

Debashis Ghosh\* and Efrén Cruz Cortés

# A Gaussian Process Framework for Overlap and Causal Effect Estimation with High-Dimensional Covariates

<https://doi.org/10.1515/jci-2018-0024>

Received September 5, 2018; revised June 21, 2019; accepted June 27, 2019

**Abstract:** A powerful tool for the analysis of nonrandomized observational studies has been the potential outcomes model. Utilization of this framework allows analysts to estimate average treatment effects. This article considers the situation in which high-dimensional covariates are present and revisits the standard assumptions made in causal inference. We show that by employing a flexible Gaussian process framework, the assumption of strict overlap leads to very restrictive assumptions about the distribution of covariates, results for which can be characterized using classical results from Gaussian random measures as well as reproducing kernel Hilbert space theory. In addition, we propose a strategy for data-adaptive causal effect estimation that does not rely on the strict overlap assumption. These findings reveal under a focused framework the stringency that accompanies the use of the treatment positivity assumption in high-dimensional settings.

**Keywords:** Average causal effect, Covariate balance, Functional data, Machine learning, Positivity

## 1 Introduction

The availability of high-dimensional covariates in administrative databases and electronic health records has led to increasing scientific focus on attempting to evaluate and develop methods for causal inference with these data structures. There has been a concomitant focus in the statistics and econometrics literature towards the use of machine learning-based methods for performing causal inference with high-dimensional data (e. g., [1, 2, 3, 4, 5]).

In light of this work, it is worth revisiting the standard assumptions necessary for performing causal inference in the potential outcomes framework of Rubin [6] and Holland [7]. A key assumption that is needed for proper definition of a causal estimand is the unconfoundedness assumption, which states that the treatment is independent of the potential outcomes conditional on confounders. Part of the interest in observational studies with high-dimensional covariates is the belief that sufficiently rich sets of covariates will render the unconfoundedness assumption more plausible.

Another assumption, which is the focus of the current paper, is the treatment positivity assumption, which states that the probability of treatment given covariates is strictly between zero and one for values of the covariate vectors. This is related to the notion of covariate/confounder overlap between the treatment groups. This is taken virtually as a given in most causal analyses, but the emergence of clinical decision support systems, deterministic treatment rules and high-dimensional covariates raises the possibility of this assumption being violated. A simple example of such a violation occurring is a situation in which treatment assignment is made deterministically in a medical setting based on the patient presenting with certain risk factors and comorbidities. In such a case, the treatment positivity assumption would be violated. More generally, Robins and Ritov [8] have shown that for estimators of the average causal effect to potentially be semiparametrically

---

\*Corresponding author: Debashis Ghosh, Department of Biostatistics and Informatics, Colorado School of Public Health, 80045 Aurora, CO, United States, e-mail: debashis.ghosh@ucdenver.edu

Efrén Cruz Cortés, Department of Biostatistics and Informatics, Colorado School of Public Health, 80045 Aurora, CO, United States, e-mail: efren.cruzortes@gmail.com

efficient, the treatment positivity assumption has to be strengthened to the propensity score being uniformly bounded away from zero and one.

Methods for performing causal inference with violations of treatment positivity have been limited in the literature. Crump et al. [9] characterized its effects in a setting with limited numbers of covariates and developed a simple rule to exclude subjects based on the propensity score. The remaining subjects would be those for whom there is sufficient covariate overlap by which one could make valid causal inference. Traskin and Small [23] use classification and regression trees (CART) to model group labels derived from the Crump et al. [9] definition of a study population with sufficient overlap to identify factors by which one can define a study population for which one can make causal inferences about. A recent proposal from Ghosh [10] suggests defining study populations based on the margin from machine learning algorithms. Another practical approach is to delete observations with extreme propensities close to zero or one, which is known as propensity score trimming. Note that these approaches all estimate a causal parameter that is effectively data-dependent; see Ghosh [10] for further discussion.

In recent work, D'Amour et al. [11] studied the notion of strict overlap in the high-dimensional case and described a concept termed strict covariate overlap. This represents one way of extending the treatment positivity assumption into the high-dimensional setting. For this situation, they showed that their covariate overlap assumption implied a bound on the discrepancy between the joint distributions of confounders among the treatment and control groups. The assumption thus places an immediate bound on the difference in joint distributions between treatment and control population. In order to avoid such a restrictive assumption, D'Amour et al. [11] suggested that one could make sparsity assumptions of several types. First, one could assume sparsity at the level of the propensity score. Second, one could assume the existence of a latent variable that renders treatment independent of the potential outcomes, much as in the standard unconfoundedness assumption. Finally, one could assume the existence of a low-dimensional subspace such that confounders are independent of the potential outcomes conditional on subspace.

Central to the notion of strict overlap is that of bounded likelihood ratio, results for which have been developed by Rukhin [12, 13] and exploited by D'Amour et al. [11]. In the current paper, we study the overlap phenomenon using Gaussian processes, which is related to but different from what was done in D'Amour et al. [11]. In particular, we will use the theory of random probability measures, and in particular, random Gaussian measures [14, 15] to characterize implications of overlap. The line of research we use enjoys a long history in statistical theory, dating back to the results on random measures based on Wiener processes developed by Cameron and Martin [16, 17]. We show that by assuming a Gaussian process framework, the bounded likelihood ratio result presented in D'Amour et al. [11] leads to an asymptotic implication of overlap that can be characterized by equivalence and orthogonality of Gaussian measures. These results have been applied to a variety of problems in statistics, including spatial statistics [18] and more recently, classification with functional data [19, 20]. The applications of these results to the causal inference setting are new.

Next, we show that for a specific Gaussian process model, a phase transition phenomenon exists with respect to the overlap assumption that can be characterized in terms of the component eigenvalues and eigenfunctions. This will entail the use of and summarizing results from Delaigle and Hall [19] and Berrendero et al. [20]. The practical implication of these results are manifold:

1. For causal inference with high-dimensional covariates, the Gaussian process framework reveals that assuming treatment positivity is equivalent to assuming equivalence of the probability laws of the confounders given treatment groups;
2. If equivalence does not hold, then one has that the probability laws of the confounders given treatment groups are orthogonal measures, which represents a complete violation of overlap;
3. We show how one can develop a causal effect estimation strategy that is less reliant on the covariate overlap assumption. It extends the work of Ghosh [10] for estimating data-adaptive causal estimands based on the margin.

The structure of this paper is as follows. In Section 2, we review the potential outcomes model, previous work on covariate overlap as well as provide a review on Gaussian processes. The latter will prove useful in the framework that we describe in Section 3. In Section 4, some practical implications for causal modelling

approaches are provided, along with an extension of the strategy from Ghosh [10] to incorporate covariance structures. Some discussion concludes the paper in the last section.

## 2 Background

### 2.1 Preliminaries and causal inference assumptions

In this paper, we will employ the potential outcomes framework of Rubin [6], Holland [7], which has been widely used in causal modelling. Throughout, we denote the set  $\{1, \dots, n\}$  by  $[n]$ . We start by assuming an underlying probability space  $(\Omega, \Sigma, \mathbb{P})$ , and random variables  $Z : (\Omega, \Sigma, \mathbb{P}) \rightarrow (\mathcal{Z}, \Sigma_{\mathcal{Z}})$ ,  $Y : (\Omega, \Sigma, \mathbb{P}) \rightarrow (\mathcal{Y}, \Sigma_{\mathcal{Y}})$ , and  $T : (\Omega, \Sigma, \mathbb{P}) \rightarrow (\mathbb{B}, \Sigma_{\mathbb{B}})$ , where  $\mathbb{B} := \{0, 1\}$ , and  $\mathcal{Z}, \mathcal{Y}$  are the confounder and response spaces, respectively. The derived distribution of  $Z$  is defined as  $P_Z := \mathbb{P} \circ Z^{-1}$ . We define  $P_Y$  and  $P_T$  analogously.  $Y$  denotes the response of interest and  $Z$  is a  $p$ -dimensional vector of confounders.  $T$  is a binary indicator of treatment exposure, where  $T = 1$  if treated and  $T = 0$  if control. We denote their joint distribution by  $P$ , and assume the observed data, represented as  $\{(Y_i, Z_i, T_i)\}_{i \in [n]}$ , is drawn from  $P$ . Note that, in this section, we will assume the confounders are simply a finite-dimensional vector.

Let  $\{Y_i(0), Y_i(1)\}$  be the potential outcomes for subject  $i$ ,  $i \in [n]$ , where  $Y_i(0)$  refers to the potential outcome under control and  $Y_i(1)$  to that under treatment. We further assume our observation  $Y_i$  has the form  $Y_i := Y_i(T_i) = Y_i(1)T_i + Y_i(0)(1 - T_i)$ , which is commonly referred to as the consistency assumption in the causal inference literature.

In an observational study, the vector of covariates  $Z$  could be related to both the outcome and the treatment assignment. Since both  $T$  and the potential outcomes  $\{Y(0), Y(1)\}$  are affected by  $Z$ ,  $T \perp \{Y(0), Y(1)\}$  will not hold. To enable causal inference in this scenario, we make the following further assumptions.

1. Strongly Ignorable Treatment Assumption (SITA):  $\{Y(1), Y(0)\}$  is independent of  $T$  given  $Z$ .
2. Stable Unit Treatment Value Assumption (SUTVA): The potential outcomes for subject  $i$  is statistically independent of the potential outcomes for all subjects  $j \neq i$ , for  $i, j \in [n]$ .
3. Treatment Positivity Assumption (TP):  $P_{T|Z}(T = 1|Z) > 0$  for all values of  $Z$ .

Taking these assumptions in order, SITA means that by conditioning on  $Z$ , the observed outcomes can be treated as if they come from a randomized complete block design. Rosenbaum and Rubin [21] show that if SITA holds, then the treatment is independent of the potential outcomes given the propensity score, defined as

$$e(Z) := P_{T|Z}(T = 1|Z).$$

Robins and Ritov [8] uses the terminology ‘no unmeasured confounders’ in lieu of SITA. The TP assumption was described in the Introduction and will be considered further in the next section.

In passing, we mention that the typical parameter of focus in causal analyses is the average causal effect (ACE), defined as

$$ACE := E(Y(1) - Y(0)), \quad (1)$$

where the expectation is taken with respect to the joint distribution of  $Y_i(0)$  and  $Y_i(1)$ . The use of propensity score modelling [21] in conjunction with outcome regression modelling leads to a variety of approaches to average causal effect modelling, a comprehensive overview on the topic being Imbens and Rubin [22].

### 2.2 Review of previous work regarding treatment positivity and covariate overlap

We focus on the treatment positivity assumption and describe previous work in this area. Crump et al. [9] noted the possibility that treatment positivity could be violated and instead defined a subpopulation causal

effect using the propensity score. Let  $\mathbf{I}(C)$  denote the indicator function for the event  $C$ . Define the region  $\mathcal{A} := \{Z : c \leq e(Z) \leq 1 - c\}$  for some  $c > 0$ . Note we suppress the explicit dependence of  $\mathcal{A}$  on  $c$ . Crump et al. [9] define the subpopulation average causal effect as

$$ACE_{\mathcal{A}} := \frac{\sum_{i=1}^n \mathbf{I}(\{e(Z_i) \in \mathcal{A}\}) (Y_i(1) - Y_i(0))}{\sum_{i=1}^n \mathbf{I}(\{e(Z_i) \in \mathcal{A}\})}.$$

Note the dependence of  $ACE_{\mathcal{A}}$  on the region of the propensity scores that is in  $\mathcal{A}$ . In practice,  $\mathcal{A}$  must be estimated from the data, and this is done by estimating the propensity score to obtain  $\hat{e}(Z_i)$ ,  $i = 1, \dots, n$ . By plugging in  $\hat{e}(Z_i)$  instead of  $e(Z_i)$  in the definitions of  $\mathcal{A}$  and  $ACE_{\mathcal{A}}$ , we obtain  $\widehat{\mathcal{A}}$  and  $\widehat{ACE}_{\mathcal{A}}$ , respectively. We finally employ  $\widehat{ACE}_{\mathcal{A}}$  as our final estimator. Based on the variability of the estimated subpopulation average causal effect, Crump et al. [9] proposed an optimization criterion for determining an optimal cutoff value in the definition of  $\mathcal{A}$  and demonstrated under some mild assumptions that an optimal  $c^*$  exists. The optimal value depends only on the marginal distribution of the propensity scores.

Traskin and Small [23] developed a meta-approach for characterizing treatment positivity using classification and regression trees [24]. In general, a tree classifier works as follows: beginning with a training data set  $\{(Z_i, T_i)\}_{i \in [n]}$ , drawn from the joint distribution of  $Z$  and  $T$ , a tree classifier repeatedly splits nodes based on one of the covariates in  $Z$ , until it stops splitting by some stopping criteria (for example, the terminal node only contains training data from one class). Each terminal node is then assigned a class label by the majority of subjects that falls in that terminal node. Once a testing data point with a covariate vector  $Z$  is introduced, the data point is run from the top of the tree until it reaches one of the terminal nodes. The prediction then will be made by the class label of that terminal node. Compared to parametric algorithms, tree-based algorithms have several advantages. There is no need to assume any parametric model for a tree; the algorithm for its constructions only requires a criterion for splitting a node and a criterion for when to stop splitting [24]. Traskin and Small [23] propose a modelling approach in which one develops a class label for subject  $i$  depending on whether or not  $\hat{e}(Z_i) \in \widehat{\mathcal{A}}$ ,  $i \in [n]$ . Define  $Y_i^* := \mathbf{I}(\hat{e}(Z_i) \in \widehat{\mathcal{A}})$ ,  $i \in [n]$ . Traskin and Small [23] then propose fitting a tree model for  $Y_i^*$  on  $Z_i$  to determine the covariates that explain being in the overlap set of Crump et al. [9]. We term this a ‘meta-approach’ because the modeling is being based on a response variable that is derived using the estimated propensity score.

On the theoretical front, Khan and Tamer [25] demonstrated that if the TP assumption is violated, then irregularities regarding identification and inference about causal effects can occur. This has echoes in the work of Robins and Ritov [8], who show that to have regular semiparametric estimators for average causal effects in the high-dimensional case, the model classes for the propensity score and outcome models have to be well-behaved. Thus, violations in standard overlap assumptions lead to irregularities in estimation and inference. This was noted in Luo et al. [26], who found that by assuming a weaker covariate overlap assumption, one could derive an estimator that exhibited super-efficiency (i.e., having an information bound that is smaller than the classical semiparametric information bound for regular estimators). This phenomenon also occurs in the collaborative targeted maximum likelihood estimator of van der Laan and Gruber [1]. The problem with superefficient estimators, as noted by D’Amour et al. [11], is that for model directions where the relaxed assumptions do not hold, the estimators can have unbounded loss functions.

D’Amour et al. [11] consider the problem of covariate overlap in one high-dimensional setting. Define the conditional probabilities  $P_{\ell}(Z) := P_{Z|T}(Z|T = \ell)$ ,  $\ell \in \mathbb{B}$ , and denote their Radon-Nikodym derivatives with respect to  $P_Z$  as  $f_{\ell} := dP_{\ell}/dP_Z$ . Let  $\alpha = P_T(T = 1)$  and note  $P_Z = \alpha P_1 + (1 - \alpha) P_0$ . They note that the strict overlap assumption:

$$\eta \leq e(Z) \leq 1 - \eta \quad \text{almost everywhere } P_Z \tag{2}$$

for some  $\eta \in (0, 0.5)$ , is equivalent to the following bounded likelihood ratio assumption: for  $P_Z$ -almost all  $z \in \mathcal{Z}$ :

$$\frac{\eta}{1 - \eta} \leq \frac{\alpha}{1 - \alpha} LR(z) \leq \frac{1 - \eta}{\eta} \tag{3}$$

where  $LR(z) := f_1(z)/f_0(z)$  is the likelihood ratio of the treatment to the control populations. For the sake of completeness, we provide a straightforward proof in the appendix, since none was given in the original paper. We note assumption (3) is a bounded likelihood ratio assumption. This allows D'Amour et al. [11] to exploit results from information theory [12, 13] to show that (3) imposes limits on the rate of growth of discriminatory information between the joint distributions of confounders in the treatment and control populations. Interestingly, these bounds are independent of  $p$ , the number of confounders. From an intuitive point of view, this makes sense, for when the number of covariates increase, one would expect that the probability of random classifiers to perfectly separate the data between the treatment and control populations would increase. This proves problematic for the causal inference problem, as it leads to a violation of the TP assumption.

## 2.3 Review of Gaussian processes

We now consider the case where  $Z$  is a Gaussian process instead of a finite-dimensional Gaussian variable.  $Y$  and  $T$  remain as in the previous section but now  $Z : (\Omega, \Sigma, \mathbb{P}) \rightarrow (\mathcal{Z}(\mathcal{T}), \Sigma_{\mathcal{Z}})$  where  $\mathcal{Z}(\mathcal{T})$  is a suitable function space on  $\mathcal{T}$ . To simplify things we can assume these functions are real-valued. It will be more useful to understand  $Z$  in the following way: write  $Z(\omega, \cdot)$  for the elements of  $\mathcal{Z}(\mathcal{T})$ , and define the maps  $\pi_t(Z) := Z(\cdot, t) \forall t \in \mathcal{T}$ . Then,  $Z_t := \pi_t \circ Z$  are functions from  $\Omega$  to  $\mathbb{R}$ . Indeed, each of these functions is a random variable if and only if  $Z$  is also a random variable. Indeed  $Z$  is fully determined by  $\{Z_t\}_{t \in \mathcal{T}}$  and, furthermore, its distribution is fully determined by the finite dimensional joint distributions  $\{Z_{t_i}\}_{i \in n}$  for all collections  $\{t_i\}_{i \in n}$  and all  $n$ . When such distributions are all Gaussian,  $Z$  is a Gaussian process. If  $Z_t$  is square-integrable for all  $t \in \mathcal{T}$ , we call  $Z$  a second-order process, throughout the paper we assume  $Z$  is second-order. Comprehensive reviews of Gaussian processes can be found in Neveu [14] and Jansson [15].

### 2.3.1 Associated Hilbert space

For simplicity, we assume here that  $E(Z_t) = 0$  for all  $t \in \mathcal{T}$ , although this assumption will be relaxed later in the paper. Gaussian Processes can be characterized by their covariance function, given by

$$k(s, t) = \text{Cov}\{Z_s, Z_t\}, \quad (s, t) \in \mathcal{T} \times \mathcal{T}.$$

For a second-order Gaussian Process, the covariance function  $k$  will be a symmetric and positive definite function on  $\mathcal{T} \times \mathcal{T}$ . Well known results from functional analysis ensure each function  $k$  of such type is in one-to-one correspondence to a Hilbert space  $\mathcal{H}_k$  with the following reproducing property:

$$\langle f, k(\cdot, t) \rangle_{\mathcal{H}_k} = f(t), \quad \forall f \in \mathcal{H}_k.$$

In lieu of the reproducing property, We call  $\mathcal{H}_k$  the reproducing kernel Hilbert space (RKHS) of  $k$ , and  $k$  the reproducing kernel of  $\mathcal{H}_k$ . For details see Berlinet and Thomas-Agnan [27], Aronszajn [28].

### 2.3.2 Equivalence of distributions

**Definition.** Two measures  $\mu_0$  and  $\mu_1$  defined on a probability space  $(X, \mathcal{X})$  are said to be mutually singular (or orthogonal), denoted by  $\mu_0 \perp \mu_1$ , if there exists a set  $A \in \mathcal{X}$  such that  $\mu_0(A) = 0$  and  $\mu_1(X \setminus A) = 0$ .  $\mu_0$  and  $\mu_1$  are said to be mutually equivalent (or equivalent), denoted by  $\mu_0 \sim \mu_1$ , if for all  $A \in \mathcal{X}$ ,  $\mu_0(A) = 0 \Leftrightarrow \mu_1(A) = 0$ .

It is well-known from probability theory that two Gaussian measures defined on the same measure space will either be orthogonal or equivalent [29, 30]. Assume  $\mu$  to be a measure that dominates both  $\mu_0$  and  $\mu_1$  (e.g.,  $\mu = \mu_0 + \mu_1$ ). Define  $d\mu_0/d\mu$  and  $d\mu_1/d\mu$  to be the Radon-Nikodym derivatives associated with  $\mu_0$  and  $\mu_1$ ,

respectively. Furthermore, define

$$B := \int \sqrt{\frac{d\mu_1}{d\mu} \frac{d\mu_0}{d\mu}} d\mu$$

and

$$J := \int \left( \frac{d\mu_0}{d\mu} - \frac{d\mu_1}{d\mu} \right) \log \left( \frac{d\mu_0/d\mu}{d\mu_1/d\mu} \right) d\mu.$$

Note that  $B$  is the Bhattacharya coefficient between two probability measures and  $J$  is the relative entropy. Several authors have developed characterization results for orthogonality and equivalence of Gaussian measures in terms of  $B$  and  $J$  [29, 31, 32]. We summarize them here using Theorem 1 from Shepp [32].

**Theorem 1** (Theorem 1 of Shepp [32]).

- (a)  $\mu_0 \perp \mu_1$  if and only if  $B = 0$  or  $J = \infty$ .
- (b)  $\mu_0 \sim \mu_1$  if and only if  $B > 0$  or  $J < \infty$ .

If  $\{\mu_0^{(n)}\}$  and  $\{\mu_1^{(n)}\}$  are sequences of Gaussian measures with corresponding Bhattacharya and relative entropy sequences  $\{B_n\}$  and  $\{J_n\}$  converging to  $B$  and  $J$ , then Theorem 1 suggests there are two asymptotic scenarios for Gaussian measures as  $n$  approaches infinity. They are either asymptotically orthogonal (part (a)) or equivalent (part (b)).

While the Gaussian process assumption may seem restrictive at first glance, we note that they in fact represent a very flexible class of probability models to fit to data. They have been widely used in Haran [33] for spatial statistics, Kennedy and O'Hagan [34] for designs of computer experiments and by Williams and Rasmussen [35] for machine learning.

### 3 Proposed framework

#### 3.1 Overlap assumptions and Gaussian processes

As before, we represent the random measures for  $Z|T = 1$  and  $Z|T = 0$  by  $P_1$  and  $P_0$ , respectively. We now reconsider the assumptions of causal inference for  $Z$  a Gaussian process. The functional version of unconfoundedness can be given by

$$T \perp \{Y(1), Y(0)\} | \mathcal{F}_Z,$$

where  $\mathcal{F}_Z := \sigma(Z_t : t \in \mathcal{T})$  denotes the smallest  $\sigma$ -algebra for which every  $Z_t$  is measurable.

The assumption of SUTVA remains the same as described in Section 2. Finally, the treatment positivity assumption can be written as

$$0 < P(T = 1|Z) < 1 \text{ almost everywhere } P_Z.$$

Analogously to (3) we define strict functional overlap to be

$$\frac{\eta}{1-\eta} < \frac{\alpha}{1-\alpha} \frac{dP_1/dP(z)}{dP_0/dP(z)} < \frac{1-\eta}{\eta}, \quad (4)$$

for almost-all  $z \in \mathcal{Z}(\mathcal{T})$  and for  $\eta \in (0, 5)$ . This is analogous to the bounded likelihood ratio assumption of D'Amour et al. [11].

Define  $\{P_0^{(n)}\}$  and  $\{P_1^{(n)}\}$  to be the restrictions of  $P_0$  and  $P_1$ , respectively, to the smallest  $\sigma$ -algebra for which  $\{Z_1, \dots, Z_n\}$  is measurable. Let the means be denoted as  $\{m_0^{(n)}\}$  and  $\{m_1^{(n)}\}$ . We have the following result:

**Theorem 2.** *Assumption (4) implies that  $P_0 \sim P_1$ .*

*Proof.* Using integration, one can show that

$$B(P_0^{(n)}, P_1^{(n)}) = \frac{1}{4}L_n + \frac{1}{8}D_n^2,$$

where

$$L_n = 2\log |\Lambda_n| - \log |\Lambda_{0n}| - \log |\Lambda_{1n}|,$$

with  $\Lambda_{jn}$  denoting the covariance operator in  $n$ -dimensional space for group  $j$  ( $j = 0, 1$ ) and  $\Lambda_n = (\Lambda_{0n} + \Lambda_{1n})/2$ . The term  $D_n^2 := (m_1^{(n)} - m_0^{(n)})' \Lambda_n^{-1} (m_1^{(n)} - m_0^{(n)})$  represents the Mahalanobis distance in  $n$ -dimensional space. By Theorem 3.2 of Rao and Varadarajan [31], we have that

$$\begin{aligned} B(P_0, P_1) &= \lim_{n \rightarrow \infty} B(P_0^{(n)}, P_1^{(n)}) \\ &= \lim_{n \rightarrow \infty} \frac{1}{4}L_n + \frac{1}{8}D_n^2. \end{aligned}$$

As noted in Rao and Varadarajan [31],  $L_n/4$  represents the Hellinger distance between two zero-mean multivariate normal distributions, one with covariance  $\Lambda_{0n}$ , the other with covariance  $\Lambda_{1n}$ . Similarly,  $D_n^2/8$  is the Hellinger distance between two  $n$ -dimensional multivariate normal distributions with means  $m_0^{(n)}$  and  $m_1^{(n)}$  and common covariance  $\Lambda_n$ . Now, note that  $L_n$  and  $D_n$  are both greater than equal to zero for all  $n$  and are increasing with  $n$ . The assumption (4), which is a bounded likelihood ratio assumption implies that  $L_n$  and  $D_n^2$  are uniformly bounded away from zero and infinity for all  $n$ . Thus, their limits will also be greater than zero so that  $B(P_0, P_1) > 0$ . We can thus use Theorem 1(b) to conclude equivalence. This concludes the proof of Theorem 2.  $\square$

**Remark.** Note that Theorems 1 and 2 apply to the situation where we assume that the joint distributions of confounders are exactly assumed to be Gaussian processes. Thus, they represent exact results in a statistical sense. However, in practice we could consider an alternate assumption like

$$\frac{\eta^*}{1 - \eta^*} < \frac{\alpha^*}{1 - \alpha^*} \frac{dP_1^{(n)}/dP^{(n)}(z)}{dP_0^{(n)}/dP^{(n)}(z)} < \frac{1 - \eta^*}{\eta^*}, \quad (5)$$

for all  $n$ , almost-all  $z \in \mathcal{Z}(\mathcal{T})$  and for  $\eta \in (0, .5)$ . In (5), the superscript  $n$  refers to the size of the dataset, so we are now in a situation where as  $n$  approaches infinity, we get to a Gaussian process in the limit. In this case, we can modify the proof of Theorem 2 to show the following:

**Corollary 1.** *Assume that (5) holds. Let  $\{P_0^{(n)}\}$  and  $\{P_1^{(n)}\}$  denote sequences of measures that converge weakly to  $P_0$  and  $P_1$ , respectively. Then as  $n \rightarrow \infty$ ,  $P_0 \sim P_1$ .*

Examples for weak convergence of data-dependent probability measures converging to Gaussian processes can be found in Billingsley [36].

### 3.2 Reinterpretation using RKHS theory and phase transition

We showed in Section 2.2 how a Gaussian stochastic process can be used to define an RKHS. In this section, we will assume that  $P_0$  and  $P_1$  represent the probability laws of Gaussian processes with mean functions  $m_0$  and  $m_1$ , respectively, and shared covariance function  $k$ . We further assume  $m_0 = 0$  and  $m_1$ , and  $k$  are continuous in their respective domains. If  $\mathcal{T}$  is compact, then  $k$  is a Mercer kernel and it admits an expansion of the form

$$k(s, t) = \sum_{j=1}^{\infty} c_j \psi_j(s) \psi_j(t), \quad (6)$$

where for  $j = 1, 2, \dots$ ,  $\{c_j\}$  and  $\{\psi_j\}$  denote the eigenvalues and eigenfunctions of the integral operator associated to  $k$ . For details on such operators, refer to Steinwart et al. [37]. Based on the decomposition (6), we can decompose the signal  $m_1$  as

$$m_1(t) = \sum_{j=1}^{\infty} a_j \psi_j(t).$$

Recently, Delaigle and Hall [19] and Berrendero et al. [20] have studied and developed this as a model for classification with functional data. These authors referred to the problem as one of ‘near-perfect’ classification of functional data. Of relevance is a quote by Delaigle and Hall [19]: ‘these (functional classification) problems have unusual, and fascinating properties that set them apart from their finite-dimensional counterparts. In particular, we show that in many quite standard settings, the performance of simple (linear) classifiers constructed from training samples becomes perfect as the sizes of the samples diverge....That property never holds for finite-dimensional data, except in pathological cases.’

Delaigle and Hall [19] and Berrendero et al. [20] developed technical characterizations of the ‘near-perfect’ classification phenomenon. In the context of our discussion, the work of Berrendero et al. [20] is perhaps the most relevant. We have the following theorem.

**Theorem 3.**

- (a) *If (4) holds, then  $\sum_{j=1}^{\infty} c_j^{-1} a_j^2 < \infty$ .*
- (b) *If  $\sum_{j=1}^{\infty} c_j^{-1} a_j^2 = \infty$ , then (4) cannot hold.*

*Proof.* We first begin with the proof of Theorem 3a. By Theorem 2, if (4) holds, then  $P_1 \sim P_0$ . By Theorem 5a of Berrendero et al. [20], this is identical to

$$\sum_{j=1}^{\infty} c_j^{-1} a_j^2 < \infty.$$

For the proof of Theorem 3b, by Theorem 4b of Berrendero et al. [20],  $\sum_{j=1}^{\infty} c_j^{-1} a_j^2 = \infty$  is identical to  $P_1 \perp P_0$ . Note that for probability measures  $P_1 \perp P_0$  implies  $P_0$  and  $P_1$  are not mutually equivalent. Hence, by contraposition of Theorem 2, Assumption (4) cannot hold.  $\square$

Theorem 3 captures a certain type of phase transition in the functional covariate overlap behavior based on the properties of the operator associated to  $k$ . In particular, the squared eigenvalues of the mean signal divided by the eigenvalues for the kernel operator, characterize strict functional overlap. If strict functional overlap holds, the summation converges (Theorem 3a), while divergence of the series (Theorem 3b) implies strict functional overlap not holding. The results of Theorem 3 are usable in situations where we can derive an explicit Mercer expansion for  $k$  in the form of (6). If the  $\{\psi_n\}$  represent a countably infinite set of orthonormal basis functions, then  $\{a_n\}$  ( $n \geq 0$ ) corresponds to the spectrum of  $k$ , and if  $k$  in (6) is symmetric and positive definite, then  $a_n \geq 0$  for all  $n$ . We now give some examples on the application of Theorem 3.

**Example.** Suppose we set  $k$  to be the Gaussian kernel in one dimension:

$$k(x, y) = \exp\{-(x - y)^2/2\sigma^2\},$$

where  $x$  and  $y$  are scalars,  $\sigma^2$  represents a scale parameter. In Steinwart et al. [37], the authors show how to represent the Gaussian kernel in the form (6), with

$$c_n = \sqrt{\frac{2\sigma^{2n}}{n}}, \quad \psi_n(x) = x^n \exp(-\sigma^2 x^2).$$

By Theorem 3a, if (4) holds, then

$$\sum_{j=1}^{\infty} c_j^{-1} a_j^2 < \infty.$$

In addition, if

$$\sum_{j=1}^{\infty} c_j^{-1} a_j^2 = \infty,$$

then (4) cannot hold. Noting that the summation will converge if  $c_j^{-1} a_j^2 = j^{-(1+\epsilon)}$  for some  $\epsilon > 0$ , some algebra yields that this is equivalent to

$$a_j = \sigma^{-j/2} j^{-1/2(1/2+\epsilon)}$$

for  $j \geq 0$ . Recalling the definition of  $a_j$ , this places an immediate constraint on the mean function  $m_1(t)$ .

## 4 Practical implications for causal inference

While the results presented so far have been of a relatively theoretical nature, one practical implication of the results in Section 3 is that by assuming a particular Gaussian process with functionally structured covariates, there is a strong tension and interplay between (a) maintaining a treatment positivity assumption and (b) having the ability to compute causal effect estimators that are semiparametrically regular in the sense of Robins and Ritov [8]. One common practice for causal effect estimation is the following strategy:

1. Fit a propensity score model for treatment given covariates;
2. Evaluate covariate balance between the treatment groups, taking into account the propensity scores. If balance does not hold, modify step 1.
3. Estimate the causal effect of interest incorporating the propensity score.

If a propensity score model can perfectly predict treatment, then we are in a situation where the treatment positivity assumption is violated. Thus, for causal modelling to proceed, one would need stronger assumptions to hold. For example, targeted learning approaches also fit a model for the potential outcomes, and if that model is correctly specified, then this would overcome the issue of violation of treatment positivity. More fundamentally, the violations in overlap necessitate the use of data-adaptive causal effect estimation approaches, which we describe in the next section.

### 4.1 Data-adaptive causal effect estimation

A recent innovation for attempting to relax some of the overlap assumptions is the margin-based approach of Ghosh [10]. The intuitive idea of that work is to first identify the margin and then conditional on the margin, develop causal effect estimation and inference procedures. This is related to, but different from, the finding in D'Amour et al. [11] that a set of observations where propensity score model misclassifications occur (i. e., the predicted treatment assignment is not concordant with the observed treatment assignment) represent a set of individuals which might be concordant with the treatment positivity assumption.

We consider the binary treatment scenario and recode treatment as  $-1$  and  $1$  and use support vector machines (SVM; [38]) to fit a propensity score model for treatment. The objective of SVM is to find a hyperplane, in an appropriate high-dimensional space, that maximizes the margin between the populations defined by  $T = 1$  and  $T = -1$ . One approach to mapping data into a high-dimensional space is through use of the so-called 'kernel trick', which corresponds to using RKHS to define support vector machines. Using a loss function framework, we can characterize the optimization problem as finding to  $f$  to minimize

$$\sum_{i=1}^n |1 - T_i f(Z_i)|_+ + \lambda \|f\|_K^2,$$

where  $|a|_+ = \max(a, 0)$ ,  $\lambda > 0$  denotes a smoothing parameter and  $\|f\|_K$  denotes the norm of a function  $f$  in an RKHS with associated kernel  $K$ . At a high level, the approach being proposed here amounts to the following:

1. Fit a support vector machine with kernel  $K$  to the data  $(T_i, \mathbf{Z}_i)$ ,  $i = 1, \dots, n$ ;
2. Determine the observations that are in the margin; denote these as  $\mathcal{M}$ ;
3. Estimate the causal effect of interest using  $(Y_i, T_i, \mathbf{Z}_i)$ ,  $i \in \mathcal{M}$ ;

Extending the arguments of Theorem 1 in Ghosh [10], the margin for nonlinear support vector machines based on a kernel  $K$  can be viewed as satisfying a functional version of relaxed covariate overlap. The relaxed covariate overlap notion is detailed more clearly in Ghosh [10] and as the name suggests, is meant to be a relaxation of the standard covariate overlap assumption in causal inference.

## 4.2 Numerical example

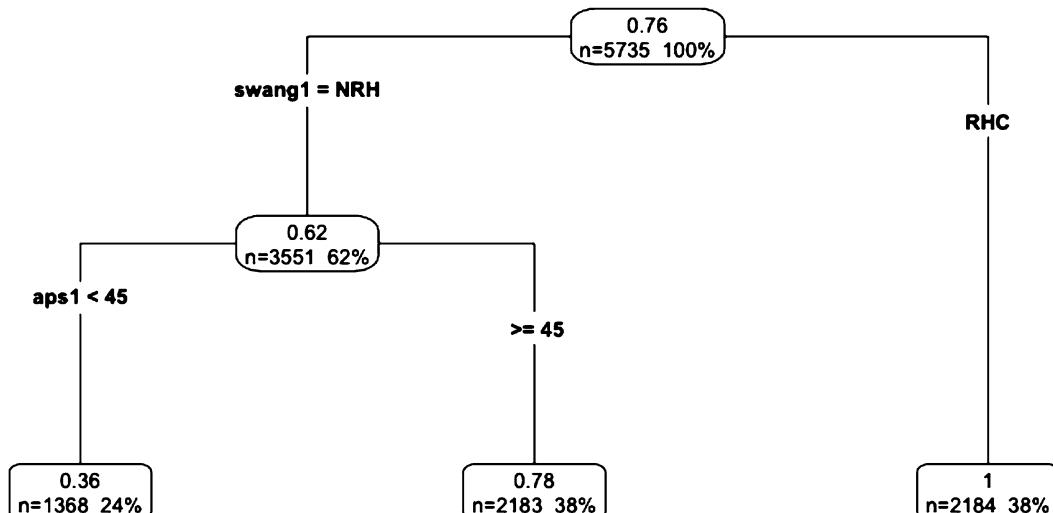
As an example, we consider the SUPPORT study from Connors et al. [39], a version of which was previously analyzed in Ghosh et al. [40]. The causal effect of interest is the effect of right heart catheterization (RHC) on 30-day survival (dead/alive at 30 days). The dataset contains information on 5735 patients, 2184 of whom received RHC. Note that the covariates are a mixture of continuous as well as discrete variables, so the assumption of a Gaussian process for the data may not be terribly realistic. Determining the robustness of results to this assumption is an important topic for future work.

We employed the implementation of support vector machines as available in the R e1071 package based on the Gaussian kernel:

$$K(x, y) = \exp(-\|x - y\|^2/2\sigma^2),$$

where  $x$  and  $y$  are  $p$ -dimensional vectors, and  $\sigma > 0$  is a scale parameter. Out of the original 5735 subjects, 3663 are selected to be in the margin, and based on these subjects, the average causal effect is estimated to be 0.049 with an associated standard error of 0.016, which yields a highly significant effect. Thus, the use of RHC is associated with decreased probability of 30-day survival, or conversely, increased risk of death within 30 days. Because of the sensitivity of  $K$  to the choice of  $\sigma^2$ , we performed sensitivity analyses in which we varied  $\sigma^2$  and reran the analysis. In all instances, we found a positive average causal effect. This aligns with the findings in Connors et al. [39].

We also used the methodology of Traskin and Small [23] to better understand how the margin population was selected. The plot of the tree is given in Figure 1. Based on Figure 1, we find that all the subjects treated



**Figure 1:** Tree plot from modelling being in the margin as a function of the covariates in the SUPPORT study. The tree was pruned using a cost-complexity parameter of 0.1.

with right-heart catheterization are included in the margin. In addition, patients with higher APACHE scores, which correlate with higher mortality, are much more likely to be included in the margin.

## 5 Discussion

In this article, we have studied the covariate overlap assumption using Gaussian processes. The results here complement those in D'Amour et al. [11]. The way in which high-dimensional covariates are considered is different from that paper. There, they consider the situation of  $p$ , the number of confounders, approaching infinity. By contrast, here, we use a Gaussian random measure to represent a theoretically infinite-dimensional confounder. An implication of the analysis here is that the treatment positivity assumption is not as innocuous as it seems, which is similar to the conclusion reached by D'Amour et al. [11]. As noted by a reviewer, while assuming exact Gaussian process models for the confounders given treatment groups is restrictive in that it leads to a strict dichotomy of overlap versus not, we also described how one could view these results in an asymptotic sense at the end of Section 3.1. Making the Gaussian process assumption allows one to mathematically characterize the interplay of the mean function with the covariance of the noise with respect to overlap as provided in Theorem 3.

Another take-home message is that in high dimensions, there is an inherent tension between covariate overlap and regular asymptotics for causal effect estimators. Virtually all proposals that are available for causal effect estimation assume a strict covariate overlap assumption or equivalently, the propensity score being bounded away from zero and one uniformly in the confounders. Letting  $n$  denote the sample size and  $p$  the dimension of the confounders, this could be in a 'large  $n$ , small  $p$ ' case (e. g., [41]) or in a 'large  $p$ , small  $n$ ' setup (e. g., [4]). The Gaussian process framework leads to the most extreme situation in a sense: one either has full covariate overlap (i. e., equivalence) or completely disjoint covariate overlap (i. e., orthogonality). For the former situation, regular asymptotics for causal effect estimators holds trivially. Under the latter situation of orthogonality, this effectively renders impossible the possibility of average causal effect estimators achieving  $n^{1/2}$  convergence [8]. In their words, relaxing the assumption of strict covariate overlap in the high-dimensional case allows for 'pathological' distributions to be considered as part of the model class, and these distributions have the possibility of leading to causal effect estimators with irregular behavior (e. g., superefficiency). The estimators in Luo et al. [26] and van der Laan and Gruber [1] fall into this class.

This article synthesized several results on Gaussian measures that have been in the literature over the last 70 years. One of the main limitations of this article is that the types of covariates that are treated here as being functional are indexed by a one-dimensional set (e. g., time). An example of biomedical data that could be treated as functional in nature are longitudinal data in electronic medical records, but many covariates in practice might not fit this structure. Understanding the robustness of the results presented in this paper to violations of the Gaussian process assumption are beyond the scope of the current manuscript and a topic for future research.

The paper addresses the issue of equivalence versus orthogonality with Gaussian measures for causal inference problems. In practice, it is difficult to verify if the strict overlap criteria in this article actually hold. This suggests that approaches that can accommodate violations in strict overlap should be considered. Petersen et al. [42] discussed possible violations of the treatment positivity assumption along with potential causes and tools to evaluate their effects in terms of biases of causal effect estimates. Suggestions there include restricting the space of treatments, redefining the causal estimand, and using alternative projection functions. We have proposed redefining the estimand in Section 4.1. Furthering work in these areas will be necessary in order to advance the use of these analytic approaches for causal effect estimation with high-dimensional confounders. It would also be of interest to see if any of these results could be extended to non-Gaussian processes.

The recent use of deep learning techniques for learning representations in causal inference has also been proposed by Johansson et al. [43]. Extension of the results in Section 3 to the deep learning case would also be of major interest. Dunlop et al. [44] have studied 'deep' Gaussian processes, which consist of a recursive

definition for the Gaussian process based on a Markov chain. It would be of interest to extend equivalence results to that setting as well.

**Funding:** This research is supported by a pilot grant from the Data Science to Patient Value (D2V) initiative from the University of Colorado.

## Appendix

### Proof of Equation (3)

We will prove the right hand side inequality only, as the left hand side inequality is proven analogously. The strict overlap assumption is:  $P_{T|Z}(T = 1|Z) \leq \eta$  almost everywhere  $P_Z$ . Hence, by Bayes rule:

$$\frac{P_{Z|T}(Z|T = 1)P_T(T = 1)}{P_Z(Z)} \leq \eta,$$

or, using the notation around equation (3):

$$\begin{aligned} \frac{\alpha P_1(Z)}{\alpha P_1(Z) + (1 - \alpha)P_0(Z)} \leq \eta &\implies \\ \alpha P_1(Z) \leq \eta(\alpha P_1(Z) + (1 - \alpha)P_0(Z)) &\implies \\ \frac{\alpha P_1(Z)}{(1 - \alpha)P_0(Z)} \leq \eta \left( \frac{\alpha P_1(Z)}{(1 - \alpha)P_0(Z)} + 1 \right) &\implies \\ \frac{\alpha P_1(Z)}{(1 - \alpha)P_0(Z)}(1 - \eta) \leq \eta &\implies \\ \frac{\alpha P_1(Z)}{(1 - \alpha)P_0(Z)} \leq \frac{\eta}{1 - \eta}. & \end{aligned}$$

Notice that if a statement holds almost-everywhere  $P_Z$  then it also holds almost-everywhere  $P_0$  and  $P_1$  since these are absolutely continuous with respect to  $P_Z$ .

Now, let  $b := \frac{1-\alpha}{\alpha} \frac{\eta}{1-\eta}$ . If  $P_1(Z) \leq bP_0(Z)$ , almost everywhere  $P_Z$ , then  $f_1(z) \leq bf_0(z)$  for  $P_Z$ -almost all  $z \in \mathcal{Z}$ . Otherwise, assuming a set  $B$  with positive probability for which  $f_1 > bf_0$ , integrating over  $B$  yields  $P_1(Z \in B) > bP_0(Z \in B)$ , a contradiction.

Conversely, if  $f_1(z) \leq bf_0(z)$  for  $P_Z$ -almost all  $z \in \mathcal{Z}$ , we just need to integrate to obtain  $P_1(Z) \leq bP_0(Z)$  almost everywhere  $P_Z$ . Hence,  $P_1 \leq bP_0$  almost everywhere  $\implies f_1 \leq bf_0$  almost everywhere.  $\square$

## References

1. van der Laan MJ, Gruber S. Collaborative double robust targeted maximum likelihood estimation. *Int J Biostat.* 2010;6(1).
2. van der Laan MJ, Rose S. Targeted learning: causal inference for observational and experimental data. Springer; 2011.
3. Athey S, Imbens G. Recursive partitioning for heterogeneous causal effects. *Proc Natl Acad Sci.* 2016;113(27):7353–60.
4. Athey S, Imbens G, Wager S. Approximate residual balancing: De-biased inference of average treatment effects in high dimensions. 1604.07125 (2018).
5. Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, Robins J. Double/debiased machine learning for treatment and structural parameters. *Econom J.* 2018. <https://doi.org/10.1111/ectj.12097>. Accepted Author Manuscript.
6. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educ Psychol.* 1974;66(5):688.
7. Holland P. Statistics and causal inference (with discussion). *J Am Stat Assoc.* 1986;81(396):945–70.
8. Robins JM, Ritov Y. Toward a curse of dimensionality appropriate (coda) asymptotic theory for semi-parametric models. *Stat Med.* 1997;16(3):285–319.
9. Crump RK, Hotz VJ, Imbens GW, Mitnik OA. Dealing with limited overlap in estimation of average treatment effects. *Biometrika.* 2009;96(1):187–99.

10. Ghosh D. Relaxed covariate overlap and margin-based causal effect estimation. *Stat Med.* 2018;37(28):4252–65. 1801.00816.
11. D'Amour A, Ding P, Feller A, Lei L, Sekhon J. Overlap in observational studies with high-dimensional covariates. 1711.02582 (2017).
12. Rukhin AL. Lower bound on the error probability for families with bounded likelihood ratios. *Proc Am Math Soc.* 1993;119(4):1307–14.
13. Rukhin A. Information-type divergence when the likelihood ratios are bounded. *Appl Math.* 1997;24:415–23.
14. Neveu J. Processus aléatoire gaussien. Séminaire de mathématiques supérieures. 1968.
15. Jansson S. Gaussian Hilbert Spaces. Cambridge Tracts in Mathematics. Cambridge; 1997.
16. Cameron RH, Martin WT. Transformations of wiener integrals under translations. *Ann Math.* 1944;386–96.
17. Cameron R, Martin W. Transformations of wiener integrals under a general class of linear transformations. *Trans Am Math Soc.* 1945;58(2):184–219.
18. Stein ML. Interpolation of spatial data: some theory for kriging. Springer; 2012.
19. Delaigle A, Hall P. Achieving near perfect classification for functional data. *J R Stat Soc, Ser B, Stat Methodol.* 2012;74(2):267–86.
20. Berrendero JR, Cuevas A, Torrecilla JL. On the use of reproducing kernel Hilbert spaces in functional classification. *J Am Stat Assoc.* 2018;113(523):1210–8.
21. Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika.* 1983;70(1):41–55.
22. Imbens GW, Rubin DB. Causal inference in statistics, social, and biomedical sciences. Cambridge University Press; 2015.
23. Traskin M, Small DS. Defining the study population for an observational study to ensure sufficient overlap: a tree approach. *Stat Biosci.* 2011;3(1):94–118.
24. Breiman L, Friedman J, Olshen R, Stone C. Classification and regression trees. Statistics/Probability Series. Wadsworth & Brooks/Cole; 1984.
25. Khan S, Tamer E. Irregular identification, support conditions and inverse weight estimation. *Econometrica.* 2010;78(6):2021–42.
26. Luo W, Zhu Y, Ghosh D. On estimating regression causal effects using sufficient dimension reduction. *Biometrika.* 2017;104(1):51–65.
27. Berlinet A, Thomas-Agnan C. Reproducing kernel Hilbert spaces in probability and statistics. Springer Science & Business Media; 2011.
28. Aronszajn N. Theory of reproducing kernels. *Trans Am Math Soc.* 1950;68(3):337–404.
29. Hajek J. A property of  $j$ -divergences of marginal probability distributions. *Czechoslov Math J.* 1958;8(3):460–2.
30. Feldman J. Equivalence and perpendicularity of Gaussian processes. *Pac J Math.* 1958;8(5):699–708.
31. Rao CR, Varadarajan V. Discrimination of Gaussian processes. *Sankhya, Ser A.* 1963;303–30.
32. Shepp L. Gaussian measures in function space. *Pac J Math.* 1966;17(1):167–73.
33. Haran M. Gaussian random field models for spatial data. In: Brooks S, Gelman A, Jones G, Meng X, editors. *Handbook of Markov chain Monte Carlo.* Springer; 2011.
34. Kennedy M, O'Hagan A. Bayesian calibration of computer models. *J R Stat Soc, Ser B, Stat Methodol.* 2001;63(3):425–64.
35. Williams CK, Rasmussen CE. Gaussian processes for machine learning. Cambridge, MA: MIT Press; 2006.
36. Billingsley P. Probability and measure. John Wiley & Sons; 2008.
37. Steinwart I, Hush D, Scovel C. An explicit description of the reproducing kernel Hilbert spaces of Gaussian rbf kernels. *IEEE Trans Inf Theory.* 2006;52(10):4635–43.
38. Cristianini N, Shawe-Taylor J, et al. An introduction to support vector machines and other kernel-based learning methods. Cambridge University Press; 2000.
39. Connors AF, Speroff T, Dawson NV, Thomas C, Harrell FE, Wagner D, Desbiens N, Goldman L, Wu AW, Califf RM, et al. The effectiveness of right heart catheterization in the initial care of critically ill patients. *JAMA.* 1996;276(11):889–97.
40. Ghosh D, Zhu Y, Coffman DL. Penalized regression procedures for variable selection in the potential outcomes framework. *Stat Med.* 2015;34(10):1645–58.
41. Chan KCG, Yam SCP, Zhang Z. Globally efficient non-parametric inference of average treatment effects by empirical balancing calibration weighting. *J R Stat Soc, Ser B, Stat Methodol.* 2016;78(3):673–700.
42. Petersen ML, Porter KE, Gruber S, Wang Y, van der Laan MJ. Diagnosing and responding to violations in the positivity assumption. *Stat Methods Med Res.* 2012;21(1):31–54.
43. Johansson F, Shalit U, Sontag D. Learning representations for counterfactual inference. *ICML.* 2016;3020–9.
44. Dunlop MM, Girolami MA, Stuart AM, Teckentrup AL. How deep are deep Gaussian processes? *J Mach Learn Res.* 2018;19(54):1–46. <http://jmlr.org/papers/v19/18-015.html>.