Bernoulli **31**(4), 2025, 3013–3033
https://doi.org/10.3150/24-BEJ1836

Sequential kernel embedding for mediated and time-varying dose response curves

RAHUL SINGH[1,a], LIYUAN XU[2,b] and ARTHUR GRETTON[2,c]

[1]*Society of Fellows and Department of Economics, Harvard University, Cambridge, MA, USA,* [a]*rahul_singh@fas.harvard.edu*
[2]*Gatsby Computational Neuroscience Unit, University College London, London, UK,* [b]*liyuan.jo.19@ucl.ac.uk,* [c]*arthur.gretton@gmail.com*

We propose simple nonparametric estimators for mediated and time-varying dose response curves based on kernel ridge regression. By embedding Pearl's mediation formula and Robins' g-formula with kernels, we allow treatments, mediators, and covariates to be continuous in general spaces, and also allow for nonlinear treatment-confounder feedback. Our key innovation is a reproducing kernel Hilbert space technique called sequential kernel embedding, which we use to construct simple estimators that account for complex feedback. Our estimators preserve the generality of classic identification while also achieving nonasymptotic uniform rates. In nonlinear simulations with many covariates, we demonstrate strong performance. We estimate mediated and time-varying dose response curves of the US Job Corps, and clean data that may serve as a benchmark in future work. We extend our results to mediated and time-varying treatment effects and counterfactual distributions, verifying semiparametric efficiency and weak convergence.

Keywords: Continuous treatment; reproducing kernel Hilbert space; treatment-confounder feedback

1. Introduction

We study mediation analysis and time-varying treatment effects with possibly continuous treatments. Mediation analysis asks, how much of the total effect of the treatment D on the outcome Y is mediated by a particular mechanism M that takes place between the treatment and outcome? Time-varying analysis asks, what would be the effect of a sequence of treatments $D_{1:T}$ on the outcome Y, even when that sequence may not have been implemented? We consider nonparametric causal functions of continuous treatments. For example, the time-varying dose response curve of two continuous treatments is the function $\theta_0^{GF}(d_1, d_2) := E\{Y^{(d_1,d_2)}\}$, which may refer to medical dosages, lifestyle habits, occupational exposures, or training durations.

The time-varying dose response curve arises when evaluating social programs from several rounds of surveys. For example, the National Job Corps Study randomized access to a large scale job training program in the US and collected several rounds of surveys (Schochet, Burghardt and McConnell, 2008). Individuals could decide whether to participate and for how many hours, possibly over multiple years. A natural question is: how much would the average individual benefit from a certain number of class hours in year one and a possibly different number of class hours in year two? This quantity is an example of a time-varying dose response curve, where the number of class hours is the time-varying dose, and the expected benefit is the response.

The difficulty in estimating time-varying dose response curves is the complex feedback loop resulting from the initial dose. Formally, we model class hours in different years as a sequence of continuous treatments subject to treatment-confounder feedback. In other words, we consider the possibility that class hours in one year may affect health behaviors such as drug use in a subsequent year, which may then affect subsequent class hours. Though several rounds of Job Corps surveys were collected,

1350-7265 This research was funded, in whole or in part, by Sainsbury Wellcome Trust, GAT 3850 and GAT 3528.

economists typically study only the initial survey due to a concern for treatment-confounder feedback, and a lack of simple yet flexible estimators that can adjust for it, while allowing treatments to be continuous. We propose such an estimator using kernels, establish its properties, share a cleaned data set that includes the additional surveys, and carry out empirical analysis on the cleaned data as well as on simulated data. Previous methods using kernels for continuous treatments (Singh, Xu and Gretton, 2024) do not handle treatment-confounder feedback, and therefore cannot analyze the additional surveys. This paper's technical innovation is a sequential kernel embedding to do so.

While mediated and time-varying causal functions are identified in theory, they are challenging to estimate in practice. For example, under standard assumptions on time-varying covariates $X_{1:T}$, the time-varying dose response is identified as the g-formula, i.e. the sequential integral $\theta_0^{GF}(d_1,d_2) = \int \gamma_0(d_1,d_2,x_1,x_2)\mathrm{d}P(x_2|d_1,x_1)\mathrm{d}P(x_1)$, where $\gamma_0(d_1,d_2,x_1,x_2) = E(Y|D_1 = d_1, D_2 = d_2, X_1 = x_1, X_2 = x_2)$ (Robins, 1986). When treatments are continuous, the functional $\gamma \mapsto \int \gamma(d_1,d_2,x_1,x_2)\ \mathrm{d}P(x_2|d_1,x_1)\ \mathrm{d}P(x_1)$ is generally not bounded (van der Vaart, 1991, Newey, 1994) or pathwise differentiable (Bickel et al., 1993, ch. 3 and 5). Popular estimators restrict attention to a binary treatment, parametric models, Markov simplifications, or constrained effect modification for tractability, and may even redefine the estimand. See Vansteelandt and Joffe (2014) for a review. Each of these restrictions simplifies the sequential integral in order to simplify estimation. Our research question is: can we devise simple machine learning estimators for causal functions that preserve the richness of the sequential integral, and therefore the generality of treatment-confounder feedback in classic identification, while also achieving nonasymptotic uniform rates?

In this paper, we match the generality of mediated and time-varying identification with the flexibility and simplicity of kernel ridge regression estimation. We propose a new family of nonparametric estimators for causal inference over short horizons. Our algorithms combine kernel ridge regressions, so they inherit the practical and theoretical virtues that make kernel ridge regression widely used. Crucially, we preserve the nonlinearity, dependence, and effect modification of identification theory with time-varying confounders. Our contribution has three aspects.

First, we introduce an algorithmic technique that appears to be an innovation in the reproducing kernel Hilbert space (RKHS) literature: sequential kernel embeddings, i.e. RKHS representations of mediator and covariate conditional distributions given a hypothetical treatment sequence, which account for treatment-confounder feedback. For example, we introduce the sequential embedding $\mu_{x_1,x_2}(d_1)$ such that the inner product $\langle f, \mu_{x_1,x_2}(d_1)\rangle$ in an appropriately defined Hilbert space equals the g-formula's sequential integral $\int f(x_1,x_2)\mathrm{d}P(x_2|d_1,x_1)\mathrm{d}P(x_1)$. We prove that the sequential kernel embedding exists because the RKHS restores boundedness of the g-formula's functional $\gamma \mapsto \int \gamma(d_1,d_2,x_1,x_2)$ $\mathrm{d}P(x_2|d_1,x_1)\mathrm{d}P(x_1)$, even when the treatments are continuous.

Second, we use our new technique to derive estimators with simple closed forms that combine kernel ridge regressions, extending the regression product (Baron and Kenny, 1986) and recursive regression (Bang and Robins, 2005) insights to machine learning. We use sequential embeddings to propose uniformly consistent machine learning estimators of time-varying dose response curves without restrictive linearity, Markov, or no-effect-modification assumptions, which to our knowledge is new. As extensions, we propose what may be the first unrestricted incremental response curves and counterfactual distributions for time-varying treatments, relaxing the restrictions of the structural nested distribution model (Robins, 1992). In Section 9 of Singh, Xu and Gretton (2025), for discrete treatments, we use sequential embeddings to propose simpler nuisance parameter estimators for known inferential procedures. In particular, we avoid multiple levels of sample splitting and iterative fitting.

Third, we prove that our simple estimators based on sequential embedding achieve nonasymptotic uniform rates for causal functions. Specifically, for the continuous treatment case, we prove uniform consistency with finite sample rates that combine minimax rates for kernel ridge regression (Caponnetto and De Vito, 2007, Fischer and Steinwart, 2020, Li et al., 2022). The rates do not directly depend on the data dimension, but rather smoothness and the effective dimension, generalizing Sobolev rates. We extend these results to incremental response curves and counterfactual outcome distributions. We relate our results to semiparametric analysis in Section 9 of Singh, Xu and Gretton (2025). In particular, for the discrete treatment case, we verify $n^{-1/2}$ Gaussian approximation and semiparametric efficiency, articulating a double spectral robustness whereby some kernels may have higher effective dimensions as long as others have sufficiently low effective dimensions.

We illustrate the practicality of our approach by conducting comparative simulations and estimating the mediated and time-varying response curves of the Jobs Corps. In nonlinear simulations over short horizons, the algorithms reliably outperform some state of the art alternatives. Under standard identifying assumptions, our direct and indirect dose response curve estimates suggest that job training reduces arrests via social mechanisms besides employment. By allowing for continuous treatments and treatment-confounder feedback, our time-varying dose response curve estimates suggest that relatively few class hours in the first and second year confer most of the benefit to counterfactual employment. Of independent interest, we clean and share a version of the Job Corps data that may serve as a benchmark for new approaches to time-varying estimation.

2. Related work

In seminal works, Robins and Greenland (1992), Pearl (2001), Imai, Keele and Yamamoto (2010) and Robins (1986) rigorously develop identification theory for mediation analysis and time-varying treatment effects, respectively. A rich class of mediated and time-varying causal functions are estimable in principle if the analyst has a sufficiently rich set of covariates over time $X_{1:T}$. Treatments may be continuous; relationships among the outcome, treatments, and covariates may be nonlinear; and dependences may include treatment-confounder feedback and effect modification by time-varying confounders (Gill and Robins, 2001, VanderWeele and Vansteelandt, 2009).

For continuous treatments, nonparametric estimators for mediated response curves of Huber et al. (2020), Ghassami et al. (2021) use density estimation, which can be challenging as dimension increases. Machine learning estimators for time-varying dose response curves of Lewis and Syrgkanis (2021) rely on restrictive linearity, Markov, and no-effect-modification assumptions, which imply additive effects of time-varying treatments.

For binary treatments, a rich literature provides abstract conditions for $n^{-1/2}$ semiparametric estimation, to which we relate our results in Section 9 of Singh, Xu and Gretton (2025); see e.g. Scharfstein, Rotnitzky and Robins (1999), van der Laan and Rubin (2006), Zheng and van der Laan (2011), van der Laan and Gruber (2012), Tchetgen Tchetgen and Shpitser (2012), Petersen et al. (2014), Molina et al. (2017), Luedtke et al. (2017), Rotnitzky, Robins and Babino (2017), Chernozhukov et al. (2018), Farbmacher et al. (2022), Bodory, Huber and Lafférs (2022), Singh (2021a) and references therein. Still, estimators that preserve the full generality of identification for binary treatments are not widely used in empirical research (Vansteelandt and Joffe, 2014), perhaps due to the complexity of nuisance parameter estimation.

Unlike previous work that incorporates the RKHS into causal inference, we provide a framework for mediated and time-varying estimands. Previous work incorporates the RKHS into time-fixed causal inference. Nie and Wager (2021), Foster and Syrgkanis (2023), Kennedy (2023) propose methods based on orthogonal loss minimization for heterogeneous treatment effects, and Wong and Chan (2018), Zhao (2019), Kallus (2020), Hirshberg, Maleki and Zubizarreta (2019), Singh (2021b) propose methods based on balancing weights for average treatment effects. Muandet et al. (2021) propose counterfactual distributions for a binary treatment, while Singh, Xu and Gretton (2024) propose dose responses and counterfactual distributions for a continuous treatment. Whereas previous work studies the time-fixed setting, we study sequential settings and prove strong results despite the additional challenges of treatment-confounder feedback and effect modification by time-varying confounders.

This paper subsumes our earlier draft Singh, Xu and Gretton (2021), which subsumed Singh, Xu and Gretton (2020, Sections B and C).

3. RKHS assumptions

We summarize RKHS notation, interpretation, and assumptions that we use in this paper. Let $k : \mathcal{W} \times \mathcal{W} \to \mathbb{R}$ be a function that is continuous, symmetric, and positive definite. We call k the kernel, and we call $\phi : w \mapsto k(w, \cdot)$ the feature map. The kernel is the inner product of features $k(w, w') = \langle \phi(w), \phi(w') \rangle_{\mathcal{H}}$, and we formally define the inner product below. The RKHS is the closure of the span of the features $\{\phi(w)\}_{w \in \mathcal{W}}$. As such, the features are interpretable as the dictionary of basis functions for the RKHS: for $f \in \mathcal{H}$, we have that $f(w) = \langle f, \phi(w) \rangle_{\mathcal{H}}$.

Kernel ridge regression uses the RKHS $\mathcal{H}$ as the hypothesis space in an infinite dimensional optimization problem with a ridge penalty, and it has a well known closed form solution:

$$\hat{f} := \operatorname*{argmin}_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n} \{Y_i - f(W_i)\}^2 + \lambda \|f\|_{\mathcal{H}}^2; \quad \hat{f}(w) = Y^\top (K_{WW} + n\lambda I)^{-1} K_{Ww}, \tag{1}$$

where $K_{WW} \in \mathbb{R}^{n \times n}$ is the kernel matrix with (i, j)th entry $k(W_i, W_j)$ and $K_{Ww} \in \mathbb{R}^n$ is the kernel vector with ith entry $k(W_i, w)$. To tune the ridge penalty hyperparameter λ, both generalized cross validation and leave one out cross validation have closed form solutions, and the former is asymptotically optimal (Li, 1986). To analyze the bias and variance of kernel ridge regression, the statistical learning literature places assumptions on the smoothness of f_0 and the effective dimension of $\mathcal{H}$. Both assumptions describe the kernel's spectrum, which we now define.

Denote by $\mathbb{L}_\nu^2(\mathcal{W})$ the space of square integrable functions with respect to the measure ν. Consider the convolution operator $L : \mathbb{L}_\nu^2(\mathcal{W}) \to \mathbb{L}_\nu^2(\mathcal{W})$, $f \mapsto \int k(\cdot, w) f(w) \mathrm{d}\nu(w)$. By the spectral theorem, we define its spectrum, where (η_j) are weakly decreasing eigenvalues and (φ_j) are orthonormal eigenfunctions that form a basis of $\mathbb{L}_\nu^2(\mathcal{W})$. As such, $Lf = \sum_{j=1}^\infty \eta_j \langle \varphi_j, f \rangle_{\mathbb{L}_\nu^2(\mathcal{W})} \cdot \varphi_j$.

Remark 3.1 (RKHS versus $\mathbb{L}^2$ inner product). To interpret how the RKHS $\mathcal{H}$ compares to $\mathbb{L}_\nu^2(\mathcal{W})$, we express both function spaces in terms of the orthonormal basis (φ_j). In other words, we present the spectral view of the RKHS. For any $f, g \in \mathbb{L}_\nu^2(\mathcal{W})$, write $f = \sum_{j=1}^\infty f_j \varphi_j$ and $g = \sum_{j=1}^\infty g_j \varphi_j$. Then

$$\mathbb{L}_\nu^2(\mathcal{W}) = \left(f = \sum_{j=1}^\infty f_j \varphi_j : \sum_{j=1}^\infty f_j^2 < \infty \right), \quad \langle f, g \rangle_{\mathbb{L}_\nu^2(\mathcal{W})} = \sum_{j=1}^\infty f_j g_j;$$

$$\mathcal{H} = \left(f = \sum_{j=1}^\infty f_j \varphi_j : \sum_{j=1}^\infty \frac{f_j^2}{\eta_j} < \infty \right), \quad \langle f, g \rangle_{\mathcal{H}} = \sum_{j=1}^\infty \frac{f_j g_j}{\eta_j}.$$

The space $\mathcal{H}$ is the subset of $\mathbb{L}_\nu^2(\mathcal{W})$ for which higher order terms in (φ_j) have a smaller contribution, subject to ν satisfying the conditions of Mercer's theorem (Steinwart and Scovel, 2012). Under those conditions, $k(w, w') = \sum_{j=1}^\infty \eta_j \varphi_j(w) \varphi_j(w')$; (η_j) and (φ_j) describe the kernel's spectrum.

An analyst can place smoothness and effective dimension assumptions on the spectral properties of a statistical target f_0 estimated in the RKHS $\mathcal{H}$ (Caponnetto and De Vito, 2007, Fischer and Steinwart, 2020). These assumptions are formalized by parameters (b, c):

$$f_0 \in \mathcal{H}^c := \left(f = \sum_{j=1}^{\infty} f_j \varphi_j : \sum_{j=1}^{\infty} \frac{f_j^2}{\eta_j^c} < \infty \right) \subset \mathcal{H}, \quad c \in (1,2]; \quad \eta_j \leq C j^{-b}, \quad b \geq 1. \tag{2}$$

The value c quantifies how well the leading terms in (φ_j) approximate f_0; a larger value of c corresponds to a smoother target f_0. A larger value of b corresponds to a faster rate of spectral decay and therefore a lower effective dimension. Both (b, c) are joint assumptions on the kernel and data distribution. Correct specification implies $c \geq 1$ and a bounded kernel implies $b \geq 1$ (Fischer and Steinwart, 2020, Lemma 10). Minimax optimal rates for regression are governed by (b, c) (Caponnetto and De Vito, 2007, Fischer and Steinwart, 2020), with faster rates corresponding to higher values. In our analysis of causal estimands, we obtain nonparametric rates and semiparametric rate conditions that combine regression rates in terms of (b, c).

Spectral assumptions are easy to interpret in Sobolev spaces (Fischer and Steinwart, 2020). Let $\mathcal{W} \subset \mathbb{R}^p$. Denote by $\mathbb{H}_2^s$ the Sobolev space with $s > p/2$ square integrable derivatives, which can be generated by the Matèrn kernel. Suppose $\mathcal{H} = \mathbb{H}_2^s$ is chosen as the RKHS for estimation. If $f_0 \in \mathbb{H}_2^{s_0}$, then $c = s_0/s$; c quantifies the additional smoothness of f_0 relative to $\mathcal{H}$. In this Sobolev space, $b = 2s/p > 1$. The effective dimension is increasing in the original dimension p and decreasing in the degree of smoothness s. The minimax optimal regression rates are

$$n^{-\frac{1}{2}\frac{c}{c+1/b}} = n^{-\frac{s_0}{2s_0+p}} \text{ in } \mathbb{L}^2 \text{ norm}, \quad n^{-\frac{1}{2}\frac{c-1}{c+1/b}} = n^{-\frac{s_0-s}{2s_0+p}} \text{ in Sobolev norm}, \tag{3}$$

and both are achieved by kernel ridge regression with $\lambda = n^{-1/(c+1/b)} = n^{-2s/(2s_0+p)}$.

We place five types of assumptions in this paper, generalizing the standard RKHS learning theory assumptions from kernel ridge regression to mediated and time-varying causal inference: identification, RKHS regularity, original space regularity, smoothness, and effective dimension. We formally instantiate these assumptions for mediated responses in Section 4, time-varying responses in Section 5, and counterfactual distributions in Section 10 of Singh, Xu and Gretton (2025). We uncover a double spectral robustness in related semiparametric inferential theory in Section 9 of Singh, Xu and Gretton (2025): some kernels may have higher effective dimensions, as long as other kernels have lower effective dimensions.

4. Mediated response curves

4.1. Pearl's mediation formula

Mediation analysis decomposes the total effect of a treatment D on an outcome Y into the direct effect versus the indirect effect mediated via the mechanism M. The problem is sequential since D causes M and Y, then M also causes Y. We denote the counterfactual mediator $M^{(d)}$ given a hypothetical intervention on the treatment $D = d$. We denote the counterfactual outcome $Y^{(d,m)}$ given a hypothetical intervention on the treatment $D = d$ and the mediator $M = m$.

Definition 4.1 (Pure mediated response curves (Robins and Greenland, 1992)). Suppose the treatment D is continuous.

1. The total response $\theta_0^{TE}(d, d') = E[Y^{\{d',M^{(d')}\}} - Y^{\{d,M^{(d)}\}}]$ is the total effect of a new treatment value d' compared to an old value d.
2. The indirect response $\theta_0^{IE}(d, d') = E[Y^{\{d',M^{(d')}\}} - Y^{\{d',M^{(d)}\}}]$ is the component of the total response mediated by M.
3. The direct response $\theta_0^{DE}(d, d') = E[Y^{\{d',M^{(d)}\}} - Y^{\{d,M^{(d)}\}}]$ is the component of the total response that is not mediated by M.
4. The mediated response $\theta_0^{ME}(d, d') = E[Y^{\{d',M^{(d)}\}}]$ is the counterfactual mean outcome in the thought experiment that the treatment is set at a new value $D = d'$ but the mediator M follows the distribution it would have followed if the treatment were set at its old value $D = d$.

Likewise we define incremental response curves, e.g. $\theta_0^{ME,\nabla}(d, d') = E[\nabla_{d'} Y^{\{d',M^{(d)}\}}]$.

Remark 4.1 (Interventional mediated response curves). Definition 4.1 considers cross world counterfactuals that involve different treatment values for the potential outcomes and potential mediators (Robins and Greenland, 1992, Pearl, 2001, Imai, Keele and Yamamoto, 2010). An alternative paradigm instead considers interventional counterfactuals; see e.g. Robins and Richardson (2011), Richardson and Robins (2013), Robins, Richardson and Shpitser (2022) and references therein. The alternative view avoids defining potential mediators and instead supposes that the treatment can be decomposed into multiple separable components. Though these two paradigms define different mediated response curves, when they are identified, their identifying formulae coincide as Pearl's mediation formula, which we quote in Lemma 4.1 as the starting point for our analysis. As such, our estimation results apply to both important paradigms for mediation analysis.

$\theta_0^{TE}(d, d')$ generalizes average total effects. The average total effect of a binary treatment is $E[Y^{\{1,M^{(1)}\}} - Y^{\{0,M^{(0)}\}}]$. For a continuous treatment, the function $d \mapsto E[Y^{\{d,M^{(d)}\}}]$ may be infinite dimensional, which makes this problem fully nonparametric.

An analyst may wish to measure how much of the total effect is indirect: how much of the total effect would be achieved by simply intervening on the distribution of the mediator M? For example, in Section 6, we investigate the extent to which employment mediates the effect of job training on arrests. With a binary treatment, the indirect effect is $E[Y^{\{1,M^{(1)}\}} - Y^{\{1,M^{(0)}\}}]$. In the former term, the mediator follows the counterfactual distribution under the intervention $D = 1$, and in the latter, it follows the counterfactual distribution under the intervention $D = 0$.

The remaining component of the total effect is the direct effect: if the mediator were held at the original distribution corresponding to $D = d$, what would be the impact of the treatment $D = d'$? For example, in Section 6, we investigate the effect of job training on arrests holding employment at the original distribution. With a binary treatment, the direct effect is $E[Y^{\{1,M^{(0)}\}} - Y^{\{0,M^{(0)}\}}]$.

The final target parameter is $\theta_0^{ME}(d, d')$. It is useful because $\theta_0^{TE}(d, d')$, $\theta_0^{IE}(d, d')$, and $\theta_0^{DE}(d, d')$ can be expressed in terms of $\theta_0^{ME}(d, d')$. With a binary treatment, this quantity is a matrix in $\mathbb{R}^{2\times 2}$. With a continuous treatment, it is a surface over $\mathcal{D} \times \mathcal{D}$.

Proposition 4.1 (Convenient expressions). *Mediated response curves can be expressed in terms of $\theta_0^{ME}(d, d')$:*

1. $\theta_0^{TE}(d, d') = \theta_0^{DE}(d, d') + \theta_0^{IE}(d, d')$;
2. $\theta_0^{IE}(d, d') = \theta_0^{ME}(d', d') - \theta_0^{ME}(d, d')$;
3. $\theta_0^{DE}(d, d') = \theta_0^{ME}(d, d') - \theta_0^{ME}(d, d)$.

In seminal works, Robins and Greenland (1992), Pearl (2001), Imai, Keele and Yamamoto (2010) state sufficient conditions under which the mediated response curves can be measured from the outcome Y, treatment D, mediator M, and covariates X, which we call selection on observables for mediation. We modestly extend the classic identification result from $\theta_0^{ME}(d,d')$ to its incremental version $\theta_0^{ME,\nabla}(d,d')$. Let $\gamma_0(d,m,x) = E(Y|D=d, M=m, X=x)$.

Lemma 4.1 (Pearl's mediation formula). *Under selection on observables for mediation,*

1. $\theta_0^{ME}(d,d') = \int \gamma_0(d',m,x)\mathrm{d}P(m|d,x)\mathrm{d}P(x)$ *(Robins and Greenland, 1992, Pearl, 2001, Imai, Keele and Yamamoto, 2010) and*
2. $\theta_0^{ME,\nabla}(d,d') = \int \nabla_{d'}\gamma_0(d',m,x)\mathrm{d}P(m|d,x)\mathrm{d}P(x)$.

See Section 12 of Singh, Xu and Gretton (2025) for the identifying assumptions and proof of our extension. Proposition 4.1 identifies the other quantities in Definition 4.1. For subsequent analysis, it helps to define $\omega_0(d,d';x) = \int \gamma_0(d',m,x)\mathrm{d}P(m|d,x)$, so that $\theta_0^{ME}(d,d') = \int \omega_0(d,d';x)\mathrm{d}P(x)$.

Remark 4.2 (Pearl's mediation formula is unbounded over $\mathbb{L}^2$ when the treatment is continuous). Define the functional $F : \mathbb{L}^2 \to \mathbb{R}$, $\gamma \mapsto \int \gamma(d',m,x)\mathrm{d}P(m|d,x)\mathrm{d}P(x)$. When the treatment is continuous, F is generally unbounded, i.e. there does not exists some $C < \infty$ such that $F(\gamma) \le C\|\gamma\|_{\mathbb{L}^2}$ for all $\gamma \in \mathbb{L}^2$ (van der Vaart, 1991, Newey, 1994). This technical challenge is well documented in the causal inference literature; see van der Laan, Bibaut and Luedtke (2018) for references.

Remark 4.3 (Mediational g-formula). A rich literature defines mediated response curves in the context of time-varying treatments; see e.g. VanderWeele and Tchetgen Tchetgen (2017), Malinsky, Shpitser and Richardson (2019) and references therein. The mediational g-formula synthesizes Pearl's mediation formula in Lemma 4.1 and Robins' g-formula in Lemma 5.1. Our framework generalizes to these more complex causal functions, using the techniques in Section 11 of Singh, Xu and Gretton (2025).

4.2. Sequential kernel embedding

Lemma 4.1 makes precise how each mediated response curve is identified as a sequential integral of the form $\int \gamma_0(d',m,x)\mathrm{d}Q$ for the distribution $Q = P(m|d,x)P(x)$. Since x appears in $\gamma_0(d',m,x)$, $P(m|d,x)$, and $P(x)$, the sequential integral is coupled and therefore challenging to estimate. We prove that, with the appropriate RKHS construction, the components $\gamma_0(d',m,x)$, $P(m|d,x)$, and $P(x)$ can be decoupled. Moreover, the sequential distribution Q can be encoded by a sequential kernel embedding, which is our key innovation. We use these techniques to reduce sequential causal inference into the combination of kernel ridge regressions, which then allows us to propose simple estimators with closed form solutions.

To begin, we construct the appropriate RKHS for γ_0. In our construction, we define RKHSs for the treatment D, mediator M, and covariates X, then assume that the regression is an element of a certain composite space. To lighten notation, we will suppress subscripts when arguments are provided. We assume the regression γ_0 is an element of the RKHS $\mathcal{H}$ with the kernel $k(d,m,x;d',m',x') = k_{\mathcal{D}}(d,d')k_{\mathcal{M}}(m,m')k_{\mathcal{X}}(x,x')$. Formally, this choice of kernel corresponds to the tensor product: $\gamma_0 \in \mathcal{H} = \mathcal{H}_{\mathcal{D}} \otimes \mathcal{H}_{\mathcal{M}} \otimes \mathcal{H}_{\mathcal{X}}$, with the tensor product dictionary of basis functions $\phi(d)\otimes\phi(m)\otimes\phi(x)$. As such, $\gamma_0(d,m,x) = \langle \gamma_0, \phi(d)\otimes\phi(m)\otimes\phi(x)\rangle_{\mathcal{H}}$ and $\|\phi(d)\otimes\phi(m)\otimes\phi(x)\|_{\mathcal{H}} = \|\phi(d)\|_{\mathcal{H}_{\mathcal{D}}}\|\phi(m)\|_{\mathcal{H}_{\mathcal{M}}}\|\phi(x)\|_{\mathcal{H}_{\mathcal{X}}}$. We place regularity conditions on this RKHS construction in order to prove our decoupling result. Anticipating Section 10 of Singh, Xu and Gretton (2025), we include conditions for an outcome RKHS in parentheses.

Assumption 4.1 (RKHS regularity conditions). Assume

1. $k_{\mathcal{D}}$, $k_{\mathcal{M}}$, $k_{\mathcal{X}}$ (and $k_{\mathcal{Y}}$) are continuous and bounded, i.e. $\sup_{d\in\mathcal{D}}\|\phi(d)\|_{\mathcal{H}_{\mathcal{D}}} \le \kappa_d$, $\sup_{m\in\mathcal{M}}\|\phi(m)\|_{\mathcal{H}_{\mathcal{M}}} \le \kappa_m$, $\sup_{x\in\mathcal{X}}\|\phi(x)\|_{\mathcal{H}_{\mathcal{X}}} \le \kappa_x$ {and $\sup_{y\in\mathcal{Y}}\|\phi(y)\|_{\mathcal{H}_{\mathcal{Y}}} \le \kappa_y$};
2. $\phi(d)$, $\phi(m)$, $\phi(x)$ {and $\phi(y)$} are measurable;
3. $k_{\mathcal{M}}$, $k_{\mathcal{X}}$ (and $k_{\mathcal{Y}}$) are characteristic (Sriperumbudur, Fukumizu and Lanckriet, 2010).

For incremental responses, further assume $\mathcal{D}\subset\mathbb{R}$ is an open set and $\nabla_d\nabla_{d'}k_{\mathcal{D}}(d,d')$ exists and is continuous, hence $\sup_{d\in\mathcal{D}}\|\nabla_d\phi(d)\|_{\mathcal{H}} \le \kappa'_d$.

Commonly used kernels are continuous and bounded. Measurability is a similarly weak condition. The characteristic property means that different distributions will have different embeddings in the RKHS. For example, the indicator kernel is characteristic over a discrete domain, while the exponentiated quadratic kernel is characteristic over a continuous domain.

Theorem 4.1 (Decoupling via sequential kernel embeddings). *Suppose the conditions of Lemma 4.1 hold. Further suppose Assumption 4.1 holds and $\gamma_0\in\mathcal{H}$. Then*

1. $\omega_0(d,d';x) = \langle\gamma_0, \phi(d')\otimes\mu_m(d,x)\otimes\phi(x)\rangle_{\mathcal{H}}$,
2. $\theta_0^{ME}(d,d') = \langle\gamma_0, \phi(d')\otimes\mu_{m,x}(d)\rangle_{\mathcal{H}}$, *and*
3. $\theta_0^{ME,\nabla}(d,d') = \langle\gamma_0, \nabla_{d'}\phi(d')\otimes\mu_{m,x}(d)\rangle_{\mathcal{H}}$,

where $\mu_m(d,x) = \int\phi(m)\mathrm{d}P(m|d,x)$ *and* $\mu_{m,x}(d) = \int\{\mu_m(d,x)\otimes\phi(x)\}\mathrm{d}P(x)$.

Proof sketch. Consider $\omega_0(d,d';x) = \int\gamma_0(d',m,x)\mathrm{d}P(m|d,x)$. We show

$$\omega_0(d,d';x) = \int\langle\gamma_0, \phi(d')\otimes\phi(m)\otimes\phi(x)\rangle_{\mathcal{H}}\mathrm{d}P(m|d,x) = \langle\gamma_0, \phi(d')\otimes\mu_m(d,x)\otimes\phi(x)\rangle_{\mathcal{H}}.$$

Sequentially repeating this technique for $\theta_0^{ME}(d,d') = \int\omega_0(d,d';x)\mathrm{d}P(x)$,

$$\theta_0^{ME}(d,d') = \int\langle\gamma_0, \phi(d')\otimes\mu_m(d,x)\otimes\phi(x)\rangle_{\mathcal{H}}\mathrm{d}P(x) = \langle\gamma_0, \phi(d')\otimes\mu_{m,x}(d)\rangle_{\mathcal{H}}. \qquad \square$$

See Section 13 of Singh, Xu and Gretton (2025) for the full proof. The quantity $\mu_m(d,x) = \int\phi(m)\mathrm{d}P(m|d,x)$ embeds the conditional distribution $P(m|d,x)$ as an element of the RKHS $\mathcal{H}_{\mathcal{M}}$, which is a popular technique in the RKHS literature. It satisfies, for $f\in\mathcal{H}_{\mathcal{M}}$, $\langle f, \mu_m(d,x)\rangle_{\mathcal{H}_{\mathcal{M}}} = \int f(m)\mathrm{d}P(m|d,x)$.

Remark 4.4 (Key innovation). The quantity $\mu_{m,x}(d)$ is a sequential kernel embedding that encodes the counterfactual distribution of the mediator M and covariates X when the counterfactual treatment value is $D = d$. It is our key innovation. It has the property that, for $f\in\mathcal{H}_{\mathcal{M}}\otimes\mathcal{H}_{\mathcal{X}}$, $\langle f, \mu_{m,x}(d)\rangle_{\mathcal{H}_{\mathcal{M}}\otimes\mathcal{H}_{\mathcal{X}}} = \int f(m,x)\mathrm{d}P(m|d,x)\mathrm{d}P(x)$, implementing Pearl's mediation formula.

Remark 4.5 (Pearl's mediation formula is bounded over $\mathcal{H}$ when the treatment is continuous). Define the functional $F:\mathcal{H}\to\mathbb{R}, \gamma\mapsto\int\gamma(d',m,x)\mathrm{d}P(m|d,x)\mathrm{d}P(x)$. Under Assumption 4.1, we show that F is bounded even when the treatment is continuous, i.e. there exists some $C<\infty$ such that $F(\gamma)\le C\|\gamma\|_{\mathcal{H}}$ for all $\gamma\in\mathcal{H}$. This observation generalizes the observation of Singh, Xu and Gretton (2024) for time-fixed dose response curves. It follows immediately from the definition of the RKHS as the space of functions for which the functional $\gamma\mapsto\gamma(d',m,x)$ is bounded (Berlinet and Thomas-Agnan, 2004).

Boundedness of the functional is what guarantees the existence of the sequential kernel embeddings in Theorem 4.1.

Theorem 4.1 decouples $\gamma_0(d',m,x)$, $P(m|d,x)$, and $P(x)$, providing a blueprint for estimation that avoids density estimation and sequential integration. Our estimator will be $\hat{\theta}^{ME}(d,d') = \langle \hat{\gamma}, \phi(d') \otimes \hat{\mu}_{m,x}(d)\rangle_{\mathcal{H}}$ where $\hat{\mu}_{m,x}(d) = n^{-1}\sum_{i=1}^{n}\{\hat{\mu}_m(d,X_i)\otimes\phi(X_i)\}$. The estimator $\hat{\gamma}$ is a standard kernel ridge regression. The estimator $\hat{\mu}_m(d,x)$ is an appropriately generalized kernel ridge regression. We combine them by averaging and taking the product.

Algorithm 4.1 (Nonparametric estimation of mediated response curves). Denote the kernel matrices by K_{DD}, K_{MM}, $K_{XX} \in \mathbb{R}^{n\times n}$. Let $\odot$ be the elementwise product. Mediated response curves have closed form solutions:

1. $\hat{\omega}(d,d';x) = Y^\top(K_{DD}\odot K_{MM}\odot K_{XX} + n\lambda I)^{-1}[K_{Dd'}\odot\{K_{MM}(K_{DD}\odot K_{XX}+n\lambda_1 I)^{-1}(K_{Dd}\odot K_{Xx})\}\odot K_{Xx}]$ and
2. $\hat{\theta}^{ME}(d,d') = n^{-1}\sum_{i=1}^{n}\hat{\omega}(d,d';X_i)$,

where (λ,λ_1) are ridge regression penalty parameters. For mediated incremental response curve estimators, we replace $K_{Dd'}$ with $\nabla_{d'}K_{Dd'}$ where $(\nabla_{d'}K_{Dd'})_i = \nabla_{d'}k(D_i,d')$.

Derivation sketch. Consider $\omega_0(d,d';x)$. Analogously to (1), the kernel ridge regression estimators of the regression γ_0 and the conditional kernel embedding $\mu_m(d,x)$ are

$$\hat{\gamma} = \operatorname*{argmin}_{\gamma\in\mathcal{H}} \frac{1}{n}\sum_{i=1}^{n}[Y_i - \langle\gamma, \phi(D_i)\otimes\phi(M_i)\otimes\phi(X_i)\rangle_{\mathcal{H}}]^2 + \lambda\|\gamma\|^2_{\mathcal{H}},$$

$$\hat{E} = \operatorname*{argmin}_{E\in\mathcal{L}_2(\mathcal{H}_{\mathcal{M}},\mathcal{H}_{\mathcal{D}}\otimes\mathcal{H}_{\mathcal{X}})} \frac{1}{n}\sum_{i=1}^{n}[\phi(M_i) - E^*\{\phi(D_i)\otimes\phi(X_i)\}]^2 + \lambda_1\|E\|^2_{\mathcal{L}_2(\mathcal{H}_{\mathcal{M}},\mathcal{H}_{\mathcal{D}}\otimes\mathcal{H}_{\mathcal{X}})},$$

where $\hat{\mu}_m(d,x) = \hat{E}^*\{\phi(d)\otimes\phi(x)\}$ and E^* is the adjoint of E. The closed forms are

$$\hat{\gamma}(d',\cdot,x) = Y^\top(K_{DD}\odot K_{MM}\odot K_{XX}+n\lambda I)^{-1}(K_{Dd'}\odot K_{M\cdot}\odot K_{Xx}),$$

$$[\hat{\mu}_m(d,x)](\cdot) = K_{\cdot M}(K_{DD}\odot K_{XX}+n\lambda_1 I)^{-1}(K_{Dd}\odot K_{Xx}).$$

To arrive at the main result, match the empty arguments $(\cdot)$ of the kernel ridge regressions. □

We derive this algorithm in Section 13 of Singh, Xu and Gretton (2025). We give theoretical values for (λ,λ_1) that optimally balance bias and variance in Theorem 4.2 below. Section 16 gives practical tuning procedures with closed form solutions to empirically balance bias and variance, one of which is asymptotically optimal. We formally define the operator space $\mathcal{L}_2(\mathcal{H}_{\mathcal{M}},\mathcal{H}_{\mathcal{D}}\otimes\mathcal{H}_{\mathcal{X}})$ below.

4.3. Uniform consistency with finite sample rates

Towards a guarantee of uniform consistency, we place regularity conditions on the original spaces. Anticipating Section 10 of Singh, Xu and Gretton (2025), we include conditions for the outcome in parentheses.

Assumption 4.2 (Original space regularity conditions). Assume (i) $\mathcal{D}$, $\mathcal{M}$, $\mathcal{X}$ (and $\mathcal{Y}$) are Polish spaces; (ii) $\mathcal{Y}\subset\mathbb{R}$ and $|Y|\le C$ almost surely.

A Polish space is a separable and completely metrizable topological space. Random variables supported in a Polish space may be discrete or continuous and even texts, graphs, or images. Boundedness of the outcome Y can be relaxed. Next, we place assumptions on the smoothness of the regression γ_0 and the effective dimension of $\mathcal{H}$ the sense of (2).

Assumption 4.3 (Smoothness and effective dimension of the regression). Assume $\gamma_0 \in \mathcal{H}^c$ with $c \in (1,2]$, and $\eta_j(\mathcal{H}) \le Cj^{-b}$ with $b \ge 1$.

See Section 14 of Singh, Xu and Gretton (2025) for alternative ways of writing and interpreting Assumption 4.3 in the tensor product space $\mathcal{H}$. We place similar conditions on the conditional kernel embedding $\mu_m(d,x)$, which is a generalized regression. We articulate this assumption abstractly for the conditional kernel embedding $\mu_a(b) = \int \phi(a)\mathrm{d}P(a|b)$ where $a \in \mathcal{A}_\ell$ and $b \in \mathcal{B}_\ell$. As such, all one has to do is specify $\mathcal{A}_\ell$ and $\mathcal{B}_\ell$ to specialize the assumption. For $\mu_m(d,x)$, $\mathcal{A}_1 = \mathcal{M}$ and $\mathcal{B}_1 = \mathcal{D} \times \mathcal{X}$. We parametrize the effective dimension and smoothness of $\mu_a(b)$ by (b_ℓ, c_ℓ).

Formally, define the conditional expectation operator $E_\ell : \mathcal{H}_{\mathcal{A}_\ell} \to \mathcal{H}_{\mathcal{B}_\ell}$, $f(\cdot) \mapsto E\{f(A_\ell)|B_\ell = \cdot\}$. By construction, E_ℓ encodes the same information as $\mu_a(b)$ since

$$\mu_a(b) = \int \phi(a)\mathrm{d}P(a|b) = [E_\ell\{\phi(\cdot)\}](b) = [E_\ell^*\{\phi(b)\}](\cdot), \quad a \in \mathcal{A}_\ell, \quad b \in \mathcal{B}_\ell,$$

where E_ℓ^* is the adjoint of E_ℓ. We denote the space of Hilbert–Schmidt operators between $\mathcal{H}_{\mathcal{A}_\ell}$ and $\mathcal{H}_{\mathcal{B}_\ell}$ by $\mathcal{L}_2(\mathcal{H}_{\mathcal{A}_\ell}, \mathcal{H}_{\mathcal{B}_\ell})$. Here, $\mathcal{L}_2(\mathcal{H}_{\mathcal{A}_\ell}, \mathcal{H}_{\mathcal{B}_\ell})$ is an RKHS in its own right, for which we place smoothness and effective dimension assumptions in the sense of (2). See Grünewälder, Gretton and Shawe-Taylor (2013, Appendix B) and Singh, Sahani and Gretton (2019, Appendix A.3) for details.

Assumption 4.4 (Smoothness and effective dimension of a conditional kernel embedding). Assume $E_\ell \in \mathcal{L}_2(\mathcal{H}_{\mathcal{A}_\ell}, \mathcal{H}_{\mathcal{B}_\ell}^{c_\ell})$ with $c_\ell \in (1,2]$, and $\eta_j(\mathcal{H}_{\mathcal{B}_\ell}) \le Cj^{-b_\ell}$ with $b_\ell \ge 1$.

Just as we place approximation assumptions for γ_0 in terms of $\mathcal{H}$, which provides the features onto which we project Y, we place approximation assumptions for E_ℓ in terms of $\mathcal{H}_{\mathcal{B}_\ell}$, which provides the features $\phi(B_\ell)$ onto which we project $\phi(A_\ell)$. Under these conditions, we arrive at our first main theoretical guarantee.

Theorem 4.2 (Uniform consistency of mediated response curves). *Suppose the conditions of Theorem 4.1 hold, as well as Assumptions 4.2, 4.3, and 4.4 with $\mathcal{A}_1 = \mathcal{M}$ and $\mathcal{B}_1 = \mathcal{D} \times \mathcal{X}$. Set $(\lambda, \lambda_1) = \{n^{-1/(c+1/b)}, n^{-1/(c_1+1/b_1)}\}$, which is rate optimal regularization. Then*

1. $\|\hat{\theta}^{ME} - \theta_0^{ME}\|_\infty = O_p\left[n^{-(c-1)/\{2(c+1/b)\}} + n^{-(c_1-1)/\{2(c_1+1/b_1)\}}\right]$ *and*
2. $\|\hat{\theta}^{ME,\nabla} - \theta_0^{ME,\nabla}\|_\infty = O_p\left[n^{-(c-1)/\{2(c+1/b)\}} + n^{-(c_1-1)/\{2(c_1+1/b_1)\}}\right]$.

See Section 14 of Singh, Xu and Gretton (2025) for the proof. By Proposition 4.1, the quantities in Definition 4.1 are uniformly consistent with the same rate, which combines optimal rates for standard (Fischer and Steinwart, 2020, Theorem 2) and generalized (Li et al., 2022, Theorem 3) kernel ridge regression in RKHS norm. Section 14 gives the exact finite sample rate and the explicit specialization of Assumption 4.4. The rate is at best $n^{-1/4}$ when $(c, c_1) = 2$ and $(b, b_1) \to \infty$, i.e. when (γ_0, μ_m) are very smooth with finite effective dimensions. The rates reflect the challenge of a sup norm guarantee, which is much stronger than an $\mathbb{L}^2$ norm guarantee and is useful for policymakers who may be concerned about each treatment value. See (3) to specialize these rates for Sobolev spaces.

Remark 4.6 (Technical contribution). The technical contribution underlying our theoretical guarantee is an RKHS norm rate for the sequential kernel embedding. In particular, Proposition 14.1 in Section 14 of Singh, Xu and Gretton (2025) derives a nonasymptotic bound on $\sup_{d\in\mathcal{D}} \|\hat{\mu}_{m,x}(d) - \mu_{m,x}(d)\|_{\mathcal{H}_{\mathcal{M}}\otimes\mathcal{H}_{\mathcal{X}}}$, as a stepping stone to Theorem 4.2. This intermediate result is not contained in Singh, Xu and Gretton (2024) nor in other previous works on kernel methods for causal inference.

Remark 4.7 (Rate improvements). For the time-fixed dose response curve, Kennedy et al. (2017) prove pointwise rates, assuming smoothness of the dose response as well as finite uniform entropy integrals for the regression estimator and for the density estimator of treatment given covariates. Later works use sample splitting to relax the entropy conditions to rate conditions, e.g. Semenova and Chernozhukov (2021), Colangelo and Lee (2020).

If the regression and density have sufficiently fast rates, then pointwise rate improvements are possible, reflecting the smoothness and lower dimension of the dose response. These are called oracle rates with second order dependence on the regression and density. See e.g. Nie and Wager (2021), Foster and Syrgkanis (2023), Kennedy (2023) for time-fixed heterogeneous treatment effects.

For our RKHS estimator of the mediated response curve, we prove uniform rates, under smoothness assumptions on the regression function and a conditional expectation operator. To achieve sup norm rate improvements for our RKHS estimator, future work may place an additional smoothness assumption on the mediated response curve and rate conditions on the regression, conditional expectation operator, and appropriate conditional densities.

5. Time-varying response curves

5.1. Robins' g-formula

So far we have considered the effect of a single treatment D. Next, we consider the effect of a sequence of time-varying treatments $D_{1:T} = d_{1:T}$ on the counterfactual outcome $Y^{(d_{1:T})}$. If the sequence of treatment values $d_{1:T}$ is observed in the data, this problem may be called on-policy planning; if not, it may be called off-policy planning.

Definition 5.1 (Time-varying dose response curves (Robins, 1986)). Suppose the treatments $D_{1:T}$ are continuous.

1. The time-varying response $\theta_0^{GF}(d_{1:T}) = E\{Y^{(d_{1:T})}\}$ is the counterfactual mean outcome given the interventions $D_{1:T} = d_{1:T}$ for the entire population.
2. With distribution shift, $\theta_0^{DS}(d_{1:T}, \tilde{P}) = E_{\tilde{P}}\{Y^{(d_{1:T})}\}$ is the counterfactual mean outcome given the interventions $D_{1:T} = d_{1:T}$ for an alternative population with the data distribution $\tilde{P}$.

Likewise we define incremental response curves, e.g. $\theta_0^{GF,\nabla}(d_{1:T}) = E\{\nabla_{d_T} Y^{(d_{1:T})}\}$.

Remark 5.1 (Randomized dynamic policies). For clarity, we focus on the deterministic, static counterfactual policy $d_{1:T}$. It is deterministic in that it is nonrandom. It is static in that it does not depend on the observed sequence of covariates $X_{1:T}$. The time-varying treatment effect literature extends to policies that may be randomized and dynamic (Robins, 1986). Our approach extends to randomized and dynamic policies with additional notation; see Remark 5.3.

Whereas much of the semiparametric literature restricts d_t to be discrete, we allow d_t to be continuous and consider a nonparametric approach to time-varying response curves (Gill and Robins, 2001).

For example, in Section 6, we estimate the effect of d_1 class hours in year one, and d_2 class hours in year two, on counterfactual employment, with treatment-confounder feedback. In the spirit of off-policy planning, we consider a distribution shift from P to $\tilde{P}$.

In seminal work, Robins (1986) states sufficient conditions under which time-varying responses can be measured from the outcome Y, treatments $D_{1:T}$, and covariates $X_{1:T}$. We refer to this collection of conditions as sequential selection on observables. We modestly extend the classic identification result by considering incremental responses.

Lemma 5.1 (Robins' g-formula). *Under sequential selection on observables and a distribution shift condition,*

1. $\theta_0^{GF}(d_{1:T}) = \int \gamma_0(d_{1:T}, x_{1:T}) \mathrm{d}P(x_1) \prod_{t=2}^{T} \mathrm{d}P\{x_t | d_{1:(t-1)}, x_{1:(t-1)}\}$ *(Robins, 1986) and*
2. $\theta_0^{DS}(d_{1:T}, \tilde{P}) = \int \gamma_0(d_{1:T}, x_{1:T}) \mathrm{d}\tilde{P}(x_1) \prod_{t=2}^{T} \mathrm{d}\tilde{P}\{x_t | d_{1:(t-1)}, x_{1:(t-1)}\}$ *(Pearl and Bareinboim, 2014).*
3. *For incremental response curves, we replace* $\gamma_0(d_{1:T}, x_{1:T})$ *with* $\nabla_{d_T} \gamma_0(d_{1:T}, x_{1:T})$.

See Section 12 of Singh, Xu and Gretton (2025) for the identifying assumptions and proof of our extension. We consider a fully nonparametric g-formula with possibly continuous treatments that allows for distribution shift. Lemma 5.1 handles auxiliary Markov restrictions as special cases, e.g. if covariates follow a Markov process, then θ_0^{GF} simplifies by setting $P\{x_t | d_{1:(t-1)}, x_{1:(t-1)}\} = P(x_t | d_{t-1}, x_{t-1})$.

Remark 5.2 (Robins' g-formula is unbounded over $\mathbb{L}^2$ when the treatments are continuous). Define the functional $F : \mathbb{L}^2 \to \mathbb{R}$, $\gamma \mapsto \int \gamma(d_{1:T}, x_{1:T}) \mathrm{d}P(x_1) \prod_{t=2}^{T} \mathrm{d}P\{x_t | d_{1:(t-1)}, x_{1:(t-1)}\}$. When the treatments are continuous, F is generally unbounded in the sense of Remark 4.2 (van der Vaart, 1991, Newey, 1994, van der Laan, Bibaut and Luedtke, 2018).

Remark 5.3 (Randomized dynamic policies). To accommodate randomized dynamic policies, the product in Lemma 5.1 will include factors for the conditional distributions of treatments. Specifically, the integral will replace d_t with $g_t : \{d_{1:(t-1)}, x_{1:t}\} \mapsto G\{d_t | d_{1:(t-1)}, x_{1:t}\}$ where G is the distribution induced by the randomized dynamic policy (g_t).

5.2. Sequential kernel embedding

Similar to Lemma 4.1, Lemma 5.1 identifies each time-varying response as a sequential integral of the form $\int \gamma_0(d_{1:T}, x_{1:T}) \mathrm{d}Q$ for the distribution $Q = P(x_1) \prod_{t=2}^{T} P\{x_t | d_{1:(t-1)}, x_{1:(t-1)}\}$ or $Q = \tilde{P}(x_1) \prod_{t=2}^{T} \tilde{P}\{x_t | d_{1:(t-1)}, x_{1:(t-1)}\}$. As before, components of the sequential integral are coupled and therefore challenging to estimate, since x_1 appears in $\gamma_0(d_{1:T}, x_{1:T})$, $P\{x_t | d_{1:(t-1)}, x_{1:(t-1)}\}$, and $P(x_1)$. As in Section 4, we construct an appropriate RKHS to decouple these components, then to encode Q by a sequential kernel embedding. With these techniques, we again reduce sequential causal inference into the combination of kernel ridge regressions. For clarity, we present the algorithm with $T = 2$, and we define $\omega_0(d_1, d_2; x_1) = \int \gamma_0(d_1, d_2, x_1, x_2) \mathrm{d}P(x_2 | d_1, x_1)$ so that $\theta_0^{GF}(d_1, d_2) = \int \omega_0(d_1, d_2; x_1) \mathrm{d}P(x_1)$. We consider $T > 2$ in Section 11 of Singh, Xu and Gretton (2025), which also showcases the role of Markov assumptions.

To construct the RKHS for γ_0, we define RKHSs for each treatment D_t and each covariate X_t. Using identical notation as Section 4, we assume the regression γ_0 is an element of the RKHS $\mathcal{H}$ with the kernel $k(d_1, d_2, x_1, x_2; d_1', d_2', x_1', x_2') = k_{\mathcal{D}}(d_1, d_1') k_{\mathcal{D}}(d_2, d_2') k_{\mathcal{X}}(x_1, x_1') k_{\mathcal{X}}(x_2, x_2')$, i.e. $\gamma_0 \in$

$\mathcal{H} = \mathcal{H}_{\mathcal{D}} \otimes \mathcal{H}_{\mathcal{D}} \otimes \mathcal{H}_{\mathcal{X}} \otimes \mathcal{H}_{\mathcal{X}}$. As such, $\gamma_0(d_1, d_2, x_1, x_2) = \langle \gamma_0, \phi(d_1) \otimes \phi(d_2) \otimes \phi(x_1) \otimes \phi(x_2) \rangle_{\mathcal{H}}$ and $\|\phi(d_1) \otimes \phi(d_2) \otimes \phi(x_1) \otimes \phi(x_2)\|_{\mathcal{H}} = \|\phi(d_1)\|_{\mathcal{H}_{\mathcal{D}}} \|\phi(d_2)\|_{\mathcal{H}_{\mathcal{D}}} \|\phi(x_1)\|_{\mathcal{H}_{\mathcal{X}}} \|\phi(x_2)\|_{\mathcal{H}_{\mathcal{X}}}$. Under regularity conditions on this RKHS construction, we prove an analogous decoupling result.

Theorem 5.1 (Decoupling via sequential kernel embeddings). *Suppose the conditions of Lemma 5.1 hold. Further suppose Assumption 4.1 holds and $\gamma_0 \in \mathcal{H}$. Then*

1. $\omega_0(d_1, d_2; x_1) = \langle \gamma_0, \phi(d_1) \otimes \phi(d_2) \otimes \phi(x_1) \otimes \mu_{x_2}(d_1, x_1) \rangle_{\mathcal{H}}$;
2. $\theta_0^{GF}(d_1, d_2) = \langle \gamma_0, \phi(d_1) \otimes \phi(d_2) \otimes \mu_{x_1,x_2}(d_1) \rangle_{\mathcal{H}}$ *where* $\mu_{x_2}(d_1, x_1) = \int \phi(x_2) \mathrm{d}P(x_2|d_1, x_1)$ *and* $\mu_{x_1,x_2}(d_1) = \int \{\phi(x_1) \otimes \mu_{x_2}(d_1, x_1)\} \mathrm{d}P(x_1)$;
3. $\theta_0^{DS}(d_1, d_2; \tilde{P}) = \langle \gamma_0, \phi(d_1) \otimes \phi(d_2) \otimes \nu_{x_1,x_2}(d_1) \rangle_{\mathcal{H}}$ *where* $\nu_{x_2}(d_1, x_1) = \int \phi(x_2) \mathrm{d}\tilde{P}(x_2|d_1, x_1)$ *and* $\nu_{x_1,x_2}(d_1) = \int \{\phi(x_1) \otimes \nu_{x_2}(d_1, x_1)\} \mathrm{d}\tilde{P}(x_1)$.

For incremental responses, we replace $\phi(d_2)$ with $\nabla_{d_2} \phi(d_2)$.

See Section 13 of Singh, Xu and Gretton (2025) for the proof. In θ_0^{GF}, the conditional kernel embedding of $P(x_2|d_1, x_1)$ is $\mu_{x_2}(d_1, x_1) = \int \phi(x_2) \mathrm{d}P(x_2|d_1, x_1)$, and it satisfies $\langle f, \mu_{x_2}(d_1, x_1) \rangle_{\mathcal{H}_{\mathcal{X}}} = \int f(x_2) \mathrm{d}P(x_2|d_1, x_1)$.

Remark 5.4 (Key innovation). Here, $\mu_{x_1,x_2}(d_1)$ is a sequential kernel embedding that encodes the counterfactual distribution of the covariates (X_1, X_2) when the initial, counterfactual treatment value is $D_1 = d_1$. It satisfies $\langle f, \mu_{x_1,x_2}(d_1) \rangle_{\mathcal{H}_{\mathcal{X}} \otimes \mathcal{H}_{\mathcal{X}}} = \int f(x_1, x_2) \mathrm{d}P(x_2|d_1, x_1) \mathrm{d}P(x_1)$, implementing Robins' g-formula. It is our key innovation, and it accounts for treatment-confounder feedback in this setting.

Remark 5.5 (Robins' g-formula is bounded over $\mathcal{H}$ when the treatments are continuous). Define the functional $F : \mathcal{H} \to \mathbb{R}$, $\gamma \mapsto \int \gamma(d_{1:T}, x_{1:T}) \mathrm{d}P(x_1) \prod_{t=2}^{T} \mathrm{d}P\{x_t | d_{1:(t-1)}, x_{1:(t-1)}\}$. Similar to Remark 4.5, when the treatments are continuous and when Assumption 4.1 holds, we show that F is bounded by appealing to the definition of the RKHS. Boundedness of F guarantees the existence of the sequential kernel embeddings in Theorem 5.1.

As before, this decoupling is a blueprint for estimation. For example, our estimator will be $\hat{\theta}^{GF}(d_1, d_2) = \langle \hat{\gamma}, \phi(d_1) \otimes \phi(d_2) \otimes \hat{\mu}_{x_1,x_2}(d_1) \rangle_{\mathcal{H}}$ where $\hat{\mu}_{x_1,x_2}(d_1) = n^{-1} \sum_{i=1}^{n} \{\phi(X_{1i}) \otimes \hat{\mu}_{x_2}(d_1, X_{1i})$. Here, $\hat{\gamma}$ is a kernel ridge regression, $\hat{\mu}_{x_2}(d_1, x_1)$ is a generalized kernel ridge regression, and we combine them by averaging and taking the product.

Algorithm 5.1 (Nonparametric estimation of time-varying response curves). Denote the kernel matrices $K_{D_1D_1}, K_{D_2D_2}, K_{X_1X_1}, K_{X_2X_2} \in \mathbb{R}^{n \times n}$ calculated from the population P. Denote the kernel matrices $K_{\tilde{D}_1\tilde{D}_1}, K_{\tilde{D}_2\tilde{D}_2}, K_{\tilde{X}_1\tilde{X}_1}, K_{\tilde{X}_2\tilde{X}_2} \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$ calculated from the population $\tilde{P}$. Let $\odot$ be the elementwise product. Time-varying dose response curves have closed form solutions:

1. $\hat{\omega}(d_1, d_2; x_1) = Y^{\top}(K_{D_1D_1} \odot K_{D_2D_2} \odot K_{X_1X_1} \odot K_{X_2X_2} + n\lambda I)^{-1}$
   $[K_{D_1d_1} \odot K_{D_2d_2} \odot K_{X_1x_1} \odot \{K_{X_2X_2}(K_{D_1D_1} \odot K_{X_1X_1} + n\lambda_4 I)^{-1}(K_{D_1d_1} \odot K_{X_1x_1})\}]$;
2. $\hat{\theta}^{GF}(d_1, d_2) = n^{-1} \sum_{i=1}^{n} \hat{\omega}(d_1, d_2; X_{1i})$;
3. $\hat{\theta}^{DS}(d_1, d_2; \tilde{P}) = \tilde{n}^{-1} \sum_{i=1}^{\tilde{n}} Y^{\top}(K_{D_1D_1} \odot K_{D_2D_2} \odot K_{X_1X_1} \odot K_{X_2X_2} + n\lambda I)^{-1}$
   $[K_{D_1d_1} \odot K_{D_2d_2} \odot K_{X_1\tilde{x}_{1i}} \odot \{K_{X_2\tilde{X}_2}(K_{\tilde{D}_1\tilde{D}_1} \odot K_{\tilde{X}_1\tilde{X}_1} + \tilde{n}\lambda_5 I)^{-1}(K_{\tilde{D}_1d_1} \odot K_{\tilde{X}_1\tilde{x}_{1i}})\}]$,

where $(\lambda, \lambda_4, \lambda_5)$ are ridge regression penalty parameters. For incremental responses, we replace $K_{D_2d_2}$ with $\nabla_{d_2} K_{D_2d_2}$ where $(\nabla_{d_2} K_{D_2d_2})_i = \nabla_{d_2} k(D_{2i}, d_2)$.

We derive these algorithms in Section 13 of Singh, Xu and Gretton (2025). We give theoretical values for $(\lambda, \lambda_4, \lambda_5)$ that optimally balance bias and variance in Theorem 5.2 below. Section 16 gives practical tuning procedures with closed form solutions to empirically balance bias and variance, one of which is asymptotically optimal. Note that $\hat{\theta}^{DS}$ requires observations of the treatments and covariates from the alternative population $\tilde{P}$.

5.3. Uniform consistency with finite sample rates

Towards a guarantee of uniform consistency, we place regularity conditions on the RKHSs and original spaces via Assumptions 4.1 and 4.2. We also assume the regression γ_0 is smooth and quantify the effective dimension of $\mathcal{H}$ via Assumption 4.3. For the conditional kernel embeddings $\mu_{x_2}(d_1, x_1)$ and $\nu_{x_2}(d_1, x_1)$, we place further smoothness and effective dimension conditions via Assumption 4.4. With these assumptions, we arrive at our next main result.

Theorem 5.2 (Uniform consistency of time-varying response curves). *Suppose the conditions of Theorem 5.1 hold, as well as Assumptions 4.2 and 4.3. Set* $(\lambda, \lambda_4, \lambda_5) = \{n^{-1/(c+1/b)}, n^{-1/(c_4+1/b_4)}, \tilde{n}^{-1/(c_5+1/b_5)}\}$*, which is rate optimal regularization.*

1. *If in addition Assumption 4.4 holds with* $\mathcal{A}_4 = \mathcal{X}$ *and* $\mathcal{B}_4 = \mathcal{D} \times \mathcal{X}$*, then* $\|\hat{\theta}^{GF} - \theta_0^{GF}\|_\infty = O_p\left[n^{-(c-1)/\{2(c+1/b)\}} + n^{-(c_4-1)/\{2(c_4+1/b_4)\}}\right]$.
2. *If in addition Assumption 4.4 holds with* $\mathcal{A}_5 = \mathcal{X}$ *and* $\mathcal{B}_5 = \mathcal{D} \times \mathcal{X}$*, then* $\|\hat{\theta}^{DS} - \theta_0^{DS}\|_\infty = O_p\left[n^{-(c-1)/\{2(c+1/b)\}} + \tilde{n}^{-(c_5-1)/\{2(c_5+1/b_5)\}}\right]$.

Likewise for the incremental responses. For example, $\|\hat{\theta}^{GF,\nabla} - \theta_0^{GF,\nabla}\|_\infty = O_p\left[n^{-(c-1)/\{2(c+1/b)\}} + n^{-(c_4-1)/\{2(c_4+1/b_4)\}}\right]$.

See Section 14 of Singh, Xu and Gretton (2025) for the proof, exact finite sample rates, and explicit specializations of Assumption 4.4. As before, these rates are at best $n^{-1/4}$ when $(c, c_4, c_5) = 2$ and $(b, b_4, b_5) \to \infty$. See (3) for the Sobolev special case.

Remark 5.6 (Technical contribution). As before, the technical contribution underlying this theoretical guarantee is an RKHS norm rate for the sequential kernel embedding. In particular, Proposition 14.2 in Section 14 of Singh, Xu and Gretton (2025) derives nonasymptotic bounds on $\sup_{d_1 \in \mathcal{D}} \|\hat{\mu}_{x_1,x_2}(d_1) - \mu_{x_1,x_2}(d_1)\|_{\mathcal{H}_\mathcal{X} \otimes \mathcal{H}_\mathcal{X}}$ and $\sup_{d_1 \in \mathcal{D}} \|\hat{\nu}_{x_1,x_2}(d_1) - \nu_{x_1,x_2}(d_1)\|_{\mathcal{H}_\mathcal{X} \otimes \mathcal{H}_\mathcal{X}}$, as a stepping stone to Theorem 5.2. These intermediate results appear to be new.

Remark 5.7 (No effect). Consider the scenario when there is no effect of the second dose, i.e. $E\{Y^{(d_1,d_2)}\} = E\{Y^{(d_1)}\}$. As argued in Section 12 of Singh, Xu and Gretton (2025), under this additional restriction, $\gamma_0(d_1, d_2, x_1, x_2) = E(Y|D = d_1, X_1 = x_1, X_2 = x_2)$. In the proof technique of Section 14 of Singh, Xu and Gretton (2025), if the kernel ridge regression estimator $\hat{\gamma}$ is consistent for a function γ_0 that is constant in d_2, our rates remain valid.

In particular, our RKHS estimator for the time-varying dose response remains uniformly consistent as long as $\mathcal{H}_\mathcal{D}$ contains constant functions. While the RKHS with exponentiated quadratic kernel does not satisfy this property (Steinwart, Hush and Scovel, 2006), other RKHSs do. Another option is to augment an RKHS that does not contain constant functions with constant functions.

Interestingly, in this scenario Robins' g-formula simplifies to $\theta_0^{GF}(d_1) = \int E(Y|D = d_1, X = x_1) \mathrm{d}P(x_1)$ by the law of iterated expectations. When neither dose has any effect, a similar argument yields $\theta_0^{GF} = E(Y)$. Simplifying the g-formula from a surface to a curve to a scalar suggests that rate improvements may be possible. Remark 4.7 discusses possible directions for future work.

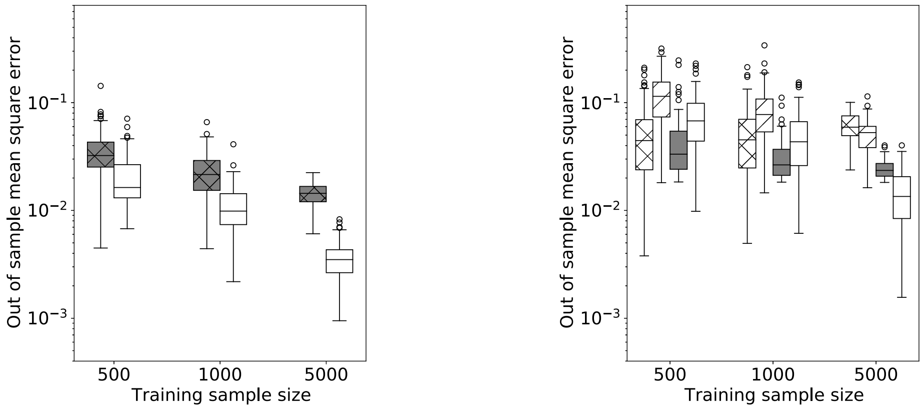

(a) Mediated response.

(b) Time-varying dose response.

Figure 1. Nonparametric response simulations. For the mediated response, we implement two estimators: Huber et al. (2020) (`IPW`, checkered gray) and our own (`RKHS`, white). For the time-varying dose response, we implement four estimators. From left to right, these are Singh, Xu and Gretton (2024) {`RKHS(ATE)`, checkered white}, Singh, Xu and Gretton (2024) {`RKHS(CATE)`, lined white}, Lewis and Syrgkanis (2021) (`SNMM`, gray), and our own {`RKHS(GF)`, white}.

6. Simulations and application

6.1. Simulations

We evaluate our estimators on various nonparametric designs. For each nonparametric design and sample size, we implement 100 simulations and calculate mean square error (MSE) with respect to the true causal function. Figure 1 visualizes results, where a lower MSE is desirable. See Section 17 of Singh, Xu and Gretton (2025) for the data generating processes and implementation details.

The mediated response design (Huber et al., 2020) involves learning the nonlinear causal function $\theta_0^{ME}(d,d') = 0.3d' + 0.09d + 0.15dd' + 0.25(d')^3$. A single observation is a tuple (Y,D,M,X) for the outcome, treatment, mediator, and covariates where $Y,D,M,X \in \mathbb{R}$. In addition to our estimator (`RKHS`, white), we implement the estimator of Huber et al. (2020) (`IPW`, checkered gray), which involves Nadaraya–Watson density estimation en route to generalized inverse propensity weighting. By the Wilcoxon rank sum test, `RKHS` significantly outperforms `IPW` at all sample sizes, with p values below 10^{-3}.

Next, we consider a time-varying dose response design, extending a time-fixed design (Colangelo and Lee, 2020). The nonlinear causal function is $\theta_0^{GF}(d_1,d_2) = 0.6d_1 + 0.5d_1^2 + 1.2d_2 + d_2^2$. A single observation is a tuple $(Y,D_{1:2},X_{1:2})$ for the outcome, treatments, and covariates where $Y,D_t \in \mathbb{R}$ and $X_t \in \mathbb{R}^{100}$. See Section 17 of Singh, Xu and Gretton (2025) for low and moderate dimensional settings. Our machine learning approach for time-varying response curves is uniformly consistent and allows for nonlinearity, dependence over time, and effect modification, which appears to be new.

To illustrate why treatment-confounder feedback and effect modification matter, we compare $\hat{\theta}^{GF}(d_1,d_2)$ {`RKHS(GF)`, white} with estimators that ignore these complexities to various degrees. Using the dose response estimator of Singh, Xu and Gretton (2024) {`RKHS(ATE)`, checkered white}, we take D_2 to be the treatment and misclassify D_1 as a covariate. Using the heterogeneous response estimator of Singh, Xu and Gretton (2024) {`RKHS(CATE)`, lined white}, we take D_2 to be the treatment

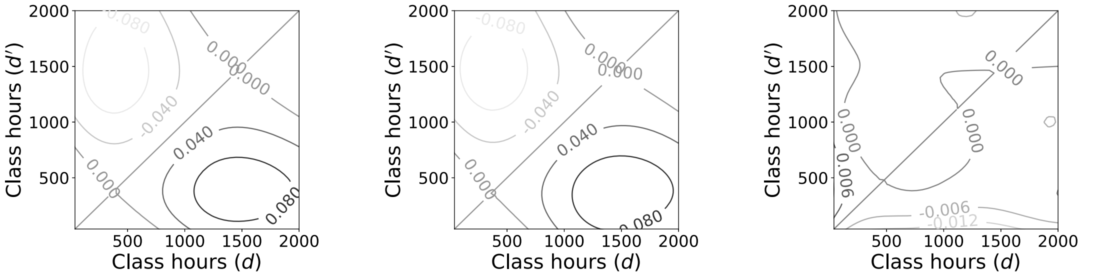

(a) Total response. (b) Direct response. (c) Indirect response.

Figure 2. Effect of job training on arrests. We implement our estimators for total, direct, and indirect response curves (RKHS, solid).

and misclassify D_1 as the subcovariate with meaningful heterogeneity. We also implement the estimator of Lewis and Syrgkanis (2021) (SNMM, gray), which is a machine learning approach with linearity, Markov, and no-effect-modification assumptions that do not hold in this setting. By the the Wilcoxon rank sum test, RKHS(GF) significantly outperforms the alternatives at $n = 5000$ with p value below 10^{-3}. The ability of RKHS(GF) to capture treatment-confounder feedback, and effect modification by time-varying confounders, helps when the sample size is large enough.

6.2. Application: US Job Corps

We estimate the mediated and time-varying responses of the Job Corps, the largest job training program for disadvantaged youth in the US. The Job Corps serves about 50,000 participants annually, and it is free for individuals who meet low income requirements. Access to the program was randomized from November 1994 to February 1996; see Schochet, Burghardt and McConnell (2008) for details. Though access to the program was randomized, individuals could decide whether to participate and for how many hours over multiple years. We assume that, conditional on the observed covariates, those decisions were as good as random in the sense formalized in Section 12 of Singh, Xu and Gretton (2025).

First, we consider employment to be a possible mechanism through which class hours affect arrests, under the identifying assumptions of Flores et al. (2012), Huber et al. (2020). The covariates $X \in \mathbb{R}^{40}$ are measured at baseline; the treatment $D \in \mathbb{R}$ is total hours spent in academic or vocational classes in the first year after randomization; the mediator $M \in \mathbb{R}$ is the proportion of weeks employed in the second year after randomization; and the outcome $Y \in \mathbb{R}$ is the number of times an individual is arrested by police in the fourth year after randomization. We use the same covariates $X \in \mathbb{R}^{40}$ and sample as Colangelo and Lee (2020), with $n = 2,913$ observations. In Figure 2, $\theta_0^{TE}(d, d')$, $\theta_0^{DE}(d, d')$, and $\theta_0^{IE}(d, d')$ are the total, direct, and indirect responses, respectively, of d' class hours relative to d class hours on arrests. In particular, $\theta_0^{IE}(d, d')$ estimates the extent to which the response is mediated by the mechanism of employment.

At best, the total response of receiving 1,600 class hours (40 weeks) versus 480 class hours (12 weeks) may be a reduction of about 0.1 arrests. The direct response estimate, of class hours on arrests, mirrors the total response estimate. Our indirect response estimate of class hours on arrests, as mediated through employment, is essentially zero. Our results extend the findings of Huber et al. (2020), allowing both (d, d') to vary. It appears that the effect of class hours on arrests is direct; there may be benefits of

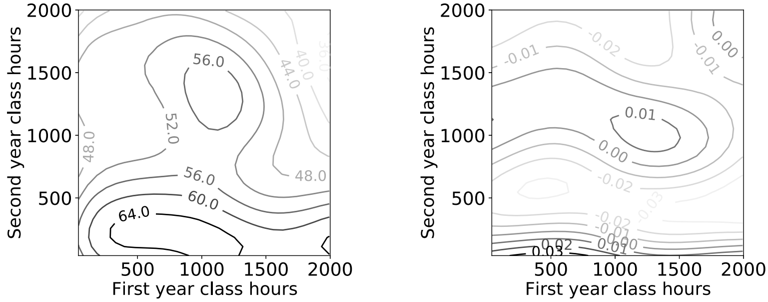

(a) Time-varying dose response. (b) Time-varying incremental response.

Figure 3. Effect of job training on employment. We implement our estimators for time-varying dose and incremental response curves {`RKHS(GF)`, solid}.

the training program that are not explained by employment alone. These benefits, however, may require many class hours.

Next, we evaluate the time-varying response of job training on employment. Here, $X_1 \in \mathbb{R}^{65}$ are covariates at baseline; $D_1 \in \mathbb{R}$ is the total class hours in the first year; $X_2 \in \mathbb{R}^{30}$ are covariates observed at the end of the first year; $D_2 \in \mathbb{R}$ is the total class hours in the second year; and $Y \in \mathbb{R}$ is the proportion of weeks employed in the fourth year. The covariates and the sample of $n = 3,141$ observations we use are similar to Colangelo and Lee (2020). The time-varying response $\theta_0^{GF}(d_1, d_2)$ is the counterfactual mean employment given d_1 class hours in year one and d_2 class hours in year two; $\theta_0^{GF,\nabla}(d_1, d_2)$ is the increment of counterfactual mean employment given d_1 class hours in year one and incrementally more than d_2 class hours in year two. Figure 3 visualizes the time-varying response estimate and its derivative with respect to the second dose.

The effect of training on employment appears to be positive when the duration of training is relatively brief. At best, the response to receiving job training appears to be 64% employment, compared to receiving no class hours at all which gives 56% employment. The maximum response is achieved by 480–1280 class hours (12–32 weeks) in year one and 0–480 (0–12 weeks) in year two. There is another local maximum of counterfactual employment achieved by 1200 class hours (30 weeks) in both years. Class hours in year one and year two may be complementary at low levels, as visualized by the incremental response. The large plateau in counterfactual employment suggests that a successful yet cost effective policy may be 480 class hours (12 weeks) in the first year and an optional, brief follow up in the second year.

In summary, under standard identifying assumptions, we find that the US Job Corps may provide two distinct benefits: reducing arrests and increasing employment, under different durations of class hours. Many class hours in the first year may directly decrease arrests in the fourth year, while few class hours in the first and second years may significantly increase employment in the fourth year. Section 18 of Singh, Xu and Gretton (2025) provides implementation details and verifies that our results are robust to the sample choice.

7. Discussion

Previous methods using kernels for continuous treatments (Singh, Xu and Gretton, 2024) do not handle treatment-confounder feedback, and therefore cannot analyze the later rounds of Job Corps surveys. We propose the sequential kernel embedding to do so. Whereas survey data for mediation analysis were previously available (Huber et al., 2020), we clean additional survey data for time-varying analysis from raw files (Schochet, Burghardt and McConnell, 2008, Section III.A). By providing clean data, we enable empirical analysis of the later rounds of Job Corps surveys, where class hours in different years may be viewed as a sequence of time-varying continuous treatments. Future work may apply the sequential embedding in dynamic programming for optimal policy estimation.

Acknowledgments

The authors would like to thank Neil Shephard for helpful comments. William Liu, Miriam Nelson, and Ikenna Ogbogu provided excellent research assistance.

Funding

The first author was supported by the Hausman Dissertation Fellowship and Simons-Berkeley Research Fellowship. The second and third authors were supported by the Gatsby Charitable Foundation, GAT 3850 and GAT 3528.

Supplementary Material

Supplement to "Sequential kernel embedding for mediated and time-varying dose response curves" (DOI: 10.3150/24-BEJ1836SUPPA; .pdf). Supplementary material includes an overview of results; connections to semiparametric efficiency, counterfactual distributions, and longer horizons; standard identifying assumptions; proofs; and implementation details.

Replication package for "Sequential kernel embedding for mediated and time-varying dose response curves" (DOI: 10.3150/24-BEJ1836SUPPB; .zip). The replication package includes data and code.

References

Bang, H. and Robins, J.M. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics* **61** 962–972. MR2216189 https://doi.org/10.1111/j.1541-0420.2005.00377.x

Baron, R.M. and Kenny, D.A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *J. Pers. Soc. Psychol.* **51** 1173–1182. http://doi.org/10.1037/0022-3514.51.6.1173

Berlinet, A. and Thomas-Agnan, C. (2004). *Reproducing Kernel Hilbert Spaces in Probability and Statistics*, 1st ed. Springer, New York. https://doi.org/10.1007/978-1-4419-9096-9

Bickel, P.J., Klaassen, C.A.J., Ritov, Y. and Wellner, J.A. (1993). *Efficient and Adaptive Estimation for Semiparametric Models*. *Johns Hopkins Series in the Mathematical Sciences*. Baltimore, MD: Johns Hopkins Univ. Press. MR1245941

Bodory, H., Huber, M. and Laffèrs, L. (2022). Evaluating (weighted) dynamic treatment effects by double machine learning. *Econom. J.* **25** 628–648. MR4530103 https://doi.org/10.1093/ectj/utac018

Caponnetto, A. and De Vito, E. (2007). Optimal rates for the regularized least-squares algorithm. *Found. Comput. Math.* **7** 331–368. MR2335249 https://doi.org/10.1007/s10208-006-0196-8

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econom. J.* **21** C1–C68. MR3769544 https://doi.org/10.1111/ectj.12097

Colangelo, K. and Lee, Y.-Y. (2020). Double debiased machine learning nonparametric inference with continuous treatments. Available at arXiv:2004.03036. https://doi.org/10.48550/arXiv.2004.03036

Farbmacher, H., Huber, M., Lafférs, L., Langen, H. and Spindler, M. (2022). Causal mediation analysis with double machine learning. *Econom. J.* **25** 277–300. MR4423402 https://doi.org/10.1093/ectj/utac003

Fischer, S. and Steinwart, I. (2020). Sobolev norm learning rates for regularized least-squares algorithms. *J. Mach. Learn. Res.* **21** Paper No. 205. MR4209491

Flores, C.A., Flores-Lagunes, A., Gonzalez, A. and Neumann, T.C. (2012). Estimating the effects of length of exposure to instruction in a training program: The case of Job Corps. *Rev. Econ. Stat.* **94** 153–171. http://doi.org/10.1162/REST_a_00177

Foster, D.J. and Syrgkanis, V. (2023). Orthogonal statistical learning. *Ann. Statist.* **51** 879–908. MR4630373 https://doi.org/10.1214/23-AOS2258

Ghassami, A., Sani, N., Xu, Y. and Shpitser, I. (2021). Multiply robust causal mediation analysis with continuous treatments. Available at arXiv:2105.09254. https://doi.org/10.48550/arXiv.2105.09254

Gill, R.D. and Robins, J.M. (2001). Causal inference for complex longitudinal data: The continuous case. *Ann. Statist.* **29** 1785–1811. MR1891746 https://doi.org/10.1214/aos/1015345962

Grünewälder, S., Gretton, A. and Shawe-Taylor, J. (2013). Smooth operators. In *International Conference on Machine Learning* **28** 1184–1192. PMLR.

Hirshberg, D.A., Maleki, A. and Zubizarreta, J.R. (2019). Minimax linear estimation of the retargeted mean. Available at arXiv:1901.10296. https://doi.org/10.48550/arXiv.1901.10296

Huber, M., Hsu, Y.-C., Lee, Y.-Y. and Lettry, L. (2020). Direct and indirect effects of continuous treatments based on generalized propensity score weighting. *J. Appl. Econometrics* **35** 814–840. MR4186786 https://doi.org/10.1002/jae.2765

Imai, K., Keele, L. and Yamamoto, T. (2010). Identification, inference and sensitivity analysis for causal mediation effects. *Statist. Sci.* **25** 51–71. MR2741814 https://doi.org/10.1214/10-STS321

Kallus, N. (2020). Generalized optimal matching methods for causal inference. *J. Mach. Learn. Res.* **21** Paper No. 62. MR4095341

Kennedy, E.H. (2023). Towards optimal doubly robust estimation of heterogeneous causal effects. *Electron. J. Stat.* **17** 3008–3049. MR4667730 https://doi.org/10.1214/23-ejs2157

Kennedy, E.H., Ma, Z., McHugh, M.D. and Small, D.S. (2017). Non-parametric methods for doubly robust estimation of continuous treatment effects. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 1229–1245. MR3689316 https://doi.org/10.1111/rssb.12212

Lewis, G. and Syrgkanis, V. (2021). Double/debiased machine learning for dynamic treatment effects. In *Conference on Neural Information Processing Systems* **34** 22695–22707. Curran Associates. https://doi.org/10.48550/arXiv.2002.07285

Li, K.-C. (1986). Asymptotic optimality of C_L and generalized cross-validation in ridge regression with application to spline smoothing. *Ann. Statist.* **14** 1101–1112. MR0856808 https://doi.org/10.1214/aos/1176350052

Li, Z., Meunier, D., Mollenhauer, M. and Gretton, A. (2022). Optimal rates for regularized conditional mean embedding learning. In *Conference on Neural Information Processing Systems* **35** 4433–4445. Curran Associates. https://doi.org/10.48550/arXiv.2208.01711

Luedtke, A.R., Sofrygin, O., van der Laan, M.J. and Carone, M. (2017). Sequential double robustness in right-censored longitudinal models. Available at arXiv:1705.02459. https://doi.org/10.48550/arXiv.1705.02459

Malinsky, D., Shpitser, I. and Richardson, T. (2019). A potential outcomes calculus for identifying conditional path-specific effects. In *Artificial Intelligence and Statistics* **89** 3080–3088. PMLR. https://doi.org/10.48550/arXiv.1903.03662

Molina, J., Rotnitzky, A., Sued, M. and Robins, J.M. (2017). Multiple robustness in factorized likelihood models. *Biometrika* **104** 561–581. MR3694583 https://doi.org/10.1093/biomet/asx027

Muandet, K., Kanagawa, M., Saengkyongam, S. and Marukatat, S. (2021). Counterfactual mean embeddings. *J. Mach. Learn. Res.* **22** Paper No. 162. MR4318518

Newey, W.K. (1994). The asymptotic variance of semiparametric estimators. *Econometrica* **62** 1349–1382. MR1303237 https://doi.org/10.2307/2951752

Nie, X. and Wager, S. (2021). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* **108** 299–319. MR4259133 https://doi.org/10.1093/biomet/asaa076

Pearl, J. (2001). Direct and indirect effects. In *Uncertainty in Artificial Intelligence* 411–420. Association for Computing Machinery. http://doi.org/10.1145/3501714.3501736

Pearl, J. and Bareinboim, E. (2014). External validity: From do-calculus to transportability across populations. *Statist. Sci.* **29** 579–595. MR3300360 https://doi.org/10.1214/14-STS486

Petersen, M., Schwab, J., Gruber, S., Blaser, N., Schomaker, M. and van der Laan, M. (2014). Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *J. Causal Inference* **2** 147–185. MR4289419 https://doi.org/10.1515/jci-2013-0007

Richardson, T.S. and Robins, J.M. (2013). Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. University of Washington Center for Statistics and the Social Sciences Working Paper No. 128.

Robins, J.M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math. Model.* **7** 1393–1512. https://doi.org/10.1016/0270-0255(86)90088-6

Robins, J. (1992). Estimation of the time-dependent accelerated failure time model in the presence of confounding factors. *Biometrika* **79** 321–334. MR1185134 https://doi.org/10.1093/biomet/79.2.321

Robins, J.M. and Greenland, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology* **3** 143–155. http://doi.org/10.1097/00001648-199203000-00013

Robins, J.M. and Richardson, T.S. (2011). Alternative graphical causal models and the identification of direct effects. In *Causality and Psychopathology: Finding the Determinants of Disorders and Their Cures* 103–158. Oxford Univ. Press. http://doi.org/10.1093/oso/9780199754649.003.0011

Robins, J.M., Richardson, T.S. and Shpitser, I. (2022). An interventionist approach to mediation analysis. In *Probabilistic and Causal Inference: The Works of Judea Pearl* 713–764. Association for Computing Machinery. http://doi.org/10.1145/3501714.3501754

Rotnitzky, A., Robins, J.M. and Babino, L. (2017). On the multiply robust estimation of the mean of the g-functional. Available at arXiv:1705.08582. https://doi.org/10.48550/arXiv.1705.08582

Scharfstein, D.O., Rotnitzky, A. and Robins, J.M. (1999). Adjusting for nonignorable drop-out using semiparametric nonresponse models. *J. Amer. Statist. Assoc.* **94** 1096–1146. MR1731478 https://doi.org/10.2307/2669923

Schochet, P.Z., Burghardt, J. and McConnell, S. (2008). Does Job Corps work? Impact findings from the national Job Corps study. *Amer. Econ. Rev.* **98** 1864–1886. http://doi.org/10.1257/aer.98.5.1864

Semenova, V. and Chernozhukov, V. (2021). Debiased machine learning of conditional average treatment effects and other causal functions. *Econom. J.* **24** 264–289. MR4281225 https://doi.org/10.1093/ectj/utaa027

Singh, R. (2021a). A finite sample theorem for longitudinal causal inference with machine learning: Long term, dynamic, and mediated effects. Available at arXiv:2112.14249. https://doi.org/10.48550/arXiv.2112.14249

Singh, R. (2021b). Debiased kernel methods. Available at arXiv:2102.11076. https://doi.org/10.48550/arXiv.2102.11076

Singh, R., Sahani, M. and Gretton, A. (2019). Kernel instrumental variable regression. In *Conference on Neural Information Processing Systems* **32** 4595–4607. Curran Associates. https://doi.org/10.48550/arXiv.1906.00232

Singh, R., Xu, L. and Gretton, A. (2020). Kernel methods for policy evaluation: Treatment effects, mediation analysis, and off-policy planning. In *Conference on Neural Information Processing Systems Workshop on ML for Economic Policy*. https://doi.org/10.48550/arXiv.2010.04855

Singh, R., Xu, L. and Gretton, A. (2021). Kernel methods for multistage causal inference: Mediation analysis and dynamic treatment effects. Available at arXiv:2111.03950. http://doi.org/10.48550/arXiv.2111.03950

Singh, R., Xu, L. and Gretton, A. (2024). Kernel methods for causal functions: Dose, heterogeneous and incremental response curves. *Biometrika* **111** 497–516. MR4745578 https://doi.org/10.1093/biomet/asad042

Singh, R., Xu, L. and Gretton, A. (2025). Supplement to "Sequential kernel embedding for mediated and time-varying dose response curves." https://doi.org/10.3150/24-BEJ1836SUPPA, https://doi.org/10.3150/24-BEJ1836SUPPB

Sriperumbudur, B., Fukumizu, K. and Lanckriet, G. (2010). On the relation between universality, characteristic kernels and RKHS embedding of measures. In *Artificial Intelligence and Statistics* **9** 773–780. PMLR.

Steinwart, I., Hush, D. and Scovel, C. (2006). An explicit description of the reproducing kernel Hilbert spaces of Gaussian RBF kernels. *IEEE Trans. Inf. Theory* **52** 4635–4643. MR2300845 https://doi.org/10.1109/TIT.2006.881713

Steinwart, I. and Scovel, C. (2012). Mercer's theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constr. Approx.* **35** 363–417. MR2914365 https://doi.org/10.1007/s00365-012-9153-3

Tchetgen Tchetgen, E.J. and Shpitser, I. (2012). Semiparametric theory for causal mediation analysis: Efficiency bounds, multiple robustness and sensitivity analysis. *Ann. Statist.* **40** 1816–1845. MR3015045 https://doi.org/10.1214/12-AOS990

VanderWeele, T.J. and Tchetgen Tchetgen, E.J. (2017). Mediation analysis with time varying exposures and mediators. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 917–938. MR3641414 https://doi.org/10.1111/rssb.12194

VanderWeele, T.J. and Vansteelandt, S. (2009). Conceptual issues concerning mediation, interventions and composition. *Stat. Interface* **2** 457–468. MR2576399 https://doi.org/10.4310/SII.2009.v2.n4.a7

Vansteelandt, S. and Joffe, M. (2014). Structural nested models and G-estimation: The partially realized promise. *Statist. Sci.* **29** 707–731. MR3300367 https://doi.org/10.1214/14-STS493

van der Laan, M.J., Bibaut, A. and Luedtke, A.R. (2018). CV-TMLE for nonpathwise differentiable target parameters. In *Targeted Learning in Data Science*. *Springer Ser. Statist.* 455–481. Cham: Springer. MR3820740

van der Laan, M.J. and Gruber, S. (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *Int. J. Biostat.* **8** Art. 9. MR2923282 https://doi.org/10.1515/1557-4679.1370

van der Laan, M.J. and Rubin, D. (2006). Targeted maximum likelihood learning. *Int. J. Biostat.* **2** Art. 11. MR2306500 https://doi.org/10.2202/1557-4679.1043

van der Vaart, A. (1991). On differentiable functionals. *Ann. Statist.* **19** 178–204. MR1091845 https://doi.org/10.1214/aos/1176347976

Wong, R.K.W. and Chan, K.C.G. (2018). Kernel-based covariate functional balancing for observational studies. *Biometrika* **105** 199–213. MR3768874 https://doi.org/10.1093/biomet/asx069

Zhao, Q. (2019). Covariate balancing propensity score by tailored loss functions. *Ann. Statist.* **47** 965–993. MR3909957 https://doi.org/10.1214/18-AOS1698

Zheng, W. and van der Laan, M.J. (2011). Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning*. *Springer Ser. Statist.* 459–474. New York: Springer. MR2867139 https://doi.org/10.1007/978-1-4419-9782-1_27

Received September 2023 and revised November 2024