



Published in final edited form as:

J Comput Graph Stat. 2023 ; 32(3): 1036–1045. doi:10.1080/10618600.2022.2141752.

Transfer learning of individualized treatment rules from experimental to real-world data

Lili Wu,

Shu Yang

Department of Statistics, North Carolina State University

Abstract

Individualized treatment effect lies at the heart of precision medicine. Interpretable individualized treatment rules (ITRs) are desirable for clinicians or policymakers due to their intuitive appeal and transparency. The gold-standard approach to estimating the ITRs is randomized experiments, where subjects are randomized to different treatment groups and the confounding bias is minimized to the extent possible. However, experimental studies are limited in external validity because of their selection restrictions, and therefore the underlying study population is not representative of the target real-world population. Conventional learning methods of optimal interpretable ITRs for a target population based only on experimental data are biased. On the other hand, real-world data (RWD) are becoming popular and provide a representative sample of the target population. To learn the generalizable optimal interpretable ITRs, we propose an integrative transfer learning method based on weighting schemes to calibrate the covariate distribution of the experiment to that of the RWD. Theoretically, we establish the risk consistency for the proposed ITR estimator. Empirically, we evaluate the finite-sample performance of the transfer learner through simulations and apply it to a real data application of a job training program.

Keywords

Augmented inverse probability weighting; Calibration weighting; Classification error; Covariate shift; Cross-validation; Transportability

1 Introduction

A personalized recommendation, tailored to individual characteristics, is becoming increasingly popular in real life such as healthcare, education and e-commerce. An optimal *individualized treatment rule* (ITR) is defined as maximizing *the mean value function* of an outcome of interest over the target population had all individuals in the whole population follow the given rule. Many machine learning approaches to estimating the optimal ITR are available, such as outcome-weighted learning with support vector machine (Zhao et al. 2012; Zhou et al. 2017), regression-based methods with Adaboost (Kang et al. 2014), regularized linear basis function (Qian and Murphy 2011), K-nearest neighbor (Zhou and Kosorok 2017) and generalized additive model (Moodie et al. 2014). However, the

ITRs derived by machine learning methods can be too complex to extract practically meaningful insights for policymakers. For example, in healthcare, clinicians need to scrutinize the estimated ITRs for scientific validity, but the black-box rules conceal the clear relationship between patient characteristics and treatment recommendations. In these cases, parsimonious and interpretable ITRs are desirable. Interpretable ITRs usually refer to the ones where clinicians can explicitly obtain the information of how the decision rules recommend a certain treatment/intervention given a patient's covariate information instead of some complex black-box approaches such as random forest and neural network. We consider the simple forms: linear rules, which can help clinicians to see how each covariate leads to the final treatment recommendations. Other examples of interpretable decision rules include the decision trees (Laber and Zhao 2015; Tao et al. 2018) and "if-then-else" decision lists (Gail and Simon 1985; Zhang et al. 2015; Lakkaraju and Rudin 2017).

It is well known that randomized experiments are the gold standard for learning the optimal ITRs because randomization of treatments in the design stage can ensure no unmeasured confounding (Greenland 1990). However, randomized experiments might lack external validity or representativeness (Rothwell 2005) because of the selective criteria for individual eligibility in the experiments. Consequently, the distribution of the covariates in the randomized experiments might differ from that in the target population. Thus, a parsimonious ITR constructed from randomized experiments cannot directly generalize to a target population (Zhao et al. 2019) due to the covariate shift. Alternatively, real-world data (RWD) such as population-based survey data usually involve a large and representative sample of the target population. Figure 1 presents a simple example to illustrate the bias of the estimated optimal linear ITR based only on the experiment. A large body of literature in statistics has focused on transporting or generalizing the average treatment effects from the experiments to a target population or a different study population (e.g., Cole and Stuart 2010; Hartman et al. 2015). One exception for the ITRs is Zhao et al. (2019) and Mo et al. (2020) which estimates the optimal linear ITR from the experimental data by optimizing the worst-case quality assessment among all covariate distributions in the target population satisfying some moment conditions or distributional closeness, while we address the problem in a different situation where we already have a representative sample at hand.

In this article, we propose a general framework for estimating a generalizable optimal interpretable ITR from randomized experiments to a target population represented by the RWD. We propose *transfer weighting* that shifts the covariate distribution in the randomized experiments to that in the RWD. To correctly estimate the population optimal ITR, transfer weights are used in two places. First, the empirical estimator of the population value function is biased when considering only the experimental sample, and hence the transfer weights are used to correct the bias of the estimator of the value function. Second, the nuisance functions in the value function such as the propensity score and the outcome mean also require transfer weighting to estimate the population counterparts. Then we can obtain the weighted interpretable optimal ITR by reformulating the problem of maximizing the weighted value to that of minimizing the weighted classification error with flexible loss function choices. We consider different parametric and nonparametric approaches to learning transfer weights. Moreover, a practical issue arises regarding the choice of transfer weights. Weighting can correct for the bias but also reduce the effective sample size and

increase the variance of the weighted value estimator. To address the bias-variance tradeoff, we propose a cross-validation procedure that maximizes the estimated population value function to choose among the weighting schemes. Continuing with the simple example in Figure 1, the optimal linear ITR from the proposed transfer learner becomes closer to the population optimal linear ITR than the biased experimental estimator. Lastly, we provide the theoretical guarantee of the risk consistency of the proposed transfer weighted ITRs.

The remaining of the article is organized as follows. We introduce our proposed method in Section 3. In Section 4, we study the asymptotic properties of the proposed estimator to guarantee its performance in the target population theoretically. In Section 5, we use various simulation settings to illustrate how this approach works. In Section 6, we apply the proposed method to the National Supported Work program (LaLonde 1986) for improving the decisions of whether to recommend the job training program based on the individual characteristics of the target population. We conclude the article in Section 7.

2 Basic Setup

2.1 Notation

Let $X \in \mathcal{X} \subseteq \mathbb{R}^p$ be a vector of pre-treatment covariates, $A \in \mathcal{A} = \{0, 1\}$ the binary treatment, and $Y \in \mathbb{R}$ the outcome of interest. Under the potential outcomes framework (Rubin 1974), let $Y^*(a)$ denote the potential outcome had the individual received treatment $a \in \mathcal{A}$. Define a treatment regime as a mapping $d: \mathcal{X} \rightarrow \mathcal{A}$ from the space of individual characteristic to the treatment. Then we define the potential outcome under regime d as $Y^*(d) = \sum_{a \in \mathcal{A}} Y^*(a) \mathbf{1}\{d(X) = a\}$, the value of regime d as $V(d) = \mathbb{E}\{Y^*(d)\}$, where the expectation is taken over the distribution in the target population, and the optimal regime d^{opt} as the one satisfying $V(d^{\text{opt}}) \geq V(d), \forall d$. As discussed in the introduction, in many applications especially the high-stake ones such as supporting clinicians and judges to make decisions (Lakkaraju and Rudin 2017), interpretable decision rules are desirable since they might bring more domain insights and be more trustworthy to decision-makers compared with the complex black-box regimes. Thus, it is valuable to learn an optimal interpretable regime. Specifically, we focus on the regimes with linear forms, i.e, $d(X; \eta) = \mathbf{1}(\eta_0 + \eta_1^\top X > 0)$, denoted by d_η for simplicity.

We aim at learning ITRs which lead to the highest benefits for the target population with size N , $\{X_i, Y_i^*(0), Y_i^*(1)\}_{i=1}^N$. Suppose we have access to two independent data sources: the experiment data with $\{(X_i, A_i, Y_i)\}_{i \in \mathcal{F}_n}$ and the RWD with $\{X_i\}_{i \in \mathcal{F}_m}$, where, $\mathcal{F}_n, \mathcal{F}_m \subseteq \{1, \dots, N\}$ are samples' index sets with size n, m , respectively. In the experimental data, individuals are selected based on certain criteria and then are randomized to receive treatments. The experimental data might induce selection bias which leads to a different distribution of X in experiments from that in the target population. On the other hand, we assume that individuals in the RWD are a random sample from the target population, thus the distribution of the covariates in the RWD can characterize that of the target population. Let S_i be the binary indicator of the i th individual participates in the randomized experiment: $S_i = 1$ if $i \in \mathcal{F}_n$ and 0 if $i \in \mathcal{F}_m$. We denote $\pi_A(X)$ as the propensity score of receiving the

active treatment $P(A = 1 | \mathbf{X})$, $Q(\mathbf{X}, a)$ as the conditional expectation of potential outcomes given covariates $\mathbb{E}\{Y^*(a) | \mathbf{X}\}$, and $\tau(\mathbf{X})$ as the contrast function $Q(\mathbf{X}, 1) - Q(\mathbf{X}, 0)$.

We make the following assumptions, which are standard and well-studied assumptions in the causal inference and treatment regime literature (Chakraborty 2013).

Assumption 1.—*a. $A \perp \perp \{Y^*(0), Y^*(1)\} | (\mathbf{X}, S = 1)$. b: $Y = Y^*(A)$. c: There exist $c_1, c_2 \in (0, 1)$ such that $\pi_A(\mathbf{X}) \in [c_1, c_2]$ for all $\mathbf{X} \in \mathcal{X}$.*

Assumption 1.a holds by the design of a completely randomized experiment, where the treatment is independent of the potential outcomes and covariates. It also holds by the design of a conditional randomized experiment, where the treatment is independent of the potential outcomes within each level of \mathbf{X} . Besides, it also holds in observational studies without unmeasured confounders. In general, our framework is applicable to generalizing the findings from the observational studies to the target population.

2.2. Existing learning methods for ITRs

Identifying the optimal individualized treatment rules has been studied in a large body of statistical literature. One category is regression-based by directly estimating the conditional expectation of the outcome given patients' characteristics and treatments received (Murphy 2005). Both parametric and nonparametric models can be used to approximate $Q(\mathbf{X}, A)$. To encourage interpretability, one strategy is to posit a parametric model as $Q(\mathbf{X}, A; \beta)$, and the estimated optimal ITR within the class indexed by β is defined as $\mathbf{1}\{\tau(\mathbf{X}; \hat{\beta}) > 0\}$.

Another category is value-based by first estimating the value of an ITR and then searching the ITR that renders the highest value. One strategy is to estimate $\pi_A(\mathbf{X})$ first, where one common way is to posit a parametric model such as logistic regression or a constant, and then use the inverse probability weighting (IPW; Horvitz and Thompson 1952) to estimate the value of an ITR indexed by η . This estimator is consistent to the true value only if $\pi_A(\mathbf{X}_i)$ is correctly specified and can be unstable if the estimate $\hat{\pi}_A(\mathbf{X}_i)$ is close to zero. The augmented inverse probability weighting (AIPW; Zhang, Tsiatis, Laber and Davidian (2012)) approach is introduced to combine the IPW and outcome regression estimator $\hat{Q}(\mathbf{X}_i, A)$ to estimate the value

$$\hat{V}_{aw}(d_\eta; \mathcal{F}_n) = \frac{1}{n} \sum_{i \in \mathcal{F}_n} \left(\left[\frac{A_i d(\mathbf{X}_i; \eta)}{\hat{\pi}_A(\mathbf{X}_i)} + \frac{(1 - A_i)\{1 - d(\mathbf{X}_i; \eta)\}}{1 - \hat{\pi}_A(\mathbf{X}_i)} \right] [Y_i - \hat{Q}\{\mathbf{X}_i, d(\mathbf{X}_i; \eta)\}] - \hat{Q}\{\mathbf{X}_i, d(\mathbf{X}_i; \eta)\} \right). \tag{1}$$

The AIPW estimator is doubly robust in the sense that it will be consistent to the value of the ITR d_η if either the model for $\pi_A(\mathbf{X})$ or $Q(\mathbf{X}, A)$ is correctly specified, but not necessarily both.

Learning the optimal ITR by optimizing the value can also be formulated as a classification problem (Zhang, Tsiatis, Davidian, Zhang and Laber 2012). The classification perspective

has equivalent forms to all the regression-based and value-based methods above, and can also bring us additional optimization advantages over directly value searching.

Lemma 1.—*The optimal treatment rule by maximizing $V(d_\eta)$ in η is equivalent to that by minimizing the risk*

$$\mathbb{E}\left(\left|\tau(\mathbf{X})\right|[\mathbf{1}\{\tau(\mathbf{X}) > 0\} - d(\mathbf{X}; \eta)]^2\right) \tag{2}$$

in η (Zhang, Tsiatis, Davidian, Zhang and Laber 2012).

Lemma 1 reformulates the problem of estimating an optimal treatment rule as a weighted classification problem, where $|\tau(\mathbf{X})|$ is regarded as the weight, $\mathbf{1}\{\tau(\mathbf{X}) > 0\}$ the binary label, and $d(\mathbf{X}; \eta)$ the classification rule of interest. Let l_{0-1} be a 0–1 loss, i.e., $l_{0-1}(u) = \mathbf{1}(u \leq 0)$. The objective function (2) can also be rewritten as

$$\mathcal{R}_{\mathcal{F}}(\tau; \eta, l_{0-1}) = \mathbb{E}\{\mathcal{L}_{(\tau, \eta, l_{0-1})}, \mathcal{L}_{(\tau, \eta, l_{0-1})} = |\tau(\mathbf{X})|l_{0-1}\{[2\mathbf{1}\{\tau(\mathbf{X}) > 0\} - 1] f(\mathbf{X}; \eta)\}, \tag{3}$$

Where $f(\mathbf{X}; \eta) = \eta^\top \mathbf{X}$, $d(\mathbf{X}; \eta) = \mathbf{1}\{f(\mathbf{X}; \eta) > 0\}$, \mathcal{F} is a certain class covering the parameters of the ITRs, i.e., $\eta \in \mathcal{F}$.

Let $\hat{\tau}_i$ be a predictor of $\tau(\mathbf{X}_i)$. From Lemma 1, to estimate the optimal ITR, one can minimize the following empirical objective function:

$$\begin{aligned} n^{-1} \sum_{i \in \mathcal{F}_n} |\hat{\tau}_i| \{ \mathbf{1}\{\hat{\tau}_i > 0\} - d(\mathbf{X}_i; \eta) \}^2 &= n^{-1} \sum_{i \in \mathcal{F}_n} |\hat{\tau}_i| \mathbf{1}\{ \mathbf{1}\{\hat{\tau}_i > 0\} \neq d(\mathbf{X}_i; \eta) \} \\ &= n^{-1} \sum_{i \in \mathcal{F}_n} |\hat{\tau}_i| l_{0-1} \{ [2\mathbf{1}\{\hat{\tau}_i > 0\} - 1] f(\mathbf{X}_i; \eta) \} \\ &= n^{-1} \sum_{i \in \mathcal{F}_n} \mathcal{L}(\hat{\tau}_i; \eta, l_{0-1}). \end{aligned} \tag{4}$$

In this classification framework, we can also consider the three different approaches to estimating $\tau(\mathbf{X}_i)$, which have the equivalent forms and enjoy the same properties with the regression- and value-based estimator. For example, the AIPW-type predictor of $\tau(\mathbf{X}_i)$ is

$$\hat{\tau}_{aw,i} = \left[\frac{\mathbf{1}(A_i = 1) \{ Y_i - \hat{Q}(\mathbf{X}_i, 1) \}}{\hat{\pi}_A(\mathbf{X}_i)} + \hat{Q}(\mathbf{X}_i, 1) \right] - \left[\frac{\mathbf{1}(A_i = 0) \{ Y_i - \hat{Q}(\mathbf{X}_i, 0) \}}{1 - \hat{\pi}_A(\mathbf{X}_i)} + \hat{Q}(\mathbf{X}_i, 0) \right].$$

The solution obtained by minimizing (4) in η with replacing $\hat{\tau}_i$ with $\hat{\tau}_{aw,i}$ is equivalent to that obtained by maximizing the AIPW value estimator $\hat{V}_{aw}(d_\eta; \mathcal{F}_n)$ in (1) and enjoys the doubly robust property.

When the experimental sample is representative of the target population, the above ITR learning methods can estimate the population optimal ITRs that lead to the highest value. To correct for the selection bias of the experimental sample, we propose an integrative strategy that leverages the randomization in the experimental sample and the representativeness

of the RWD. As pointed out by Zhao et al. (2019), the optimal rule obtained from the experimental sample is also asymptotically optimal for the target population if no restriction to the class of d , but such optimal rules are too complex. Therefore, we focus on learning interpretable and parsimonious ITRs, and thus integrative methods are required to correct for the selection bias of the experimental sample.

3. Integrative transfer learning of ITRs

3.1 Assumptions for identifying population optimal ITRs

To identify population optimal ITRs, we assume the covariates \mathbf{X} can fully explain the selection mechanism of inclusion into the randomized experiment and the covariate distribution in the RWD is representative to the target population.

Assumption 2.—*a*: $\{Y^*(0), Y^*(1)\} \perp \perp S \mid \mathbf{X}$. *b*: $f(\mathbf{X} \mid S = 0) = f(\mathbf{X})$. *c*: There exist $c_3, c_4 \in (0, 1)$ such that $\pi_s(\mathbf{X}) \in [c_3, c_4]$ for all $\mathbf{X} \in \mathcal{X}$.

Assumption 2.a is commonly assumed in the generalizability literature (see Stuart et al. (2011); Buchanan et al. (2018)), which requires that the observed covariates \mathbf{X} are rich enough such that it can decide whether or not the subject is selected into the experimental data. Weaker versions of the assumption include the mean exchangeability, i.e., $\mathbb{E}\{Y^*(a) \mid \mathbf{X}, S = 0\} = \mathbb{E}\{Y^*(a) \mid \mathbf{X}, S = 1\}$, $a = 0, 1$ (Zhao et al. 2019; Dahabreh et al. 2019).

Assumption 2.b requires that covariate distribution in the RWD is the same as that in the target population so that we can leverage the representativeness of the RWD. As a result, it implies that $m^{-1} \sum_{i \in \mathcal{I}_m} g(\mathbf{X}_i)$ would be an unbiased estimator of $\mathbb{E}\{g(\mathbf{X})\}$, the covariate expectation in the target population for any g . Assumption 2.a and Assumption 2.b are also adopted in the covariate shift problem in machine learning, a sub-area in domain adaptation, where standard classifiers do not perform well when the independent variables in the training data have a different covariate distribution from that in the testing data. Re-weighting individuals to correct the over- or under-representation is one way to cope with the covariate shift in classification problems (Kouw and Loog 2018).

Assumption 2.c means that all subjects have nonzero probabilities of being selected into the experiment which require no existence of covariates that prohibit to participate the experiment, which is used to make $1/\pi_s(\mathbf{X})$ well-defined. If Assumption 2.c is violated, the support of covariates in two populations has limited overlap, so generalization can only be made to a restricted population without extrapolation (Yang and Ding 2018). This is also a common assumption in covariate shift literature (Uehara et al. 2020).

The following proposition identifies the population value function using the inverse of sampling score as the transfer weight, and leads to an equivalent weighted classification loss function to identify the population optimal ITR.

Proposition 1.—Let the transfer weight be $w = \pi_s^{-1}(\mathbf{X})$. Under Assumption 1.b, Assumption 2.a and Assumption 2.c, the population value of an ITR d_η is identified by $V(d_\eta) = \mathbb{E}[\omega SY^* \{d_\eta(\mathbf{X})\}]$. Then the population optimal ITR can be obtained by minimizing

$$\mathbb{E}(wS \mid \tau(\mathbf{X}) \mid [\mathbf{1}\{\tau(\mathbf{X}) > 0\} - d_\eta(\mathbf{X})]^2) \text{ in } \eta.$$

Two challenges arise for using Proposition 1 to estimate the optimal ITR: (i) the weights ω are unknown and (ii) the estimators of $\tau(\mathbf{X})$ discussed so far such as $\hat{\tau}_{aw}$ are biased based only on the experimental sample. In the following subsections, we present the estimation of the transfer weights and unbiased estimators of $\tau(\mathbf{X})$.

3.2 Estimation of transfer weights

We consider various parametric and nonparametric methods to estimate the transfer weights ω . Each method has its own benefits: parametric methods are easy to implement and efficient if the models are correctly specified, while nonparametric methods are more robust and less sensitive to model misspecification.

3.2.1 Parametric approach—For parametric approaches, we assume the target population size N is known. Similar to modeling propensity score $\pi_A(\mathbf{X})$, we can posit a logistic regression model for $\pi_s(\mathbf{X})$; i.e., $\text{logit}\{\pi_s(\mathbf{X}; \alpha)\} = \alpha_0 + \alpha_1^\top \mathbf{X}$, where $\alpha = (\alpha_0, \alpha_1^\top)^\top \in \mathbb{R}^{p+1}$. The typical maximum likelihood estimator of α can be obtained by

$$\begin{aligned} \hat{\alpha} &= \underset{\alpha}{\text{argmax}} \frac{1}{N} \sum_{i=1}^N [S_i \log \pi_s(\mathbf{X}_i; \alpha) + (1 - S_i) \log \{1 - \pi_s(\mathbf{X}_i; \alpha)\}] \\ &= \underset{\alpha}{\text{argmax}} \frac{1}{N} \sum_{i=1}^N [S_i (\alpha_0 + \alpha_1^\top \mathbf{X}_i) - \log \{1 + \exp(\alpha_0 + \alpha_1^\top \mathbf{X}_i)\}]. \end{aligned} \quad (5)$$

However, the second term $N^{-1} \sum_{i=1}^N \log \{1 + \exp(\alpha_0 + \alpha_1^\top \mathbf{X}_i)\}$ in (5) is not feasible to calculate because we can not observe \mathbf{X} for all individuals in the population. The key insight is that this term can be estimated by $m^{-1} \sum_{i \in \mathcal{J}_m} \log \{1 + \exp(\alpha_0 + \alpha_1^\top \mathbf{X}_i)\}$ based on the RWD. This strategy leads to our modified maximum likelihood estimator

$$\hat{\alpha} = \underset{\alpha}{\text{argmax}} \left[\frac{1}{N} \sum_{i=1}^N S_i (\alpha_0 + \alpha_1^\top \mathbf{X}_i) - \frac{1}{m} \sum_{i=1}^m \log \{1 + \exp(\alpha_0 + \alpha_1^\top \mathbf{X}_i)\} \right].$$

Alternatively, Assumption 2.b leads to unbiased estimating equations of α

$$\mathbb{E} \left\{ \frac{Sg(\mathbf{X})}{\pi_s(\mathbf{X}; \alpha)} \right\} = \mathbb{E} \{ g(\mathbf{X}) \},$$

for any g . Therefore, we can use the following estimating equations

$$\frac{1}{N} \sum_{i=1}^N \frac{S_i g(\mathbf{X}_i)}{\pi_S(\mathbf{X}_i; \alpha)} = \frac{1}{m} \sum_{i \in \mathcal{J}_m} g(\mathbf{X}_i)$$

for α , where $g(\mathbf{X}) \in \mathbb{R}^{p+1}$. For simplicity, one can take $g(\mathbf{X})$ as $\partial \log \{\pi_S(\mathbf{X}; \alpha)\} / \partial \alpha$.

3.2.2 Nonparametric approach—The above parametric approaches require $\pi_S(\mathbf{X}; \alpha)$ to be correctly specified and the population size N to be known. These requirements may be stringent, because the sampling mechanism into the randomized experiment is unknown and the population size is difficult to obtain. We now consider the constrained optimization algorithm of Wang and Zubizarreta (2020) to estimate the transfer weights by $\min_w \sum_{i=1}^N S_i w_i \log w_i$ subject to

$$\left| \sum_{i=1}^N w_i S_i g_k(\mathbf{X}_i) - \frac{1}{m} \sum_{i \in \mathcal{J}_m} g_k(\mathbf{X}_i) \right| \leq \sigma_k (k = 1, \dots, K), \tag{6}$$

Where $\sigma_k \geq 0$; $w_i \in [0, 1]$, $i \in \mathcal{J}_m$, $\sum_{i \in \mathcal{J}_m} w_i = 1$, and $\{g_1, \dots, g_K\}$ can be chosen as the first-, second- and higher order moments of the covariate distribution. The constants σ_k are the tolerate limits of the imbalances in g_k . When all the σ_k 's are taken 0, (6) becomes the exact balance, and the solution w becomes the entropy balancing weights (Hainmueller 2012). The choice of σ_k is related to the standard bias-variance tradeoff. On the one hand, if σ_k is too small, the weight distribution can have large variability in order to satisfy the stringent constraints, and in some extreme scenarios, the weights may not exist. On the other hand, if σ_k is too large, the covariate imbalances remain and therefore the resulting weights are not sufficient to correct for the selection bias. Following Wang and Zubizarreta (2020), we choose σ_k from a pre-specified set by the tuning algorithm described in Supplementary Materials S1.

3.3 Estimation of $\tau(\mathbf{X})$

The transfer weights are also required to estimate $\tau(\mathbf{X})$ from the experimental data. In parallel to Section 2.2, we can obtain the weighted estimators of $\tau(\mathbf{X})$. For example, the transfer AIPW-type predictor of $\tau(\mathbf{X}_i)$

$$\hat{\tau}_{aw,i}^w = \left[\frac{\mathbf{1}(A_i = 1) \{Y_i - \hat{Q}_w(\mathbf{X}_i, 1)\}}{\hat{\pi}_A^w(\mathbf{X}_i)} + \hat{Q}_w(\mathbf{X}_i, 1) \right] - \left[\frac{\mathbf{1}(A_i = 0) \{Y_i - \hat{Q}_w(\mathbf{X}_i, 0)\}}{1 - \hat{\pi}_A^w(\mathbf{X}_i)} + \hat{Q}_w(\mathbf{X}_i, 0) \right],$$

where $\hat{Q}_w(\cdot, a)$ can be estimated by weighted least square for linear Q functions with the estimated transfer weights, and $\hat{\pi}_A^w(\mathbf{X}_i) = \hat{P}_w(A_i = 1 | \mathbf{X}_i)$ is the weighted propensity score which can be obtained by the weighted regression model with the estimated transfer weights using the experimental data.

Let $\hat{\tau}_i^w$ be a generic weighted predictor of $\tau(\mathbf{X})$. Proposition 1 implies an empirical risk function to estimate η

$$\frac{1}{n} \sum_{i \in \mathcal{J}_n} \hat{w}_i |\hat{\tau}_i^w| \{ \mathbf{1}(\hat{\tau}_i^w > 0) - d(\mathbf{X}_i; \eta) \}^2 = \frac{1}{n} \sum_{i \in \mathcal{J}_n} \hat{w}_i |\hat{\tau}_i^w| l_{0-1} [\{ 2\mathbf{1}(\hat{\tau}_i^w > 0) - 1 \} f(\mathbf{X}_i; \eta)]. \tag{7}$$

It is challenging to optimize the objective function (7) since it involves a non-smooth non-convex objective function of η . One way is to optimize it directly using grid search or the genetic algorithm in Zhang, Tsiatis, Laber and Davidian (2012) for learning linear decision rules (implemented with rgenoud in R), but grid search becomes untenable for a higher dimension of η , and the genetic algorithm cannot guarantee a unique solution. Alternatively, one can replace l_{0-1} with a surrogate loss function, such as a hinge loss used in the support vector machine (Vapnik 1998) and outcome weighted learning (Zhao et al. 2012), or a ramp loss (Collobert et al. 2006) which truncates the unbounded hinge loss by trading convexity for robustness to outliers (Wu and Liu 2007). We adopt the smoothed ramp loss function proposed in Zhou et al. (2017)

$$l(u) = 0 \text{ if } u \geq 1, (1 - u)^2 \text{ if } 0 \leq u < 1, 2 - (1 + u)^2 \text{ if } -1 \leq u < 0, 2 \text{ if } u \leq -1. \tag{8}$$

which retains the robustness and gains computational advantages for being smooth everywhere. $l(u)$ is the difference of two smooth convex functions, $l(u) = l_1(u) - l_0(u)$, where

$$l_1(u) = 0 \text{ if } u \geq s; (s - u)^2 \text{ if } s - 1 \leq u < s; 2s - 2u - 1 \text{ if } u < s - 1. \tag{9}$$

Similarly to Zhou et al. (2017), we can apply the d.c. algorithm (Le Thi Hoai and Tao 1997) to solve the non-convex minimization problem by first decomposing the original objective function into the summation of a convex component and a concave component and then minimizing a sequence of convex subproblems. Applying (9) to (7), our empirical risk objective becomes

$$\underbrace{\frac{1}{n} \sum_{i \in \mathcal{J}_n} \hat{w}_i |\hat{\tau}_i^w| l_1(u_i)}_{\text{a convex function of } \eta} + \underbrace{\frac{1}{n} \sum_{i \in \mathcal{J}_n} \{-\hat{w}_i |\hat{\tau}_i^w| l_0(u_i)\}}_{\text{a concave function of } \eta}, \tag{10}$$

where $u_i = \{ 2\mathbf{1}(\hat{\tau}_i^w > 0) - 1 \} f(\mathbf{X}_i; \eta)$. Let $C_i^{\text{cav}} = -\hat{w}_i |\hat{\tau}_i^w| l_0(u_i)$ and

$$\xi_i = \frac{\partial C_i^{\text{cav}}}{\partial u_i} = -\hat{w}_i |\hat{\tau}_i^w| \frac{\partial l_0(u_i)}{\partial u_i}, \tag{11}$$

for $i \in \mathcal{J}_n$. Therefore, the convex subproblem at $(t + 1)$ th iteration in the d.c. algorithm is to

$$\min_{\eta} \frac{1}{n} \sum_{i \in \mathcal{J}_n} \hat{w}_i |\hat{\tau}_i^w| l_1(u_i) + \frac{1}{n} \sum_{i \in \mathcal{J}_n} \xi_i^{(t)} u_i, \tag{12}$$

which is a smooth unconstrained optimization problem and can be solved with many efficient algorithms such as L-BFGS (Nocedal 1980). Algorithm 2 in the Supplementary Material summarizes the proposed procedure to obtain the minimizer $\hat{\eta}$.

3.4 Cross-validation

We have discussed various approaches for the estimation of transfer weights and the individual contrast function. Although transfer weights can correct the selection bias of the experimental data, they also increase the variance of the estimator compared to the unweighted counterpart. The effect of weighting on the precision of estimators can be quantified by the effective sample size originated in survey sampling (Kish and Stat 1992). If the true contrast function can be approximated well by linear models, the estimated ITR with transfer weights may not outperform the one without weighting. As a result, for a given application, it is critical to developing a data-adaptive procedure to choose the best method among the aforementioned weighted estimators and whether or not to use transfer weights.

Toward this end, we propose a cross-validation procedure. To proceed, we split the sample into two parts, one for learning the ITRs and the other for evaluating the estimated ITRs. We also adopt multiple sample splits to increase the robustness since a single split might have dramatic varying results due to randomness. For evaluation of an ITR d_η , we estimate the value for the given ITR by the weighted AIPW estimator in (13) with the estimated nonparametric transfer weights \hat{w} ,

$$\hat{V}(d_\eta; \hat{w}, \mathcal{F}_n) = \sum_{i \in \mathcal{F}_n} \hat{w}_i \left(\left[\frac{A_i d(\mathbf{X}_i; \eta)}{\hat{\pi}_A^w(\mathbf{X}_i)} + \frac{(1 - A_i) \{1 - d(\mathbf{X}_i; \eta)\}}{1 - \hat{\pi}_A^w(\mathbf{X}_i)} \right] \right. \\ \left. [Y_i - \hat{Q}_w\{\mathbf{X}_i, d(\mathbf{X}_i; \eta)\}] - \hat{Q}_w\{\mathbf{X}_i, d(\mathbf{X}_i; \eta)\} \right). \tag{13}$$

Since we are more interested in the comparisons among different weighting and unweighting methods instead of different contrast function estimation methods, thus we fix an estimation method for the contrast function at first, and then use cross-validation to choose among the set of candidate methods \mathcal{S} . This \mathcal{S} contains the proposed learning method with nonparametric transfer weights described in Section 3.2.2, and unweighted estimation method. If the target population size N is available, we can also include the weighted estimators with the two parametric approaches introduced in Section 3.2.1 into \mathcal{S} . Then we can get the value estimates averaged over all the sample splits for each learning method in \mathcal{S} , and return the method with the highest value estimates. The more explicit description of the procedure is presented in Algorithm 3 in the Supplementary Material.

4 Asymptotic Properties

In this section, we provide a theoretical guarantee that the estimated ITR within a certain class \mathcal{F} is consistent with the true population optimal ITR within \mathcal{F} . Toward this end, we show that the expected risk of the estimated ITR is consistent with that of the true optimal ITR within \mathcal{F} . In particular, we consider the ITRs learned by the nonparametric transfer weights, weighted AIPW estimator $\hat{\tau}_{aw}^w$ with the logistic regression model $\pi_A(\mathbf{X}; \gamma)$

and the linear Q function. The extensions to other cases with parametric transfer weights and regression and IPW estimators are straightforward.

We introduce more notations. Let $\theta = (\gamma^\top, \beta^\top)^\top$ and let the contrast function

$$\tau_{aw}^\theta(\mathbf{X}) = \frac{\mathbf{1}(A = 1)\{Y - Q(\mathbf{X}, 1; \beta)\}}{\pi_A(\mathbf{X}; \gamma)} + Q(\mathbf{X}, 1; \beta) - \frac{\mathbf{1}(A = 0)\{Y - Q(\mathbf{X}, 0; \beta)\}}{1 - \pi_A(\mathbf{X}; \gamma)} - Q(\mathbf{X}, 0; \beta). \tag{14}$$

Denote the general smooth ramp loss function $l^\zeta(u) = l(\zeta_1 u)/\zeta_2$, where $\zeta_1, \zeta_2 > 0$, l is defined in (8), so it is easy to see that when $\zeta = (\zeta_1, \zeta_2) \rightarrow \zeta^* \triangleq (+\infty, 2)$, the general smooth ramp loss function will converge to the 0–1 loss function, i.e., $l^\zeta(u) \rightarrow l_{0-1}(u), \forall u$. Note that the general smooth ramp loss function can also be written as the difference of two convex functions, $l^\zeta(u) = l_1(\zeta_1 u)/\zeta_2 - l_0(\zeta_1 u)/\zeta_2$, thus the d.c. algorithm in Algorithm 2 can be easily generalized to $l^\zeta(u)$. Recalling the loss function $\mathcal{L}(\tau; \eta, l_{0-1})$ and the corresponding \mathcal{F} -risk $\mathcal{R}_{\mathcal{F}}(\tau; \eta, l_{0-1})$ defined in (3), we replace τ with τ_{aw}^θ and obtain:

$$\mathcal{R}_{\mathcal{F}}(\tau_{aw}^\theta; \eta, l_{0-1}) = \mathbb{P}\{\mathcal{L}(\tau_{aw}^\theta; \eta, l_{0-1})\}, \mathcal{L}(\tau_{aw}^\theta; \eta, l_{0-1}) = |\tau_{aw}^\theta(\mathbf{X})|_{l_{0-1}}([\mathbf{21}\{\tau_{aw}^\theta(\mathbf{X}) > 0\} - 1]f(\mathbf{X}; \eta)), \forall \eta \in \mathcal{F},$$

denoted as $\mathcal{R}_{\mathcal{F}}(\eta)$ and $\mathcal{L}(\eta, \theta)$ respectively for simplicity. We define the minimal \mathcal{F} -risk as $\mathcal{R}_{\mathcal{F}}^* = \inf_{\eta \in \mathcal{F}} \mathcal{R}_{\mathcal{F}}(\eta)$. Similarly, replacing l_{0-1} in $\mathcal{R}_{\mathcal{F}}(\eta)$ and $\mathcal{L}(\eta, \theta)$ with the general smooth ramp loss l^ζ , we can obtain $\mathcal{R}_{\mathcal{F}}(\tau_{aw}^\theta; \eta, l^\zeta)$ and $\mathcal{L}(\tau_{aw}^\theta; \eta, l^\zeta)$, denoted as $\mathcal{R}_{\zeta\mathcal{F}}(\eta)$ and $L^\zeta(\eta, \theta)$, respectively. Likewise, the minimal \mathcal{F} -risk with l^ζ is defined as $\mathcal{R}_{\zeta\mathcal{F}}^* = \inf_{\eta \in \mathcal{F}} \mathcal{R}_{\zeta\mathcal{F}}(\eta)$.

The goal is to show that the decision rule is \mathcal{F} -consistency (Bartlett et al.

2006), i.e., $\lim_{N \rightarrow \infty} \mathcal{R}_{\mathcal{F}}(\hat{\eta}) = \mathcal{R}_{\mathcal{F}}^*$ Where $\hat{\eta} = \operatorname{argmin}_{\eta \in \mathcal{F}} \sum_{i=1}^N S_i \hat{w}_i L^\zeta(\eta, \hat{\theta}_w)$ with $\hat{\theta}_w = (\hat{\gamma}_w^\top, \hat{\beta}_w^\top)^\top$.

In particular, $\hat{\gamma}_w = \operatorname{argmax}_{\gamma} \sum_{i=1}^N \hat{w}_i S_i [A_i \log \pi_A(\mathbf{X}_i; \gamma) + (1 - A_i) \log \{1 - \pi_A(\mathbf{X}_i; \gamma)\}]$ converges in probability to $\operatorname{argmax}_{\gamma} \mathbb{E}[A \log \pi_A(\mathbf{X}; \gamma) + (1 - A) \log \{1 - \pi_A(\mathbf{X}; \gamma)\}] \triangleq \gamma^*$, $\hat{\beta}_w$ converges in probability to $\operatorname{argmin}_{\beta} \sum_{i: A_i = a | S_i = 1} \hat{w}_i S_i (Y_i - \beta^\top \mathbf{X}_i)^2$

probability to $\{\mathbb{E}(\mathbf{X}\mathbf{X}^\top)\}^{-1} \times \mathbb{E}\{\mathbf{X}Y^*(a)\} \triangleq \beta_a^*$, for $a = 0, 1$, and thus $\hat{\beta}_w = (\hat{\beta}_{w,1}^\top, \hat{\beta}_{w,0}^\top)^\top$ converges to $\beta^* = (\beta_1^{*\top}, \beta_0^{*\top})^\top$. The convergences of $\hat{\gamma}_w$ and $\hat{\beta}_w$ are implied from the weak law of large numbers and the result $\max_i |N\hat{w}_i - 1/\pi_S(\mathbf{X}_i)| = o_p(1)$ proved in Theorem 2 in Wang and Zubizarreta (2020). Let $\theta^* = (\gamma^{*\top}, \beta^{*\top})^\top$, thus we have

$$\hat{\theta}_w \xrightarrow{p} \theta^*. \tag{15}$$

It is easy to see that $\mathcal{R}_{\zeta\mathcal{F}}(\eta)$ will converge to $\mathcal{R}_{\mathcal{F}}(\eta)$ when $\zeta \rightarrow \zeta^*$.

Therefore, to obtain \mathcal{F} -consistency, it suffices to show that $\lim_{N \rightarrow \infty} \mathcal{R}_{\zeta\mathcal{F}}(\hat{\eta}) = \mathcal{R}_{\zeta\mathcal{F}}^*$

which is stated in the theorem under the regularity conditions below. Define $m^*(\mathbf{x}) = (\rho')^{-1}[1/\{N\pi(\mathbf{x})\}]$, $h(\mathbf{x}) = \mathbb{E}\{L_i^\zeta(\eta, \theta^*) \mid \mathbf{x}\}$ and denote $\pi_i(\mathbf{X}_i)$ as π_i for simplicity. The following regularity conditions follow from Wang and Zubizarreta (2020) to ensure the asymptotic properties of nonparametric weights.

Assumption 3.

a: There exists the limit of $\hat{\eta}$, denoted as η^ . b: $|\partial L_i^\zeta(\eta, \theta)/\partial \theta|_{\theta=\theta^*} \leq M_1 < \infty$, $|\partial L_i^\zeta(\eta, \theta)/\partial \eta|_{\eta=\eta^*} \leq M_2 < \infty$. c: $\mathbb{P}L^\zeta(\eta, \theta^*)$ is continuous in η . d: $\mathbb{E}|L_i^\zeta(\eta, \theta^*) - \mathbb{E}\{L_i^\zeta(\eta, \theta^*) \mid \mathbf{X}\}| < \infty$. e: $\rho'(\cdot)$ is bounded away from zero and Lipschitz. f: $m^*(\cdot) \in \mathcal{M}$ and $h(\cdot) \in \mathcal{H}$, where \mathcal{M} and \mathcal{H} are two sets of smooth functions satisfying $\log N_{[]}(\epsilon, \mathcal{M}, L_2(P)) \leq c(1/\epsilon)^{1/k_1}$ and $\log N_{[]}(\epsilon, \mathcal{H}, L_2(P)) \leq c(1/\epsilon)^{1/k_2}$ for a positive constant c and $k_1, k_2 > 1/2$, with $N_{[]}(\epsilon, \mathcal{S}, L_2(P))$ denoting the covering number of the set \mathcal{S} by ϵ -brackets. g: There exist $r_x > 1/2$ and λ_2^* such that $h(\mathbf{X})$ satisfies $\sup_{\mathbf{X} \in \mathcal{X}} |h(\mathbf{X}) - B(\mathbf{X})^\top \lambda_2^*| = O(K^{-r_x})$. h: $K^{-r_x - r_y + 1/2} = o(1)$. i: $\|\sigma\|^2 = o_p(1)$. j: $m^{-1} \sum_{i \in \mathcal{J}_m} B(\mathbf{X}_i) - N^{-1} \sum_{i=1}^N B(\mathbf{X}_i) = o_p(1)$.*

Theorem 1.

Under Assumptions 2–3, we have $\lim_{N \rightarrow \infty} \mathcal{R}_{\zeta, \mathcal{F}}(\hat{\eta}) = \mathcal{R}_{\zeta, \mathcal{F}}^$.*

Our asymptotic results require n and m both to go to infinity. If $m \nrightarrow \infty$, $m^{-1} \sum_{i \in \mathcal{J}_m} g(\mathbf{X}_i)$ for any function $g(\cdot)$ will not converge to $\mathbb{E}\{g(\mathbf{X})\}$. If $n \nrightarrow \infty$, the estimated transferring weights will not converge to the true ones. In either case, the final estimated optimal ITRs will not be consistent for the ones of the target population.

5 Simulation study

We evaluate the finite-sample performances of the proposed estimators with a set of simulation studies. We first generate a target population of size $N = 10^6$. The covariates are generated by $\mathbf{X}_i = (X_{i1}, X_{i2})^\top \sim N(\mathbf{1}_{2 \times 1}, I_{2 \times 2})$, and the potential outcome is generated by $Y_i^*(a) \mid \mathbf{X}_i = 1 + 2X_{i1} + 3X_{i2} + a\tau(\mathbf{X}_i) + \epsilon_i(a)$, where $\epsilon_i(a)$ i.i.d. follows $N(0, 0.5^2)$ for $a = 0, 1$. We consider three contrast functions: (I). $\tau(\mathbf{X}_i) = \arctan\{\exp(1 + X_{i1}) - 3X_{i2} - 5\}$; (II). $\tau(\mathbf{X}_i) = \cos(1) + \cos(X_{i1}) + \cos(X_{i2}) - 3/2$; (III). $\tau(\mathbf{X}_i) = 1 + 2X_{i1} + 3X_{i2}$.

To evaluate the necessity and performance of weighting methods, we consider two categories of data-generating models. One is the scenarios where the true optimal policies are not in linear forms, i.e., the settings (I) and (II), while the other one is considering the true optimal policy is in a linear form as the setting (III). We start with the first two scenarios. From the data visualization, we observed that the covariates shifted between the experimental data and RWD to a varying extent in the two generative models: the covariates generated by (I) have similar distributions between the experimental data and RWD, while those generated by (II) have very different ones. Then we generate the

RWD by random sample $m = 5000$ subjects from the target population $\{\mathbf{X}_i\}_{i=1}^N$, denoted as $\{\mathbf{X}_i\}_{i \in \mathcal{J}_m}$. To form the experimental sample, we generate the indicator of selection according to $S_i | \mathbf{X}_i \sim \text{Bernoulli}[\exp(\alpha_0 + \alpha_1^\top \mathbf{X}_i) / \{1 + \exp(\alpha_0 + \alpha_1^\top \mathbf{X}_i)\}]$, where $\alpha_1 = (1, -2)^\top$ and $\alpha_0 = -8$ which gives the average sample sizes of the experimental data over the simulation replicates round 1386, respectively. In the experimental data denoted as \mathcal{J}_n , the treatments are randomly assigned according to $A_i \sim \text{Bernoulli}(0.5)$, and then the actual observed outcomes $Y_i = Y_i^*(A_i), i \in \mathcal{J}_n$.

For each setting, we compare the following estimators:

1. “ W_1 ”: our proposed weighted estimator with parametric weights learned by the modified maximum likelihood introduced in Section 3.2.1.
2. “ W_2 ”: our proposed weighted estimator with parametric weights learned by the estimating equations introduced in Section 3.2.1.
3. “ cv ”: our proposed weighted estimator using the cross-validation procedure introduced in Section 3.4, and we set the number of sample splits as 10 in all the simulation studies.
4. “ np ”: our proposed weighted estimator with nonparametric weights.
5. “ $unweight$ ”: only using the experimental data to learn the optimal linear ITRs.
6. “ bm ”: using the weighted estimator but we replace the estimated transfer weights \hat{w}_i with the normalized true weights $w_i = \{\pi_s(\mathbf{X}_i)\}^{-1} / \sum_{i \in \mathcal{J}_n} \{\pi_s(\mathbf{X}_i)\}^{-1}$, and replace the estimated contrast functions $\hat{\tau}_i$ with the true contrast values $Y_i^*(1) - Y_i^*(0)$.
7. “ $Imai$ ”: the method proposed in Imai et al. (2013). They construct the transfer weights by fitting a Bayesian additive regression trees (BART) model using covariates as predictors and whether to be in the experimental data as outcomes, and then use a variant of SVM to estimate the conditional treatment effect (implemented in the R package FindIt). Then given \mathbf{X} , we can assign treatment 1 if its estimated conditional treatment effect is positive, and 0 otherwise.
8. “ DBN ”: the method proposed in Mo et al. (2020), where they propose to use distributionally robust ITRs, that is, value functions are evaluated under all testing covariate distributions that are “close” to the training distribution, and the worse-case takes a minimal one, and then the ITRs are learned by maximizing the defined worst-case value function.

For the estimators 1. – 5. in the above, we estimate the contrast function using τ_{aw} with linear Q functions and constant propensity scores estimated via the proportion of $A = 1$ in the experimental data, $n^{-1} \sum_{i \in \mathcal{J}_n} A_i$.

To study the impact of model misspecification of $\pi_s(\mathbf{X})$, we consider to estimate α by fitting logistic regression wrongly: $\pi_s(\mathbf{X}; \alpha) = \exp(\alpha_0 + \alpha_1^\top \mathbf{X}) / \{1 + \exp(\alpha_0 + \alpha_1^\top \mathbf{X})\}$, where $\mathbf{X} = (\mathbf{X}_1^\top, \mathbf{X}_2^\top)^\top$. And as described in the proposed methods, only estimators 1. – 3. need to specify the

sampling model. To measure the performance of the different approaches, we use the mean value MSE defined as $N^{-1} \sum_{i=1}^N (Y_i^* \{d_i(\mathbf{X}_i)\} - Y_i^* [\mathbf{1}\{\tau(\mathbf{X}_i) > 0\}])^2$. Thus the lower the MSE, the better performance the approach has.

The results are summarized in Figure 2. For the generative model (I) shown in Figure 2, unweighted and DBN rules perform worst in all the cases. The second parametric weighting method is relatively robust compared with the first parametric one when the sampling model is misspecified; this is reasonable since the first one uses maximum likelihood estimation which is more reliable on the correctness of the parametric model specification. When the training data and testing data are similarly distributed, the performances of the parametric, nonparametric weighting methods and the cross-validation procedure are similar; in this case, sampling model misspecification does not have much influence on the rule learning.

For the generative model (II) shown in Figure 2, among the performances of weighted and unweighted estimators, we can see the weighted estimators outperform the unweighted estimators when the model correctly specified; the nonparametric weighted estimator outperforms other estimators and the performance of the parametric weighted estimators has dramatically decreased when the sampling model was misspecified. These findings make sense, because when the more different the training and testing data are, the more crucial the bias correction is. Thus sampling model misspecification might have a huge influence on the learned rules.

From both generative models, we can see that the nonparametric weighting and cross-validation procedure are more robust when the sampling model is unknown and recommended in practice use. As for DBN estimator, it performs badly in terms of MSE magnitude or variability, one possible reason might be that the tuning parameters when measuring the "closeness" of two distributions need to be chosen more carefully in practice (Mo et al. 2020). As for the comparison Imai also has relatively good performance among all estimators and has relatively small variability.

As for setting (III), we consider the correctly specified sample score model. As shown in Figure 3, the unweighted method performs the best, close to the benchmark. This demonstrates that when the true optimal policies are in linear forms, then learning with only the experimental data performs the best. Thus weighting is not a necessity, and it even will induce the variance of the estimators. The other competitive estimator is using cross-validation. Therefore, when we have no idea of the true data generating models, using the CV is a good choice in practice.

6 Real Data Application

We apply the proposed methods to personalized recommendations of a job training program on post-market earning. Our analysis is based on two data sources (downloaded from <https://users.nber.org/~rdehejia/data/.nswdata2.html>): an experimental dataset from the National Supported Work (NSW) program (LaLonde 1986) and a non-experimental dataset from the Current Population Survey (CPS; LaLonde (1986)). The NSW dataset includes 185 individuals who received the job training program ($A = 1$) and 260 individuals who did not

receive the job training program ($A = 0$), and the CPS dataset includes 15992 individuals without the job training program. Let \mathbf{X} be a vector including an intercept 1 and the pre-intervention covariates: age, education, Black (1 if black, 0 otherwise), Hispanic (1 if Hispanic, 0 otherwise), married (1 if married, 0 otherwise), whether having high school degree (1 if no degree, 0 otherwise), earning in 1974, and earning in 1975. The outcome variable is the earning in 1978. We transfer the earnings in the year 1974, 1975, and 1978 by taking the logarithm of the earnings plus one, denoted as “log.Re74”, “log.Re75” and “log.Re78”. We aim to learn a linear rule $d(\mathbf{X}; \eta) = \mathbf{1}(\mathbf{X}^\top \eta > 0)$ tailoring the recommendation of the job training program to an individual with characteristics \mathbf{X} .

The CPS sample is generally representative of the target population; while the participants in the NSW program may not represent the target population. Table 1 contrasts sample averages of each pre-intervention covariate in the NSW and CPS samples and shows that the covariates are highly imbalanced between the two samples. On average, the NSW sample is younger, less educated, more black people, fewer married, more without high school degree, fewer earnings in 1974 and 1975. Because of imbalance, ITRs cannot be learned based only on the NSW sample, thus we are motivated to use the proposed weighted transfer learning approaches to learn optimal linear ITRs which benefit the target population most.

We consider the unweighted and weighted transfer learning methods. For the weighted transfer learning method, we consider the nonparametric weighting approach, because the parametric approach requires the target population size N which is unknown in this application, and use the AIPW estimator $\hat{\tau}_{aw}$ with linear Q functions. We first normalize all the covariates in the combined two data sets for a better learning process. We also implement the cross-validation procedure to choose between the unweighted and weighted transfer learning methods.

Table 1 presents the estimated linear coefficients of each corresponding variable in the linear decision rule from the weighted transfer and unweighted learning methods, respectively. Moreover, from the cross-validation procedure, the weighted transfer learning method produces a smaller population risk than the unweighted learning method and therefore is chosen as the final result. Compared with the learned decision rule with only the experiment data NSW which oversamples the low-income individuals, younger, less educational or black individuals are more likely to have higher outcomes if offered the job training program in the target population.

To assess the performances of the learned rules for the target population, following Kallus (2017), we create a test sample that is representative of the target population. We sample 100 individuals at random from the CPS sample to form a control group and then find their closest matches in the NSW treated sample to form the treatment group (use `pairmatch` in R package `optmatch`). The test sample thus contains the 100 matched pairs, mimicking an experimental sample where binary actions are randomly assigned with equal probabilities. We use random forest on the test data to estimate Q function $\mathbb{E}(Y | \mathbf{X}, a)$, denoted the estimator as $\hat{Q}_{\text{test}}(\cdot, a)$, $a = 0, 1$, then we can use $\sum_{i \in \text{CPS}} \hat{Q}_{\text{test}}\{\mathbf{X}_i, d(\mathbf{X}_i; \eta)\} / |\text{CPS}|$ to estimate the value of an ITR $d(\cdot; \eta)$, where $|\text{CPS}|$ is the sample size of CPS. We replicate this

sampling test data procedure 100 times, estimate values of the three ITRs for each replicate, and get the averaged estimated values (standard error) as follows: 6.73 (0.01) and 6.39 (0.01) for $d(\cdot; \hat{\eta}_w)$ and $d(\cdot; \hat{\eta}_a)$, thus the transfer weighted rule has the higher value estimate. Although this evaluation method needs assumption of no unmeasured confounders for the non-experimental data CPS, which is not needed in our learning procedure, it can still bring us confidence, to some degree, that the weighted linear rule has better generalizability to the target population.

7 Discussion

We develop a general framework to learn the optimal interpretable ITRs for the target population by leveraging no unmeasured confounding in the experimental data and the representativeness of covariates in the RWD. The proposed procedure is easy to implement. By constructing transfer weights to correct the distribution of covariates in the experimental data, we can learn the optimal interpretable ITRs for the target population. Besides, we provide a data-adaptive procedure to choose among different weighting methods by modified cross-validation. Moreover, we establish the theoretical guarantee for the consistency of the proposed estimator of ITRs.

For data integration problems, there are generally two study designs: nested and non-nested ones (Dahabreh et al. 2019). For the nested designs, the experimental data is a subsample from the RWD (e.g., comprehensive cohort studies or when the study is nested in a larger healthcare system that forms the target population), while in the non-nested designs, the experimental data and the RWD are separate (e.g., different geographic locations) (see Dahabreh et al. 2021; Steingrímsson et al. 2021; Lee et al. 2021, for details). Although our simulation and data application consider non-nested studies, our proposed method is equally applicable to nested studies. The key assumption is that the domain of covariates in the target and experimental populations is the same. In the nested designs, the target population becomes the population that the RWD represents, which can be different from the experimental population.

Some future directions worth considering. One is to quantify the uncertainty of the estimated ITRs and therefore to conduct statistical inference. Moreover, in some scenarios, beside the covariate information, comparable treatment and outcome information are also available in the RWD. Unlike the randomized experiment, the treatment assignment in the RWD is determined by the preference of physicians and patients. In this case, an interesting topic is to develop integrative methods that can leverage the strengths of both data sources for efficient estimation of the ITRs. Besides, we only consider binary treatments in this paper to illustrate the proposed framework. However, it can be extended to the case with multiple treatments. To do this, we can still use the same transferring weights and apply them to the value function for the multiple treatments (Zhang et al.; 2012). The additional difficulty to overcome is how to optimize this nonsmooth objective function. One possible direction is to reformulate the problem as a weighted multi-classification problem, which can be an interesting future research topic. Finally, Assumption 2.a is plausible if \mathbf{X} captures all variables that related to the experiment participation and outcome (Buchanan et al. 2018) or the sampling mechanisms of the study samples depend only on \mathbf{X} . Similar to the no

unmeasured confounders assumption in causal inference from observational studies, this assumption is generally not testable empirically but needs domain knowledge to evaluate its plausibility. In the future, we will develop sensitivity analysis to assess the robustness of the study conclusion to the unverifiable assumption.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors gratefully acknowledge *NSF grant DMS 1811245, NCI grant P01 CA142538, NIA grant 1R01AG066883, and NIEHS grant 1R01ES031651.*

References

- Bartlett PL, Jordan MI and McAuliffe JD (2006). Convexity, classification, and risk bounds, *Journal of the American Statistical Association* 101(473): 138–156.
- Buchanan AL, Hudgens MG, Cole SR, Mollan KR, Sax PE, Daar ES, Adimora AA, Eron JJ and Mugavero MJ (2018). Generalizing evidence from randomized trials using inverse probability of sampling weights, *Journal of the Royal Statistical Society. Series A (Statistics in Society)* 181(4): 1193. [PubMed: 30555215]
- Chakraborty B. (2013). *Statistical methods for dynamic treatment regimes*, Springer.
- Cole SR and Stuart EA (2010). Generalizing evidence from randomized clinical trials to target populations: the actg 320 trial, *American journal of epidemiology* 172(1): 107–115. [PubMed: 20547574]
- Collobert R, Sinz F, Weston J. and Bottou L. (2006). Trading convexity for scalability, *Proceedings of the 23rd international conference on Machine learning*, pp. 201–208.
- Dahabreh IJ, Haneuse SJA, Robins JM, Robertson SE, Buchanan AL, Stuart EA and Hernán MA (2021). Study designs for extending causal inferences from a randomized trial to a target population, *American journal of epidemiology* 190(8): 1632–1642. [PubMed: 33324969]
- Dahabreh IJ, Robertson SE, Tchetgen EJ, Stuart EA and Hernán MA (2019). Generalizing causal inferences from individuals in randomized trials to all trial-eligible individuals, *Biometrics* 75(2): 685–694. [PubMed: 30488513]
- Gail M. and Simon R. (1985). Testing for qualitative interactions between treatment effects and patient subsets., *Biometrics* 41(2): 361–372. [PubMed: 4027319]
- Greenland S. (1990). Randomization, statistics, and causal inference, *Epidemiology* pp. 421–429. [PubMed: 2090279]
- Hainmueller J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies, *Political analysis* pp. 25–46.
- Hartman E, Grieve R, Ramsahai R. and Sekhon JS (2015). From sample average treatment effect to population average treatment effect on the treated: combining experimental with observational studies to estimate population treatment effects, *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 178(3): 757–778.
- Horvitz DG and Thompson DJ (1952). A generalization of sampling without replacement from a finite universe, *Journal of the American statistical Association* 47(260): 663–685.
- Imai K, Ratkovic M. et al. (2013). Estimating treatment effect heterogeneity in randomized program evaluation, *The Annals of Applied Statistics* 7(1): 443–470.
- Kallus N. (2017). Recursive partitioning for personalization using observational data, *Proceedings of the 34th International Conference on Machine Learning-Volume 70, JMLR. org*, pp. 1789–1798.
- Kang C, Janes H. and Huang Y. (2014). Combining biomarkers to optimize patient treatment recommendations, *Biometrics* 70(3): 695–707. [PubMed: 24889663]

- Kish L. and Stat JO (1992). Weighting for unequal pi.
- Kouw WM and Loog M. (2018). An introduction to domain adaptation and transfer learning, arXiv preprint arXiv:1812.11806.
- Laber E. and Zhao Y. (2015). Tree-based methods for individualized treatment regimes, *Biometrika* 102(3): 501–514. [PubMed: 26893526]
- Lakkaraju H. and Rudin C. (2017). Learning cost-effective and interpretable treatment regimes, *Artificial Intelligence and Statistics*, pp. 166–175.
- LaLonde RJ (1986). Evaluating the econometric evaluations of training programs with experimental data, *The American economic review* pp. 604–620.
- Le Thi Hoai A. and Tao PD (1997). Solving a class of linearly constrained indefinite quadratic problems by dc algorithms, *Journal of global optimization* 11(3): 253–285.
- Lee D, Yang S, Dong L, Wang X, Zeng D. and Cai J. (2021). Improving trial generalizability using observational studies, *Biometrics*.
- Mo W, Qi Z. and Liu Y. (2020). Learning optimal distributionally robust individualized treatment rules, *Journal of the American Statistical Association* pp. 1–16.
- Moodie EE, Dean N. and Sun YR (2014). Q-learning: Flexible learning about useful utilities, *Statistics in Biosciences* 6(2): 223–243.
- Murphy SA (2005). A generalization error for q-learning, *Journal of Machine Learning Research* 6(Jul): 1073–1097. [PubMed: 16763665]
- Nocedal J. (1980). Updating quasi-newton matrices with limited storage, *Mathematics of computation* 35(151): 773–782.
- Qian M. and Murphy SA (2011). Performance guarantees for individualized treatment rules, *Annals of statistics* 39(2): 1180. [PubMed: 21666835]
- Rothwell PM (2005). External validity of randomised controlled trials: “to whom do the results of this trial apply?”, *The Lancet* 365(9453): 82–93.
- Rubin DB (1974). Estimating causal effects of treatments in randomized and nonrandomized studies., *Journal of educational Psychology* 66(5): 688.
- Steingrimsson JA, Gatsonis C. and Dahabreh IJ (2021). Transporting a prediction model for use in a new target population, arXiv preprint arXiv:2101.11182.
- Stuart EA, Cole SR, Bradshaw CP and Leaf PJ (2011). The use of propensity scores to assess the generalizability of results from randomized trials, *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 174(2): 369–386.
- Tao Y, Wang L. and Almirall D. (2018). Tree-based reinforcement learning for estimating optimal dynamic treatment regimes, *The Annals of Applied Statistics* 12(3): 1914–1938. [PubMed: 30984321]
- Uehara M, Kato M. and Yasui S. (2020). Off-policy evaluation and learning for external validity under a covariate shift, *Advances in Neural Information Processing Systems* 33: 49–61.
- Vapnik VN (1998). *Statistical Learning Theory*, Wiley-Interscience.
- Wang Y. and Zubizarreta JR (2020). Minimal dispersion approximately balancing weights: asymptotic properties and practical considerations, *Biometrika* 107(1): 93–105.
- Wu Y. and Liu Y. (2007). Robust truncated hinge loss support vector machines, *Journal of the American Statistical Association* 102(479): 974–983.
- Yang S. and Ding P. (2018). Asymptotic inference of causal effects with observational studies trimmed by the estimated propensity scores, *Biometrika* 105(2): 487–493.
- Zhang B, Tsiatis AA, Davidian M, Zhang M. and Laber E. (2012). Estimating optimal treatment regimes from a classification perspective, *Stat* 1(1): 103–114. [PubMed: 23645940]
- Zhang B, Tsiatis AA, Laber EB and Davidian M. (2012). A robust method for estimating optimal treatment regimes, *Biometrics* 68(4): 1010–1018. [PubMed: 22550953]
- Zhang Y, Laber EB, Tsiatis A. and Davidian M. (2015). Using decision lists to construct interpretable and parsimonious treatment regimes, *Biometrics* 71(4): 895–904. [PubMed: 26193819]
- Zhao Y-Q, Zeng D, Tangen CM and LeBlanc ML (2019). Robustifying trial-derived optimal treatment rules for a target population, *Electronic journal of statistics* 13(1): 1717–1743. [PubMed: 31440323]

- Zhao Y, Zeng D, Rush AJ and Kosorok MR (2012). Estimating individualized treatment rules using outcome weighted learning, *Journal of the American Statistical Association* 107(499): 1106–1118. [PubMed: 23630406]
- Zhou X. and Kosorok MR (2017). Causal nearest neighbor rules for optimal treatment regimes, arXiv preprint arXiv:1711.08451.
- Zhou X, Mayer-Hamblett N, Khan U. and Kosorok MR (2017). Residual weighted learning for estimating individualized treatment rules, *Journal of the American Statistical Association* 112(517): 169–187. [PubMed: 28943682]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Optimal individualized treatment rules

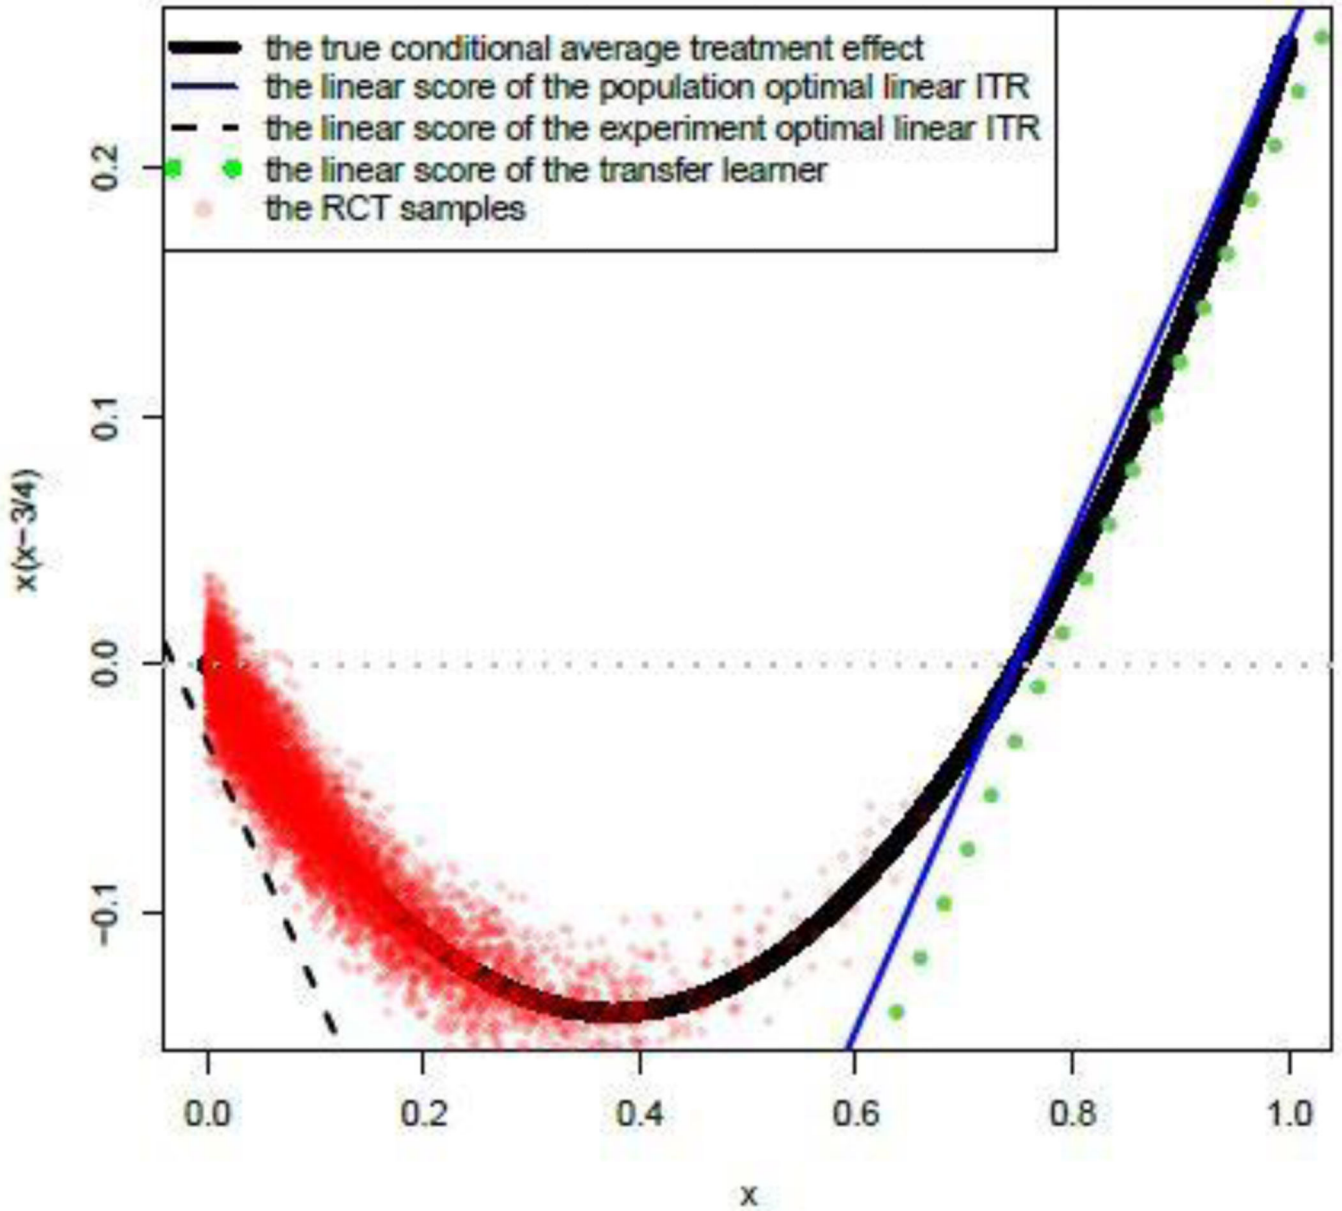


Fig. 1. Linear scores from the population, experiment, and transfer-learned optimal linear ITRs. The population X follows $\text{Unif}^{[0,1]}$ (the x -axis), and the conditional average treatment effect is $\tau(x) = x(x - 3/4)$ (the solid black line). The experimental sample of size 6227 (red dots) is selected based on a Bernoulli trial with selection probability $\exp(-11x)/\{1 + \exp(-11x)\}$, and the RWD sample is a random sample of 8000 from the population. The plot shows that the linear score of the optimal linear ITR based only on the experiment (the dash black line) is far from the population counterpart (the solid blue line); while the one from the transfer

learner (the dot green line) becomes closer by utilizing the representativeness of the RWD sample (Section 3).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

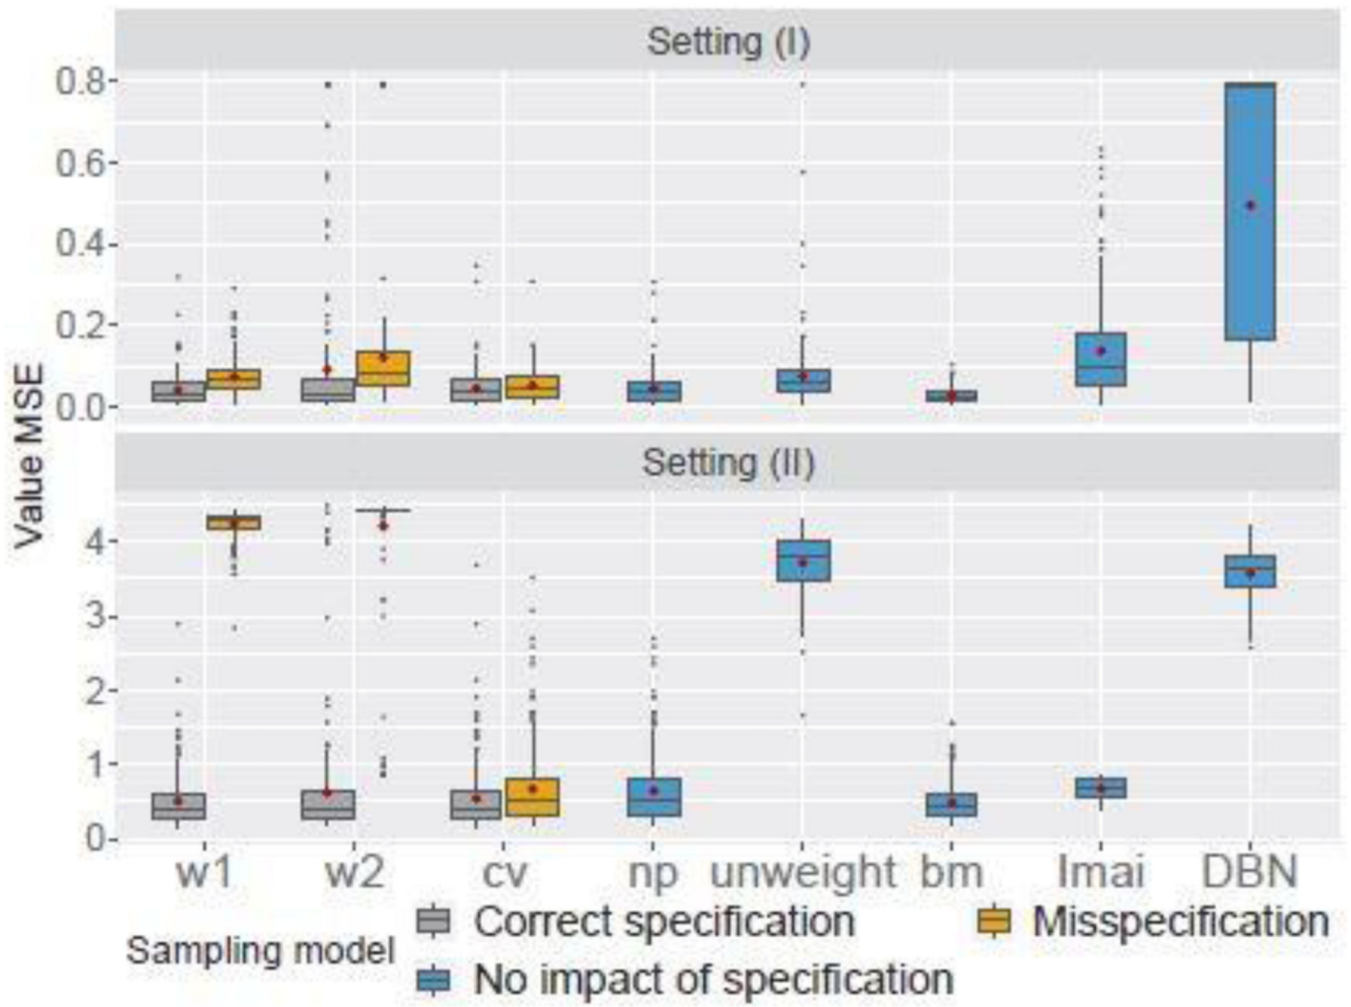


Fig. 2. Value MSE boxplots of different methods for the generative model (I) and (II). The red point in each boxplot is the mean value MSE averaged over 200 Monte Carlo replicates.

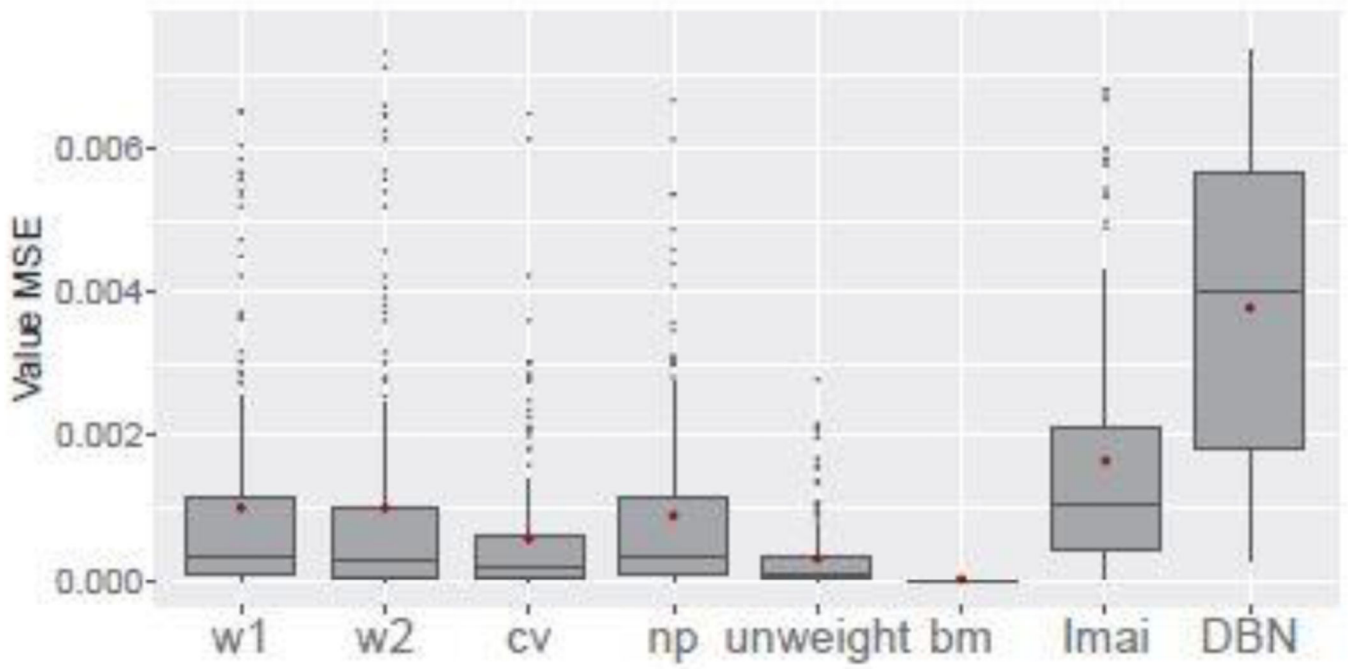


Fig. 3. Value MSE boxplots of different methods for the generative model (III). The red point in each boxplot is the mean value MSE averaged over 200 Monte Carlo replicates.

Table 1

Sample averages of pre-intervention variables in the NSW and CPS samples and results from the weighted transfer and unweighted learning methods: The first two rows are sample averages of each pre-intervention covariate in the NSW and CPS samples, respectively. The last two rows are estimated linear coefficients of each corresponding covariate in the linear decision rule from the weighted transfer and unweighted learning methods after normalizing the covariates, respectively.

	intercept	Age	Educ	Black	Hisp	Married	Nodegr	log.Re74	log.Re75
Experiment NSW	1	25.37	10.20	0.83	0.09	0.17	0.78	2.27	2.70
Non-experiment CPS	1	33.23	12.03	0.07	0.07	0.71	0.30	8.23	8.29
	η_0	η_1	η_2	η_3	η_4	η_5	η_6	η_7	η_8
Unweighted learning $\hat{\eta}_u$	1	0.11	0.16	-0.01	0.04	0.67	-0.67	-3.20	2.53
Weighted transfer learning $\hat{\eta}_w$	1	-0.62	-0.21	0.13	0.10	0.52	-0.85	-5.14	4.02