

Causal Effect Estimation Using Random Hyperplane Tessellations

Abhishek Dalvi

Neil Ashtekar

Vasant Honavar

The Pennsylvania State University, University Park, PA, USA.

ADALVI@PSU.EDU

AN2@SAMPLE.COM

AN3@SAMPLE.COM

Editors: Francesco Locatello and Vanessa Didelez

Abstract

Matching is one of the simplest approaches for estimating causal effects from observational data. Matching techniques compare the observed outcomes across pairs of individuals with similar covariate values but different treatment statuses in order to estimate causal effects. However, traditional matching techniques are unreliable given high-dimensional covariates due to the infamous curse of dimensionality. To overcome this challenge, we propose a simple, fast, yet highly effective approach to matching using Random Hyperplane Tessellations (RHPT). First, we prove that the RHPT representation is an approximate balancing score – thus maintaining the strong ignorability assumption – and provide empirical evidence for this claim. Second, we report results of extensive experiments showing that matching using RHPT outperforms traditional matching techniques and is competitive with state-of-the-art deep learning methods for causal effect estimation. In addition, RHPT avoids the need for computationally expensive training of deep neural networks.

Keywords: Matching, Randomized algorithms, High-Dimensional data

1. Introduction

Randomized experiments, whenever feasible, provide the most reliable estimates of causal effects. In a randomized experiment, each individual is randomly assigned to either the treatment or control group, thereby ensuring covariate balance: the two groups will be only *randomly different* from each other with respect to all covariates, both observed and unobserved. In this case, it is straightforward to show that the average treatment effect $ATE = \mathbb{E}[Y|T = 1] - \mathbb{E}[Y|T = 0]$, where Y is the observed outcome and T is the treatment status. In other words, the ATE is simply the difference in average *observed* outcomes across the treatment and control groups. In the case of a randomized experiment, association is causation.

However, randomized experiments are not always feasible due to their high costs, ethical concerns, or other fundamental limitations. As such, there is high interest in methods capable of reliably estimating causal effects from observational data. The following assumptions suffice for estimating causal effects from observational data: *Ignorability*, *Positivity*, *Non-interference* and *Consistency* (Rosenbaum and Rubin, 1983; Cox, 1958; Rubin, 1980). Under these conditions, it can be shown that:

$$ATE = \mathbb{E}_{\mathbf{X}}[\mathbb{E}[Y|T = 1, \mathbf{X} = \mathbf{x}] - \mathbb{E}[Y|T = 0, \mathbf{X} = \mathbf{x}]]$$

In observational data the treatment assignment is not randomized, it is possible that some of the covariates influence both the treatment assignment and the outcome; thus, introducing confounding bias while estimating the causal effect. In the presence of confounding, causation is *not* association. One way to eliminate confounding bias under the standard assumptions is to *simulate* a randomized experiment such that the treatment and control groups have similar covariate distributions.

This can be accomplished through matching individuals with similar covariates across the treated and control groups, thereby reducing bias in treatment assignment (Hernan and Robins, 2023; Stuart, 2010). Note that matching on the observed covariates also matches on the unobserved covariates, in so much as they are correlated with those that are observed¹.

Comparing the observed outcome for an individual with covariates x_i with treatment status $T_i = 1$ against the individual they are matched with, $m(x_i)$ with treatment status $T_{m(x_i)} = 0$ yields an estimate for the individual treatment effect $\text{ITE}(x_i)$. Matching techniques have been extensively studied in the literature on causal inference – see Stuart (2010) for a review.

Because perfect matches can seldom be found, matching in practice entails nearest neighbor search. However, matching techniques typically perform poorly when the number of covariates is large, since all points tend to appear dissimilar in high-dimensional spaces. This phenomenon is known as the curse of dimensionality (Beyer et al., 1999). Empirical results from Beyer et al. (1999); Aggarwal et al. (2001) suggest that distance metrics become unreliable given as few as 10-20 dimensions.

A simple, popular workaround is to match individuals based on their propensity scores – an individual’s probability of being treated – and to rely on the theoretical result that once conditioned on the propensity scores, treatment becomes independent of the covariates, thereby satisfying the identifiability conditions. In practice, matching on propensity scores seldom achieves its intended objective (King and Nielsen, 2019). Recently there has been increased interest in deep neural networks that learn low-dimensional, “treatment-agnostic” representations of high-dimensional data using non-linear mappings (Johansson et al., 2016; Shalit et al., 2017; Louizos et al., 2017). While deep learning approaches offer strong performance, they require large amounts of training data, incur substantial computational overhead for training, and require expensive hyper-parameter tuning.

In this work, we explore Random Hyperplane Tessellations (RHPT) an alternative approach which overcomes the limitations of both existing matching techniques as well as deep neural networks for causal effect estimation. RHPT involves mapping covariates to a high-dimensional, binary space based on the relationship between covariate values and randomly chosen hyperplanes (Plan and Vershynin, 2014; Dirksen et al., 2022). The RHPT representation is attractive for causal effect estimation because it offers a compromise between information-poor (and potentially inaccurate) predicted propensity scores and information-rich (but potentially noisy) matching on covariates. RHPT can be viewed as a denoising, information-reduction step performed prior to matching.

The key contributions of this paper are as follows:

- We present a novel application of Random Hyperplane Tessellations for estimating causal effects from observational data and explain via theory why it works.
- We show that our method preserves strong ignorability of treatment assignments as our RHPT representations are approximate balancing scores.
- We present experiments comparing treatment effect estimates obtained using our proposed method to estimates obtained using traditional matching approaches as well as state-of-the-art deep learning techniques. The results of these experiments on several standard benchmark data sets show that our proposed method outperforms traditional matching approaches and is

1. Bias introduced by unobserved covariates that are uncorrelated with the observed covariates is assessed using sensitivity analyses (Hernan and Robins, 2023).

competitive with the state-of-the-art deep neural networks in terms of accuracy while avoiding the substantial computational cost of training deep neural networks.

2. Preliminaries

2.1. Problem Definition

We are given data on n individuals each represented as a tuple (\mathbf{x}_i, T_i, Y_i) . \mathbf{x}_i denotes the (pre-treatment) covariates, T_i denotes treatment status ($T_i = 1$ indicating treatment and $T_i = 0$ indicating control), and Y_i denotes the observed outcome. Formally, $\mathbf{x}_i \in \mathcal{X}$ with $\mathcal{X} \subseteq \mathbb{R}^d$, where d is the dimensionality of the covariates, $T_i \in \{0, 1\}$, and $Y_i \in \mathcal{Y}$ with $\mathcal{Y} \subseteq \mathbb{R}$. Our goal is to estimate the ATE across all individuals and the ITE_i for each individuals i under the strong ignorability, consistency, and non-interference assumptions.

2.2. Locality Sensitive hashing

Locality-Sensitive Hashing (LSH) (Indyk and Motwani, 1998) was one of the first works related to hyperplane tessellations and is a widely used approach for approximately finding nearest neighbors in high-dimensional spaces. A locality sensitive hashing scheme is a distribution on a family \mathcal{H} of hash functions² h operating on a collection of objects, such that for two objects \mathbf{x}_i and \mathbf{x}_j , $\Pr_{h \in \mathcal{H}}[h(\mathbf{x}_i) = h(\mathbf{x}_j)] = \text{sim}(\mathbf{x}_i, \mathbf{x}_j) \in [0, 1]$, where $\text{sim}(\mathbf{x}_i, \mathbf{x}_j)$ is a similarity function. Such a scheme leads to a compact representation of objects such that the similarity of objects can be estimated from their compact sketches.

Given covariates $\mathbf{x}_i \in \mathbb{R}^d$, SimHash (Charikar, 2002) generates a random hyperplane $\mathbf{a}^\perp \in \mathbb{R}^d$ using a standard Gaussian random vector $\mathbf{a} \sim \mathcal{N}(0, I_d)$. Using random vector \mathbf{a} and function $\text{sign}(\cdot) : \mathbb{R} \rightarrow \{0, 1\}$, we generate a hash value for \mathbf{x}_i :

$$h_{\mathbf{a}}(\mathbf{x}_i) = \text{sign}(\mathbf{a} \cdot \mathbf{x}_i)$$

with $\Pr[h_{\mathbf{r}}(\mathbf{x}_i) = h_{\mathbf{r}}(\mathbf{x}_j)] = 1 - \frac{\theta_{ij}}{\pi}$ where θ_{ij} is the angle between \mathbf{x}_i and \mathbf{x}_j . A binary sketch can be obtained by performing hashing β times. Let $\mathbf{A} : \mathbb{R}^d \rightarrow \mathbb{R}^\beta$ be a matrix with each row $\mathbf{a}_1^T, \dots, \mathbf{a}_\beta^T$ drawn from $\mathcal{N}(0, I_d)$. For conciseness, denote $\mathbf{A} \sim \mathcal{N}^\beta(0, I_d)$ and apply the sign function elementwise, i.e. $\text{sign}(\cdot) : \mathbb{R}^\beta \rightarrow \{0, 1\}^\beta$. The binary sketch obtained using SimHash is $h_{\mathbf{A}}(\mathbf{x}_i) = \text{sign}(\mathbf{A}\mathbf{x}_i)$.

2.3. Random Hyper-Plane Tessellations (RHPT)

Plan and Vershynin (2014) showed that Simhash results in a representation where angular distances are preserved with error $\leq \delta$; given that the points are subset of a unit Euclidean sphere S^{d-1} . Formally, we are given $\mathcal{X} \subseteq S^{d-1}$, positive constants c_1 and c_2 , $\ell_*(\mathcal{X}) = \mathbb{E} \sup_{\mathbf{x} \in \mathcal{X}} |\langle \mathbf{a}, \mathbf{x} \rangle|$, and $\delta > 0$.

If the following holds:

$$\beta \geq c_1 \frac{\ell_*^2(\mathcal{X})}{\delta^6}$$

2. A hash function maps input samples to fixed-sized output hash codes. We consider Locality-sensitive hashing (LSH); a fuzzy hashing technique that hashes similar input items into the same to similar hash codes.

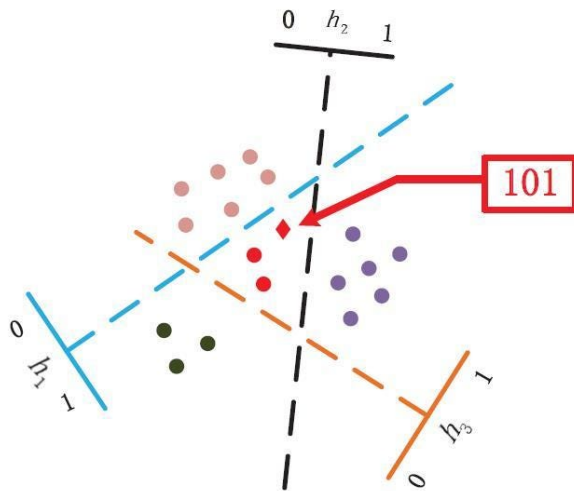


Figure 1: An example of a Random Hyperplane Tessellation (RHPT) for a set of points in a two-dimensional space from Poli (2021). Three hyperplanes – indicated by dashed lines – are used to construct a three-dimensional binary hash code for each point. The color of each point represents its binary hash code: red represents 101, black represents 100, etc.

then with a probability of at least $1 - 2 \exp(-c_2 \beta \delta^2)$, for every $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$:

$$\left| \frac{1}{\beta} d_H[\text{sign}(\mathbf{A}\mathbf{x}_i), \text{sign}(\mathbf{A}\mathbf{x}_j)] - d_{S^{d-1}}[\mathbf{x}_i, \mathbf{x}_j] \right| \leq \delta \quad (1)$$

where d_H is the Hamming distance and $d_{S^{d-1}}[\mathbf{x}_i, \mathbf{x}_j] = \frac{1}{\pi} \arccos \frac{\langle \mathbf{x}_i, \mathbf{x}_j \rangle}{\|\mathbf{x}_i\|_2 \|\mathbf{x}_j\|_2}$ is the angular distance. This is consistent with the findings in Charikar (2002); Pratap et al. (2020) which state that:

$$\cos(\theta_{ij}) \approx \cos\left(\frac{\pi}{\beta} d_H[\text{sign}(\mathbf{A}\mathbf{x}_i), \text{sign}(\mathbf{A}\mathbf{x}_j)]\right)$$

To summarize, $\delta \rightarrow 0$ with probability $\rightarrow 1$ when $\beta \rightarrow \infty$. In other words, distances are approximately preserved as the dimensionality of the embedding is increased.

Dirksen et al. (2022) extend the results from Plan and Vershynin (2014) beyond the unit Euclidean sphere S^{d-1} to an arbitrary set \mathbb{R}^d by introducing shifts on the hyperplane represented by γ . Let $\mathbf{A} \sim \mathcal{N}^\beta(0, I_d)$ and γ be uniformly distributed on $[-\lambda, \lambda]$. For conciseness, we represent β samples of γ as $\Gamma \sim [-\lambda, \lambda]^\beta$. Using \mathbf{A} and Γ , a binary sketch $h_{\mathbf{A}, \Gamma}(\mathbf{x}_i)$ is obtained where $h_{\mathbf{A}, \Gamma}(\mathbf{x}_i) = \text{sign}(\mathbf{A}\mathbf{x}_i + \Gamma)$. Given sufficiently large values of λ and β , the following holds probabilistically:

$$\sup_{\mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}} \left| \frac{\sqrt{2\pi}\lambda}{\beta} d_H[h_{\mathbf{A}, \Gamma}(\mathbf{x}_i), h_{\mathbf{A}, \Gamma}(\mathbf{x}_j)] - \|\mathbf{x}_i - \mathbf{x}_j\|_2 \right| \leq \delta \quad (2)$$

See Theorem 1.7 of Dirksen et al. (2022) for details. Similar to inequality 1, distances are preserved as the dimensionality of the embedding increases, assuming an appropriate choice of hyperplane

shift. Note that [Plan and Vershynin \(2014\)](#) showed only an existence of such hyperplane over arbitrary sets in \mathbb{R}^d but [Dirksen et al. \(2022\)](#) showed an estimate with a bound.

2.4. Controlling for Confounding Covariates

As noted earlier, matching offers a way to reduce the bias introduced by confounding when estimating causal effects from observational data ([Rubin, 1980](#)). By comparing the outcomes of pairs of treated and untreated individuals that are nearly identical (i.e. matched), we can eliminate any bias arising from confounding due to any observed covariates (as well as any unobserved covariates that are sufficiently correlated with the observed covariates). The estimated individual treatment effect $\widehat{\text{ITE}}(\mathbf{x}_i)$ for individual i is given by:

$$\widehat{\text{ITE}}(\mathbf{x}_i) = \begin{cases} Y_i - Y_{\text{Match}_0(\mathbf{x}_i)} & \text{if } T_i = 1 \\ Y_{\text{Match}_1(\mathbf{x}_i)} - Y_i & \text{if } T_i = 0 \end{cases}$$

where $\text{Match}_0(\mathbf{x}_i)$ and $\text{Match}_1(\mathbf{x}_i)$ denote the indices of untreated and treated individuals that match \mathbf{x}_i . The estimated unobserved outcome is denoted by $Y_{\text{Match}_T(\mathbf{x}_i)}$. If the covariates are perfectly matched, then the ignorability condition clearly holds i.e. $(Y_i^{T=0}, Y_i^{T=1}) \perp\!\!\!\perp T_i | \mathbf{x}_i$ ([Rosenbaum and Rubin, 1983](#)).

In practice, matches are seldom perfect and we must settle for nearest neighbor matching. In the context of this paper, matching is done on the β -dimensional hashed sketches of the covariate vectors, i.e., their RHPT representations.

3. Causal Effect Estimation Using Random Hyperplane Tessellations

We project the original d -dimensional covariates \mathbf{x}_i into the Hamming space to get our β -dimensional hashed sketch. In our experiments, we use the [Dirksen et al. \(2022\)](#) and [Plan and Vershynin \(2014\)](#) methods in order to approximately preserve both the angle and the Euclidean distance between vectors. We concatenate the results from both methods: $h_{\mathbf{A}, \Gamma}(\mathbf{x}_i) = \text{sign}(\mathbf{A}\mathbf{x}_i + \Gamma)$ and $h_{\mathbf{A}}(\mathbf{x}_i) = \text{sign}(\mathbf{A}\mathbf{x}_i)$. For brevity, from this point onward the concatenated embedding is represented by $h_{\beta}(\mathbf{x}_i)$. To summarize, we first use RHPT to obtain a binary representation of the covariates, and then use nearest-neighbor matching in the binary space in order to estimate causal effects. Specifically, we use 1-Nearest Neighbor (1-NN) matching with replacement and use Hamming distance to calculate nearest neighbors.

Matching is a form of covariate adjustment, and covariate adjustment eliminates confounding under the standard assumptions stated in Section 2. To justify matching using RHPT representations of covariate vectors, it suffices to show that the RHPT embeddings achieve *covariate balance*. We discuss this property of the RHPT embedding in the following section.

3.1. Approximate Balancing Scores using RHPT

A function $b(\mathbf{x}_i)$ is a balancing score if there exists a mapping from $b(\mathbf{x}_i)$ to the propensity score $e(\mathbf{x}_i) = \mathbb{E}[T_i | \mathbf{x}_i]$; $\forall \mathbf{x}_i \in \mathcal{X}$. Let this mapping be $g(\cdot)$, where $g(b(\mathbf{x}_i)) = e(\mathbf{x}_i)$. Matching can be performed on the balancing scores rather than the raw covariates as implied by Theorems 2 and 3 from [Rosenbaum and Rubin \(1983\)](#), i.e. $\mathbf{x}_i \perp\!\!\!\perp T_i | b(\mathbf{x}_i)$ and $Y_i(1), Y_i(0) \perp\!\!\!\perp T_i | b(\mathbf{x}_i)$, assuming the *Ignorability*, *Positivity*, *Non-interference* and *Consistency* assumptions hold on the original data.

If the embedding $b(x_i)$ is a one-to-one mapping from \mathcal{X} , and there is a mapping from the original covariates to the propensity scores, then there must exist a mapping from the embedding to the propensity scores $g(b(\mathbf{x}_i)) = e(\mathbf{x}_i)$. Because the function $b(\cdot)$ is one-to-one mapping, an inverse function exists, therefore, we can express the function $g(b(\mathbf{x}_i))$ as $e(b^{-1}(b(\mathbf{x}_i)))$.

Unfortunately, for our method which uses RHPT embeddings of covariate vectors, we cannot guarantee one-to-one mapping using $h_\beta(\cdot) : \mathcal{X} \rightarrow \{0, 1\}^\beta$. Let $\mathbf{x}_i \in \mathcal{X}$ and $\mathbf{x}_j \in \mathcal{X}$, be two covariate vector that are mapped to the same binary representation using RHPT. In this case, the Hamming distance $d_H[h_\beta(\mathbf{x}_i), h_\beta(\mathbf{x}_j)] = 0$. Substituting this term into inequalities 1 and 2, we get:

$$d_{S^{d-1}}[\mathbf{x}_i, \mathbf{x}_j] \leq \delta \quad \text{and} \quad \|\mathbf{x}_i - \mathbf{x}_j\|_2 \leq \delta$$

Therefore, two covariate vectors that are within distance δ (angular and L2) from each other may be mapped to the same Hamming code. Because $h_\beta(\cdot)$ is not a one-to-one mapping, we cannot exactly recover these covariate vectors by ‘‘inverting’’ their representations. Particularly, if $d_H[h_\beta(\mathbf{x}_i), h_\beta(\mathbf{x}_j)] = 0$, then the ‘‘reconstruction’’ $h^{-1}(h_\beta(\mathbf{x}_i))$ could be \mathbf{x}_i or \mathbf{x}_j . For simplicity we say that the reconstruction maps to only one of the arbitrarily chosen points with the same binary code (i.e. either \mathbf{x}_i or \mathbf{x}_j).

The results in section 2 state that as β increases, $\delta \rightarrow 0$ with high probability. This means that if the Hamming codes are high-dimensional, and the Hamming codes of two points are the same, then the original points will be similar with high probability. In such a scenario, the propensity scores predicted using the reconstruction will be approximately equal if the propensity score function is assumed to be *smooth*. Formally, $e(h^{-1}(h_\beta(\mathbf{x}_i))) = e(h^{-1}(h_\beta(\mathbf{x}_j))) \approx e(\mathbf{x}_i) \approx e(\mathbf{x}_j)$ if $e(\cdot)$ is a smooth function i.e. $e(\mathbf{x}_i) \approx e(\mathbf{x}_i \pm \epsilon)$, where ϵ is a small perturbation close to 0. It is reasonable to assume that the propensity score function of real-world data is smooth – this means that the probability of treatment does not change sharply across individuals with slightly different covariates vectors.

Therefore, RHPT embeddings are approximate balancing scores, given that β is sufficiently large and $e(\cdot)$ is a smooth function. Note that there is no way of knowing the value of β such that $h_\beta(\mathbf{x}_i)$ is an approximate balancing score, since calculating the value of such a β involves dealing with unknown constants mentioned in section 2.

3.2. Empirical evidence that RHPT yields approximate balancing scores

An empirical approach used to verify if a representation is as an approximate balancing score involves fitting a function to predict propensity score given the representation, then comparing this predicted score to the true propensity score. In real-world data sets, we do not have access to true propensity scores. To address this limitation, we generate a semi-synthetic data set using the method proposed by [Jesson et al. \(2021\)](#) for the MNIST data set. For the purpose of evaluating the proposed method, we obtain the true propensity scores by simulating the data-generation process.

To assess if our RHPT embeddings are balancing scores, we predict propensity scores using logistic regression learned on our embeddings and compare our predictions to the true propensity scores. Specifically, we calculate the mean absolute difference between the the true propensity scores and our predicted propensity scores, given by $\psi = \frac{1}{n} \sum_{i=1}^n |e(\mathbf{x}_i) - \hat{e}(h_\beta(\mathbf{x}_i))|$ where $\hat{e}(h_\beta(\mathbf{x}_i))$ is the logistic regression model trained on $h_\beta(\mathbf{x}_i)$. Low values of ψ indicate that our embeddings $h_\beta(\mathbf{x}_i)$ are an approximate balancing scores.




Figure 2: Error of propensity score predictions ψ across various RHPT embedding dimensionalities. Means and confidence intervals are computed over 100 instantiations of RHPT on a single draw of the synthetic data. Lower ψ favors ignorability.

Figure 2 illustrates our empirical results across varying embedding dimensions β . We see that as β increases, the mean absolute difference between predicted and true propensity scores decreases. This evaluation was performed in order to complement our theoretical arguments and provide empirical evidence that RHPT embeddings are (approximate) balancing scores. As discussed in Section 3.1, causal effects can be estimated by matching on balancing scores rather than by matching on the raw covariates. The prediction error in Figure 2 is relatively low, and decreases as the RHPT dimensionality increases. These results suggest that RHPT embeddings are approximate balancing scores and that matching using RHPT is a valid method of estimating causal effects.

3.3. Uncertainty in causal estimates using RHPT matching

Our proposed method makes use of a randomized algorithm to map covariates to our RHPT representation. As such, conducting experiments with different random seeds will yield different results. Specifically, different random mappings will result in different nearest-neighbor matches, thus affecting ATE and ITE predictions. Given the theoretical results from section 2, we would expect more variation in our results for lower-dimensional RHPT embeddings. This is because distances δ are approximately preserved when the embedding dimensionality β increases.

To address this, we perform experiments over 100 randomized runs of RHPT with various choice of β , with results and variability illustrated in Figure 3. We observe that as we increase the value of β , the variability of the ATE decreases until it eventually levels off when β reaches 10000.

4. Experiments

We conduct experiments to compare the performance of RHPT with both classic matching techniques as well as the state-of-the-art deep learning techniques. Baseline 1-Nearest Neighbor (1-NN) matching techniques include matching on the raw covariates \mathbf{X} and matching on the PCA-transformed covariates \mathbf{Z} . For matching on PCA-transformed covariates we use 5 principal compo-

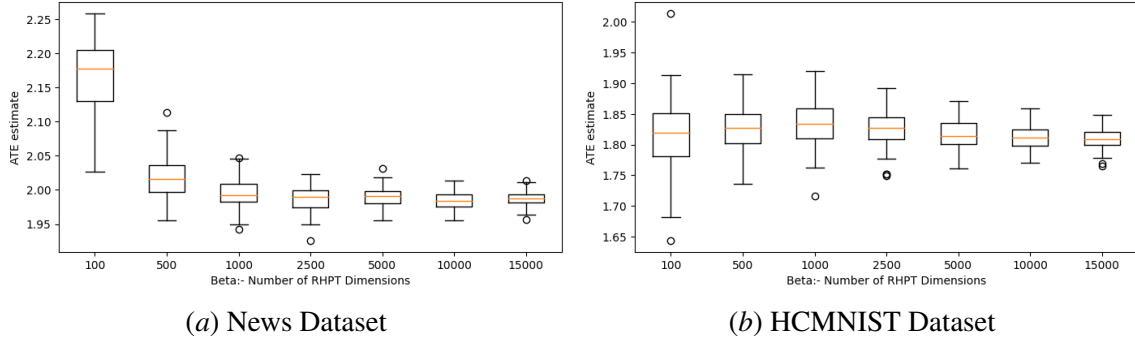


Figure 3: Estimated ATE using RHPT on one instantiation of the synthetic dataset with various hash dimensions, each over 100 randomized runs. Higher dimensionality embeddings result in more reliable estimates.

nents following the experiments of [Beyer et al. \(1999\)](#), which showed that nearest neighbor queries become unreliable when the number of principal components used exceeds 10. We also benchmark against a propensity scores which are calculated using logistic regression over the raw covariate space \mathbf{X} and PCA features \mathbf{Z} , which we refer to as $\hat{e}(\mathbf{X})$ and $\hat{e}(\mathbf{Z})$, respectively. In addition, we consider matching uniformly at random (random matching) as a simple baseline. Finally, we compare our results with those obtained using the state-of-the-art deep learning techniques: CFRNet with Maximum Mean Discrepancy (MMD) ([Johansson et al., 2016](#)) and DragonNet with Targeted Regularization ([Shi et al., 2019](#)).

4.1. Evaluation Metrics

We compare RHPT with other methods using two different classes of performance metrics: (1) Error of estimated causal effects (relative to their true value); and (2) Computational cost of causal effect estimation (training as well as inference).

We evaluate the error of causal effect estimates in two different settings: (1) *within-sample*, where the task is to estimate the causal effect where the factual outcome is observed, and (2) *out-of-sample*, where the goal is to estimate the causal effect where the factual outcome is unobserved. For within-sample evaluation we estimate the true $\text{ITE}(\mathbf{x}_i)$ by direct modelling. Given an individual with observation (\mathbf{x}_i, T_i, Y_i) , the transductive ITE [Johansson et al. \(2016\)](#) is given by:

$$\widehat{\text{ITE}}_{in}(\mathbf{x}_i) = \begin{cases} Y_i - g(\mathbf{x}_i, 0) & \text{if } T_i = 1 \\ g(\mathbf{x}_i, 1) - Y_i & \text{if } T_i = 0 \end{cases} \quad (3)$$

where the function $g(\mathbf{x}_i, 1)$ and $g(\mathbf{x}_i, 0)$ are the estimate of the counterfactual for individual i . For CFRnet and DragonNet, $g(\mathbf{x}_i, 1)$ and $g(\mathbf{x}_i, 0)$ are neural networks, while for matching techniques they are $\hat{Y}_{NN_1(\mathbf{x}_i)}$ and $\hat{Y}_{NN_0(\mathbf{x}_i)}$, respectively.

We introduce two within-sample metrics for evaluation. The first is within-sample ϵ_{ATE} , which is the mean absolute difference between the true and the predicted ATE:

$$\epsilon_{\text{ATE}} = \left| \left(\frac{1}{n} \sum_{i=0}^{n-1} \widehat{\text{ITE}}_{in}(\mathbf{x}_i) \right) - \left(\frac{1}{n} \sum_{i=0}^{n-1} \text{ITE}(\mathbf{x}_i) \right) \right|$$

Note that the true effect i.e. $\text{ITE}(\mathbf{x}_i)$ considers noiseless outcomes, while the factual outcomes Y_i seen in the data-set contain noise in them. We also use the RMSE error for ITE for within-sample evaluation:

$$\epsilon_{\text{ITE}} = \sqrt{\frac{1}{n} \sum_{i=0}^{n-1} (\widehat{\text{ITE}}_{in}(\mathbf{x}_i) - \text{ITE}(\mathbf{x}_i))^2}$$

For Out-of-sample metrics, we use inductive ITE: $\widehat{\text{ITE}}_{out}(\mathbf{x}_i) = g(\mathbf{x}_i, 1) - g(\mathbf{x}_i, 0)$, where both outcomes are predicted. For matching techniques, $g(\mathbf{x}_i, 1)$ and $g(\mathbf{x}_i, 0)$ are 1-NN matches from the within-sample individuals. For deep learning techniques, CFRNet and DragonNet, $g(\mathbf{x}_i, 1)$ and $g(\mathbf{x}_i, 0)$ are predicted using neural networks. In this setting we also use out-of-sample ϵ_{ATE} – the absolute error in the ATE estimation. The final metric we use is Precision in Estimation of Heterogeneous Effect (PEHE) from Hill (2011), given by:

$$\epsilon_{\text{PEHE}} = \sqrt{\frac{1}{n} \sum_{i=0}^{n-1} (\widehat{\text{ITE}}_{out}(\mathbf{x}_i) - \text{ITE}(\mathbf{x}_i))^2}$$

	Within-sample metrics		Out-of-sample metrics		Wall-Time	
	ϵ_{ATE}	ϵ_{ITE}	ϵ_{ATE}	ϵ_{PEHE}	GPU	CPU
Random	0.87 ± 0.08	4.08 ± 0.20	0.84 ± 0.09	5.72 ± 0.30	-	0.63
X	0.42 ± 0.06	3.04 ± 0.11	0.75 ± 0.11	4.22 ± 0.18	-	1.82
Z	0.34 ± 0.03	2.77 ± 0.07	0.36 ± 0.04	3.86 ± 0.15	-	1.67
$\hat{e}(X)$	0.48 ± 0.09	2.88 ± 0.13	0.66 ± 0.13	5.00 ± 0.24	-	1.59
$\hat{e}(Z)$	0.44 ± 0.03	3.19 ± 0.09	0.47 ± 0.05	4.49 ± 0.18	-	1.66
RHPT ¹⁰⁰	0.45 ± 0.04	3.11 ± 0.12	0.50 ± 0.06	4.23 ± 0.19	-	1.37
RHPT ⁵⁰⁰	0.22 ± 0.02	2.51 ± 0.06	0.27 ± 0.03	3.37 ± 0.12	-	1.93
RHPT ^{1k}	0.19 ± 0.02	2.39 ± 0.06	0.25 ± 0.03	3.18 ± 0.12	-	2.14
RHPT ^{2.5k}	0.18 ± 0.02	2.36 ± 0.06	0.25 ± 0.03	3.18 ± 0.12	-	5.16
RHPT ^{5k}	0.17 ± 0.02	2.35 ± 0.05	0.25 ± 0.04	3.15 ± 0.13	-	11.09
RHPT ^{10k}	0.16 ± 0.02	2.34 ± 0.05	0.25 ± 0.04	3.14 ± 0.13	-	16.00
CFRNet ^{MMD}	0.16 ± 0.01	2.03 ± 0.03	0.22 ± 0.02	2.63 ± 0.07	38.84	≥ 2000
DragonNet ^{T-reg}	0.13 ± 0.01	1.92 ± 0.04	0.16 ± 0.02	2.25 ± 0.07	39.26	≥ 2000

Table 1: Mean errors (and standard errors) with Wall-times (minutes) across 50 simulations of News Dataset. Rows indicate different causal effect estimation methods (randomly matching, matching on covariates X, matching on PCA-transformed covariates Z, etc).

We compare the computational cost of different methods using run-time (training and inference) for each method. The deep learning methods are run on two types of machines: (1) using an NVIDIA A-40 GPU and (2) using the free version of Google Colab in CPU mode (with no GPU). Matching techniques are not evaluated using a GPU since matching does not require computationally intensive training. Note that for the free version of Google Colab in CPU mode, there is a limit

	Within-Sample Metrics		Out-of-Sample Metrics		Wall-Time	
	ϵ_{ATE}	ϵ_{ITE}	ϵ_{ATE}	ϵ_{PEHE}	GPU	CPU
Random	0.26 ± 0.01	4.04 ± 0.01	0.32 ± 0.02	4.94 ± 0.02	-	1.86
X	0.26 ± 0.01	3.59 ± 0.01	0.32 ± 0.02	3.71 ± 0.02	-	1.98
Z	0.24 ± 0.01	3.70 ± 0.01	0.30 ± 0.02	4.06 ± 0.02	-	1.92
$\hat{e}(X)$	0.25 ± 0.02	4.03 ± 0.01	0.39 ± 0.03	4.97 ± 0.02	-	8.38
$\hat{e}(Z)$	0.23 ± 0.01	3.98 ± 0.01	0.31 ± 0.02	4.64 ± 0.02	-	1.97
RHPT ¹⁰⁰	0.23 ± 0.01	3.68 ± 0.01	0.29 ± 0.02	3.95 ± 0.02	-	1.96
RHPT ⁵⁰⁰	0.21 ± 0.01	3.63 ± 0.01	0.24 ± 0.02	3.80 ± 0.02	-	2.55
RHPT ^{1k}	0.20 ± 0.01	3.61 ± 0.01	0.24 ± 0.02	3.77 ± 0.02	-	3.08
RHPT ^{2.5k}	0.20 ± 0.01	3.61 ± 0.01	0.23 ± 0.02	3.75 ± 0.02	-	3.59
RHPT ^{5k}	0.20 ± 0.01	3.60 ± 0.01	0.23 ± 0.02	3.75 ± 0.02	-	6.18
RHPT ^{10k}	0.19 ± 0.01	3.58 ± 0.01	0.23 ± 0.02	3.73 ± 0.02	-	8.30
CFRNet ^{MMD}	0.21 ± 0.01	3.31 ± 0.01	0.25 ± 0.02	3.09 ± 0.02	56.34	181.43
DragonNet ^{T-reg}	0.23 ± 0.01	2.84 ± 0.00	0.24 ± 0.02	2.16 ± 0.01	21.01	59.41

Table 2: Mean errors (and standard errors) with wall-times in minutes across 100 simulations of HCMNIST Dataset. Rows indicate different causal effect estimation methods (randomly matching, matching on covariates X, matching on PCA-transformed covariates Z, etc).

on the of execution time. If that limit is exceeded during execution, we report a run-time ≥ 2000 minutes.

4.2. Data sets

Data sets used for evaluation are semi-synthetic, meaning that the covariates come from real-world data while the treatments and outcomes were generated through a user-specified data generating process. For each data set, we report the average and standard error of the within-sample and out-of-sample errors over different instantiations of the data generating process. The wall-time is reported as the total time required to run each method on of all samples from the data generating process. For all data sets, 10% of the samples are used for out-of-sample evaluation. For deep learning models, 20% of in-sample data points are used as validation set to avoid over-fitting.

4.2.1. NEWS

Johansson et al. (2016) used a simulated data set to replicate how opinions and perceptions toward news articles could differ based on whether they accessed the content on a desktop (T=0) or a mobile device (T=1). In each draw of the data generating process, 5000 news articles are randomly sampled from the NY Times corpus (Newman, 2008). The covariates x_i of each news article is a Bag of Words with dimensionality 3477. Using the covariates, treatment assignments (desktop (T=0) or mobile device (T=1)) and outcomes (opinions) for each news articles are generated. We use 50 draws of the Data Generating Process provided by Johansson et al. (2016).

4.2.2. HC-MNIST

Jesson et al. (2021) introduced a semi-synthetic data set that based on the MNIST dataset LeCun (1998). The covariates are the 28×28 pixel values of the images with total dimensionality of 784. The representation of the images are used to generate binary treatment and continuous outcomes. The data generating process contains $\Gamma \geq 1$ factor which controls the degree of hidden confounding with $\Gamma = 1$ representing no unobserved confounding. For each instance/draw out of the 100 draws of the simulation, the treatment and outcomes are generated according to Jesson et al. (2021) method and then 3000 images are randomly selected. Only 3000 samples/images are selected out of the complete data set (42000 images) to introduce sparsity in the data space so that we can emulate the curse of dimensionality. We set $\Gamma = 1.1$ to mimic a real world scenario. (Full unconfoundedness is not plausible in the real world).

In addition to News and HC-MNIST dataset, we also perform experiment on 100 draws of Infant Health and Development Program (IHDP) dataset by Hill (2011) which includes covariates from a randomized controlled trial devised to assess how cognitive test scores in premature infants are impacted by visits from specialized doctors. IHDP data set is a low dimensional dataset with 25 covariates of which 16 are binary and 9 are continuous. While this dataset falls outside the primary scope of our algorithm due to its low dimensionality, we have included its results in the appendix to demonstrate that our method remains functional even for low-dimensional data.

4.3. Hyperparameters for Deep learning models

The hyperparameters for deep learning models are presented in Table 3. The hyperparameters are found using the method described in Shalit et al. (2017). The loss functions of these deep learning models consist of multiple terms. All these terms are weighted equally i.e. weights/imbalance parameter are 1. In all the models, the initial two layers don't have L2 regularization; every other layer has L2 regularization weighted by 0.01. One significant observation is that the hyper-parameters for early stopping and learning rate were more important as compared to the neural network architecture for achieving strong performance. For training CFRNet, we use Adam optimizer Kingma and Ba (2014) and for DragonNet we use Nesterov accelerated gradient descent (NAG) with momentum.

Dataset	Model	Layers	Early Stopping Patience	Starting Learning Rate (LR)	LR Reduction Patience
News	CFRNET	[400,400,[200,200]]	40	2.5e-3	10
	Dragonnet	[200,200,[100,100]]	40	1e-7	5
HCMNIST	CFRNET	[512,256,[256,128]]	30	2.5e-3	10
	Dragonnet	[512,256,128,[64,32]]	20	1e-7	5

Table 3: Hyperparameters for the deep learning models

4.4. Results

For the News dataset, we observe that RHPT matching significantly outperforms traditional matching techniques on all metrics. Additionally, it is noteworthy that the Within-sample ϵ_{ATE} error for

both RHPT matching and deep learning models is almost the same, while the ϵ_{ITE} error and the Out-of-sample ϵ_{ATE} error is slightly worse for RHPT, but still within the same ballpark as the deep learning models. There is a significant difference between RHPT and the deep learning models for the ϵ_{PEHE} error. RHPT matching outperforms CFRNet and DragonNet in terms of the wall-time needed for estimating the causal effect. Due to high dimensionality a large sample size of the News dataset, it takes a significant amount of time to execute on a CPU when using deep learning models. Even when a GPU is used, the wall time for RHPT matching is less than half of the GPU wall time of the deep learning models.

As seen from Table 2, HCMNIST dataset appears to be more challenging for causal effect estimation than the News dataset, as the improvement from traditional matching techniques to RHPT and deep learning models is not as pronounced as it is for the News dataset. The In-sample ϵ_{ATE} error for all methods is nearly the same, with RHPT performing the best. The out-of-sample ϵ_{ATE} error is significantly better for RHPT matching when compared to traditional matching techniques and on par with the deep learning models. The ϵ_{PEHE} error is significantly better for RHPT matching than traditional matching techniques, but deep learning methods outperform any of the matching techniques.

RHPT outperforms traditional matching techniques in terms of both in-sample and out-of-sample metrics when the dimensionality of the RHPT representation is sufficiently large ($\geq 10k$) with only modest increase in computational cost. We see that our matching technique significantly and consistently outperforms all traditional matching techniques when it comes to the ϵ_{ATE} , ϵ_{ITE} and ϵ_{PEHE} and is competitive with CFRNet and DragonNet. Deep learning methods slightly outperform RHPT in terms of out-of-sample ϵ_{PEHE} error; and RHPT is competitive with deep learning methods in terms of within-sample error. In both cases, RHPT matching incurs computational cost that is an order or two magnitude lower than deep learning methods.

5. Conclusion and Discussion

We have introduced a simple, fast, yet effective approach to estimating causal effects from high-dimensional observational data. The proposed method makes use of Random Hyperplane Tessellations (RHPT) to map covariates into a high-dimensional, binary space based on the relationship between covariate values and randomly chosen hyperplanes. We show that RHPT ensures approximate covariate balance, a theoretical requirement for accurate causal effect estimation from observational data. We report results of extensive experiments showing that matching using RHPT outperforms traditional matching techniques and is competitive with state-of-the-art deep learning methods for causal effect estimation. Matching using RHPT also provides several important benefits not offered by deep learning methods. Namely, RHPT avoids the need for computationally expensive training, does not require large amounts of data, and requires little hyperparameter tuning.

A promising direction for future research include extensions of the proposed method to causal effect estimation in settings where the observations are not independent and identically distributed, e.g., when individuals are linked by social or other ties. Another interesting direction is deeper theoretical analyses of the conditions under which RHPT or similar methods can be expected to yield accurate causal effect estimates. Finally, causal inference in the online learning and continual learning settings remains relatively unexplored. Matching-based approaches such as RHPT are well-suited for this setting, as they do not require iterative training and can be easily updated as new data becomes available.

Acknowledgments

This work was funded in part by grants from the National Science Foundation (2226025, 2041759), the National Center for Advancing Translational Sciences, and the National Institutes of Health (UL1 TR002014).

References

- Charu C. Aggarwal, Alexander Hinneburg, and Daniel A. Keim. On the surprising behavior of distance metrics in high dimensional spaces. In *International Conference on Database Theory*, 2001.
- Kevin S. Beyer, Jonathan Goldstein, Raghu Ramakrishnan, and Uri Shaft. When is "nearest neighbor" meaningful? In *International Conference on Database Theory*, 1999.
- Moses S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing, STOC '02*, page 380–388, New York, NY, USA, 2002. Association for Computing Machinery. ISBN 1581134959. doi: 10.1145/509907.509965.
- David Roxbee Cox. Planning of experiments. 1958.
- Sjoerd Dirksen, Shahar Mendelson, and Alexander Stollenwerk. Sharp estimates on random hyperplane tessellations. *SIAM Journal on Mathematics of Data Science*, 4(4):1396–1419, 2022. doi: 10.1137/22M1485826.
- Miguel Hernan and James Robins. *Causal Inference - What if*. Taylor and Francis, 2023.
- Jennifer L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20:217 – 240, 2011.
- Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing, STOC '98*, page 604–613, New York, NY, USA, 1998. Association for Computing Machinery. ISBN 0897919629. doi: 10.1145/276698.276876. URL <https://doi.org/10.1145/276698.276876>.
- Andrew Jesson, Sören Mindermann, Yarin Gal, and Uri Shalit. Quantifying ignorance in individual-level causal-effect estimates under hidden confounding. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4829–4838. PMLR, 18–24 Jul 2021.
- Fredrik D. Johansson, Uri Shalit, and David A. Sontag. Learning representations for counterfactual inference. In *International Conference on Machine Learning*, 2016.
- Gary King and Richard Nielsen. Why propensity scores should not be used for matching. *Political analysis*, 27(4):435–454, 2019.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.

- Yann LeCun. The mnist database of handwritten digits, 1998. URL <http://yann.lecun.com/exdb/mnist/>.
- Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. *Advances in neural information processing systems*, 30, 2017.
- David Newman. Bag of words data set, 2008. URL <https://archive.ics.uci.edu/ml/datasets/Bag+of+Words>.
- Yaniv Plan and Roman Vershynin. Dimension reduction by random hyperplane tessellations. *Discrete Comput. Geom.*, 51(2):438–461, mar 2014. doi: 10.1007/s00454-013-9561-6.
- Iacopo Poli. Article on efficient neural networks training through locality sensitive hashing. *medium.com*, 2021.
- Rameshwar Pratap, Anup Deshmukh, Pratheeksha Nair, and Anirudh Ravi. Scaling up simhash. In Sinno Jialin Pan and Masashi Sugiyama, editors, *Proceedings of The 12th Asian Conference on Machine Learning*, volume 129 of *Proceedings of Machine Learning Research*, pages 705–720. PMLR, 18–20 Nov 2020.
- Paul R. Rosenbaum and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 04 1983. ISSN 0006-3444. doi: 10.1093/biomet/70.1.41.
- Donald B. Rubin. Bias reduction using mahalanobis-metric matching. *Biometrics*, 36(2):293–298, 1980.
- Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pages 3076–3085. PMLR, 2017.
- Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems*, 32, 2019.
- Elizabeth A. Stuart. Matching methods for causal inference: A review and a look forward. *Statistical science : a review journal of the Institute of Mathematical Statistics*, 25 1:1–21, 2010.

Appendix A. Additional experimental results

Results for the IHDP dataset are provided in Table 4. RHPT matching provides similar performance to matching on the raw covariates. Since the IHDP dataset is relatively low-dimensional with only 25 features, the curse-of-dimensionality does not hinder matching performance, therefore RHPT matching does not provide much added benefit.

Appendix B. Implementation Details

The code is implemented in Python 3 using TensorFlow 2. The codebase for the experiments can be found at <https://ANONYMOUS.com>.

	Within-Sample Metrics		Out-of-Sample Metrics		Wall-Time	
	ϵ_{ATE}	ϵ_{ITE}	ϵ_{ATE}	ϵ_{PEHE}	GPU	CPU
Random	0.30 ± 0.04	4.23 ± 0.50	0.94 ± 0.18	8.47 ± 1.16	-	0.07
X	0.14 ± 0.02	2.64 ± 0.25	0.67 ± 0.13	4.85 ± 0.63	-	0.12
Z	0.17 ± 0.02	3.0 ± 0.30	0.70 ± 0.16	5.63 ± 0.74	-	0.12
e(X)	0.22 ± 0.02	4.16 ± 0.5	0.79 ± 0.14	8.23 ± 1.12	-	0.13
e(Z)	0.20 ± 0.02	4.02 ± 0.47	0.8 ± 0.12	7.93 ± 1.08	-	0.20
RHPT ¹⁰⁰	0.14 ± 0.02	2.98 ± 0.31	0.57 ± 0.10	5.75 ± 0.77	-	0.19
RHPT ⁵⁰⁰	0.14 ± 0.02	2.84 ± 0.28	0.55 ± 0.08	5.20 ± 0.71	-	0.19
RHPT ^{1k}	0.15 ± 0.02	2.79 ± 0.25	0.53 ± 0.07	5.11 ± 0.69	-	0.19
RHPT ^{2.5k}	0.14 ± 0.01	2.76 ± 0.27	0.53 ± 0.07	4.98 ± 0.63	-	0.32
RHPT ^{5k}	0.14 ± 0.01	2.76 ± 0.28	0.49 ± 0.07	4.91 ± 0.60	-	0.45
RHPT ^{10k}	0.14 ± 0.02	2.75 ± 0.27	0.48 ± 0.10	4.92 ± 0.65	-	0.72
CFRNet ^{MMD}	0.10 ± 0.01	1.21 ± 0.04	0.19 ± 0.03	1.15 ± 0.17	19.58	117.76
DragonNet ^{T-reg}	0.11 ± 0.00	1.17 ± 0.02	0.16 ± 0.02	0.96 ± 0.12	47.25	513.32

Table 4: Mean errors (and standard errors) with Wall-times in minutes across 100 simulations of IHDP Dataset. Rows indicate different causal effect estimation methods (randomly matching, matching on covariates X, matching on PCA-transformed covariates Z, etc).