_{tr}^{(C)}} \sum_{s_l \in \mathcal{S}_{tr}^{(C)}} \text{loss}[\mathcal{M}, s_i, s_l] = \frac{L_{emp}(\mathcal{M}, \mathcal{S}_{tr}^{(C)})}{\exp(-\frac{g_0 \mathbf{C}_x^2}{c})}. \quad (7)$$

Similarly for the treatment units, we have

$$\Delta_{\mathcal{S}_{tr}}^{(T)}(\mathcal{M}) \leq \frac{L_{emp}(\mathcal{M}, \mathcal{S}_{tr}^{(T)})}{\exp(-\frac{g_0 \mathbf{C}_x^2}{c})}. \quad (8)$$

Now, we define a few concepts important for our results including covering number, multi-robustness, and generalizability. The following definitions and results closely align with the theoretical guarantees of distance metric learning algorithms in Bellet and Habrard (2015) and Xu and Mannor (2012). Our work extends these results to learn a distance metric for causal inference.

Definition 2 (Covering Number) Let $(\mathcal{U}, \mathbf{d})$ be a metric space. Consider a subset \mathcal{V} of \mathcal{U} , then $\hat{\mathcal{V}} \subset \mathcal{V}$ is called a γ -cover of \mathcal{V} if for any $v \in \mathcal{V}$, we can always find a $\hat{v} \in \hat{\mathcal{V}}$ such that $\mathbf{d}(v, \hat{v}) \leq \gamma$. Further, the γ -covering-number of \mathcal{V} under the distance metric \mathbf{d} is defined by $\mathbf{N}(\gamma, \mathcal{V}, \mathbf{d}) := \min \{|\hat{\mathcal{V}}| : \hat{\mathcal{V}} \text{ is a } \gamma\text{-cover of } \mathcal{V}\}$.

Note that $\mathbf{N}(\gamma, \mathcal{V}, \mathbf{d})$ is finite if \mathcal{U} is compact.

Definition 3 (Robustness) A learned distance metric $\mathcal{M}(\cdot)$ is $(K, \epsilon(\cdot))$ -robust for a given K and $\epsilon(\cdot) : (\mathcal{Z} \times \mathcal{Z})^n \rightarrow \mathbb{R}$, if we can partition \mathcal{X} into K disjoint sets $\{C_i\}_{i=1}^K$ such that for any subsample \mathcal{S}_{tr} and its corresponding pair set $\mathcal{S}_{tr}^2 := \mathcal{S}_{tr} \times \mathcal{S}_{tr}$, we have for any pair of training units $(s_1 = (\mathbf{x}_1, y_1, t_1), s_2 = (\mathbf{x}_2, y_2, t_2)) \in \mathcal{S}_{tr}^2$, and for any pair of units in the support $(z_1 = (\mathbf{x}'_1, y'_1, t'_1), z_2 = (\mathbf{x}'_2, y'_2, t'_2)) \in \mathcal{Z}^2$, $\forall i, l \in \{1, \dots, K\}$,

if $\mathbf{x}_1, \mathbf{x}'_1 \in C_i$ and $\mathbf{x}_2, \mathbf{x}'_2 \in C_l$ such that $t_1 = t'_1 = t_2 = t'_2$ then

$$\left| \text{loss}[\mathcal{M}(\mathcal{S}_{tr}), s_1, s_2] - \text{loss}[\mathcal{M}(\mathcal{S}_{tr}), z_1, z_2] \right| \leq \epsilon(\mathcal{S}_{tr}).$$

Intuitively, *robustness* means that for any possible unit in the support, the loss is not far away from the loss of nearby units in the training set, should some training units exist nearby. (This terminology is aligned with the distance metric learning literature, e.g., Bellet and Habrard 2015, Xu and Mannor 2012, and it is different from robustness to model misspecification that frequently appears in the causal inference literature in terms such as “doubly robust estimator.”)

Definition 4 (Multi-Robustness) A learned distance metric $\mathcal{M}(\cdot)$ is $(K, \epsilon(\cdot))$ -multirobust for a given K and $\epsilon(\cdot) : \mathcal{Z}^n \rightarrow \mathbb{R}$, if we can partition \mathcal{X} into K disjoint sets $\mathbf{C} = \{C_i\}_{i=1}^K$ such that for any subsample \mathcal{S}_n and its corresponding pair set $\mathcal{S}_n^2 := \mathcal{S}_n \times \mathcal{S}_n$, we have $\forall (s_1 = (x_1, y_1, t_1), s_2 = (x_2, y_2, t_2)) \in \mathcal{S}_n^2, \forall (z_1 = (x'_1, y'_1, t'_1), z_2 = (x'_2, y'_2, t'_2)) \in \mathcal{Z}^2, \forall i, l \in \{1, \dots, K\}$,

$$\text{given } \widehat{\text{loss}}[\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')}] := \frac{1}{|C_i^{(t')}| |C_l^{(t')}|} \sum_{(s_i, s_l) \in C_i^{(t')} \times C_l^{(t')}} \text{loss}[\mathcal{M}(\mathcal{S}_n), s_1, s_2]$$

$$\text{and } \overline{\text{loss}}[\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')}] := \mathbb{E}[\text{loss}(\mathcal{M}, Z_i, Z_l) \mid X'_i \in C_i^{(t')}, X'_l \in C_l^{(t')}]$$

$$\forall C_i, C_l \in \mathbf{C}, \quad \left| \widehat{\text{loss}}[\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')}] - \overline{\text{loss}}[\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')}] \right| \leq \epsilon(\mathcal{S}_n).$$

Intuitively, *multi-robustness* means that for any possible pair of points from any two partitions of \mathcal{X} , the empirical average loss over training points is not far away from the population average loss. As the training procedure aims at minimizing the total loss, we can safely say that a multi-robust method will not perform poorly out of sample.

Definition 5 (Generalizability) A learned distance metric $\mathcal{M}(\cdot)$ is said to generalize with respect to the given training sample \mathcal{S}_n if

$$P_{\mathcal{S}_n} \left(\sum_{t' \in \mathcal{T}} \left| L_{\text{pop}}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{\text{emp}}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')}) \right| \geq \epsilon \right) \leq \delta_\epsilon$$

where δ_ϵ is a decreasing function of ϵ with zero-intercept.

Definition 6 (Asymptotic Generalizability) A learned distance metric $\mathcal{M}(\cdot)$ is said to asymptotically generalize with respect to the given training sample \mathcal{S}_n if

$$\lim_{n \rightarrow \infty} \sum_{t' \in \mathcal{T}} \left| L_{\text{pop}}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{\text{emp}}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')}) \right| = 0$$

Given these definitions, we first show that the distance metric learned using MALTS is robust in Theorem 2 and we extend the argument to show that it is also generalizable in Theorem 4.

Theorem 2 (MALTS' learned distance metric is multi-robust) With probability greater than $\left(1 - \exp \left(-\frac{\beta^2 (\rho_\gamma^{(t')})^2}{n^{(t')} B^2} \right) \right)$, the distance metric $\mathcal{M}(\cdot)$ learned using MALTS

is $\left(\mathbf{N}(\gamma, \mathcal{X}, \|\cdot\|_2), \beta \right)$ -multirobust for arbitrary chosen values of $\gamma > 0$ and $\beta \geq 0$, where B is $\max_{z_1, z_2} \text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2)$, $\{C_i\}_{i=1}^K$ is the partition of \mathcal{X} into non-empty sets C_i 's such that K is the γ -covering number of \mathcal{X} , $C_i^{(t')} = \{z_j = (\mathbf{x}_j, y_j, t_j) : t_j = t', \mathbf{x}_j \in C_i\}$ and $\rho_\gamma^{(t')} = \min_i |C_i^{(t')}|$.

Proof (Theorem 2). Given $\mathcal{Z} = \mathcal{X} \times \mathcal{Y} \times \mathcal{T}$, we consider the following definition of a minimum sized γ -cover $\hat{\mathcal{V}}$ of the set \mathcal{X} under the distance metric $\|\cdot\|_2$: Partition the set into K disjoint subsets $\mathbf{C}_\gamma = \{C_i\}_{i=1}^K$ such that K is the γ -covering-number of \mathcal{X} under $\|\cdot\|_2$ (which is exactly equal to $|\hat{\mathcal{V}}|$) where each C_i is contained in the γ -neighborhood of each $\hat{v}_i \in \hat{\mathcal{V}}$ and each C_i contains at least one control and one treated sample. Note that if \mathcal{X} is a compact convex set, then such a cover and the corresponding packing \mathbf{C}_γ exists and $K = |\mathbf{C}_\gamma|$ is finite.

For any arbitrary C_i and C_l in \mathbf{C}_γ , consider the empirical average loss for all training units $s_i \in C_i$ and $s_l \in C_l$ with treatment t' :

$$\widehat{\text{loss}} \left[\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')} \right] = \frac{1}{|C_i^{(t')}| |C_l^{(t')}|} \sum_{(s_i, s_l) \in C_i^{(t')} \times C_l^{(t')}} \text{loss}[\mathcal{M}(\mathcal{S}_n), s_i, s_l]$$

and the expected loss for units Z_i and Z_l :

$$\overline{\text{loss}} \left[\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')} \right] = \mathbb{E} \left[\text{loss}(\mathcal{M}, Z_i, Z_l) \mid X_i' \in C_i^{(t')}, X_l' \in C_l^{(t')} \right].$$

Let f be a function of the set of independent random variables such that

$$f(s_1, \dots, s_{|C_i^{(t')}|}, s_{|C_i^{(t')}|+1}, \dots, s_{|C_i^{(t')}|+|C_l^{(t')}|}) = \frac{1}{|C_i^{(t')}| |C_l^{(t')}|} \sum_{j=1}^{C_i^{(t')}} \sum_{i=C_i^{(t')}+1}^{C_l^{(t')}} \text{loss}[\mathcal{M}(\mathcal{S}_n), s_i, s_l].$$

Thus, $f(s_1, \dots, s_{|C_i^{(t')}|}, s_{|C_i^{(t')}|+1}, \dots, s_{|C_i^{(t')}|+|C_l^{(t')}|}) = \widehat{\text{loss}} \left[\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')} \right]$.

Now, let $\rho_\gamma^{(t')}$ be the density of the γ -cover for treatment t' , defined as the number of units with treatment t' in the smallest partition set $\rho_\gamma^{(t')} = \min_i |C_i^{(t')}|$ and $B = \max_{z_1, z_2} \text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2)$. Now, we show that $f(\cdot)$ has bounded difference. Without loss of generality, consider an index $j \leq |C_i^{(t')}|$, then

$$\begin{aligned} & \left| f(s_1, \dots, s_j, \dots, s_{|C_i^{(t')}|+|C_l^{(t')}|}) - f(s_1, \dots, s'_j, \dots, s_{|C_i^{(t')}|+|C_l^{(t')}|}) \right| \\ &= \left| \frac{1}{|C_i^{(t')}| |C_l^{(t')}|} \sum_{i=|C_i^{(t')}|+1}^{|C_i^{(t')}|+|C_l^{(t')}|} \text{loss}[\mathcal{M}(\mathcal{S}_n), s_i, s_j] - \text{loss}[\mathcal{M}(\mathcal{S}_n), s_i, s'_j] \right| \\ &\leq \frac{1}{|C_i^{(t')}| |C_l^{(t')}|} \sum_{i=|C_i^{(t')}|+1}^{|C_i^{(t')}|+|C_l^{(t')}|} \left| \text{loss}[\mathcal{M}(\mathcal{S}_n), s_i, s_j] - \text{loss}[\mathcal{M}(\mathcal{S}_n), s_i, s'_j] \right| \\ &\leq \frac{1}{|C_i^{(t')}| |C_l^{(t')}|} \sum_{i=|C_i^{(t')}|+1}^{|C_i^{(t')}|+|C_l^{(t')}|} \left(\left| \text{loss}[\mathcal{M}(\mathcal{S}_n), s_i, s_j] \right| + \left| \text{loss}[\mathcal{M}(\mathcal{S}_n), s_i, s'_j] \right| \right) \\ &\leq \frac{|C_l^{(t')}|}{|C_i^{(t')}| |C_l^{(t')}|} B = \frac{B}{|C_i^{(t')}|} \leq \frac{B}{\rho_\gamma^{(t')}}. \end{aligned}$$

Similarly, for any $j > |C_i^{(t')}|$,

$$|f(s_1, \dots, s_j, \dots, s_{|C_i^{(t')}|+|C_l^{(t')}|}) - f(s_1, \dots, s'_j, \dots, s_{|C_i^{(t')}|+|C_l^{(t')}|})| \leq \frac{2B}{\rho_\gamma^{(t')}}.$$

As $f()$ is a function of independent $|C_i^{(t')}| + |C_l^{(t')}|$ random variables, by McDiarmid's inequality:

$$\begin{aligned} & P \left(\left| \widehat{\text{loss}} [\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')}] - \overline{\text{loss}} [\mathcal{M}(\mathcal{S}_n), C_i^{(t')}, C_l^{(t')}] \right| \geq \beta \right) \\ & \leq \exp \left(- \frac{2\beta^2}{\sum_{i=1}^{|C_i^{(t')}|+|C_l^{(t')}|} \frac{B^2}{(\rho_\gamma^{(t')})^2}} \right) = \exp \left(- \frac{2\beta^2 (\rho_\gamma^{(t')})^2}{(|C_i^{(t')}| + |C_l^{(t')}|) B^2} \right) \leq \exp \left(- \frac{\beta^2 (\rho_\gamma^{(t')})^2}{n^{(t')} B^2} \right). \end{aligned}$$

■

We will need the following lemma to prove Theorem 4. The lemma provides a bound for a particular treatment assignment, while the theorem sums over all treatment assignments.

Lemma 3 (Error Bound) *Given sample $\mathcal{S}_n \stackrel{i.i.d}{\sim} \mu(\mathcal{Z})$ where $n^{(t')}$ is the number of units with $t_i = t'$ in \mathcal{S}_n , and choosing $B > 0$ for which $\text{loss}[\cdot, z_i, z_l] \leq B \forall z_i, z_l \in \mathcal{Z}$ (B is finite because \mathcal{X} is compact and \mathcal{Y} is bounded): if a learning algorithm provides a distance metric $\mathcal{M}(\mathcal{S}_n)$ that is $(K, \epsilon(\cdot))$ -multi-robust with probability $p_{mr}(\epsilon)$, then for any $\mathcal{E} > 0$, with probability greater than or equal to $(1 - \mathcal{E})(p_{mr}(\epsilon))^{K^2}$ we have*

$$\forall t' \in \mathcal{T}, \left| L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')}) \right| \leq \epsilon(\mathcal{S}_n^{(t')}) + 2B \sqrt{\frac{2K \ln(2) + 2 \ln(1/\mathcal{E})}{n^{(t')}}}.$$

Theorem 4 (MALTS' distance metric is generalizable) *The distance metric $\mathcal{M}(\cdot)$ learned using the data \mathcal{S}_n and MALTS algorithm is generalizable and asymptotically generalizable, as follows:*

1. *Generalizability:*

With probability at least

$$(1 - \mathcal{E})^{|\mathcal{T}|} \left(1 - \exp \left(- \frac{\beta^2 (\rho_\gamma^{(t')})^2}{K^2 n^{(t')} B^2} \right) \right)^{|\mathcal{T}| K^2}$$

with respect to the random draw of data,

$$\sum_{t' \in \mathcal{T}} \left| L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')}) \right| \leq 2|\mathcal{T}|\beta + \sum_{t' \in \mathcal{T}} 2B \sqrt{\frac{2K \ln(2) + 2 \ln(1/\mathcal{E})}{n^{(t')}}}$$

for arbitrary chosen constants $\gamma > 0$, $\mathcal{E} > 0$, and $\beta \geq 0$, where B is $\max_{z_1, z_2} \text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2)$, $\{C_i\}_{i=1}^K$ is the partition of \mathcal{X} into non-empty sets C_i 's such that K is the γ -covering number of \mathcal{X} , $C_i^{(t')} = \{z_j = (\mathbf{x}_j, y_j, t_j) : t_j = t', \mathbf{x}_j \in C_i\}$, and $\rho_\gamma = \min_{i, t'} |C_i^{(t')}|$.

2. *Asymptotic Generalizability:*

$$\lim_{n \rightarrow \infty} \left(\left| L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(C)}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(C)}) \right| + \left| L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(T)}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(T)}) \right| \right) = 0$$

Now that we have theoretically proven the functionality of MALTS, we will next discuss and compare MALTS performance with other methods on different datasets.

6. Experiments

In this section, we discuss and compare the performance of MALTS with other competing methods on a few different simulation setups with continuous covariates, discrete covariates and mixed (continuous and discrete) covariates. Lastly, we demonstrate MALTS performance for estimating ATE on LaLonde’s NSW and PSID-2 data samples (LaLonde 1986, Dehejia and Wahba 1999).

MALTS performs an η -fold honest causal inference procedure with the estimator ϕ inside each matched group being linear regression. We split the observed samples \mathcal{S}_n into η equal parts such that the ratio of treated to control units in each part is similar. For each fold, we use one of the η partitions as the training set \mathcal{S}_{tr} (not used for matching) and the rest of the $\eta - 1$ partitions as the estimation set \mathcal{S}_{est} . Using the output from each of the η folds, we calculate the estimated CATE for each unit (averaged across folds), estimated distance metric (averaged across folds) and a weighted unified matched group for each unit $s_i \in \mathcal{S}_n$. The weight of each matched unit s_k corresponds to the number of times a particular unit s_k was in the matched group of unit s_i across the $\eta - 1$ constructed matched groups. Here, η was chosen to be 5 in our experiments.

For interpretability, we let \mathcal{M}_c be a diagonal matrix, which allows stretches of the continuous covariates. (Note that \mathcal{M}_d , which is the stretch matrix over discrete covariates, is always set to be diagonal.) This way, the magnitude of an entry in \mathcal{M}_c or \mathcal{M}_d provides the relative importance of the indicated covariate for the causal inference problem.

We further analyzed strategies for variance estimation for MALTS in Section 6.8, and performance under limited overlap between the covariates distribution of treated and control groups, and sensitivity to unobserved confounding. Detailed results are shown in Appendix B.

The main results of these experiments are that **MALTS’ performance is on par with existing state-of-the-art methods for causal inference**, including black box methods. **MALTS tends to have fairly consistent performance, even if the training set is fairly small or the number of irrelevant covariates is large.** Further, **MALTS provides interpretable distance metrics and matched groups** that black box machine learning methods do not provide.

6.1 Data Generation Processes

In this subsection we describe the data generation process (DGP) used in the simulation experiments. We use two main data-generation processes: The first DGP has a linear baseline with linear and quadratic treatment effects while the second DGP is the extension

of Friedman’s function introduced to test performance of prediction algorithms of Friedman (1991). This second DGP, also termed as Friedman’s DGP, has a scaled sinusoidal treatment effect.

6.1.1 QUADRATIC DGP

This simulation includes both linear and quadratic terms. Let $\mathbf{x}_{i,p} = \{\mathbf{x}_{i,p_c}, \mathbf{x}_{i,p_d}\}$ be a p -dimensional covariate vector composed of $|p_c|$ continuous covariates and $|p_d|$ discrete ones. There are $k = k_c \cup k_d$ relevant covariates and the rest of the dimensions are irrelevant. Here, $p_c, k_c, p_d,$ and k_d refer to the subsets of indices of the covariates: all continuous, relevant continuous, all discrete, and relevant discrete, respectively. \mathbf{x}_{i,k_c} and \mathbf{x}_{i,k_d} refer to the vectors of relevant continuous and discrete covariates respectively. $\mathbf{x}_{i,k}$ refers to all $|k|$ relevant covariates. $\kappa_c \subseteq k_c$ is the set of continuous covariates and $\kappa_d \subseteq k_d$ is the set of discrete which are relevant in determining the treatment choice. The potential outcomes and treatment assignment are determined as follows:

$$\begin{aligned} \mathbf{x}_{i,p_c} &\stackrel{iid}{\sim} \mathcal{N}(\mu, \Sigma), \{x_{i,j}\}_{j \in p_d} \stackrel{iid}{\sim} \text{Bernoulli}(\psi), \epsilon_{i,0}, \epsilon_{i,1} \stackrel{iid}{\sim} \mathcal{N}(0, 1), \epsilon_{i,\text{treat}} \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2) \\ s_1, \dots, s_{|k|} &\stackrel{iid}{\sim} \text{Uniform}\{-1, 1\}, \alpha_j | s_j \stackrel{iid}{\sim} \mathcal{N}(10s_j, 9), \beta_1, \dots, \beta_{|k|} \stackrel{iid}{\sim} \mathcal{N}(1, 0.25) \\ y_i^{(0)} &= \sum_{j \in k_c \cup k_d} \alpha_j x_{i,j} + \epsilon_{i,0} \\ y_i^{(1)} &= \sum_{j \in k_c \cup k_d} \alpha_j x_{i,j} + \sum_{j \in k_c \cup k_d} \beta_j x_{i,j} + \sum_{j \in k_c \cup k_d} \sum_{j' \in k_c \cup k_d} x_{i,j} x_{i,j'} + \epsilon_{i,1} \\ t_i &= \mathbb{1} \left[\text{expit} \left(\sum_{j \in \kappa_c \subseteq k_c} x_{i,j} + \sum_{j \in \kappa_d \subseteq k_d} x_{i,j} - (|\kappa_c| \mu + |\kappa_d| \psi) + \epsilon_{i,\text{treat}} \right) > 0.5 \right] \\ y_i &= t_i y_i^{(1)} + (1 - t_i) y_i^{(0)}. \end{aligned}$$

Here $\text{expit}(z) = \exp(z)/(1 + \exp(z))$. The variance of $\epsilon_{i,\text{treat}}$ determines how much confounding and overlap there is in the dataset: higher values of the variance make the dataset look like a randomized experiment with good overlap, while very small values of the variance lead to poor overlap and a very hard to analyze observational study. We explore these issues in detail in Appendix B.

6.1.2 FRIEDMAN’S DGP

The data generation process of Friedman (1991) was first proposed to assess the performance of prediction methods. We augmented Friedman’s simulation setup to evaluate causal inference methods. The potential outcome under control is Friedman’s function as provided by Friedman (1991) and Chipman et al. (2010). The expected treatment effect we study is equal to the cosine of the product of the first two covariates scaled by the third covariate.

$$\begin{aligned} x_{i,1} \dots x_{i,10} &\stackrel{iid}{\sim} \mathcal{U}(0, 1), \epsilon_{i,0}, \epsilon_{i,1} \sim \mathcal{N}(0, 1), \epsilon_{i,\text{treat}} \stackrel{iid}{\sim} \mathcal{N}(0, 1) \\ y_i^{(0)} &= 10 \sin(\pi x_{i,1} x_{i,2}) + 20 (x_{i,3} - 0.5)^2 + 10 x_{i,4} + 5 x_{i,5} + \epsilon_{i,0} \end{aligned}$$

$$\begin{aligned}
 y_i^{(1)} &= 10 \sin(\pi x_{i,1} x_{i,2}) + 20 (x_{i,3} - 0.5)^2 + 10 x_{i,4} + 5 x_{i,5} + x_{i,3} \cos(\pi x_{i,1} x_{i,2}) + \epsilon_{i,1} \\
 t_i &= \mathbb{1}[\text{expit}(x_{i,0} + x_{i,1} - 0.5 + \epsilon_{i,\text{treat}}) > 0.5] \\
 y_i &= t_i y_i^{(1)} + (1 - t_i) y_i^{(0)}.
 \end{aligned}$$

6.2 Continuous Covariates

We use the data-generation process described in Section 6.1.1 to generate 2500 units with no discrete covariates, 15 important continuous covariates and 25 irrelevant continuous covariates. Further, we set the parameters for the DGP as follows: $\mu = 1$, $\Sigma = 1.5\mathbf{I}$, $\psi = 0.5$, $\sigma^2 = 1$ and $\kappa_c = \{0, 1\}$. We estimate CATE for each unit using matching methods like propensity score matching, prognostic score matching and genetic matching, and non-matching (uninterpretable) methods like causal forest and BART. Figure 2 shows the performance of these methods. *MALTS’ performance is on par with existing state-of-the-art non-matching methods and outperforms all other matching methods for continuous covariates in the quadratic data generation process.*

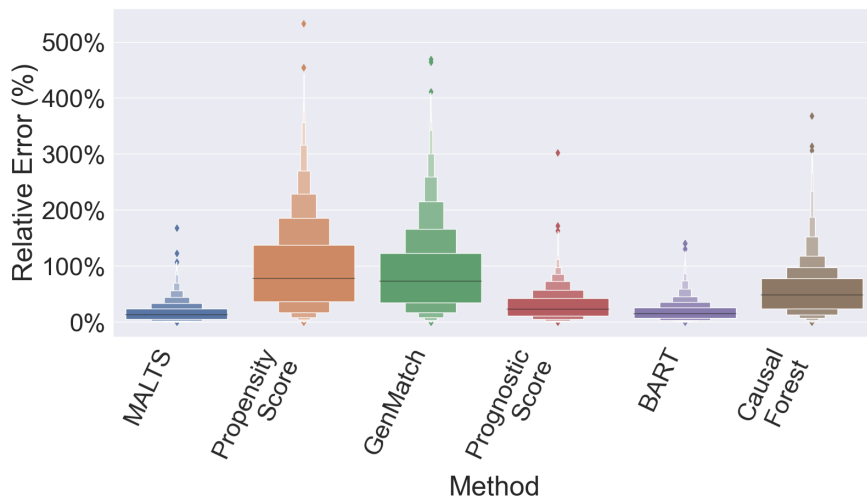


Figure 2: *MALTS performs well with respect to other methods for continuous data.* Letter-box plots of CATE Absolute Error relative to the true ATE on the test set for several methods.

6.3 Discrete Covariates

We use the data-generation process described in Section 6.1.1 to generate 2500 units with no continuous covariates, 15 important discrete covariates and 10 irrelevant discrete covariates. Further, we set the parameters of the DGP as follows: $\sigma^2 = 1$, $c = 2$ and $\kappa_d = \{0, 1\}$. We used the weighted Hamming distance metric for this experiment.

Figure 3 shows the performance comparison, again showing that MALTS’ performance is on par with existing state-of-the-art non-matching methods; it also performs better than

FLAME (a state-of-the-art matching method for discrete data) as it is able to provide additional smoothing in this relatively small- n setting. Hence, *MALTS performs well for discrete covariates in the quadratic data generation process.*

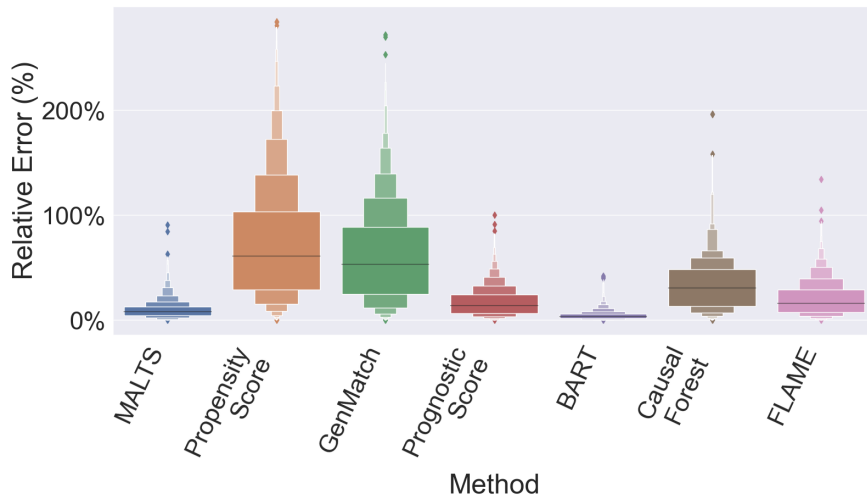


Figure 3: *MALTS performs well with respect to other methods for discrete data.* Letter-box plots of CATE Absolute Error relative to the true ATE on the test set for several methods.

6.4 Mixed Covariates

We use the data-generation process used for experiments on continuous and discrete covariates (described in Section 6.1.1) to generate 2500 units with 5 relevant continuous covariates, 15 relevant discrete covariates, 10 irrelevant continuous and 10 irrelevant discrete covariates. We used the same set of parameters for the DGP as the previous two experiments. Similar to the previous two experiments, Figure 4 shows that *MALTS performs on par with the state-of-the-art non-matching methods and outperforms all matching methods that can handle mixed covariates for the quadratic data generation process.*

6.5 Number of Covariates

We studied the performance of various causal inference methods to estimate CATEs as the number of covariates (p) changes, keeping the number of relevant covariates ($|k|$) constant and equal to 8. We simulated the data using the DGP described in Section 6.1.1. The number of units is constant ($n = 2048$) while the number of covariates (p) changes from 8 to 256. The performance of MALTS is on-par with or better than other causal inference methods as the number of irrelevant covariates increases (see Figure 5). This indicates that MALTS can be used to help reduce the effects of the curse of dimensionality.

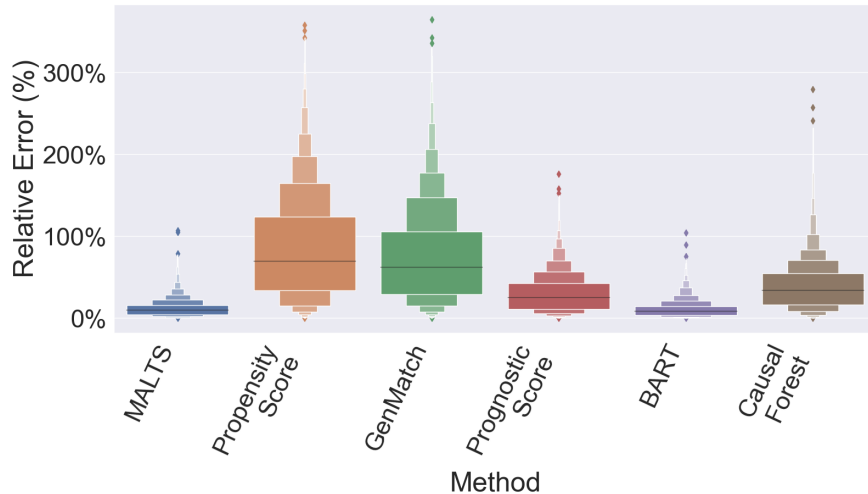


Figure 4: *MALTS performs well on data with mixed covariates.* Letter-box plots of CATE Absolute Error relative to the true ATE on the test set for several methods. MALTS performs well on the setup with mixed (continuous+discrete) covariates.

6.6 Number of Units

We studied the change in CATE estimation error-rates as the number of units in a dataset increases. We simulated the data using the DGP described in Section 6.1.1, keeping the number of covariates constant and equal to 20 (all of them are relevant in outcome determination). We changed the number of units from 2^8 to 2^{12} . *MALTS' performance is on-par with or better than BART and the error-rate is significantly lower than that of other causal inference methods* (see Figure 6).

6.7 Friedman's Setup

We further compare MALTS and other flexible methods' performance on data generated using the process described in Section 6.1.2. This DGP is particularly interesting because the potential outcomes are highly non-linear functions with trigonometric expressions.

As shown in Figure 7, we observe that *MALTS performs on par with Causal Forest while BART's error-rate is significantly higher (worse) than MALTS, for the Friedman's data generation process.*

6.8 Coverage Study

We use the DGP described in Section 6.1.1 with 2 relevant continuous covariates and no irrelevant covariates for the coverage study. Further, we set the parameters to the DGP as follows: $\mu = 1$, $\Sigma = 1.5\mathbf{I}$, $\psi = 0.5$, and $c = 2$. We selected 9 reference points in a grid from the covariate space as shown in Figure 8(b) and conducted an experiment that considered these reference points, over 100 repetitions. We compared coverage for CATEs estimated

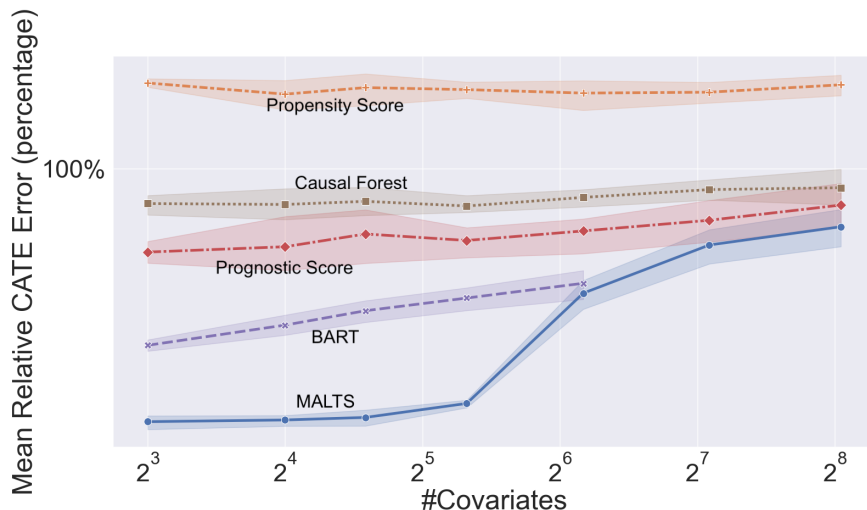


Figure 5: *MALTS performs on-par with other methods for a range of values of p . Comparative performance in estimating CATE using causal inference methods as the number of covariates increases, keeping the number of relevant covariates constant and equal to 8. The number of units is fixed: $n = 2^{11}$. (For the given n , BART does not return CATE estimates for some units when $p > 2^6$. Prognostic scores use BART for $p \leq 2^7$ and gradient boosted trees for $p > 2^6$.)*

using MALTS for different values of the variance, ranging from 1.0 to 4.0, for noise term ϵ_0 and ϵ_1 in the potential outcomes function.

Variance estimation is notoriously hard in matching problems, even for overall quantities such as the average treatment effect (Abadie and Imbens 2006). We consider both a conservative variance estimator (Wang et al. 2021) and estimators that sacrifice some interpretability for better coverage. Specifically, we consider the CATEs estimated using MALTS and study how well an uninterpretable method can predict those estimates to obtain a variance estimate. We use the predictive variance from gradient boosting regression, from gaussian process regression and from Bayesian ridge regression on the covariates, where we estimated CATEs and quantify variance of each CATE estimate.

Based on Figure 8(a), the coverage for each the nine points of interest is between 0.85 and 1 for most values of the variance using any of the three variance estimation approaches.

6.9 LaLonde Data

The LaLonde data pertain to the National Support Work Demonstration (NSW) temporary employment program and its effect on income level of the participants (LaLonde 1986). This dataset is frequently used as a benchmark for the performance of methods for observational causal inference. We employ the male sub-sample from the NSW in our analysis as well as the PSID-2 control sample of male household-heads under age 55 who did not classify themselves as retired in 1975 and who were not working when surveyed in the spring of 1976

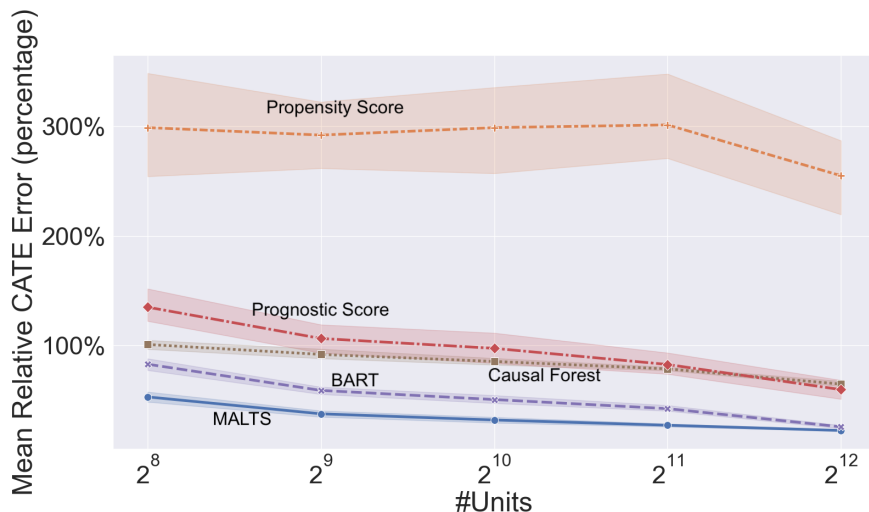


Figure 6: *MALTS* consistently performs on par with or better than non-interpretable approaches. Trend plots of average CATE Absolute Error for several methods, for different numbers of units in the datasets.

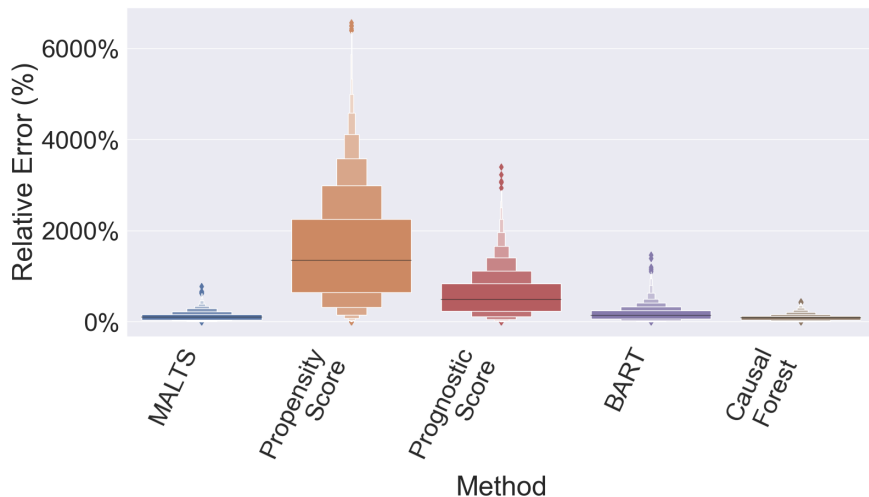
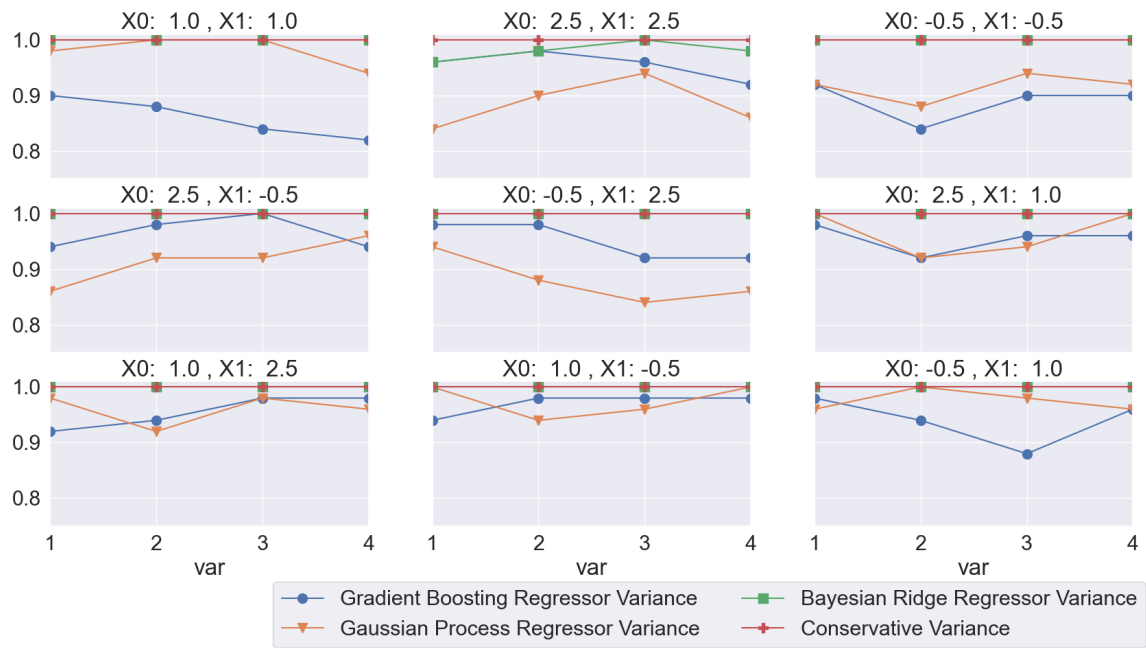


Figure 7: *MALTS* performs well on Friedman’s setup. Letter-box plots of CATE absolute error relative to true ATE for MALTS and other causal inference methods.

(Dehejia and Wahba 1999). The outcome variable for both experimental and observational analyses is earnings in 1978 and the considered variables are age, education, whether a respondent is Black, is Hispanic, is married, has a degree, and their earnings in 1975. Previously, it has been demonstrated that almost any adjustment during the analysis of the experimental and observational variants of these data (both by modeling the outcome and

MALTS



(a)



(b)

Figure 8: (a) Coverage of 95 percent confidence interval for 9 points: $(1.0,1.0)$, $(2.5,2.5)$, $(-0.5,-0.5)$, $(2.5,-0.5)$, $(-0.5,2.5)$, $(4.0,4.0)$, $(-3.0,-3.0)$, $(4.0,-3.0)$ and $(-3.0,4.0)$. (b) Covariate space showing positions of 9 points-of-interest as black-stars, with other points color-coded according to their treatment assignments.

by modeling the treatment variable) can lead to extreme bias in the estimate of average treatment effects (LaLonde 1986).

Table 2: *Estimated ATE for different methods on Lalonde’s NSW experimental dataset. The MALTS estimate of ATE is closer to the true ATE than other methods. We provide estimates for MALTS before and after pruning the matched groups with large diameters. The threshold to prune was chosen by rule of thumb on diameters of matched groups as shown in Figure 9(b).*

Method	ATE Estimate	Estimation Bias (%)
Truth	886	-
MALTS	881.67	-0.49
MALTS (pruned)	888.53	0.29
GenMatch	859.72	-2.97
Propensity Score	513.30	-42.06
Prognostic Score	943.81	6.52
BART-CV	1164.72	31.46
Causal Forest-CV	509.32	-42.51

Table 3: *Estimated ATE for different methods on Lalonde’s NSW experimental data and PSID-2 observational dataset. We provide estimates for MALTS before and after pruning the matched groups with large diameters. The threshold to prune was chosen by rule of thumb on diameters of matched groups as shown in Figure 9(b).*

Method	ATE Estimate	Estimation Bias (%)
Truth	886	-
MALTS	608.37	-31.34
MALTS (pruned)	891.75	0.65
GenMatch	549.53	-37.98
Propensity Score	513.79	-42.01
Prognostic Score	-897.76	-201.33
BART-CV	713.20	-19.50
Causal Forest-CV	-179.98	-120.31

Performance results: Tables 2 and 3 present the average treatment effect estimates based on MALTS, state-of-the-art modeling methods, and matching methods. *MALTS (after appropriately pruning low-quality matched groups) is able to achieve accurate ATE estimation on both experimental and observational datasets.*

Figure 9 illustrates how the matched groups were pruned. There was a clear visual separation between high-quality matched groups, which had low diameters, and low-quality matched groups, with larger diameters.

Model Interpretability: One difference between MALTS and the other methods is that its solution can be described concisely: MALTS produces a total of seven numbers that define the distance metric on the LaLonde data. The distribution of the learned distance metric values across folds is shown in Figure 9(a). Once the researcher has these seven numbers, along with the value of k in k -nearest neighbors used to train MALTS, they know precisely which units should be matched. In contrast, causal forest and BART require a model whose size depends on the number of trees, where each tree is several levels deep—in this case, 2000 trees and 150 trees, respectively.

Interpretability of Matched Groups: To examine the interpretability of MALTS’ matched groups, we present two of the matched groups from MALTS for the observational Lalonde dataset in Table 4, corresponding to two “query” individuals in the dataset. Query 1 is a 22 year old with no income in 1975. MALTS was able to construct a tight matched group for this individual (both in control and in treatment). In contrast, Query 2 is a 42-year-old high-income individual without a degree, which is an extremely unlikely scenario, leading to a matched group with a very large diameter, which should probably not be used during analysis. Such granular analysis is not possible for regression methods like BART and matching methods like prognostic score or propensity score matching.

This further highlights the troubleshooting capabilities of interpretable matching methods: by identifying units that are poorly matched, we know exactly which units to study in more detail. In this case, it is possible that the “degree” field might have a data error, which means it would be better not to match this unit and to potentially follow up on the veracity of responses to the survey.

7. Conclusion and Discussion

This paper introduces the MALTS algorithm, which learns a distance metric on the covariate space for use with matching. The learned metric stretches important covariates and compresses irrelevant covariates for outcome prediction in order to produce high-quality matches. Unlike other methods, MALTS can handle a large number of irrelevant covariates by compressing them to the point where they are effectively eliminated, which helps handle the curse of dimensionality. Unlike black-box machine learning methods, MALTS produces interpretable matched groups and returns the stretch matrix on covariates for counterfactual prediction. The stretch matrix is chosen here to be diagonal, so that it can be represented using only a few “stretch” numbers that determine the importance of each covariate in determining the matched groups.

Whereas deep neural networks mainly show improvements over other methods for problems that do not have natural data representations (computer vision, speech, etc.), we conjecture that the stretch/almost-exact match combination should suffice for most datasets. A natural extension, however, is to use neural networks to learn a flexible distance metric in a latent space, thus allowing us to match on medical records, images, and text documents. This will allow us to incorporate complex data structures by introducing a flexible learning framework (e.g., interpretable neural networks) for coding the data. That is, we can redefine the distance metric via

$$\mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j) = \langle \omega_{\mathcal{M}}(\mathbf{x}_i), \omega_{\mathcal{M}}(\mathbf{x}_j) \rangle \quad \text{or} \quad \mathbf{d}_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j) = (\omega_{\mathcal{M}}(\mathbf{x}_i) - \omega_{\mathcal{M}}(\mathbf{x}_j))^2,$$

Table 4: *Learned distance metric and examples of matched-groups on Lalonde Experimental treatment and Observational control datasets for two example query points drawn from the same datasets.* Query 1 represents a high quality (low diameter) matched group while Query 2 represents a poor quality (high diameter) matched group that could be discarded during analysis.

Stretch Matrix							
	Age	Education	Black	Hispanic	Married	No-Degree	Income-1975
mean($Diag(\mathcal{M})$)	0.780	1.786	1.254	1.110	1.205	1.229	1.001
std($Diag(\mathcal{M})$)	0.361	0.778	0.641	0.577	0.614	0.618	0.512

Two Matched Groups									
Unit-ID	Treated	Age	Education	Black	Hispanic	Married	No-Degree	Income-1975	Income-1978
Query-1: 1	Yes	22	9	No	Yes	No	Yes	\$0	\$3595
94	Yes	23	8	No	Yes	No	Yes	\$0	\$3881
330	No	22	8	No	Yes	No	Yes	\$0	\$9920
299	No	22	9	Yes	No	No	Yes	\$0	\$0
5	Yes	22	9	Yes	No	No	Yes	\$0	\$4056
82	Yes	21	9	Yes	No	No	Yes	\$0	\$0
416	No	22	9	Yes	No	No	Yes	\$0	\$12898
333	No	21	9	Yes	No	No	Yes	\$0	\$3343
292	Yes	20	9	Yes	No	No	Yes	\$0	\$8881
17	Yes	23	10	Yes	No	No	Yes	\$0	\$7693
116	Yes	24	10	Yes	No	No	Yes	\$0	\$0

Unit-ID	Treated	Age	Education	Black	Hispanic	Married	No-Degree	Income-1975	Income-1978
Query-2: 968	No	42	11	No	No	Yes	Yes	\$44758	\$54675
274	Yes	35	9	Yes	No	Yes	Yes	\$13830	\$12803
141	Yes	25	8	Yes	No	No	Yes	\$37431	\$2346
967	No	50	17	No	No	Yes	No	\$30435	\$25860
948	No	35	12	No	No	Yes	No	\$26854	\$29554
210	Yes	25	8	No	No	No	Yes	\$23096	\$6421
241	Yes	24	15	Yes	No	No	No	\$13008	\$14683
311	No	28	12	Yes	No	Yes	No	\$29009	\$10067
183	Yes	23	10	Yes	No	No	Yes	\$15709	\$5665
182	Yes	23	12	Yes	No	Yes	No	\$15079	\$10283

where $\omega_{\mathcal{M}}$ is a summary of relevant data features learned using a complex modeling framework.

In the future, the MALTS framework could be extended to deal with missing covariates, and can be adapted to instrumental variables.

Acknowledgments

We gratefully acknowledge funding from the National Science Foundation under grants III 1703431, CCF 1934964, IIS 2130250, IIS 2147061 (with Amazon), and CAREER DMS 2046880, and the National Institute of Health under grants NIDA DA054994 and R01EB025021. We also acknowledge funding from an Amazon Graduate fellowship.

REFERENCES

- A. Abadie and G. W. Imbens. Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1):235–267, 2006.
- A. Abadie and G. W. Imbens. Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics*, 29(1):1–11, 2011.
- A. Bellet and A. Habrard. Robustness and generalization for metric learning. *Neurocomputing*, 151:259–267, 2015.
- V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018.
- H. A. Chipman, E. I. George, and R. E. McCulloch. BART: Bayesian additive regression trees. *Annals of Applied Statistics*, pages 266–298, 2010.
- W. G. Cochran and D. B. Rubin. Controlling bias in observational studies: A review. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 417–446, 1973.
- R. H. Dehejia and S. Wahba. Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs. *Journal of the American Statistical Association*, 94(448):1053–1062, 1999.
- A. Dieng, Y. Liu, S. Roy, C. Rudin, and A. Volfovsky. Interpretable almost-exact matching for causal inference. *Proceedings of Machine Learning Research (Proceedings of AISTATS)*, 89:2445, 2019.
- V. Dorie, H. Chipman, R. McCulloch, A. Dadgar, R. C. Team, G. U. Draheim, M. Bosmans, C. Tournayre, M. Petch, R. de Lucena Valle, et al. Package ‘dbarts’. 2019.
- M. H. Farrell. Robust inference on average treatment effects with possibly more covariates than observations. *Journal of Econometrics*, 189(1):1–23, 2015.
- J. H. Friedman. Multivariate adaptive regression splines. *The Annals of Statistics*, pages 1–67, 1991.
- J. Goldberger, G. E. Hinton, S. T. Roweis, and R. R. Salakhutdinov. Neighbourhood components analysis. In *Advances in Neural Information Processing Systems*, pages 513–520, 2005.
- X. S. Gu and P. R. Rosenbaum. Comparison of multivariate matching methods: Structures, distances, and algorithms. *Journal of Computational and Graphical Statistics*, 2(4):405–420, 1993.
- P. R. Hahn, J. S. Murray, and C. M. Carvalho. Bayesian regression tree models for causal inference: regularization, confounding, and heterogeneous effects. *Bayesian Analysis*, 15(3), September 2020.
- B. B. Hansen. The prognostic analogue of the propensity score. *Biometrika*, 95(2):481–488, 2008.
- J. L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- D. E. Ho, K. Imai, G. King, and E. A. Stuart. MatchIt: Nonparametric preprocessing for parametric causal inference. *Journal of Statistical Software*, 42(8):1–28, 2011.
- S. M. Iacus, G. King, and G. Porro. Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, 20(1):1–24, 2012.
- G. W. Imbens. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and Statistics*, 86(1):4–29, 2004.
- H. Jiang. Non-asymptotic uniform rates of consistency for k-nn regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3999–4006, 2019.
- N. Kallus. A Framework for Optimal Matching for Causal Inference. In A. Singh and J. Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 372–381, Fort Lauderdale, FL, USA, 20–22 Apr 2017.

- L.-Z. Kara, A. Laksaci, M. Rachdi, and P. Vieu. Data-driven kNN estimation in nonparametric functional data analysis. *Journal of Multivariate Analysis*, 153:176–188, 2017.
- L. Keele and J. R. Zubizarreta. Optimal multilevel matching in clustered observational studies: A case study of the school voucher system in Chile. *Journal of the American Statistical Association*, 112(518):547–560, 2017.
- R. J. LaLonde. Evaluating the Econometric Evaluations of Training Programs with Experimental Data. *American Economic Review*, 76(4):604–620, September 1986.
- M. Morucci, V. Orlandi, S. Roy, C. Rudin, and A. Volfovsky. Adaptive hyper-box matching for interpretable individualized treatment effect estimation. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2020.
- M. Morucci, M. Noor-E-Alam, and C. Rudin. A robust approach to quantifying uncertainty in matching problems of causal inference. *INFORMS Journal on Data Science*, 2022. accepted.
- H. Parikh, C. Rudin, and A. Volfovsky. An application of matching after learning to stretch (MALTS) to the ACIC 2018 causal inference challenge data. *Observational Studies*, 5:118–130, 2019.
- H. Parikh, K. Hoffman, H. Sun, W. Ge, J. Jing, R. Amerineni, L. Liu, J. Sun, S. Zafar, A. Struck, et al. Why interpretable causal inference is important for high-stakes decision making for critically ill patients and how to do it. *arXiv preprint arXiv:2203.04920*, 2022.
- M. Resa and J. R. Zubizarreta. Evaluation of subset matching methods and forms of covariate balance. *Statistics in Medicine*, 2016.
- P. R. Rosenbaum. Imposing minimax and quantile constraints on optimal matching in observational studies. *Journal of Computational and Graphical Statistics*, 26(1), 2017.
- P. R. Rosenbaum and D. B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- D. B. Rubin. Matching to remove bias in observational studies. *Biometrics*, pages 159–183, 1973a.
- D. B. Rubin. The use of matched sampling and regression adjustment to remove bias in observational studies. *Biometrics*, pages 185–203, 1973b.
- D. B. Rubin. Multivariate matching methods that are equal percent bias reducing, I: Some examples. *Biometrics*, pages 109–120, 1976.
- D. B. Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100:322–331, 2005.
- C. J. Stone. Consistent nonparametric regression. *The Annals of Statistics*, pages 595–620, 1977.
- E. A. Stuart. Matching methods for causal inference: A review and a look forward. *Statistical Science*, 25(1):1, 2010.
- T. Wang, M. Morucci, M. U. Awan, Y. Liu, S. Roy, C. Rudin, and A. Volfovsky. FLAME: A fast large-scale almost matching exactly approach to causal inference. *Journal of Machine Learning Research*, 22(31):1–41, 2021.
- K. Q. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10(2), 2009.
- K. Q. Weinberger, J. Blitzer, and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. In *Advances in Neural Information Processing Systems*, pages 1473–1480, 2006.
- H. Xu and S. Mannor. Robustness and generalization. *Machine Learning*, 86(3):391–423, 2012.
- Z. Zhao. Using matching to estimate treatment effects: Data requirements, matching metrics, and monte carlo evidence. *The Review of Economics and Statistics*, 86(1):91–107, 2004.
- J. R. Zubizarreta. Using mixed integer programming for matching in an observational study of kidney failure after surgery. *Journal of the American Statistical Association*, 107(500):1360–1371, 2012.

J. R. Zubizarreta, R. D. Paredes, and P. R. Rosenbaum. Matching for balance, pairing for heterogeneity in an observational study of the effectiveness of for-profit and not-for-profit high schools in Chile. *The Annals of Applied Statistics*, 8(1):204–231, 2014.

Appendix Appendix A.

In this section we provide proofs for theorems and lemmas discussed in Section 5.

Proof (Lemma 3). If (D_1, \dots, D_K) is the multinomially distributed random vector with parameters d and p_1, \dots, p_K then, by the Bretagnolle-Huber-Carol inequality,

$$P\left(\sum_{i=1}^K \left|\frac{D_i}{d} - p_i\right| \geq \lambda\right) \leq 2^K e^{-\frac{d\lambda^2}{2}}.$$

Thus, for our case, we can consider N_i corresponding to the set of indices of units in sample $\mathcal{S}_n^{(t')}$ such that their x 's are contained in the partition \mathbf{C}_i as in Theorem 2. Hence, by the Bretagnolle-Huber-Carol inequality, we know that

$$P\left(\sum_{i=1}^K \left|\frac{|N_i|}{n^{(t')}} - \mu(\mathbf{C}_i)\right| \geq \sqrt{\frac{2K \ln(2) + 2 \ln(1/\mathcal{E})}{n^{(t')}}}\right) \leq \mathcal{E}.$$

Now, for some arbitrary $t' \in \mathcal{T}$ let us consider $|L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')})|$.

We know that

$$\begin{aligned} & \left|L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')})\right| \\ &= \left| \sum_{i,j=1}^K \left(\mathbb{E}_{z_1, z_2}[\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1 = (\mathbf{x}'_1, y'_1, t'_1), z_2 = (\mathbf{x}'_2, y'_2, t'_2)) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \mu(\mathbf{C}_i)\mu(\mathbf{C}_j)\right) \right. \\ & \quad \left. - \frac{1}{(n^{(t')})^2} \sum_{s_1, s_2 \in \mathcal{S}_n^{(t')}} \text{loss}(\mathcal{M}(\mathcal{S}_n), s_1, s_2) \right| \\ &= \left| \sum_{i,j=1}^K \left(\mathbb{E}_{z_1, z_2}[\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \mu(\mathbf{C}_i)\mu(\mathbf{C}_j)\right) \right. \\ & \quad - \sum_{i,j=1}^K \left(\mathbb{E}_{z_1, z_2}[\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \mu(\mathbf{C}_i) \frac{|N_j|}{n^{(t')}}\right) \\ & \quad + \sum_{i,j=1}^K \left(\mathbb{E}_{z_1, z_2}[\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \mu(\mathbf{C}_i) \frac{|N_j|}{n^{(t')}}\right) \\ & \quad + \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2}[\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \frac{|N_i|}{n^{(t')}} \frac{|N_j|}{n^{(t')}} \\ & \quad \left. - \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2}[\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \frac{|N_i|}{n^{(t')}} \frac{|N_j|}{n^{(t')}} \right| \end{aligned}$$

$$\begin{aligned}
 & \left| -\frac{1}{(n^{(t')})^2} \sum_{s_1, s_2 \in \mathcal{S}_n^{(t')}} \text{loss}(\mathcal{M}(\mathcal{S}_n), s_1, s_2) \right| \\
 \leq & \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2} [\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \mu(\mathbf{C}_i) \left(\mu(\mathbf{C}_j) - \frac{|N_j|}{n^{(t')}} \right) \right| \\
 & + \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2} [\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \frac{|N_j|}{n^{(t')}} \left(\mu(\mathbf{C}_i) - \frac{|N_i|}{n^{(t')}} \right) \right| \\
 & + \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2} [\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \frac{|N_i|}{n^{(t')}} \frac{|N_j|}{n^{(t')}} \right| \\
 & \left| -\frac{1}{(n^{(t')})^2} \sum_{s_1, s_2 \in \mathcal{S}_n^{(t')}} \text{loss}(\mathcal{M}(\mathcal{S}_n), s_1, s_2) \right| \\
 \leq & 2B \sum_{i=1}^K \left| \frac{|N_i|}{n^{(t')}} - \mu(\mathbf{C}_i) \right| + \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2} [\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \frac{|N_i|}{n^{(t')}} \frac{|N_j|}{n^{(t')}} \right| \\
 & \left| -\frac{1}{(n^{(t')})^2} \sum_{s_1, s_2 \in \mathcal{S}_n^{(t')}} \text{loss}(\mathcal{M}(\mathcal{S}_n), s_1, s_2) \right| \text{ where } B \text{ is } \max_{z_1, z_2} \text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2).
 \end{aligned}$$

Recall, $\mathcal{M}(\mathcal{S}_n)$ is $(K, \epsilon(\cdot))$ -multi-robust with probability $p_{mr}(\epsilon)$. Thus,

$$\begin{aligned}
 & P \left(\left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2} [\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] \frac{|N_i|}{n^{(t')}} \frac{|N_j|}{n^{(t')}} \right| \leq \epsilon(\mathcal{S}_n^{(t')}) \right) \\
 & \geq \prod_{i,j} P \left(\left| \frac{|N_i| |N_j|}{(n^{(t')})^2} \left[\mathbb{E}_{z_1, z_2} [\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j] - \frac{1}{|N_i| |N_j|} \sum_{s_1, s_2 \in \mathcal{S}_n^{(t')}} \text{loss}(\mathcal{M}(\mathcal{S}_n), s_1, s_2) \right] \right| \leq \epsilon(\mathcal{S}_n^{(t')}) / K^2 \right) \\
 & \geq \prod_{i,j} P \left(\left| \frac{\mathbb{E}_{z_1, z_2} [\text{loss}(\mathcal{M}(\mathcal{S}_n), z_1, z_2) \mid \mathbf{x}'_1 \in \mathbf{C}_i, \mathbf{x}'_2 \in \mathbf{C}_j]}{1/|N_i| |N_j|} - \sum_{(s_1, s_2) \in (\mathbf{C}_i \times \mathbf{C}_j)} \text{loss}(\mathcal{M}(\mathcal{S}_n), s_1, s_2) \right| \leq \epsilon(\mathcal{S}_n^{(t')}) / K^2 \right) \\
 & \geq (p_{mr}(\epsilon/K^2))^{K^2}
 \end{aligned}$$

Hence, by combining the above results, we can conclude for all $t' \in \mathcal{T}$ we have

$$\begin{aligned}
 P_{\mathcal{S}_n} \left(\left| L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')}) \right| \geq \epsilon(\mathcal{S}_n^{(t')}) + 2B \sqrt{\frac{2K \ln(2) + 2 \ln(1/\mathcal{E})}{n^{(t')}}} \right) \\
 \leq 1 - (1 - \mathcal{E})(p_{mr}(\epsilon/K^2))^{K^2}.
 \end{aligned}$$

■

Lemma 5 (Used for proof of Theorem 1) Let $\{\mathcal{S}_n\}_{n=1}^\infty$ be a sequence of nested datasets, each of which includes n i.i.d. samples from $\mu(\mathcal{Z})$, $n = 1.. \infty$. Given a smooth distance metric $\mathbf{d}_{\mathcal{M}}$, covariate vector \mathbf{x} , and $\alpha > 0$, if there exists a small enough value of “ a ” and

a large enough value of N such that $\mathcal{K}_n^{(t')}(x) = \{z_k : d_{\mathcal{M}}(x_k, x) < a, t_k = t', z_k \in \mathcal{S}_n\}$ is non-empty and $\alpha > \delta_{d_{\mathcal{M}}}(a)$ for all $n \geq N$ then,

$$P(|E[Y^{(t')}|x] - \widehat{Y}_x^{(t')}| > \alpha) \leq \exp(-|\mathcal{K}_n^{(t')}(x)|(\alpha - \delta_{d_{\mathcal{M}}}(a))^2/2\mathbf{C}_y)$$

where $\delta_{d_{\mathcal{M}}}(a)$ is the bound from Definition 1 (definition of smooth distance metric). As the above choice of a holds for all $n \geq N$, we have that the bound goes to zero as $n \rightarrow \infty$.

Proof (Lemma 5). $\mathcal{K}_n^{(t')}(x)$ is a matched group of nearest neighbors z_k of unit z such that $d_{\mathcal{M}}(x, x_k) < a$ and treatment indicator $t_k = t'$, i.e.,

$$\mathcal{K}_n^{(t')}(x) = \{z_k : d_{\mathcal{M}}(x, x_k) < a \text{ and } t_k = t'\}.$$

We estimate the conditional average potential outcome for treatment choice t' and $X = x$ as $\frac{1}{|\mathcal{K}_n^{(t')}(x)|} \sum_{z_k \in \mathcal{K}_n^{(t')}(x)} Y_k$. If $d_{\mathcal{M}}$ is a smooth distance metric then as all the units in $\mathcal{K}_n^{(t')}(x)$ have distance to x less than a , for every z_k in $\mathcal{K}_n^{(t')}(x)$, we have $|\mathbb{E}[Y^{(t')}|X = x] - \mathbb{E}[Y^{(t')}|X = x_k]| < \delta_{d_{\mathcal{M}}}(a)$. Consider α such that $\delta_{d_{\mathcal{M}}}(a) < \alpha$ and $\alpha \leq \mathbf{C}_y$, then

$$\begin{aligned} & P_{\mathcal{S}_n \sim \mu(\mathcal{Z}^n)} \left(\left| \mathbb{E}[Y^{(t')}|X = x] - \frac{1}{|\mathcal{K}_n^{(t')}(x)|} \sum_{z_k \in \mathcal{K}_n^{(t')}(x)} Y_k \right| > \alpha \middle| \{\mathbf{X}_i\}, \{T_i\} \right) \\ &= P \left(\left| \frac{1}{|\mathcal{K}_n^{(t')}(x)|} \sum_{z_k \in \mathcal{K}_n^{(t')}(x)} (\mathbb{E}[Y^{(t')}|X = x] - Y_k) \right| > \alpha \middle| \{\mathbf{X}_i\}, \{T_i\} \right). \end{aligned}$$

Thus,

$$\begin{aligned} & P \left(\left| \frac{1}{|\mathcal{K}_n^{(t')}(x)|} \sum_{z_k \in \mathcal{K}_n^{(t')}(x)} (\mathbb{E}[Y^{(t')}|X = x] - Y_k) \right| > \alpha \middle| \{\mathbf{X}_i\}, \{T_i\} \right) \\ &= P \left(\left| \sum_{z_k \in \mathcal{K}_n^{(t')}(x)} \mathbb{E}[Y^{(t')}|X = x] - \mathbb{E}[Y^{(t')}|X = \mathbf{X}_k] + \mathbb{E}[Y^{(t')}|X = \mathbf{X}_k] - Y_k \right| > |\mathcal{K}| \alpha \middle| \{\mathbf{X}_i\}, \{T_i\} \right) \end{aligned}$$

and using the triangle inequality,

$$\begin{aligned} & \leq P \left(\sum_{z_k \in \mathcal{K}_n^{(t')}(x)} \left| \mathbb{E}[Y^{(t')}|X = x] - \mathbb{E}[Y^{(t')}|X = \mathbf{X}_k] \right| + \right. \\ & \quad \left. \left| \sum_{z_k \in \mathcal{K}_n^{(t')}(x)} \mathbb{E}[Y^{(t')}|X = \mathbf{X}_k] - Y_k \right| > |\mathcal{K}_n^{(t')}(x)| \alpha \middle| \{\mathbf{X}_i\}, \{T_i\} \right). \end{aligned}$$

By the definition of smooth distance metric,

$$\begin{aligned} & \leq P \left(\sum_{z_k \in \mathcal{K}_n^{(t')}(x)} \delta_{d_{\mathcal{M}}}(a) + \left| \sum_{z_k \in \mathcal{K}_n^{(t')}(x)} \mathbb{E}[Y^{(t')}|X = \mathbf{X}_k] - Y_k \right| > |\mathcal{K}_n^{(t')}(x)| \alpha \middle| \{\mathbf{X}_i\}, \{T_i\} \right) \\ &= P \left(\left| \sum_{z_k \in \mathcal{K}_n^{(t')}(x)} \mathbb{E}[Y^{(t')}|X = \mathbf{X}_k] - Y_k \right| > |\mathcal{K}_n^{(t')}(x)| (\alpha - \delta_{d_{\mathcal{M}}}(a)) \middle| \{\mathbf{X}_i\}, \{T_i\} \right). \end{aligned}$$

and by Hoeffding's inequality,

$$\leq 2 \exp\left(\frac{-2|\mathcal{K}_n^{(t')}(\mathbf{x})|(\alpha - \delta_{\mathbf{d}_{\mathcal{M}}}(a))^2}{4\mathbf{C}_y}\right) = 2 \exp\left(\frac{-|\mathcal{K}_n^{(t')}(\mathbf{x})|(\alpha - \delta_{\mathbf{d}_{\mathcal{M}}}(a))^2}{2\mathbf{C}_y}\right).$$

Thus,

$$\begin{aligned} P\left(\left|\mathbb{E}[Y^{(t')}|X = \mathbf{x}] - \frac{1}{|\mathcal{K}_n^{(t')}(\mathbf{x})|} \sum_{z_k \in \mathcal{K}_n^{(t')}(\mathbf{x})} Y_k\right| > \alpha \middle| \{\mathbf{X}_i\}, \{T_i\}\right) \\ \leq 2 \exp\left(\frac{-|\mathcal{K}_n^{(t')}(\mathbf{x})|(\alpha - \delta_{\mathbf{d}_{\mathcal{M}}}(a))^2}{2\mathbf{C}_y}\right). \end{aligned}$$

As $n \rightarrow \infty$, for a constant a (and hence constant $\delta_{\mathbf{d}_{\mathcal{M}}}(a)$), a constant α , and letting the number of units matched to the target unit go to infinity: $|\mathcal{K}_n^{(t')}(\mathbf{x})| \rightarrow \infty$, we have

$$\lim_{n \rightarrow \infty} P\left(\left|\mathbb{E}[Y^{(t')}|X = \mathbf{x}] - \frac{1}{|\mathcal{K}_n^{(t')}(\mathbf{x})|} \sum_{z_k \in \mathcal{K}_n^{(t')}(\mathbf{x})} Y_k\right| > \alpha \middle| \{\mathbf{X}_i\}, \{T_i\}\right) = 0.$$

■

Lemma 6 (Also used for proof of Theorem 1) *If we can estimate the conditional average potential outcomes using a finite sample $\mathcal{S}_n \stackrel{i.i.d.}{\sim} \mu(\mathcal{Z}^n)$ such that for all t' , $\mathbb{E}[Y^{(t')}|X = \mathbf{x}]$ and the estimate, $\hat{Y}_{\mathbf{x}}^{(t')}$ are farther than ϵ' with probability less than $\delta'(\epsilon', n)$ for any given $z = (\mathbf{x}, y, t) \in \mathcal{Z}$ and $t \in \mathcal{T}$, then the estimated conditional average treatment effect $\hat{\tau}(\mathbf{x})$ using a finite sample $\mathcal{S}_n \stackrel{i.i.d.}{\sim} \mu(\mathcal{Z}^n)$ and the true conditional average treatment effect $\tau(\mathbf{x})$ are farther than ϵ with probability less than $2\delta'(\frac{\epsilon}{2}, n)$.*

$$\begin{aligned} \forall t \in \mathcal{T}, \quad P_{\mathcal{S}_n \sim \mu(\mathcal{Z}^n)}\left(\left|\hat{Y}_{\mathbf{x}}^{(t')} - \mathbb{E}[Y^{(t')}|X = \mathbf{x}]\right| \geq \epsilon'\right) &\leq \delta'(\epsilon', n) \\ \implies P_{\mathcal{S}_n \sim \mu(\mathcal{Z}^n)}\left(\left|\hat{\tau}(\mathbf{x}) - \tau(\mathbf{x})\right| \geq \epsilon\right) &\leq 2\delta'\left(\frac{\epsilon}{2}, n\right). \end{aligned}$$

Proof (Lemma 6). We are given in the statement that for any $\epsilon' > 0$, we can find a $\delta'(\epsilon', n)$ such that we can estimate outcomes well, i.e.,

$$\forall z \in \mathcal{Z}, \forall t' \in \mathcal{T}, P_{\mathcal{S}_n \sim \mu(\mathcal{Z}^n)}\left(\left|\hat{Y}_{\mathbf{x}}^{(t')} - \mathbb{E}[Y^{(t')}|X = \mathbf{x}]\right| \geq \epsilon'\right) \leq \delta'(\epsilon', n).$$

We can further deduce from the union bound that

$$P\left(\bigvee_{t' \in \mathcal{T}} \left(\left|\hat{Y}_{\mathbf{x}}^{(t')} - \mathbb{E}[Y^{(t')}|X = \mathbf{x}]\right| \geq \epsilon'\right)\right) \leq |\mathcal{T}| \delta'(\epsilon', n). \quad (9)$$

By the triangle inequality, we also know that

$$\sum_{t' \in \mathcal{T}} \left|\hat{Y}_{\mathbf{x}}^{(t')} - \mathbb{E}[Y^{(t')}|X = \mathbf{x}]\right| \geq \left|\sum_{t' \in \mathcal{T}} \left(\hat{Y}_{\mathbf{x}}^{(t')} - \mathbb{E}[Y^{(t')}|X = \mathbf{x}]\right)\right|. \quad (10)$$

From Equation 9, we have

$$P\left(\sum_{t' \in \mathcal{T}} \left(|\hat{Y}_{\mathbf{x}}^{(t')} - \mathbb{E}[Y^{(t')}|X = \mathbf{x}]|\right) \geq |\mathcal{T}|\epsilon'\right) \leq |\mathcal{T}| \delta'(\epsilon', n).$$

Applying the triangle inequality from Equation 10,

$$P\left(\left|\sum_{t' \in \mathcal{T}} \left(\hat{Y}_{\mathbf{x}}^{(t')} - \mathbb{E}[Y^{(t')}|X = \mathbf{x}]\right)\right| \geq |\mathcal{T}|\epsilon'\right) \leq |\mathcal{T}| \delta'(\epsilon', n).$$

Considering the case where $\mathcal{T} = \{0, 1\}$,

$$P\left(\left|\hat{\tau}(\mathbf{x}) - \tau(\mathbf{x})\right| \geq 2\epsilon'\right) \leq 2\delta'(\epsilon', n).$$

Hence, we can conclude that

$$P\left(\left|\hat{\tau}(\mathbf{x}) - \tau(\mathbf{x})\right| \geq \epsilon\right) \leq 2\delta'\left(\frac{\epsilon}{2}, n\right).$$

Proof (Theorem 1). The proof of Theorem 1 follows directly by substituting the result of Lemma 5 into Lemma 6. ■

Proof (Theorem 4). By Theorem 2, we know that the distance metric $\mathcal{M}(\cdot)$ learned using MALTS is $(\mathbf{N}(\gamma, \mathcal{X}, \|\cdot\|_2), \beta)$ -multirobust with probability more than $1 - \exp\left(-\frac{\beta^2 \left(\rho_\gamma^{(t')}\right)^2}{n^{(t')} B^2}\right)$. ■

Also, inferring from Lemma 3, for any arbitrary $\forall t' \in \mathcal{T}$ and $\mathcal{E} > 0$ we have that with probability at least

$$(1 - \mathcal{E}) \left(1 - \exp\left(-\frac{\beta^2 \left(\rho_\gamma^{(t')}\right)^2}{K^2 n^{(t')} B^2}\right)\right)^{K^2}$$

with respect to the random draw of data to form \mathcal{S}_n , we have

$$\left|L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')})\right| \leq \beta + 2B \sqrt{\frac{2K \ln(2) + 2 \ln(1/\mathcal{E})}{n^{(t')}}}.$$

Let $\rho_\gamma = \min_{t'} \rho_\gamma^{(t')}$. Then, summing over all possible $t' \in \mathcal{T}$ we have that with probability at least:

$$(1 - \mathcal{E})^{|\mathcal{T}|} \left(1 - \exp\left(-\frac{\beta^2 \left(\rho_\gamma^{(t')}\right)^2}{K^2 n^{(t')} B^2}\right)\right)^{|\mathcal{T}|K^2}$$

with respect to the random draw of data to form \mathcal{S}_n , we have:

$$\sum_{t' \in \mathcal{T}} \left| L_{pop}(\mathcal{M}(\mathcal{S}_n), \mathcal{Z}^{(t')}) - L_{emp}(\mathcal{M}(\mathcal{S}_n), \mathcal{S}_n^{(t')}) \right| \leq 2|\mathcal{T}|\beta + \sum_{t' \in \mathcal{T}} 2B \sqrt{\frac{2K \ln(2) + 2 \ln(1/\mathcal{E})}{n^{(t')}}}.$$

γ in Theorem 2 was arbitrary, allowing us to take it to 0 in such a way that K increases at a rate smaller than $\min_{t'} n^{(t')}$ increases, and ρ_γ increases at rate faster than or equal to $O(n)$ to ∞ as n approaches ∞ . Thus we can reduce β^2 to 0 at a rate slower than $\frac{1}{\rho_\gamma^2}$. \mathcal{E} was also set arbitrarily, allowing us to take it to 0 slowly enough such that as $n \rightarrow \infty$, each of the $n^{(t')} \rightarrow \infty$ and we thus have:

$$\lim_{n \rightarrow \infty} \left(\sum_{t' \in \mathcal{T}} \left| L_{pop}(\mathcal{M}(\mathcal{S}_{tr}), \mathcal{Z}^{(t')}) - L_{emp}(\mathcal{M}(\mathcal{S}_{tr}), \mathcal{S}_{tr}^{(t')}) \right| \right) = 0. \quad (11)$$

■

Appendix Appendix B.

Limited Overlap and Performance

We use the DGP described in Section 6.1.1 with 2 relevant continuous covariates and no irrelevant covariates for the limited overlap experiments. Further, we set the parameters of the DGP as follows: $\mu = 1$, $\Sigma = 1.5$, $\phi = 0.5$, and $c = 2$. We performed experiments on the overlap by changing the standard deviation of the noise term ϵ_{treat} from 0.001 to 100 in the treatment assignment equation of the DGP and measured CATE estimation error for MALTS in comparison with other methods for each of the scenarios. A lower variance leads to small overlap, i.e., large standardized difference of means, whereas large variance creates small standardized difference of means and high overlap. Figure 10 shows the relationship of ϵ_{treat} and standardized difference of means between the covariate set of treated and control units. Figure 11 shows the performance of MALTS in comparison with other methods in predicting CATEs for multiple dataset sizes – $n \in \{500, 2000, 4000\}$ – with $p = 20$, and different levels of overlap. MALTS’ performance is largely insensitive to limited overlap, however, the performance deteriorates if the control and the treated units are very different. The primary reason for MALTS deterioration of performance under almost no-overlap compared to BART is because matching methods like MALTS can be conceptualized as *interpolation*, unlike regression approaches that explicitly models the potential outcomes’ surfaces, which is closer to *extrapolation*.

Sensitivity Analysis

We performed a *sensitivity analysis* of MALTS on a data generation setup with a constant treatment effect, two observed relevant covariates, and an unobserved confounder affecting the probability distributions of outcome as well as the choice of treatment. The unobserved confounder has a linear relationship with the outcome, with the value of the coefficient equal to the “sensitivity parameter of the outcome” (γ_Y) and the choice of treatment with the value of the coefficient equal to the “sensitivity parameter of the treatment” (γ_T).

$$x_{i,1}, x_{i,2}, u_i \stackrel{iid}{\sim} \mathcal{N}(0, 1)$$

$$\begin{aligned}
& \epsilon_{i,0}, \epsilon_{i,1} \stackrel{iid}{\sim} \mathcal{N}(0, 1) \\
y_i^{(0)} &= x_{i,1} + x_{i,2} + \gamma_Y u_i + \epsilon_{i,0} \\
y_i^{(1)} &= x_{i,1} + x_{i,2} + \gamma_Y u_i + 1 + \epsilon_{i,1} \\
t_i &= \text{Bernoulli}(\text{expit}(x_{i,1} + x_{i,2} + \gamma_T u_i - 2)) \\
y_i &= t_i y_i^{(1)} + (1 - t_i) y_i^{(0)}
\end{aligned}$$

Figure 12 shows the contour plot of ATE estimates produced by MALTS as we change the sensitivity parameters in the data generation process. Here, the true ATE is equal to 1. The plot indicates that as long as the unmeasured confounders are approximately half as important as either of the two observed covariates (γ_T or γ_Y is below 0.5), MALTS performance is stable. At the extreme, which is when the unobserved confounder is as important as the total of the two observed covariates (γ_T or γ_Y is 2), the performance (expectedly) degrades.

Appendix Appendix C.

In this section, we discuss our implementation of existing causal inference methods like genmatch, propensity score matching, BART, causal forest, difference of random forest, prognostic score matching and FLAME. In Section 6, we compare the performance of each of these methods with MALTS.

We used MatchIt’s implementation of genmatch and propensity score matching as it is commonly used by empiricists (Ho et al. 2011). We allowed matching with replacement for creating match groups and estimating CATEs. As MatchIt returns only match groups and CATE estimates for treated units (and not control units), then in order to estimate CATEs for control units we flipped the sign of the treatment indicators and estimated negative CATEs (we have to estimate negative CATE in this case because we flipped the sign of the treatment indicator causing CATE estimates to become negative CATE estimates). We merged the CATE estimates for the treated units and control units to get the CATE estimates for every unit in the dataset.

We used the causal forest algorithm as implemented in the ‘grf’ package in R. The settings for causal forest were set to the default designed by the ‘grf’ developer with number of trees equal to 2000 and $\sqrt{p} + 20$ variables tried for each split.

We performed the same 5-fold CATE estimation procedure for causal forest, analogous to the one used for estimating CATEs using MALTS. We estimated CATEs for both the treated and control units in each estimation set.

We used Vincent Dorie’s R implementation of BART (Dorie et al. 2019). We performed the same 5-fold CATE estimation using BART that we used for MALTS. For each of the η folds, we trained two BART models, one for learning the response function for estimating the potential outcome under control and the other response function for estimating potential outcome under treatment using the training set. The CATEs were estimated by taking the difference of estimated response functions of treated and control units in the estimation set.

We also implemented a 5-fold FLAME CATE estimation procedure analogous to the one used by MALTS.

Lastly, we implemented 5-fold prognostic score matching using a random forest approach to model the prognostic score function. We fit a model for control units and a model for treated units using the data in the training set. To estimate the CATE for a treated unit in the estimation set, we found k-nearest neighbors in the control set with a similar estimated prognostic score. Analogously, we estimated the CATEs for the control units in the estimation set using the k-nearest treated units with similarity measured using the prognostic score.

Appendix Appendix D.

In this section, we show expanded matched groups (including all treatment and control units, not just control units) using propensity score matching, prognostic score matching and MALTS for unit id-1 from Table 5. MALTS and prognostic score matching are implemented as described in Appendix C with $K = 10$ nearest neighbors, where K was selected by cross-validation for MALTS and used also for the other methods. Propensity score matching was implemented using the MatchIt package and operationalized with 10 nearest neighbor matching.

Here, MALTS produces a high quality matched group, as shown in Table 4. While MALTS matches units based on a learned distance metric over the covariate space, prognostic score matching matches units based on a single prognostic value and propensity score matching matches units based on a single propensity value. We note that there is only one matched unit in common between the matched groups from prognostic score matching and MALTS (unit 116) and there are two units in common between the matched groups from propensity score matching and MALTS (units 330 and 416); the matched groups are almost entirely different between the three methods.

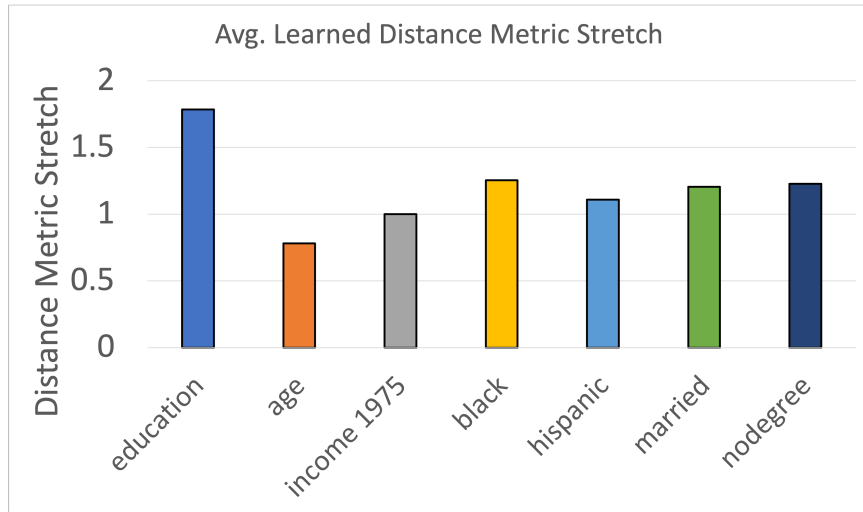
Appendix Appendix E.

In this section, we study the computation time required by MALTS to learn an optimal distance metric and estimate CATEs using matching. We study this by increasing the number of covariates from 8 to 136 (where the number of relevant covariates is 8 and others are irrelevant) keeping the number of samples constant at 2048. We compare the performance of MALTS for unrestricted stretch matrices (referred as ‘full \mathcal{M} ’) with a case when the stretch matrix is restricted to diagonal matrices (referred as ‘diagonal \mathcal{M} ’). Figure 13 shows that the difference in the computational time between the two cases increases approximately quadratically as the number of covariates increases. This is because inverting full \mathcal{M} is $O(p^2)$ more costly than inverting a diagonal \mathcal{M} .

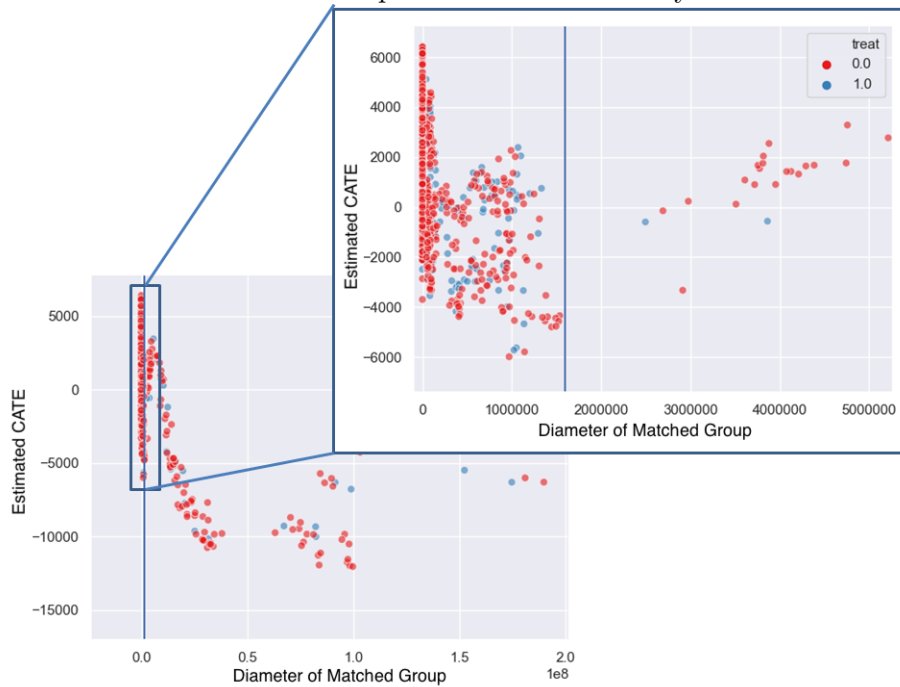
MALTS

Table 5: Example Matched Group using (a) our approach, (b) prognostic score (Hansen 2008), and (c) propensity score matching (Rosenbaum and Rubin 1983) for a query unit in the Lalonde dataset (top rows). It matched closely on almost all covariates such as age, education, marital status, whether the person had an academic degree, and income in 1975. In contrast, prognostic and propensity scores did not match closely on factors such as education, age, marital status and income. Bold is used to denote disagreement between the query unit and its matched group.

Treatment		Covariates							Outcome	
<i>Unit ID</i>	Treated	Age	Education	Black	Hispanic	Married	No-Degree	Income-1975	Income-1978	
Query: 1	Yes	22	9	No	Yes	No	Yes	\$0	\$3596	
(a) Our Approach (MALTS)										
94	Yes	23	8	No	Yes	No	Yes	\$0	\$3881	
330	No	22	8	No	Yes	No	Yes	\$0	\$9921	
299	No	22	9	Yes	No	No	Yes	\$0	\$0	
5	Yes	22	9	Yes	No	No	Yes	\$0	\$4056	
82	Yes	21	9	Yes	No	No	Yes	\$0	\$0	
416	No	22	9	Yes	No	No	Yes	\$0	\$12898	
333	No	21	9	Yes	No	No	Yes	\$0	\$3343	
292	Yes	20	9	Yes	No	No	Yes	\$0	\$8882	
17	Yes	23	10	Yes	No	No	Yes	\$0	\$7693	
116	Yes	24	10	Yes	No	No	Yes	\$0	\$0	
(b) Prognostic Scores										
154	Yes	22	10	Yes	No	No	Yes	\$1071	\$7315	
56	Yes	30	11	Yes	No	Yes	Yes	\$0	\$591	
100	Yes	17	10	Yes	No	No	Yes	\$0	\$0	
109	Yes	18	9	Yes	No	No	Yes	\$0	\$4483	
141	Yes	25	8	Yes	No	No	Yes	\$37432	\$2347	
286	Yes	23	12	No	Yes	No	No	\$1117	\$559	
338	No	44	9	Yes	No	No	Yes	\$0	\$9722	
340	No	22	12	Yes	No	No	No	\$532	\$1333	
355	No	18	10	No	Yes	No	Yes	\$0	\$1859	
116	Yes	24	10	Yes	No	No	Yes	\$0	\$0	
(c) Propensity Scores										
416	No	22	9	Yes	No	No	Yes	\$0	\$12898	
451	No	22	8	Yes	No	No	Yes	\$0	\$1391	
330	No	22	8	No	Yes	No	Yes	\$0	\$9921	
407	No	20	12	Yes	No	No	No	\$1371	\$20893	
626	No	18	10	Yes	No	No	Yes	\$2682	\$0	
774	No	21	13	No	No	No	No	\$693	\$2660	
402	No	22	11	Yes	No	Yes	Yes	\$0	\$1698	
925	No	21	12	Yes	No	No	No	\$716	\$22166	
879	No	22	12	Yes	No	Yes	No	\$0	\$665	
788	No	22	12	Yes	No	Yes	No	\$0	\$0	



(a) Distance metric learned on Lalonde data across 250 folds – 50 repeats and 5 splits within each repeat. Here, on an *average*, education is stretched more than other variables, which means it is more important to match closely on education.



(b) Criteria for pruning low-quality matched groups with large diameter.

Figure 9: (a) Box-plot of distance metric stretch values corresponding to each covariate in Lalonde data learned over 5 folds and 20 repeats. (b) Criteria to prune low-quality matched groups with large diameter from Lalonde data.

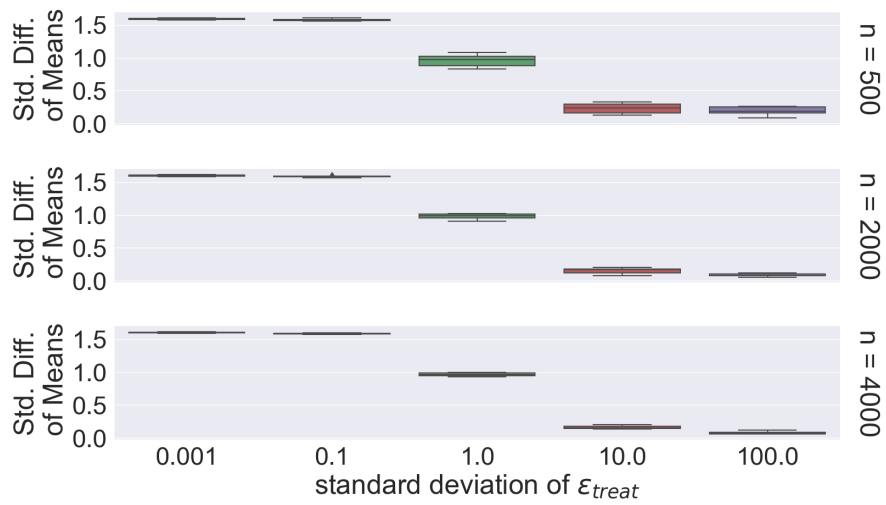
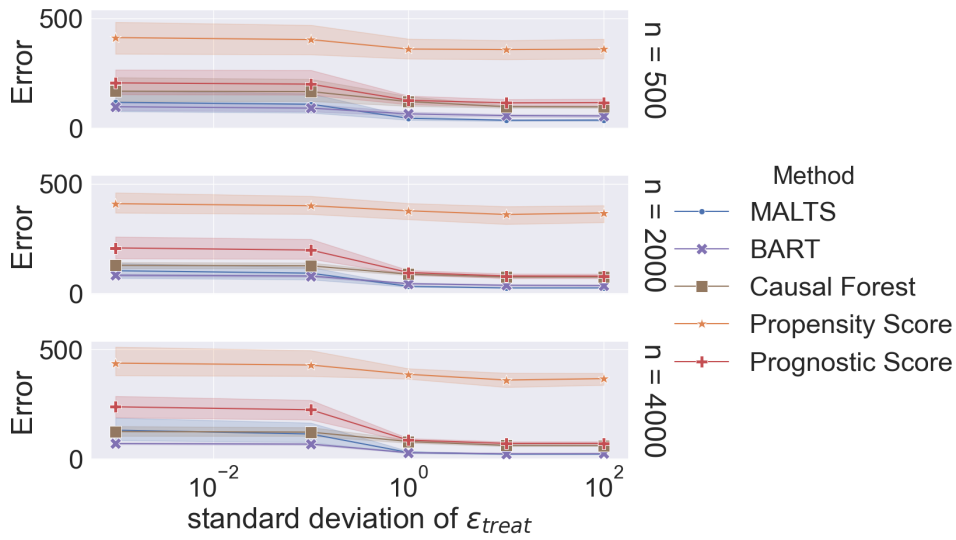
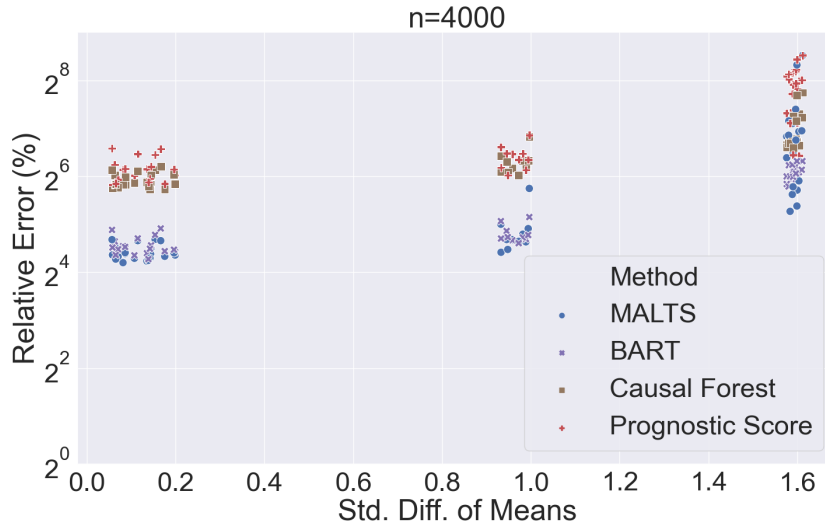


Figure 10: Standardized difference of means between covariates of treated and control units decreases as $\text{std}(\epsilon_{treat})$ increases. We increase the value of $\text{std}(\epsilon_{treat})$ in the DGP for treatment allocation which increases the overlap in the treated and control groups. We generate data with p equal to 20 covariates and for values of $n \in \{500, 2000, 4000\}$.



(a)



(b)

Figure 11: (a) Trend plot comparison of MALTS performance measured as mean relative error for CATE estimation under different levels of overlap, measured as a function of standard deviation of ϵ_{treat} (the scale of noise in treatment allocation process). Higher values of the standard deviation of ϵ_{treat} corresponds to more overlap between the control and the treated groups. (b) Scatterplot comparing MALTS' performance, as mean CATE estimation error, under different levels of overlap. Overlap is measured as standardized difference of means for $n = 4000$. Larger values of standardized difference of means corresponds to less overlap between the control and treated groups.

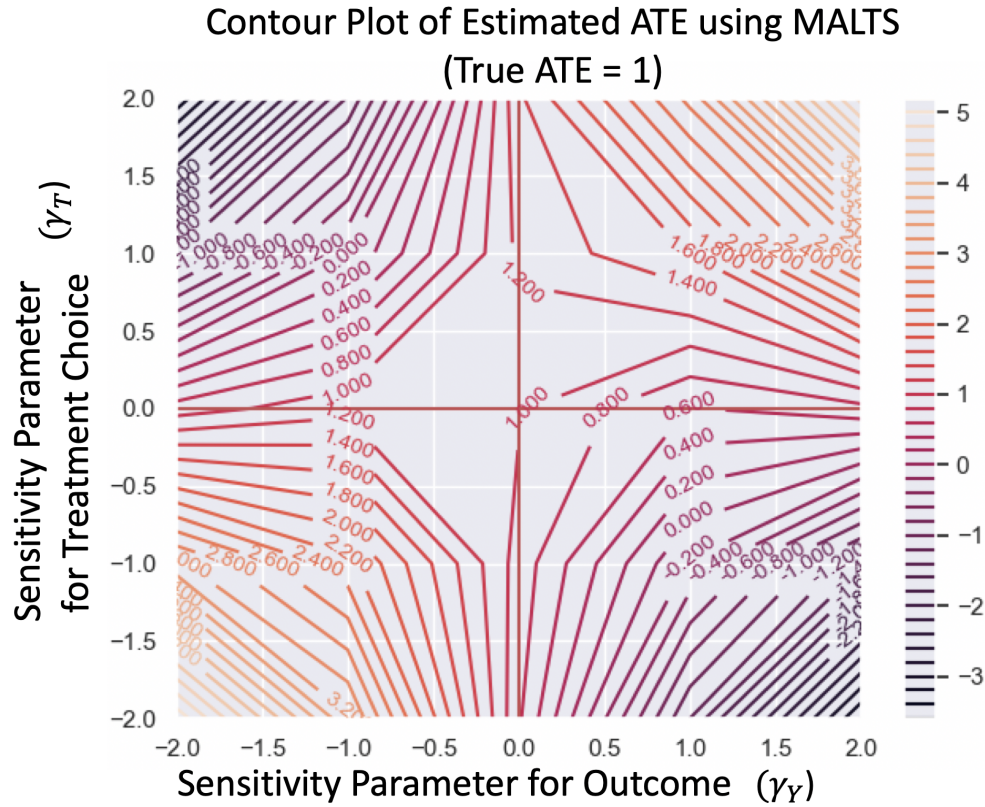


Figure 12: Sensitivity analysis contour plot of the ATE estimation using MALTS. The best performance is when there is no unmeasured confounding, which is when both sensitivity parameters are 0. When the sensitivity parameter for the outcome is 2, the unmeasured confounder is approximately as important as the total of the two known covariates and performance degrades.

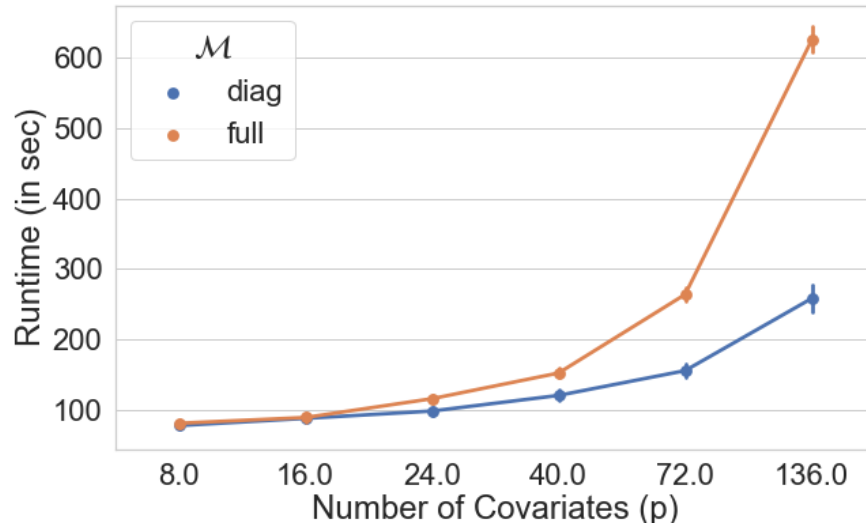


Figure 13: Run time for MALTS when distance metric is constrained to have diagonal \mathcal{M} , compared with distance metric where \mathcal{M} is a full-rank positive semi-definite matrix.