

Kernel-based covariate functional balancing for observational studies

By RAYMOND K. W. WONG

*Department of Statistics, Texas A&M University, 401E Blocker Building, 155 Ireland Street,
College Station, Texas 77843, U.S.A.*

raywong@stat.tamu.edu

AND KWUN CHUEN GARY CHAN

*Department of Biostatistics, University of Washington, 1959 NE Pacific St., Seattle,
Washington 98195, U.S.A.*

kcgchan@u.washington.edu

SUMMARY

Covariate balance is often advocated for objective causal inference since it mimics randomization in observational data. Unlike methods that balance specific moments of covariates, our proposal attains uniform approximate balance for covariate functions in a reproducing-kernel Hilbert space. The corresponding infinite-dimensional optimization problem is shown to have a finite-dimensional representation in terms of an eigenvalue optimization problem. Large-sample results are studied, and numerical examples show that the proposed method achieves better balance with smaller sampling variability than existing methods.

Some key words: Average treatment effect; Eigenvalue optimization; Reproducing-kernel Hilbert space; Sobolev space.

1. INTRODUCTION

The estimation of average treatment effects is important in the evaluation of an intervention or a treatment, but is complicated by confounding in observational studies where the treatment is not randomly assigned. When treatment assignment is unconfounded conditional on observable covariates, two popular modelling strategies are based on propensity score modelling ([Rosenbaum & Rubin, 1983](#)) and outcome regression modelling. Parametric approaches can suffer seriously from model misspecification, and there have been substantial recent efforts to construct more robust estimators within these modelling frameworks; see, for example, [Robins et al. \(1994\)](#), [Qin & Zhang \(2007\)](#), [Tan \(2010\)](#), [Graham et al. \(2012\)](#), and [Han & Wang \(2013\)](#).

Since randomization is a gold standard for identifying average treatment effects, [Rubin \(2007\)](#) advocated mimicking randomization, which balances the covariate distributions among the treated, the controls, and the combined sample, in the analysis of observational data. Based on these considerations, weighting-based covariate balancing methods have been proposed by [Qin & Zhang \(2007\)](#), [Hainmueller \(2012\)](#), [Imai & Ratkovic \(2014\)](#), [Zubizarreta \(2015\)](#) and [Chan et al. \(2016\)](#). A common feature of these methods is that a vector of user-specified functions of covariates is balanced. While balancing low-order moments of the covariates often yields good results, there is no guarantee that there will be sufficient balance over a large class of covariate functions. Matching is another general approach to attaining covariate balance. Exact matching

is not feasible for multiple continuous covariates, and a user-specified coarsening of the covariate space is needed (Iacus et al., 2011). In this paper, we shall focus on weighting-based methods.

Instead of balancing prespecified moments of covariates, we propose a method to control the covariate functional balance over a reproducing-kernel Hilbert space (Aronszajn, 1950), which can be chosen large enough to contain any functions with mild smoothness constraints, including nonlinearities and interactions. At a conceptual level, the comparison between covariate balancing with an increasing number of basis functions and kernel-based covariate functional balancing is analogous to the comparison of regression and smoothing splines in conditional mean estimation. Unlike regression splines, smoothing splines do not require preselection of the number of knots and their locations. Although achieving our goal involves a challenge due to an infinite-dimensional optimization problem, we show that it has a finite-dimensional representation and can be solved by eigenvalue optimization. Large-sample properties are derived under minimal smoothness conditions on the outcome regression model. Consistent estimation of average treatment effects is then possible without first guessing or estimating the outcome regression function, and efficient estimation can be attained when the outcome regression function is estimated. Unlike weighting methods that require stringent smoothness conditions for the propensity score function, our method does not require smoothness of the propensity score.

2. KERNEL-BASED COVARIATE FUNCTIONAL BALANCING

2.1. Preliminaries

Let $Y(1)$ and $Y(0)$ be the potential outcomes when an individual is assigned to the treatment group and control group respectively. We are interested in estimating the population average treatment effect $\tau = E\{Y(1) - Y(0)\}$. In practice, $Y(1)$ and $Y(0)$ are not both observed. With T denoting the binary treatment indicator, we can represent the observed outcome as $Y = TY(1) + (1 - T)Y(0)$. Moreover, we observe a vector of covariates $X \in \mathcal{X}$ for every individual, so the observed data are $\{(T_i, Y_i, X_i), i = 1, \dots, N\}$ where N is the sample size. We assume that $\{T_i, Y_i(1), Y_i(0), X_i\} (i = 1, \dots, N)$ are independent and identically distributed, and that T is independent of $\{Y(1), Y(0)\}$ conditional on X .

Note that τ consists of two expectations, $E\{Y(1)\}$ and $E\{Y(0)\}$. In this work, we consider weighted estimation of these expectations. Without loss of generality, we focus on $E\{Y(1)\}$. In the following, we consider a weighting estimator of $E\{Y(1)\}$ that can be represented as $N^{-1} \sum_{i=1}^N T_i w_i Y_i$. Hence, for estimation of $E\{Y(1)\}$, we only need to specify weights $w_i (i : T_i = 1)$ for individuals in the treatment group.

Let $\pi(x) = \text{pr}(T = 1 \mid X = x)$ be the propensity score. Assuming knowledge of $\pi(X_i)$ ($i : T_i = 1$), w_i can be chosen as $\{\pi(X_i)\}^{-1}$ to obtain a consistent estimator of $E\{Y(1)\}$. In practice, propensity scores are usually unknown. In such scenarios, one can estimate the propensity score function to form a plug-in estimator for $E\{Y(1)\}$. However, estimation errors and model misspecification of the propensity score function can lead to significant error in the estimation of $E\{Y(1)\}$ due to the use of inverse probability weighting. Poor finite-sample performance of such estimators has been reported in the literature (Kang & Schafer, 2007).

Due to this unsatisfactory performance, some attention has been given to choosing $w_i (i : T_i = 1)$ via covariate balancing, which mimics randomization directly. To understand this, note that

$$E \left\{ \frac{Tu(X)}{\pi(X)} \right\} = E\{u(X)\} \quad (1)$$

for any measurable function $u : \mathcal{X} \rightarrow \mathbb{R}$ such that $E\{u(X)\}$ exists and is finite. Instead of modelling the propensity function, it is therefore natural to choose weights that ensure the validity of the empirical finite-dimensional approximation of (1),

$$\frac{1}{N} \sum_{i=1}^N T_i w_i U(X_i) = \frac{1}{N} \sum_{i=1}^N U(X_i), \quad (2)$$

where $U(X) = \{u_1(X), \dots, u_L(X)\}^T$ is a L -variate function of X . Here $\text{span}\{u_1, \dots, u_L\}$ can be viewed as a finite-dimensional approximation space of functions in which the balancing is enforced. Practical considerations may suggest a choice of $\{u_1, \dots, u_L\}$. In this case, we call it parametric covariate balancing. Without assumptions on the outcome regression model, the balancing of fixed and finitely many component functions u_j in (1) may not lead to consistent estimation (Hellerstein & Imbens, 1999). To allow consistent estimation in a larger family of outcome regression functions, another possibility is to allow L to increase with N (Chan et al., 2016). This has a nonparametric flavour similar to regression splines, for which the number of knots grows with sample size. However, the choices of L and $\{u_1, \dots, u_L\}$ are not obvious. In this work, we aim to balance covariate functionals nonparametrically via reproducing-kernel Hilbert space modelling of the approximation space.

Let $m(X) = E\{Y(1) \mid X\}$ and $Y_i(1) = m(X_i) + \varepsilon_i$ for $i = 1, \dots, N$. Further assume that the ε_i are independent with $E(\varepsilon_i \mid X_i) = 0$ and $E(\varepsilon_i^2 \mid X_i) = \sigma_i^2 < \infty$. All weighting estimators of $E\{Y(1)\}$ admit the decomposition

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N T_i w_i Y_i &= \left\{ \frac{1}{N} \sum_{i=1}^N T_i w_i Y_i - \frac{1}{N} \sum_{i=1}^N m(X_i) \right\} + \left[\frac{1}{N} \sum_{i=1}^N m(X_i) - E\{Y(1)\} \right] + E\{Y(1)\} \\ &= \frac{1}{N} \sum_{i=1}^N (T_i w_i - 1) m(X_i) + \frac{1}{N} \sum_{i=1}^N T_i w_i \varepsilon_i \\ &\quad + \left[\frac{1}{N} \sum_{i=1}^N m(X_i) - E\{Y(1)\} \right] + E\{Y(1)\}, \end{aligned} \quad (3)$$

which allows a transparent understanding of the terms that have to be controlled. The first term on the right-hand side of (3) poses a challenge since the unknown outcome regression function m is intrinsically related to the outcome data, and could be complex and high-dimensional in general. To connect with covariate balancing, if $m \in \text{span}\{u_1, \dots, u_L\}$ in (2), we can control the first term. For the second term, the ε_i ($i = 1, \dots, N$) are independent of the choice of w_i ($i : T_i = 1$) if the outcome data are not used to obtain the weights. Some control over the magnitude of w_i will lead to convergence of the second term. Corresponding details will be given in § 2.4. The convergence of the third term is ensured by the law of large numbers.

2.2. Construction of the method

We consider the following empirical validity measure for any suitable function u :

$$S_N(w, u) = \left\{ \frac{1}{N} \sum_{i=1}^N (T_i w_i - 1) u(X_i) \right\}^2,$$

where $w = (w_1, \dots, w_N)^T$. In parametric covariate balancing, the weights w_i ($i : T_i = 1$) can be constructed to satisfy

$$\sup_{u \in \mathcal{U}_L} S_N(w, u) = 0,$$

where $\mathcal{U}_L = \text{span}\{u_1, \dots, u_L\}$ with u_1, \dots, u_L being suitable basis functions. In this case, the weights attain exact covariate balance as in (2) when the dimension of \mathcal{U}_L is small.

Here the overall validity of (1) is instead controlled directly on an approximation space \mathcal{H} , a reproducing-kernel Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and norm $\| \cdot \|_{\mathcal{H}}$. Ideally, one would want to pick a large enough, possibly infinite-dimensional, space \mathcal{H} to guarantee the control of $S_N(w, u)$ on a rich class of functions. Unlike sieve spaces, \mathcal{H} is specified without reference to sample size. The matching of nonlinear functions is also automatic if \mathcal{H} is large enough to contain such functions, without the need to explicitly introduce particular nonlinear basis functions in sieve spaces. For any Hilbert space \mathcal{H}_1 of functions of x_1 and any Hilbert space \mathcal{H}_2 of functions of x_2 , the tensor product space $\mathcal{H}_1 \otimes \mathcal{H}_2$ is defined as the completion of the class $\{\sum_{k=1}^{\ell} f_1(x_1)f_2(x_2) : f_1 \in \mathcal{H}_1, f_2 \in \mathcal{H}_2, \ell = 1, 2, \dots\}$ under the induced norm by \mathcal{H}_1 and \mathcal{H}_2 . A popular choice of \mathcal{H} is the tensor product reproducing-kernel Hilbert space $\mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \dots \otimes \mathcal{H}_d$ with \mathcal{H}_j being the reproducing-kernel Hilbert space of functions of the j th component of X . Suppose the support of the covariate distribution is $[0, 1]^d$ and $f^{(\ell)}$ is the ℓ th derivative of a function f . Following Wahba (1990), one can choose \mathcal{H}_j as the ℓ th-order Sobolev space $\mathcal{W}^{\ell, 2}([0, 1]) = \{f : f, f^{(1)}, \dots, f^{(\ell-1)} \text{ are absolutely continuous, } f^{(\ell)} \in L^2[0, 1]\}$ with norm

$$\|f\| = \left[\sum_{k=0}^{\ell-1} \left\{ \int_0^1 f^{(k)}(t) dt \right\}^2 + \int_0^1 \{f^{(\ell)}(t)\}^2 dt \right]^{1/2}.$$

The second-order Sobolev space is one of the most common choices in practice and will be adopted in all of our numerical illustrations. Another common choice is the space generated by the Gaussian kernel, which will also be used in our numerical studies. If it is desirable to prioritize covariates based on prior beliefs, we can raise the components to different powers to reflect their relative importance. For Gaussian kernels, this is equivalent to using different bandwidth parameters for each covariate. In cases where there are binary or categorical covariates, one can choose the corresponding \mathcal{H}_j as a reproducing-kernel Hilbert space with kernel $R(s, t) = I(s = t)$, for any levels s and t of such a covariate, as suggested by Gu (2013); here I is an indicator function.

Ideally, we want to control $\sup_{u \in \mathcal{H}} S_N$. However, there are two issues. First, that $S_N(w, cu) = c^2 S_N(w, u)$ for any $c \geq 0$ suggests a scale issue of S_N with respect to u . Therefore, in order to use $S_N(w, u)$ to determine the weights w , the magnitude of u should be standardized. To deal with this, we note that

$$S_N(w, u) = \left\{ \frac{1}{N} \sum_{i=1}^N (T_i w_i - 1) u(X_i) \right\}^2 \leq \|u\|_N^2 \left\{ \frac{1}{N} \sum_{i=1}^N (T_i w_i - 1)^2 \right\} \quad (4)$$

by the Cauchy–Schwarz inequality, where $\|u\|_N^2 = N^{-1} \sum_{i=1}^N u(X_i)^2$. In view of (4), we restrict our attention to $\tilde{\mathcal{H}}_N = \{u \in \mathcal{H} : \|u\|_N = 1\}$. Second, similar to many statistical and machine learning frameworks, the optimization of an unpenalized sample objective function will result

in overfitting. In our case, the weights become highly unstable. To alleviate this, we control $\|\cdot\|_{\mathcal{H}}$ to emphasize the balance on smoother functions. Additionally, we penalize on $V_N(w) = N^{-1} \sum_{i=1}^N T_i w_i^2$ to control the variabilities of both w and the second term in the right-hand side of (3). Overall, we consider the constrained minimization

$$\min_{w \geq 1} \left[\sup_{u \in \tilde{\mathcal{H}}_N} \{S_N(w, u) - \lambda_1 \|u\|_{\mathcal{H}}^2\} + \lambda_2 V_N(w) \right], \quad (5)$$

where $\lambda_1 > 0$ and $\lambda_2 > 0$ are tuning parameters and the above minimization is only taken over w_i ($i : T_i = 1$). The weights w_i are restricted to be greater than or equal to 1, as their counterparts, the inverse propensities, satisfy $\{\pi(X_i)\}^{-1} \geq 1$. We denote the solution of (5) by \hat{w} . Further discussion on these tuning parameters will be given in § 2.4 and § 2.5. In particular, we show that the convergence to zero of the first term of (3) can be ensured even when $\lambda_2 = 0$. This indicates that this extra tuning parameter is mostly needed for our justification of the convergence of the second term in (3).

A small number of recent papers have also considered kernel-based methods for covariate balancing. An unpublished paper by Zhao (2017) considered a dual formulation of the method of Imai & Ratkovic (2014) for the estimation of $\pi(x)$ under a logistic regression model, and generalized the linear predictor into a nonlinear one using the kernel trick. Since this method aims at estimating $\pi(x)$, it requires smoothness conditions on π and penalizes roughness of the resulting estimate. Our method does not require smoothness of π and penalizes the roughness of the balancing functions. An unpublished paper by Kallus (2017a) considered weights that minimize the dual norm of a balancing error. Given a reproducing-kernel Hilbert space, this method does not have the ability to adapt to a relevant subset of functions. An external parameter is required to index the function space, such as the dispersion parameter of a Gaussian kernel, and needs to be specified in an ad hoc manner. Due to the lack of an explicit tuning parameter, this method will not work well for the Sobolev space, which does not have extra indexing parameters. Our method works for a given reproducing-kernel Hilbert space by using data-adaptive tuning to promote balancing of smoother functions within the given space. An unpublished paper of Hazlett (2016) proposed an extension of the moment-based balancing method of Hainmueller (2012) to balance the columns of the Gram matrix. Since the Gram matrix is $N \times N$, exact balancing of N moment conditions under additional constraints on the weights is often computationally infeasible. Balancing a low-rank approximation of the Gram matrix may be an ad hoc solution but its theoretical properties have not been studied.

2.3. Finite-dimensional representation

Many common choices of reproducing-kernel Hilbert space, including Sobolev Hilbert space, are infinite-dimensional, so the inner optimization in (5) is essentially infinite-dimensional and appears impractical. Fortunately, we shall show that the solution of (5) enjoys a finite-dimensional representation. First, the inner optimization of (5) can be expressed as

$$\sup_{u \in \mathcal{H}} \left\{ \frac{S_N(w, u)}{\|u\|_N^2} - \lambda_1 \frac{\|u\|_{\mathcal{H}}^2}{\|u\|_N^2} \right\}.$$

Let K be the reproducing kernel of \mathcal{H} . By the representer theorem (Wahba, 1990), the solution lies in a finite-dimensional subspace $\text{span}\{K(X_j, \cdot) : j = 1, \dots, N\}$. Now this optimization

is equivalent to

$$\sup_{\alpha=(\alpha_1,\dots,\alpha_N)^T \in \mathbb{R}^N} \left[\frac{S_N \left\{ w, \sum_{j=1}^N \alpha_j K(X_j, \cdot) \right\}}{\alpha^T M^2 \alpha / N} - \lambda_1 \frac{\alpha^T M \alpha}{\alpha^T M^2 \alpha / N} \right], \tag{6}$$

where M is a $N \times N$ matrix with (i, j) th element $K(X_i, X_j)$. This positive semidefinite matrix is commonly known as the Gram matrix. Let the eigendecomposition of M be

$$M = (P_1 \ P_2) \begin{pmatrix} Q_1 & 0 \\ 0 & Q_2 \end{pmatrix} \begin{pmatrix} P_1^T \\ P_2^T \end{pmatrix},$$

where Q_1 and Q_2 are diagonal matrices. In particular, $Q_2 = 0$. Let r be the rank of Q_1 . We remark that P_2 and Q_2 do not exist if $r = N$, but the following derivation still holds. Moreover,

$$S_N \left\{ w, \sum_{j=1}^N \alpha_j K(X_j, \cdot) \right\} = \frac{1}{N^2} \alpha^T M A(w) M \alpha,$$

where $A(w) = a(w)a(w)^T$ with $a(w) = (T_1 w_1 - 1, \dots, T_N w_N - 1)^T$. Let $\beta = Q_1 P_1^T \alpha / N^{1/2}$. The constrained optimization (6) is then equivalent to

$$\sup_{\beta \in \mathbb{R}^r: \|\beta\| \leq 1} \beta^T \left\{ \frac{1}{N} P_1^T A(w) P_1 - N \lambda_1 Q_1^{-1} \right\} \beta.$$

Therefore, the target optimization becomes

$$\min_{w \geq 1} \left[\sigma_{\max} \left\{ \frac{1}{N} P_1^T A(w) P_1 - N \lambda_1 Q_1^{-1} \right\} + \lambda_2 V_N(w) \right], \tag{7}$$

where $\sigma_{\max}(M)$ represents the maximum eigenvalue of a matrix M . Again, the above minimization is only taken over w_i ($i : T_i = 1$). Since $P_1^T a(w)$ is an affine transformation of w and V_N is a convex function, the objective function of this minimization is convex with respect to w , due to Proposition 1, whose proof is given in the [Supplementary Material](#). Due to convexity and Slater’s condition of strict feasibility, a necessary and sufficient condition for a global minimizer of (7) is the corresponding Karush–Kuhn–Tucker condition using subdifferentials.

PROPOSITION 1. *Let $B \in \mathbb{R}^{r \times r}$ be a symmetric matrix. The function $\sigma_{\max}(v v^T + B)$ is convex with respect to $v \in \mathbb{R}^r$.*

For the computation, note that the maximum eigenvalue is evaluated at a rank-one modification of a diagonal matrix, which can be computed efficiently by solving the secular equation (O’Leary & Stewart, 1990) in a linear algebra package such as LAPACK. The objective function is second-order differentiable with respect to the w_i when the maximum eigenvalue of $P_1^T A(w) P_1 / N - N \lambda_1 Q_1^{-1}$ has multiplicity 1. Moreover, the corresponding gradient has a closed-form expression. In this case, a common and fast nonlinear optimization method such as the limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm with bound constraints can be applied. Nondifferentiability occurs when the largest two eigenvalues of $P_1^T A(w) P_1 / N - N \lambda_1 Q_1^{-1}$ coincide. To ensure validity, one could employ the following two-part computational strategy.

First, one applies the Broyden–Fletcher–Goldfarb–Shanno algorithm and checks numerically whether the maximum eigenvalue evaluated at the resulting solution is repeated. If not, the objective function is differentiable at this solution and the Karush–Kuhn–Tucker condition is satisfied. Thus, the minimizer is obtained. Otherwise, the nonlinear eigenvalue optimization method of Overton (1992, § 5), which is applicable to the scenario of repeated eigenvalues, is initialized by the former estimate and then applied. In our practical experience, the second step is seldom needed and has negligible effect on the final solution. Therefore, for fast computation, we only apply the first part in our numerical illustrations.

2.4. Theoretical properties

For notational simplicity, we shall study the theoretical properties of the proposed estimator with \mathcal{H} being the tensor product of ℓ th-order Sobolev spaces, as studied extensively in smoothing splines (Wahba, 1990; Gu, 2013). Our results can be extended to other choices of \mathcal{H} if an entropy result and a uniform boundedness condition for the unit ball $\{u \in \mathcal{H} : \|u\|_{\mathcal{H}} \leq 1\}$ are supplied; see the [Supplementary Material](#). For instance, an entropy result for Gaussian reproducing-kernel Hilbert space can be obtained from Zhou (2002). As mentioned, we concentrate on $E\{Y(1)\}$. Similar conditions are required for $E\{Y(0)\}$ to obtain results on the average treatment effect $\tau = E\{Y(1) - Y(0)\}$.

Assumption 1. The propensity $\pi(\cdot)$ is uniformly bounded away from zero. That is, there exists a constant C such that $1/\pi(x) \leq C < \infty$ for all $x \in \mathcal{X}$.

Assumption 2. The ratio d/ℓ is less than 2.

Assumption 3. The regression function $m(\cdot)$ belongs to \mathcal{H} .

Assumption 4. The errors $\{\varepsilon_i\}$ are uncorrelated, with $E(\varepsilon_i) = 0$ and $\text{var}(\varepsilon_i) = \sigma_i^2 \leq \sigma^2$ for all $i = 1, \dots, N$. Further, the $\{\varepsilon_i\}$ are independent of $\{T_i\}$ and $\{X_i\}$.

The above assumptions are very mild. Assumption 1 is the usual overlap condition required for identification. There are no additional smoothness assumptions on $\pi(\cdot)$, which would typically be required in propensity score or covariate balancing methods (Hirano et al., 2003; Chan et al., 2016). Assumption 2 corresponds to the weakest smoothness assumption on $m(\cdot)$ in smoothing spline regression. We use the notation $A_N \asymp B_N$ to represent $A_n = O(B_N)$ and $B_N = O(A_N)$ for some sequences A_N and B_N .

THEOREM 1. *Suppose Assumptions 1–3 hold. If $\lambda_1 \asymp N^{-1}$ and $\lambda_2 = O(N^{-1})$, then $S_N(\hat{w}, m) = O_p(N^{-1})\|m\|_N^2$. If $\lambda_1 \asymp N^{-1}$ and $\lambda_2 \asymp N^{-1}$, then $V_N(\hat{w}) = O_p(1)$ and there exist constants $W > 0$ and $S^2 > 0$ such that $E\{V_N(\hat{w})\} \leq W$ and $E\{NS_N(\hat{w}, m)\} \leq S^2$.*

Theorem 1 supplies the rate of convergence of the first term in (3), and boundedness of the expectation of the second term in (3). Convergence of $S_N(\hat{w}, m)$ is guaranteed even if λ_2 is chosen to be 0. However, to ensure boundedness of $E\{V_N(\hat{w})\}$, additional regularization is needed and hence $\lambda_2 > 0$ is proposed. The following theorem establishes the $N^{1/2}$ -consistency of the weighting estimator. Moreover, we show that the asymptotic distribution has finite variance.

THEOREM 2. *Suppose Assumptions 1–4 hold and that $m \in \mathcal{H}$. If $\lambda_1 \asymp N^{-1}$ and $\lambda_2 \asymp N^{-1}$, then*

$$\frac{1}{N} \sum_{i=1}^N T_i \hat{w}_i Y_i - E\{Y(1)\} = O_p(N^{-1/2}).$$

Moreover, $N^{1/2}[\sum_{i=1}^N T_i \hat{w}_i Y_i / N - E\{Y(1)\}]$ has finite asymptotic variance.

Although Theorem 2 only gives the rate of convergence of the estimator, it is stronger than recent results for other kernel-based methods for the estimation of average treatment effects. Zhao (2017) and Hazlett (2016) do not provide the rates of convergence of their estimators. To our knowledge, the only paper that contains a rate of convergence for kernel-based methods is Kallus (2017b), who showed a root- N convergence rate under the assumption that $m(X)$ is linear in X and did not give the asymptotic distribution. In fact, when linear assumptions hold, parametric covariate balancing is sufficient for estimating the average treatment effects (Qin & Zhang, 2007). When $m(\cdot)$ is a general function, the difficulty in theoretical development lies in the first term of (3), which is shown to attain the same rate of convergence as the other two terms of (3), but whose asymptotic distribution is not available. For the sieve-based method (Chan et al., 2016), the growth rate of the sieve approximation space can be carefully chosen in a range such that terms analogous to the first term of (3) have a faster convergence rate than the dominating terms. In our case, similar to nonparametric regression, there is only a particular growth rate of λ_1 such that the bias and variance of the first term of (3) are balanced. In fact, it is possible that the term has an asymptotic bias of order $N^{-1/2}$. In § 2.6, a modified estimator is studied by debiasing the first term of (3), so that its rate of convergence is faster than $N^{-1/2}$ and is dominated by the other terms. In that case, the asymptotic distribution can be derived. Further discussion of the relationship between Theorem 2 and the literature is given in Remark 3.

2.5. Tuning parameter selection

In Theorems 1 and 2, λ_1 and λ_2 are required to decrease at the same order N^{-1} , so as to achieve the desired asymptotic results. To reduce the amount of tuning, we choose $\lambda_2 = \zeta \lambda_1$ where $\zeta > 0$ is fixed. As explained above, λ_2 is chosen to be positive mostly to ensure the boundedness of $E\{V_N(\hat{w})\}$. From our practical experience, the term $V_N(\hat{w})$ is usually stable and does not take large values even if λ_2 is small. Therefore, we are inclined to choose a small ζ . In all of our numerical illustrations, ζ is fixed at 0.01. Now we focus on the choice of λ_1 . Tuning λ_1 is similar to choosing the dimension of the sieve space in Chan et al. (2016), which is a difficult and mostly unsolved problem. In this paper, we do not attempt to solve it rigorously, but just to provide a reasonable solution.

By Lagrange multipliers, the optimization $\sup_{u \in \tilde{\mathcal{H}}_N} \{S_N(w, u) - \lambda_1 \|u\|_{\mathcal{H}}^2\}$ is equivalent to $\sup_{\{u \in \tilde{\mathcal{H}}_N : \|u\|_{\mathcal{H}} \leq \gamma\}} S_N(w, u)$ for some γ , where there exists a correspondence between γ and λ_1 . Since a larger regularization parameter corresponds to a stricter constraint, γ decreases with λ_1 . We use

$$B_N(w) = \sup_{\{u \in \tilde{\mathcal{H}}_N : \|u\|_{\mathcal{H}} \leq \gamma\}} S_N(w, u) \quad (8)$$

as a measure of the balancing error over $\{u \in \tilde{\mathcal{H}}_N : \|u\|_{\mathcal{H}} \leq \gamma\}$ with respect to the weights w . Due to the large subset of functions to balance, $B_N(\hat{w})$ is large when γ is large or, equivalently, when λ_1 is small. When γ decreases, or equivalently when λ_1 increases, $B_N(\hat{w})$ typically decreases to approximately zero, as the resulting weight \hat{w} approximately balances the whole subset $\{u \in \tilde{\mathcal{H}}_N : \|u\|_{\mathcal{H}} \leq \gamma\}$. An example is given in § 3.2. When this happens, a further decrease of γ would not lead to any significant decrease in $B_N(\hat{w})$. The key idea is to choose the smallest λ_1 that achieves such approximate balancing, to ensure the largest subset of functions being well-balanced. In practice, we compute our estimator with respect to a grid of λ_1 such that $\lambda_1^{(1)} < \dots < \lambda_1^{(J)}$. Write $\hat{w}^{(j)}$ as the estimator with respect to $\lambda_1^{(j)}$. We select $\lambda_1^{(j^*)}$ as our choice of λ_1 if j^* is the

smallest j such that

$$\frac{B_N(\hat{w}^{(j+1)}) - B_N(\hat{w}^{(j)})}{\lambda_1^{(j+1)} - \lambda_1^{(j)}} \geq c,$$

where c is chosen as a negative constant of small magnitude. In the numerical illustrations, we set $c = -10^{-6}$.

2.6. An efficient modified estimator

Since the outcome regression function $m(\cdot)$ is assumed to be in a reproducing-kernel Hilbert space \mathcal{H} , a kernel-based estimator $\hat{m}(\cdot)$, such as smoothing splines (Gu, 2013), can be employed, and $N^{-1} \sum_{i=1}^N \hat{m}(X_i)$ is a natural estimator of $E\{Y(1)\} = E[E\{Y(1) | X\}] = E\{m(X)\}$. However, since randomization is administered before collecting any outcome data, Rubin (2007) advocated the estimation of treatment effects without using outcome data to avoid data snooping. On the other hand, Chernozhukov et al. (2017) and Athey et al. (2017) advocate the use of an estimated outcome regression function to improve the theoretical results in high-dimensional settings. Inspired by these results, we modify the weighting estimator by subtracting $N^{-1} \sum_{i=1}^N (T_i w_i - 1) \hat{m}(X_i)$ from both sides of (3), so that the first term in the decomposition becomes $N^{-1} \sum_{i=1}^N (T_i w_i - 1) \{m(X_i) - \hat{m}(X_i)\}$, while the remaining two terms are unchanged. It can then be shown that the first term has a rate of convergence faster than $N^{-1/2}$ under mild assumptions, and the asymptotic distribution of the resulting estimator will be derived.

The estimator takes the form

$$\frac{1}{N} \sum_{i=1}^N \{T_i w_i Y_i - (T_i w_i - 1) \hat{m}(X_i)\} = \frac{1}{N} \sum_{i=1}^N [T_i w_i \{Y_i - \hat{m}(X_i)\} + \hat{m}(X_i)],$$

which has the same form as the residual balancing estimator proposed in Athey et al. (2017). They considered a different setting of the high-dimensional linear regression model with sparsity assumptions, and showed that their estimator attains the semiparametric efficiency bound.

Our analysis requires the additional technical assumption that \hat{w} is $o_p(N^{1/2})$. To achieve this, we adopt an assumption like the one in Athey et al. (2017), that $\hat{w} \leq BN^{1/3}$ for a prespecified large positive constant B . This can be enforced in the optimization (7) easily together with the constraint $\hat{w} \geq 1$. We call this estimator $\tilde{w} = (\tilde{w}_1, \dots, \tilde{w}_N)^T$.

THEOREM 3. *Suppose Assumptions 1, 2 and 4 hold with $\sigma_i^2 \equiv \sigma^2$. Also, let $\max_i E|\varepsilon_i|^3 < \infty$. Let $h = m - \hat{m} \in \mathcal{H}$ such that $\|h\|_N = o_p(1)$ and $\|h\|_{\mathcal{H}} = O_p(1)$. Further, assume $\lambda_1 = o(N^{-1})$, $\lambda_2 \|h\|_N^2 = o_p(N^{-1})$, and $\lambda_1^{-1} = o(\lambda_2^{(2\ell-d)/d} N^{2\ell/d})$. Write*

$$J_N = N^{1/2} \left(\left[\frac{1}{N} \sum_{i=1}^N T_i \tilde{w}_i \{Y_i - \hat{m}(X_i)\} + \frac{1}{N} \sum_{i=1}^N \hat{m}(X_i) \right] - E\{Y(1)\} \right),$$

$$J_N^* = [\text{var}\{m(X_1)\}]^{1/2} F + \sigma N^{-1/2} \sum_{j=1}^N T_j \tilde{w}_j G_j,$$

where F, G_1, \dots, G_N are independent and identically distributed standard normal random variables independent of $X_1, \dots, X_N, T_1, \dots, T_N$ and $\varepsilon_1, \dots, \varepsilon_N$. Let ψ_N and ψ_N^* be the corresponding characteristic functions of J_N and J_N^* . Then as $N \rightarrow \infty$,

$$|\psi_N(t) - \psi_N^*(t)| \rightarrow 0, \quad t \in \mathbb{R},$$

where ψ_N^* is twice differentiable, and

$$\limsup_N \text{var}(J_N) \leq \text{var}\{m(X_1)\} + \sigma^2 V, \quad (9)$$

where $V = E\{1/\pi(X_1)\}$.

COROLLARY 1. Under the assumptions of Theorem 3,

$$N^{1/2}\{\sigma^2 V_N(\tilde{w})\}^{-1/2} \left[\frac{1}{N} \sum_{i=1}^N T_i \tilde{w}_i \{Y_i - \hat{m}(X_i)\} + \frac{1}{N} \sum_{i=1}^N \{\hat{m}(X_i) - m(X_i)\} \right]$$

converges in distribution to a standard normal variable as $N \rightarrow \infty$.

Remark 1. In Theorem 3, the estimand is $E\{Y(1)\}$, whereas in Corollary 1, the estimand is a finite-sample conditional average, $N^{-1} \sum_{i=1}^N E\{Y_i(1) \mid X_i\} = N^{-1} \sum_{i=1}^N m_1(X_i)$. [Athey et al. \(2017\)](#) considered a finite-sample conditional average treatment effect and obtained a result similar to Corollary 1. Normalization by $V_N(\tilde{w})$ is possible in Corollary 1 following a conditional central limit theorem, since \tilde{w} depends only on (T_i, X_i) ($i = 1, \dots, N$) and can be treated as constants upon conditioning. To derive the limiting distribution of J_N in Theorem 3, one cannot use a similar normalization because the handling of the extra term $N^{-1} \sum_{i=1}^N \{\hat{m}(X_i) - m(X_i)\}$ requires averaging across the distribution of X . If $V_N(\tilde{w})$ converges to a constant in probability, one could use Slutsky's theorem to prove asymptotic normality of J_N . Theorem 3 requires a partially conditional central limit theorem which is proven in the [Supplementary Material](#), and the distribution of J_N can be approximated by a weighted sum of independent standard normal random variables. The asymptotic variance is bounded above by the right-hand side of (9), which is the semiparametric efficiency bound ([Robins et al., 1994](#); [Hahn, 1998](#)).

Remark 2. Compared to Theorem 2, Theorem 3 requires different conditions on the orders of λ_1 and λ_2 . These order specifications, together with a diminishing $\|h\|_N$, allow a direct asymptotic comparison between $V_N(\tilde{w})$ and V , which yields $V_N(\tilde{w}) \leq V\{1 + o_p(1)\}$. This is essential for (9). To make sense of the theorem, the conditions $\lambda_1 = o(N^{-1})$ and $\lambda_1^{-1} = o\{\lambda_2^{(2\ell-d)/d} N^{2\ell/d}\}$ should not lead to a null set of λ_1 . As an illustration, suppose \hat{m} achieves the optimal rate $\|h\|_N \asymp N^{-\ell/(2\ell+d)}$; then one can take $\lambda_2 = o\{N^{-d/(2\ell+d)}\}$, which suggests $\lambda_2^{(2\ell-d)/d} N^{2\ell/d} = o\{N^{(d^2+4\ell^2)/(d^2+2\ell d)}\}$. Due to Assumption 2, $(d^2 + 4\ell^2)(d^2 + 2\ell d)^{-1} > 1$. Therefore, there exist choices of λ_1 and λ_2 that fulfil the assumptions of Theorem 3. We found in simulations that the practical performance of the modified estimator is not sensitive to λ_1 and λ_2 , and we thus use the method described in § 2.5 to obtain these tuning parameters.

Remark 3. Most existing efficient methods require explicit or implicit estimation of both $\pi(\cdot)$ and $m(\cdot)$. [Chernozhukov et al. \(2017\)](#) gave a general result on the required convergence rates of $\pi(\cdot)$ and $m(\cdot)$ for efficient estimation. Even though weighting methods do not explicitly estimate

$m(\cdot)$, estimating equation-based methods would give implicit estimators of $m(\cdot)$ that attain good rates of convergence (Hirano et al., 2003; Chan et al., 2016). However, weights constructed based on complex optimization problems may not even converge to the true inverse propensities; see Athey et al. (2017) who, under a sparse linear model assumption, proposed an efficient estimator by controlling the balancing error of linear functions and the estimation error for $m(\cdot)$. Although our modified estimator is not a direct kernel-based extension of their method, we have arrived at a similar conclusion. Our method only requires $\|\hat{m} - m\|_N = o_p(1)$ and does not require smoothness of $\pi(\cdot)$ or linearity of $m(\cdot)$, and is therefore less vulnerable to the curse of dimensionality. The weighting estimator as described in Theorem 2 corresponds to $\hat{m} = 0$, and therefore $\|\hat{m} - m\|_N$ is not $o_p(1)$ when $m(\cdot)$ is not the zero function. Theorem 2 is interesting because convergence to neither $\pi(\cdot)$ nor $m(\cdot)$ is established.

3. NUMERICAL EXAMPLES

3.1. Simulation study

Simulation studies were conducted to evaluate the finite-sample performance of the proposed estimator. We considered settings where the propensity score and outcome regression models are nonlinear functions of the observed covariates, with possibly nonsmooth propensity score functions. For each observation, we generated a ten-dimensional multivariate standard Gaussian random vector $Z = (Z_1, \dots, Z_{10})^T$. The observed covariates are $X = (X_1, \dots, X_{10})^T$ where $X_1 = \exp(Z_1/2)$, $X_2 = Z_2/\{1 + \exp(Z_1)\}$, $X_3 = (Z_1 Z_3/25 + 0.6)^3$, $X_4 = (Z_2 + Z_4 + 20)^2$ and $X_j = Z_j$ ($j = 5, \dots, 10$). Three propensity score models are studied: model 1 is $\text{pr}(T = 1 | Z) = \exp(-Z_1 - 0.1Z_4)/\{1 + \exp(-Z_1 - 0.1Z_4)\}$, model 2 is $\text{pr}(T = 1 | Z) = \exp\{-Z_1 - 0.1Z_4 + \eta_2(\tilde{Z})\}/\{1 + \exp\{-Z_1 - 0.1Z_4 + \eta_2(\tilde{Z})\}\}$, and model 3 is $\text{pr}(T = 1 | Z) = \exp\{-Z_1 - 0.1Z_4 + \eta_3(\tilde{Z})\}/\{1 + \exp\{-Z_1 - 0.1Z_4 + \eta_3(\tilde{Z})\}\}$, where $\tilde{Z} = (Z_2 + Z_4 + Z_6 + Z_8 + Z_{10})/5$, η_2 is the scaling function of the Daubechies 4-tap wavelet (Daubechies, 1992), and η_3 is the Weierstrass function with parameters $a = 2$ and $b = 13$. The functions η_2 and η_3 are chosen such that the propensity functions in models 2 and 3 are nonsmooth. Two outcome regression models are studied: model A is $Y = 210 + (1.5T - 0.5)(27.4Z_1 + 13.7Z_2 + 13.7Z_3 + 13.7Z_4) + \epsilon$, and model B is $Y = Z_1 Z_2^3 Z_3^2 Z_4 + Z_4 |Z_1|^{0.5} + \epsilon$, where ϵ follows the standard normal distribution. For each scenario, we compared the proposed weighting and modified estimators using the second-order Sobolev kernel and the Gaussian kernel with bandwidth parameter chosen via the median heuristic (Gretton et al., 2005). We also compared the Horvitz–Thompson estimator, where the weights are the inverse of propensity scores estimated by maximum likelihood under a working logistic regression model with X being the predictors; the Hájek estimator, which is a normalized version of the Horvitz–Thompson estimator with weights summing to N ; the inverse probability weighting estimator using the covariate balancing propensity score of Imai & Ratkovic (2014) or the stable balancing weights of Zubizarreta (2015); and the nonparametric covariate balancing estimator of Chan et al. (2016) with exponential weights. The first moment of X was balanced explicitly for these methods. We compared the bias, root mean squared error and covariate balance of the methods, where covariate balance is evaluated at the true conditional mean function. In particular, we calculate $S_N(\hat{w}, m)$ to evaluate the covariate balance of the treatment and the combined groups, as well as its counterpart for the controls and the combined groups, and report the sum of these two measures. The reason for comparing the covariate balance at the true conditional mean function is that it is the optimal function to balance but is unknown in practice. For each scenario, 1000 independent datasets are generated, and the results for outcome models A and B with sample size $N = 200$ are given in Tables 1 and 2 respectively.

Table 1. *Biases, root mean square errors and overall covariate balancing measures of weighting estimators for outcome model A; the numbers reported are averages obtained from 1000 simulated datasets*

	PS1			PS2			PS3		
	Bias	RMSE	Bal	Bias	RMSE	Bal	Bias	RMSE	Bal
Proposed weighting (S)	-192	470	28	-114	422	20	-86	420	21
Proposed weighting (G)	-362	599	86	-232	505	58	-201	504	67
Proposed modified (S)	150	512	28	243	486	20	255	480	21
Proposed modified (G)	-91	378	86	45	368	58	-28	370	67
Horvitz–Thompson	>9999	>9999	>9999	>9999	>9999	>9999	7298	>9999	>9999
Hájek	837	1792	275	693	1429	165	662	1480	175
Imai–Ratkovic	-527	1720	331	-187	1523	274	224	1516	253
Zubizarreta	444	715	28	389	658	23	381	630	21
Chan et al.	262	594	21	226	549	17	244	533	16

The sample sizes were $N = 200$. The values of Bias and RMSE were multiplied by 100. RMSE, the root mean squared error; Bal, an overall covariate balancing measure; S, Sobolev kernel; G, Gaussian kernel.

Table 2. *Biases, root mean square errors and overall covariate balancing measures of weighting estimators for outcome model B; the numbers reported are averages obtained from 1000 simulated datasets*

	PS1			PS2			PS3		
	Bias	RMSE	Bal	Bias	RMSE	Bal	Bias	RMSE	Bal
Proposed weighting (S)	-10	85	1	-7	81	0	-8	85	0
Proposed weighting (G)	-7	92	1	-3	88	0	-5	94	0
Proposed modified (S)	3	82	1	-8	82	0	-9	85	0
Proposed modified (G)	-6	89	1	-2	85	0	-4	92	0
Horvitz–Thompson	131	3629	1151	-1	881	66	-35	2092	451
Hájek	4	392	14	-3	221	4	-2	367	12
Imai–Ratkovic	-7	110	1	-6	97	1	-6	108	1
Zubizarreta	-8	115	1	-9	98	1	-10	108	1
Chan et al.	-8	123	1	-9	99	1	-8	111	1

The sample sizes were $N = 200$. The values of Bias and RMSE were multiplied by 100. RMSE, the root mean squared error; Bal, an overall covariate balancing measure; S, Sobolev kernel; G, Gaussian kernel.

The results show that the empirical performances of the estimators are related to the degree of covariate balancing. Without any explicit covariate balancing, the Horvitz–Thompson estimator can be highly unstable. The Hájek estimator balances the constant function; the Imai–Ratkovic estimator balances X ; the estimators of Zubizarreta and Chan et al. balance both the constant function and X . For outcome model A, the balance of both the constant function and X is important and the omission of either constraint can lead to poor performance. For outcome model B, the balance of X often implies approximate balance of the constant, and therefore the estimators of Imai and Ratkovic, Zubizarreta and Chan et al. performed similarly. However, in both cases, the proposed method outperformed the other estimators because it can also control the balance of nonlinear and higher-order moments. We attempted to compute a Horvitz–Thompson estimator using a smoothing spline logistic regression model with the same kernel as the proposed method using the R package *gss*, but the program did not converge in a reasonable time. We also tried to exactly balance the second moments in addition to the first moments of ten baseline

Table 3. Dependence of daily wage (Ugandan schillings) of males on abduction as a child, $Y(1)$, and non-abduction, $Y(0)$

	$E\{Y(1)\}$	$E\{Y(0)\}$	τ
Proposed weighting (S)	1530 (1219, 1886)	1851 (1247, 2354)	-321 (-893, 431)
Proposed weighting (G)	1516 (1212, 1822)	1671 (1231, 2204)	-156 (-809, 382)
Proposed modified (S)	1532 (1239, 1945)	1867 (1256, 2489)	-355 (-993, 444)
Proposed modified (G)	1536 (1238, 1845)	1689 (1269, 2249)	-153 (-819, 380)
Horvitz–Thompson	1573 (1234, 2033)	2135 (1478, 3075)	-562 (-1667, 242)
Hájek	1573 (1234, 2027)	2131 (1471, 3064)	-558 (-1614, 241)
Imai–Ratkovic	1599 (1256, 2062)	1998 (1381, 2857)	-399 (-1312, 365)
Zubizarreta	1591 (1253, 2073)	2165 (1492, 3046)	-574 (-1613, 229)
Chan et al.	1580 (1246, 2060)	2144 (1485, 3034)	-564 (-1576, 238)

Numbers in parentheses are 95% confidence intervals. S, Sobolev kernel; G, Gaussian kernel.

covariates in the existing methods, but the algorithms did not converge in a substantial fraction of simulations. This shortcoming of the existing methods can be circumvented by using the proposed methods.

3.2. Data analysis

We compare the proposed methods with others using a study of the impact of child abduction by a militant group on the future income of abductees who escape later (Blattman & Annan, 2010). The data were collected from 741 males in Uganda during 2005–2006, of which 462 had been abducted by militant groups before 2005 but had escaped by the time of the study. Covariates include geographical region, age in 1996, father’s education, mother’s education, whether the parents had died during or before 1996, whether the father is a farmer, and household size in 1996. The investigators chose to collect covariate values in 1996 because it predates most abductions and is also easily recalled as the year of the first election since 1980. The authors discuss the plausibility of the unconfounded treatment assignment, since abduction is mostly due to random night raids on rural homes. The outcome of interest here is the daily wage of the study participants in Ugandan shillings in 2005. We compared the estimators as in the simulation studies in § 3.1. Table 3 shows the point estimates and 95% confidence intervals based on 1000 bootstrap samples. All methods comparing the abducted to the non-abducted group give a small but nonsignificant decrease in income. However, a small difference is observed between the proposed method and other methods, indicating that a mild nonlinear effect is possibly present, especially in the non-abducted group. To further illustrate this point, we compared the maximal balancing error $B_N(w)$ as a function of λ_1 , which, as defined by (8), measures the balancing error as a function of the size of nested subsets of \mathcal{H}_N , which is chosen as the Sobolev space. The subspace is smaller with an increasing λ_1 , containing smoother functions. We standardize the comparisons by dividing by the balancing error of constant weights that are used in unweighted comparisons. As seen in Fig. 1, the proposed estimator had approximately no balancing error after reaching a data-dependent threshold, so that any smoother functions can be approximately balanced. This is not the case for the other estimators, since there will be residual imbalance for nonlinear functions of a given smoothness. The Imai–Ratkovic estimator has less balancing error than the Horvitz–Thompson estimator with maximum likelihood weights, because the former explicitly balances more moments than the latter, including linear and some nonlinear covariate functionals. The Imai–Ratkovic estimator gives the closest result to the proposed estimators.

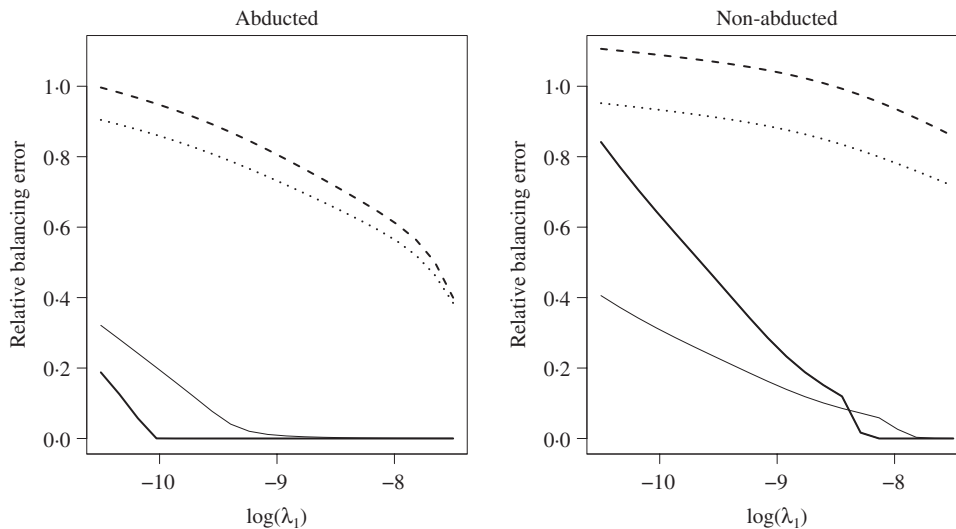


Fig. 1. Supremum balancing error for the child abduction data. Thick solid curves correspond to the proposed estimator using a Sobolev kernel, thin solid curves correspond to the proposed estimator using a Gaussian kernel, dashed curves correspond to the Horvitz–Thompson estimator, and dotted curves correspond to the Imai–Ratkovic estimator.

ACKNOWLEDGEMENT

The authors thank the editor, an associate editor and a reviewer for their helpful comments and suggestions. Wong was partially supported by the U.S. National Science Foundation. Most of this work was conducted while the first author was affiliated with Iowa State University. Chan was partially supported by the National Heart, Lung, and Blood Institute of the U.S. National Institutes of Health and by the U.S. National Science Foundation.

SUPPLEMENTARY MATERIAL

[Supplementary material](#) available at *Biometrika* online includes the proofs of Proposition 1 and Theorems 1–3.

REFERENCES

- ARONSAJN, N. (1950). Theory of reproducing kernels. *Trans. Am. Math. Soc.* **68**, 337–404.
- ATHEY, S., IMBENS, G. & WAGER, S. (2017). Approximate residual balancing: de-biased inference of average treatment effects in high dimensions. *arXiv*: 1604.07125v4.
- BLATTMAN, C. & ANNAN, J. (2010). The consequences of child soldiering. *Rev. Econ. Statist.* **92**, 882–98.
- CHAN, K. C. G., YAM, S. C. P. & ZHANG, Z. (2016). Globally efficient nonparametric inference of average treatment effects by empirical balancing calibration weighting. *J. R. Statist. Soc. B* **78**, 673–700.
- CHERNOZHUKOV, V., CHETVERIKOV, D., DEMIRER, M., DUFO, E., HANSEN, C., NEWEY, W. & ROBINS, J. (2017). Double/debiased machine learning for treatment and causal parameters. *arXiv*: 1608.0060v5.
- DAUBECHIES, I. (1992). *Ten Lectures on Wavelets*. Philadelphia: SIAM.
- GRAHAM, B. S., PINTO, C. C. D. X. & EGEL, D. (2012). Inverse probability tilting for moment condition models with missing data. *Rev. Econ. Studies* **79**, 1053–79.
- GRETTON, A., HERBRICH, R., SMOLA, A., BOUSQUET, O. & SCHÖLKOPF, B. (2005). Kernel methods for measuring independence. *J. Mach. Learn. Res.* **6**, 2075–129.
- GU, C. (2013). *Smoothing Spline ANOVA Models*. New York: Springer.

- HAHN, J. (1998). On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica* **66**, 315–32.
- HAINMUELLER, J. (2012). Entropy balancing for causal effects: a multivariate reweighting method to produce balanced samples in observational studies. *Polit. Anal.* **20**, 25–46.
- HAN, P. & WANG, L. (2013). Estimation with missing data: beyond double robustness. *Biometrika* **100**, 417–30.
- HAZLETT, C. (2016). Kernel balancing: a flexible non-parametric weighting procedure for estimating causal effects. *arXiv*: 1605.00155v1.
- HELLERSTEIN, J. K. & IMBENS, G. W. (1999). Imposing moment restrictions from auxiliary data by weighting. *Rev. Econ. Statist.* **81**, 1–14.
- HIRANO, K., IMBENS, G. & RIDDER, G. (2003). Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* **71**, 1161–89.
- IACUS, S. M., KING, G. & PORRO, G. (2011). Multivariate matching methods that are monotonic imbalance bounding. *J. Am. Statist. Assoc.* **106**, 345–61.
- IMAL, K. & RATKOVIC, M. (2014). Covariate balancing propensity score. *J. R. Statist. Soc. B* **76**, 243–63.
- KALLUS, N. (2017a). A framework for optimal matching for causal inference. *arXiv*: 1606.05188v2.
- KALLUS, N. (2017b). Generalized optimal matching methods for causal inference. *arXiv*: 1612.08321v3.
- KANG, J. D. & SCHAFER, J. L. (2007). Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data (with Discussion). *Statist. Sci.* **22**, 523–39.
- O'LEARY, D. & STEWART, G. (1990). Computing the eigenvalues and eigenvectors of symmetric arrowhead matrices. *J. Comp. Phys.* **90**, 497–505.
- OVERTON, M. L. (1992). Large-scale optimization of eigenvalues. *SIAM J. Optimiz.* **2**, 88–120.
- QIN, J. & ZHANG, B. (2007). Empirical-likelihood-based inference in missing response problems and its application in observational studies. *J. R. Statist. Soc. B* **69**, 101–22.
- ROBINS, J., ROTNITZKY, A. & ZHAO, L. (1994). Estimation of regression coefficients when some regressors are not always observed. *J. Am. Statist. Assoc.* **89**, 846–66.
- ROSENBAUM, P. R. & RUBIN, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**, 41–55.
- RUBIN, D. B. (2007). The design versus the analysis of observational studies for causal effects: parallels with the design of randomized trials. *Statist. Med.* **26**, 20–36.
- TAN, Z. (2010). Bounded, efficient and doubly robust estimation with inverse weighting. *Biometrika* **97**, 661–82.
- WAHBA, G. (1990). *Spline Models for Observational Data*. Philadelphia: SIAM.
- ZHAO, Q. (2017). Covariate balancing propensity score by tailored loss functions. *arXiv*: 1601.05890v3.
- ZHOU, D.-X. (2002). The covering number in learning theory. *J. Complex.* **18**, 739–67.
- ZUBIZARRETA, J. R. (2015). Stable weights that balance covariates for estimation with incomplete outcome data. *J. Am. Statist. Assoc.* **110**, 910–22.

[Received on 28 September 2016. Editorial decision on 20 August 2017]

