
Dual Instrumental Variable Regression

Krikamol Muandet

Max Planck Institute for Intelligent Systems
Tübingen, Germany
krikamol@tuebingen.mpg.de

Arash Mehrjou

Max Planck Institute for Intelligent Systems
ETH Zürich, Zürich, Switzerland
arash.mehrjou@inf.ethz.ch

Si Kai Lee

Booth School of Business
University of Chicago, USA
sikai.lee@chicagobooth.edu

Anant Raj

Max Planck Institute for Intelligent Systems
Tübingen, Germany
anant.raj@tuebingen.mpg.de

Abstract

We present a novel algorithm for non-linear instrumental variable (IV) regression, DualIV, which simplifies traditional two-stage methods via a dual formulation. Inspired by problems in stochastic programming, we show that two-stage procedures for non-linear IV regression can be reformulated as a convex-concave saddle-point problem. Our formulation enables us to circumvent the first-stage regression which is a potential bottleneck in real-world applications. We develop a simple kernel-based algorithm with an analytic solution based on this formulation. Empirical results show that we are competitive to existing, more complicated algorithms for non-linear instrumental variable regression.

1 Introduction

Inferring causal relationships under the influence of unobserved confounders remains one of the most challenging problems in economics, health care, and social sciences [1, 2]. A typical example in economics is the study of returns from schooling [3], which attempts to measure the causal effect of education on labor market earnings. For each individual, the treatment variable X represents the level of education and the outcome Y represents how much they earn. However, one’s level of education and income is likely confounded by the socioeconomic status or other unobserved confounding factors H [1, Ch. 4].

Since randomized control trials are often infeasible in most economic studies, economists have turned to *instrumental variables* (IVs) or *instruments* derived from naturally occurring random experiments to overcome unobserved confounding. Informally, instrumental variables Z are defined as variables that are associated with the treatment X , affect the outcome Y only through X and do not share common causes with Y . For instance, the season-of-birth was used as an instrument in [4] to estimate the impact of compulsory schooling on earnings. Because of the compulsory school attendance laws, an individual’s season-of-birth, which is likely to be random, affects how long they actually remain in school, but not their earnings. Figure 1 illustrates this example. Finding valid instruments for specific problems is an essential task in econometrics [1] and epidemiology [5].

Although IV analysis is widely used, the statistical tools employed for estimating causal effect are fairly rudimentary. Most applications of instrumental variables utilise a two-stage procedure [1, 6–8]. For instance, the two-stage least squares (2SLS) relies on the assumption that the relationship between X and Y is linear [9]. It first estimates the conditional mean $\mathbb{E}[X|Z = z]$ via linear regression and then regresses Y on the estimate of $\mathbb{E}[X|Z = z]$ to obtain an estimate of the causal effect. Since the first-stage estimate is by construction independent from confounders, the resultant




Figure 1: A data generating process (DGP) with a hidden confounder H and an instrument Z . A variation in X comes from both H and Z . Intuitively speaking, the external source of variation from Z can help improve an estimation by removing the effect of H on X .

causal estimate is therefore free from hidden confounding. In the non-linear setting, however, a poorly-fitted first-stage regression may result in inaccurate second-stage estimates [1, Ch. 4.6].

In this paper, we propose a novel procedure, DualIV, to directly estimate the structural causal function. Unlike previous works which extend 2SLS by employing non-linear models in place of their linear counterparts [7, 8], we solve the dual problem which can be expressed as a convex-concave saddle-point problem. Based on this framework, we develop a consistent reproducing kernel Hilbert spaces-based (RKHS) algorithm. Our formulation was inspired by the mathematical resemblance of non-linear IV to two-stage problems in stochastic programming [10–12].

The rest of the paper is organized in the following manner. Section 2 introduces the IV regression problem, reviews related work and identifies current limitations. We present our formulation in Section 3, followed by the kernelized estimation method in Section 4. Then, Section 5 reports empirical comparisons between DualIV and existing algorithms. Finally, we discuss the limitations of our procedure and suggest future directions in Section 6. All proofs can be found in Appendix E.

2 Instrumental variable regression

Let X , Y , and Z be treatment, outcome, and instrumental variable(s) taking values in \mathcal{X} , \mathcal{Y} , and \mathcal{Z} , respectively. In this work, we assume that $\mathcal{Y} \in \mathbb{R}$, and \mathcal{X} and \mathcal{Z} are Polish spaces. We also assume that Y is bounded, *i.e.*, $|Y| < M < \infty$ almost surely. Moreover, we denote unobserved confounder(s) by H . The underlying data generating process (DGP) is described by the causal graph in Figure 1 equipped with the following structural causal model (SCM):

$$Y = f(X) + \varepsilon, \quad \mathbb{E}[\varepsilon] = 0, \quad (1)$$

where f is an unknown, potentially non-linear continuous function and ε denotes the additive noise which depends on the hidden confounder(s) H . If $\mathbb{E}[\varepsilon | X] = 0$, we can estimate f consistently from observational data via the standard least-square regression. This allows us to identify $\mathbb{E}[Y | \text{do}(X = x)]$ where $\text{do}(X = x)$ represents an intervention on X where its value is set to x [13].

For non-expert readers, we elaborate that $\text{do}(X = x)$ here denotes a mathematical operator which simulates physical interventions by setting the value of X to x , while keeping the rest of the model unchanged [14, Sec. 3.2.1]. That is, the conditional expectation $\mathbb{E}[Y | \text{do}(X = x)]$ is computed with respect to the interventional distribution $\mathbb{P}(Y | \text{do}(X = x))$. We can estimate $\mathbb{P}(Y | \text{do}(X = x))$ if it is possible to directly manipulate X and then observe the resulting outcome Y . In Figure 1, for instance, one may assign different levels of education to people and then observe their subsequent levels of income in the labor market. Unfortunately, such experiment is not always possible and we only have access to an observational distribution $\mathbb{P}(Y | X = x)$, which can be different from $\mathbb{P}(Y | \text{do}(X = x))$. The discrepancy between interventional and observational distributions may result from the unobserved socioeconomic status, as illustrated in Figure 1.

When hidden confounders exist between X and Y , the error term ε in (1) is generally correlated with X . Hence, $\mathbb{E}[\varepsilon | X] \neq 0$ and it follows from (1) that

$$\mathbb{E}[Y | X = x] = f(x) + \mathbb{E}[\varepsilon | X = x], \quad (2)$$

which implies that $\mathbb{E}[Y | \text{do}(X = x)] \neq \mathbb{E}[Y | X = x]$. Thus, standard least-square regression no longer provides a valid estimate of f for making a prediction about the outcome of an intervention on X [7, 8, 15]. To handle hidden confounders, we assume access to an instrumental variable(s) Z which satisfies the following assumptions: (i) **Relevance**: Z has a causal influence on X . (ii) **Exclusion restriction**: Z affects Y only through X , *i.e.*, $Y \perp\!\!\!\perp Z | X, \varepsilon$. (iii) **Unconfounded instrument(s)**: Z is independent of the error, *i.e.*, $\varepsilon \perp\!\!\!\perp Z$.

The properties of Z imply that $\mathbb{E}[\varepsilon | Z] = 0$. Taking the expectation of (1) w.r.t. Y conditioned on Z yields the following integral equation

$$\mathbb{E}[Y|Z] = \int_{\mathcal{X}} f(x) d\mathbb{P}(x|Z), \quad (3)$$

which is a Fredholm integral equation of the first kind. Recent works in nonparametric IV regression have adopted this perspective [7, 8, 15, 16] despite the fact that solving (3) directly is an ill-posed problem as it involves inverting linear compact operators [15, 17, 18].

To illustrate the role of an instrument, we consider two special cases. When X is perfectly correlated with Z , the treatment is uncorrelated with the hidden confounder. In other words, we recover the strong ignorability assumption [19, 20] required for causal inference. When Z is independent of X , the instrument is useless as it has no predictive power over treatment so the structural function f is unidentifiable from the data. Therefore, the most interesting cases lie between these two extremes, especially when X and Z are weakly correlated, see, *e.g.*, [21, 22][1, pp. 205–216].

2.1 Previous work

Early applications of instrumental variables often assume linear relationships between Z and X as well as X and Y [1, 23]. When there is a single endogenous variable and instrument, the structural parameter can be estimated consistently by the instrumental variable (IV) estimator [23]. Interestingly, we can obtain this estimate using a two-stage procedure: regress X on Z using ordinary least square (OLS) to calculate the predicted value of X and used that as an explanatory variable in the structural equation to estimate the structural parameter using OLS. When there are multiple instruments, the two-stage least squares (2SLS) estimator is obtained by using all the instruments simultaneously in the first-stage regression. Wooldridge [24, Theorem 5.3] asserts that the 2SLS estimator is the most efficient IV estimator; see, *e.g.*, [1, 24] for a detailed exposition.

Recently, several extensions of 2SLS have been proposed to overcome the linearity constraint. The first line of work replaces linear regression by a linear projection onto a set of known basis functions [15, 16, 25, 26]. Chen and Christensen [27] provides a uniform convergence rate of this approach. However, there exists no principled way of choosing the appropriate set of basis functions. The second line of work replaces the first-stage regression by a conditional density estimate of $\mathbb{P}(X|Z)$ [28, 29]. Despite being more flexible, such approaches are known to suffer from the curse of dimensionality [30, Ch. 1]. Other extensions of 2SLS are DeepIV [8] and KernelIV [7] algorithms. In [8], (3) is solved by first estimating $\mathbb{P}(X|Z)$ with a mixture of deep generative models on which f is learned using another deep neural network. Instead of neural networks, Singh et al. [7] proposes to model the first-stage regression using the conditional mean embedding of $\mathbb{P}(X|Z)$ [31–33] which is then used in the second-stage kernel ridge regression.

The curse of two-stage methods. Two-stage procedures have two fundamental issues. First, such procedures violate Vapnik’s principle [34]: “[...] *when solving a problem of interest, do not solve a more general problem as an intermediate step [...]*”. Specifically, estimating the conditional density [8] or the conditional mean embedding [7] via regression in the first stage can be harder than estimating the parameter of interest in the second stage. The first stage is even referred as the “forbidden regression” in econometrics [1, Ch. 4.6]. On top of that, we usually only observe a single sample from each $\mathbb{P}(X|Z = z)$, which further increases the difficulty of the task. Second, although two-stage procedures are asymptotically consistent, the first-stage estimate creates a finite-sample bias in the second-stage estimate [1, Sec. 4.6.4]. This bias can be alleviated through sample splitting [35] which is also used in [7, 8]. Thus, two-stage procedures are less sample efficient and could yield biased estimates when run on the smaller datasets common in economics and social sciences.

The generalized method of moments (GMM) framework provides another set of popular approaches for estimating f [36, 37]. Unlike two-stage procedures, GMM-based algorithms find a function f that satisfies the orthogonality condition $\mathbb{E}[\varepsilon|Z] = 0$ directly. Specifically, if g_1, g_2, \dots, g_m are arbitrary real-valued functions, the orthogonality condition implies that $\mathbb{E}[(Y - f(X))g_j(Z)] = 0$ for $j = 1, \dots, m$. The GMM estimate of f can then be obtained by minimizing the quadratic form $\frac{1}{2} \sum_{j=1}^m \psi(f, g_j)^2$ where $\psi(f, g) := \mathbb{E}[(Y - f(X))g(Z)]$. This estimator can be interpreted as a generalization of the 2SLS estimator in the linear setting [6]. Recently, extensions of GMM-based methods where both f and g are parameterized by deep neural networks have successfully been

used to solve non-linear IV regression [38, 39]. In contrast, Muandet et al. [40] considers the set of RKHS functions which allow for an analytic formulation of the orthogonality condition.

3 Dual IV

In this section, we reformulate the integral equation (3) as an empirical risk minimization problem and present DualIV algorithm.

3.1 Empirical risk minimization

Let $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ be a proper, convex, and lower semi-continuous loss function for any value in its first argument.¹ Let \mathcal{F} be an arbitrary class of continuous functions which we assume contains f that fulfills the integral equation (3). Then, we can formulate (3) as

$$\min_{f \in \mathcal{F}} R(f) := \mathbb{E}_{YZ} [\ell(Y, \mathbb{E}_{X|Z}[f(X)])], \quad (4)$$

where $R(f)$ denotes the expected risk of f . To understand how (3) and (4) are related, let us consider the squared loss $\ell(y, y') = (y - y')^2$ and define $h(z) := \mathbb{E}_{X|z}[f(X)]$. Then, the solution to (4) is the minimum mean square error (MMSE) estimator $h^*(z) := \mathbb{E}[Y|z]$, which is exactly the LHS of (3). If there exists no $f \in \mathcal{F}$ for which $h^*(z) = \mathbb{E}[Y|z]$, we use $h^*(z)$ as the best MMSE approximation.

The key challenge here is that if f is noncontinuous in $h(z)$, it is not assured to be consistently estimated even if $h(z)$ is estimated correctly [15]. We defer further discussion to Section 3.4. In addition, it remains cumbersome to solve (4) directly because of the inner expectation. To circumvent this, we look to similar two-stage problems in stochastic programming [10, 11]. For example, in [11], the problem of learning from conditional distributions was formulated in a similar fashion to (4). Moreover, [12] proposes the deconditional mean embedding (DME) which solves the integral equation (3) by performing a closed-form ‘‘inversion’’ of the conditional mean embedding of $\mathbb{P}(X|Z)$ (see [33, 42] for a review). By contrast, we solve the initial problem in (3) by resorting to the dual formulation of (4).

3.2 Dual formulation

To derive the dual form of (4), we employ two existing results, *interchangeability* and *Fenchel duality*, which we review; see, e.g., [11, Lemma 1], [43, Ch. 14], and [10, Ch. 7] for more details.

Theorem 1 (Interchangeability). *Let ω be a random variable on Ω and, for any $\omega \in \Omega$, the function $f(\cdot, \omega) : \mathbb{R} \rightarrow (-\infty, \infty)$ is proper and upper semi-continuous concave function. Then,*

$$\mathbb{E}_\omega \left[\max_{u \in \mathbb{R}} f(u, \omega) \right] = \max_{u(\cdot) \in \mathcal{U}(\Omega)} \mathbb{E}_\omega [f(u(\omega), \omega)], \quad (5)$$

where $\mathcal{U}(\Omega) := \{u(\cdot) : \Omega \rightarrow \mathbb{R}\}$ is the entire space of functions defined on the support Ω .

Definition 2 (Fenchel duality). *Let $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ be a proper, convex, and lower semi-continuous loss function for any value in its first argument and $\ell_y^* := \ell^*(y, \cdot)$ a convex conjugate of $\ell_y := \ell(y, \cdot)$ which is also proper, convex, and lower semi-continuous w.r.t. the second argument. Then, $\ell_y(v) = \max_u \{uv - \ell_y^*(u)\}$. The maximum is achieved at $v \in \partial \ell^*(u)$, or equivalently $u \in \partial \ell(v)$.*

Applying the interchangeability and Fenchel duality to (4) yields the expected loss

$$\begin{aligned} R(f) &= \mathbb{E}_{YZ} [\max_{u \in \mathbb{R}} \{\mathbb{E}_{X|Z}[f(X)]u - \ell_Y^*(u)\}] \\ &= \max_{u \in \mathcal{U}} \mathbb{E}_{YZ} [\mathbb{E}_{X|Z}[f(X)]u(Y, Z) - \ell_Y^*(u(Y, Z))] \\ &= \max_{u \in \mathcal{U}} \mathbb{E}_{XYZ} [f(X)u(Y, Z)] - \mathbb{E}_{YZ} [\ell_Y^*(u(Y, Z))] \end{aligned}$$

where \mathcal{U} is the space of continuous functions over $\mathcal{Y} \times \mathcal{Z}$. Hence, (4) can be reformulated as

$$\min_{f \in \mathcal{F}} \max_{u \in \mathcal{U}} \mathbb{E}_{XYZ} [f(X)u(Y, Z)] - \mathbb{E}_{YZ} [\ell_Y^*(u(Y, Z))]. \quad (6)$$

¹The function f is *proper* if $\text{dom } f \neq \emptyset$ and $f(x) > -\infty, \forall x \in \mathcal{X}$. It is *lower (upper) semi-continuous* at $x_0 \in \mathcal{X}$ if for $\varepsilon > 0$ there exists a neighborhood $N(x_0)$ of x_0 such that $\varepsilon < (>)f(x) - f(x_0)$ for all $x \in N(x_0)$ [41].

Following [11], we will refer to $u \in \mathcal{U}$ as the *dual function*. Note that this function depends on only the outcome Y and the instrument Z , but not the treatment X .

The advantages of our formulation (6) over (3) and (4) are twofold. First, there is no need to estimate $\mathbb{E}_{X|Z}[f(X)]$ or $\mathbb{P}(X|Z)$ explicitly. Second, the target function f appears linearly in (6) which makes it convex in f . Since ℓ_y^* is also convex, (6) is concave in the dual function u . Hence, (6) is essentially a convex-concave saddle-point problem for which efficient solvers exist [11].

For the squared loss $\ell(y, y') = (y - y')^2$, we have $\ell_y^*(w) = wy + \frac{1}{2}w^2$ (see Appendix A for the derivation) and the saddle-point problem (6) reduces to

$$\min_{f \in \mathcal{F}} \max_{u \in \mathcal{U}} \Psi(f, u) := \mathbb{E}_{XYZ} [(f(X) - Y)u(Y, Z)] - \frac{1}{2} \mathbb{E}_{YZ} [u(Y, Z)^2]. \quad (7)$$

To solve (7), one can adopt an SGD-based algorithm developed by Dai et al. [11]. Alternatively, we propose in Section 4 a simple algorithm that can solve (7) in closed form.

3.3 Interpreting the dual function

The dual function $u(y, z)$ plays an important role in our framework. To understand its role, we consider the minimization and maximization problems in (7) separately. For any $f \in \mathcal{F}$, the maximization problem is $\max_{u \in \mathcal{U}} \mathbb{E}_{XYZ} [(f(X) - Y)u(Y, Z)] - \frac{1}{2} \|u(Y, Z)\|_{L^2(\mathbb{P}_{YZ})}^2$ where the first term can be viewed, loosely speaking, as a loss function and the second as a regularizer. Intuitively, we are seeking $u^* \in \mathcal{U}$ that is least orthogonal to the residual. Given u^* , the outer minimization problem $\min_{f \in \mathcal{F}} \mathbb{E}_{XYZ} [(f(X) - Y)u^*(Y, Z)]$ finds the function f that yields the most orthogonal residual to u^* . Our procedure clearly differs from previous two-stage methods as the minimization and maximization stages are interdependent.

From the causal inference perspective, the residual contains the variation that cannot be explained by the current estimate of f due to hidden confounding. We select u that maximally reweights the residuals according to how inconsistent they are w.r.t. the unconfounded joint distribution of Y and Z . Given u , we then select f that minimizes the inconsistencies between the residuals and u . Hence, at the equilibrium, we are left with residuals uncorrelated with Y and Z which can be attributed to noise due to unobserved confounding.

Lastly, we draw a connection between (7) and GMM. Let g_1, g_2, \dots, g_m be real-valued functions on $\mathcal{Y} \times \mathcal{Z}$ and $\psi(f, g) := \mathbb{E}[(Y - f(X))g(Y, Z)]$. When $\mathcal{U} = \text{span}\{g_1, \dots, g_m\}$, it is not difficult to show that $\max_{u \in \mathcal{U}} \Psi(f, u) = \frac{1}{2} \psi^\top \Lambda^{-1} \psi$ where $\Lambda := \mathbb{E}_{YZ}[\mathbf{g} \otimes \mathbf{g}]$ with $\mathbf{g} := (g_1(Y, Z), \dots, g_m(Y, Z))^\top$; see Appendix B. That is, minimizing the above over f yields a formulation that strongly resembles the GMM objective, with the dual function $u(Y, Z)$ playing a role similar to that of an instrument. However, we must clarify that u cannot act as an instrument since it depends on Y and thereby violates the exclusion restriction assumption. This ambiguity has been resolved in Liao et al. [44, Appendix F] by resorting to an alternative formulation similar to (4) and (7). Furthermore, we also note that AGMM [38] and DeepGMM [39] rely on minimax optimization, similar to (7), but were formulated based on the GMM framework.

3.4 Theoretical analysis

This section provides the conditions for which the true structural function f^* can be identified by the optimum of the saddle-point problem (7). We lay out the assumptions needed for the optimal dual function u^* to be unique and continuous, show that the saddle-point formulation (7) is equivalent to the problem (4) under the squared loss and prove that the solution of (7) given u^* is indeed f^* .

Assumption 1. (i) $\mathbb{P}(X|Z)$ is continuous in Z for any values of X . (ii) The function class \mathcal{F} is correctly specified, i.e., $f^* \in \mathcal{F}$.

Following [11], we define the optimal dual function for any pair $(y, z) \in \mathcal{Y} \times \mathcal{Z}$ as $u^*(y, z) \in \arg \max_{u \in \mathbb{R}} \{\mathbb{E}_{X|z}[f(X) - y]u - (1/2)u^2\}$. Since this is an unconstrained quadratic program, $u^*(y, z)$ takes the form $\mathbb{E}_{X|z}[f(X)] - y$. Given Assumption 1 and the loss function ℓ is convex and continuously differentiable, it follows from [11, Proposition 1] that u^* is unique and continuous.

Next, we shows that if (f^*, u^*) is the saddle-point of (7), f^* minimizes the original objective (4). The result follows from plugging $u^* = \mathbb{E}_{X|z}[f(X)] - y$ into the dual loss $\Psi(f, u)$ in (7); see Appendix E.1 for the detailed proof.

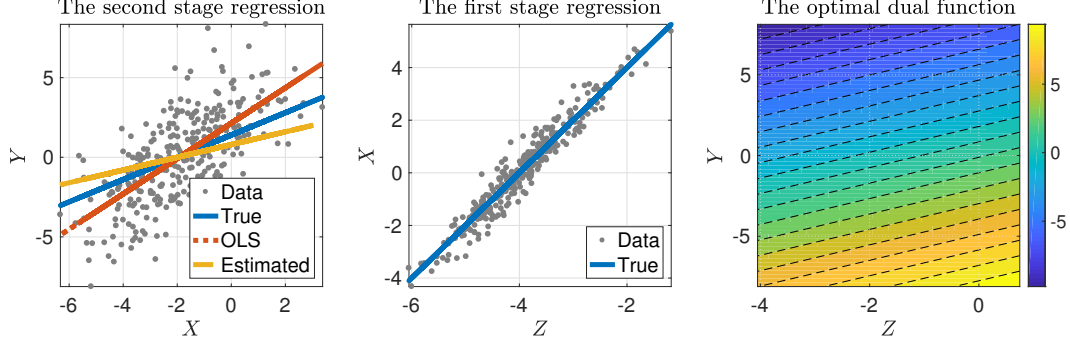


Figure 2: The dual function u w.r.t. the current estimate f in the linear setting (8). For each y and z , u can directly measure the discrepancy between y and $\mathbb{E}_{X|z}[f(X)]$.

Proposition 3. Let $\ell(y, y') = \frac{1}{2}(y - y')^2$. Then, for any fixed f , we have $R(f) = \max_u \Psi(f, u)$.

By Proposition 3 and the convexity of the loss $\ell(y, y')$, we obtain the following result.

Theorem 4. Let $\ell(y, y') = \frac{1}{2}(y - y')^2$ and assume that Assumption 1 holds. Then, (f^*, u^*) is the saddle-point of a minimax problem $\min_{f \in \mathcal{F}} \max_{u \in \mathcal{U}} \Psi(f, u)$.

By virtue of Theorem 4, we can identify the true function f^* under relatively weak assumptions. In contrast, previous work usually require stronger assumptions such as the completeness condition [7, 15] which specifies that the first-stage conditional expectation $\mathbb{E}_{X|z}[f(X)]$ is injective, or $h(z) = \mathbb{E}_{X|z}[f(X)]$ is a smooth function of z [7, 26, 27]. Since we do not perform first-stage regression, we only require $\mathbb{P}(X|Z)$ is continuous in Z for any value of X . The assumption that (4) is correctly specified, i.e., $f^* \in \mathcal{F}$, is standard in the literature [7, 15, 16].

As we can see, the optimal dual function $u^*(y, z) = \mathbb{E}_{X|z}[f(X)] - y$ acts as a residual function measuring the discrepancy between y and $\mathbb{E}_{X|z}[f(X)]$ [11]. Remarkably, this makes it possible to approximate $R(f)$ in (4) without computing the expectation $\mathbb{E}_{X|z}[f(X)]$ explicitly. We will later exploit this property in selecting hyperparameters. Moreover, since $\mathbb{E}_{X|z}[f(X)]$ allows X and Z to have a non-linear relationship, u can be non-linear even when the true structural function f^* is linear. This flexibility enables u to accommodate a larger class of functions that maps Z to X . Figure 2 illustrates this given the following generative process:

$$Y = X\beta + e + \epsilon, \quad X = (1 - \rho)Z_1 + \rho e + \eta \quad (8)$$

where $e \sim \mathcal{N}(0, 2)$, $Z_1 \sim \mathcal{N}(0, 2)$, $\epsilon \sim \mathcal{N}(0, 0.1)$, and $\eta \sim \mathcal{N}(0, 0.1)$. The parameter ρ controls the strength of the instrument w.r.t. hidden confounder e . Here, we set $n = 300$, $\beta = 0.7$, $\hat{\beta} = 0.4$, and $\rho = 0.2$ where $\hat{\beta}$ is an OLS estimate of β . Under this model, we have $u^*(y, z) \approx \hat{\beta}(1 - \rho)z - y$.

4 Kernelized DualIV

To demonstrate the effectiveness of our framework, we develop a simple kernel-based algorithm using the new formulation (7). To simplify notation, we denote by $W := (Y, Z)$ a random variable taking value in $\mathcal{W} := \mathcal{Y} \times \mathcal{Z}$. We pick \mathcal{F} and \mathcal{U} to be reproducing kernel Hilbert spaces (RKHSs) associated with positive definite kernels $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $l : \mathcal{W} \times \mathcal{W} \rightarrow \mathbb{R}$, respectively. Let $\phi : x \mapsto k(x, \cdot)$ and $\varphi : w \mapsto l(w, \cdot)$ be the canonical feature maps [45]. We assume both \mathcal{F} and \mathcal{U} are *universal* and hence are dense in the space of bounded continuous functions [46, Ch. 4].

Then, for any $f \in \mathcal{F}$ and $u \in \mathcal{U}$, we can rewrite the objective in (7) as

$$\begin{aligned} \Psi(f, u) &= \mathbb{E}_{XW}[f(X)u(W)] - \mathbb{E}_{YZ}[Yu(Y, Z)] - \frac{1}{2}\mathbb{E}_W[u(W)^2] \\ &= \langle \mathcal{C}_{WX}f - \mathbf{b}, u \rangle_{\mathcal{U}} - \frac{1}{2}\langle u, \mathcal{C}_W u \rangle_{\mathcal{U}}, \end{aligned} \quad (9)$$

where $\mathbf{b} := \mathbb{E}_{YZ}[Y\varphi(Y, Z)] \in \mathcal{U}$, $\mathcal{C}_W := \mathbb{E}_W[\varphi(W) \otimes \varphi(W)] \in \mathcal{U} \otimes \mathcal{U}$ is a covariance operator, and $\mathcal{C}_{WX} := \mathbb{E}_{WX}[\varphi(W) \otimes \phi(X)] \in \mathcal{U} \otimes \mathcal{F}$ is a cross-covariance operator [47, 48] (see Appendix

Algorithm 1 Kernelized DualIV

Input: Data $(x_i, y_i, z_i)_{i=1}^n$, kernel functions k, l , and a parameter grid Γ .

1: Compute kernel matrices $\mathbf{K}_{ij} = k(x_i, x_j)$ and $\mathbf{L}_{ij} = l((y_i, z_i), (y_j, z_j))$.

2: $(\lambda_1, \lambda_2) \leftarrow \text{SelectParams}(\mathbf{K}, \mathbf{L}, \Gamma)$.

3: $\mathbf{M} \leftarrow \mathbf{K}(\mathbf{L} + n\lambda_1 I)^{-1} \mathbf{L}$.

4: $\boldsymbol{\beta} \leftarrow (\mathbf{M}\mathbf{K} + n\lambda_2 \mathbf{K})^{-1} \mathbf{M}\mathbf{y}$.

Output: $f(x) = \sum_{i=1}^n \beta_i k(x_i, x)$.

C). Since (9) is quadratic in u , we have $\mathcal{C}_W u^* = \mathcal{C}_{WX} f - \mathbf{b}$. Substituting u^* back into (9) yields

$$f^* = \arg \min_{f \in \mathcal{F}} \frac{1}{2} \langle \mathcal{C}_{WX} f - \mathbf{b}, \mathcal{C}_W^{-1} (\mathcal{C}_{WX} f - \mathbf{b}) \rangle_{\mathcal{U}} = (\mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathcal{C}_{WX})^{-1} \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b}. \quad (10)$$

We can view (10) as a generalized least squares solution in RKHS. Since \mathcal{C}_W^{-1} and $(\mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathcal{C}_{WX})^{-1}$ may not exist in general, we replace them with regularized versions $(\mathcal{C}_W + \lambda_1 \mathcal{I})^{-1}$ and $(\mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathcal{C}_{WX} + \lambda_2 \mathcal{I})^{-1}$ where \mathcal{I} is the identity operator and $\lambda_1, \lambda_2 > 0$ are regularization parameters.

Given an i.i.d. sample $(x_i, y_i, z_i)_{i=1}^n$ from $\mathbb{P}(X, Y, Z)$, we define $\Phi := [\phi(x_1), \dots, \phi(x_n)]$, $\Upsilon := [\varphi(y_1, z_1), \dots, \varphi(y_n, z_n)]$, and $\mathbf{y} := [y_1, \dots, y_n]^\top$. Then, we can estimate $\mathbf{b}, \mathcal{C}_W$, and \mathcal{C}_{XW} with their empirical counterparts $\hat{\mathbf{b}} := n^{-1} \sum_{i=1}^n y_i \varphi(y_i, z_i) = n^{-1} \Upsilon \mathbf{y}$, $\hat{\mathcal{C}}_{XW} := n^{-1} \sum_{i=1}^n \phi(x_i) \otimes \varphi(y_i, z_i) = n^{-1} \Phi \Upsilon^\top$ and $\hat{\mathcal{C}}_W = n^{-1} \sum_{i=1}^n \varphi(y_i, z_i) \otimes \varphi(y_i, z_i) = n^{-1} \Upsilon \Upsilon^\top$. We denote the empirical version of (9) by $\hat{\Psi}(f, u)$ and the estimate of f^* by \hat{f} .

Next, we show that the representer theorem [49] for $\hat{\Psi}(f, u)$ holds for both f and u .

Lemma 5. *For any $f \in \mathcal{F}$ and $u \in \mathcal{U}$, there exist $f_\beta = \sum_{i=1}^n \beta_i k(x_i, \cdot)$ and $u_\alpha = \sum_{i=1}^n \alpha_i l(w_i, \cdot)$ for some $\alpha, \beta \in \mathbb{R}^n$ such that $\hat{\Psi}(f, u) = \hat{\Psi}(f_\beta, u_\alpha)$.*

By virtue of Lemma 5, the solution to (10) can be expressed as $f(x) = \sum_{i=1}^n \beta_i k(x_i, x)$ where the coefficients β are given by the following proposition.

Proposition 6. *Given an i.i.d. sample $(x_i, y_i, z_i)_{i=1}^n$ from $\mathbb{P}(X, Y, Z)$, let $\mathbf{K} := \Phi^\top \Phi$ and $\mathbf{L} := \Upsilon^\top \Upsilon$ be the Gram matrices such that $\mathbf{K}_{ij} = k(x_i, x_j)$ and $\mathbf{L}_{ij} = l(w_i, w_j)$ where $w_i := (y_i, z_i)$. Then, $\hat{f} = \Phi \beta$ where $\beta = (\mathbf{M}\mathbf{K} + n\lambda_2 \mathbf{K})^{-1} \mathbf{M}\mathbf{y}$ and $\mathbf{M} := \mathbf{K}(\mathbf{L} + n\lambda_1 I)^{-1} \mathbf{L}$.*

Compared to previous work which involved conditional density estimation [8, 28, 29] and vector-valued regression [7] as first-stage regression, estimating the dual function u , a real-valued function, is arguably easier. This is especially so when X and Z are high-dimensional.

Hyperparameter selection. Our estimator depends on two hyper-parameters, λ_1 and λ_2 . Given a dataset $(x_i, y_i, z_i)_{i=1}^{2n}$ of size $2n$, we provide a simple heuristic to determine the values of (λ_1, λ_2) . Ideally, if we know the optimal dual function u^* , we can interpret $u^*(y, z)^2$ as a loss function of f^* at (y, z) , as discussed in Section 3.4. To this end, we first estimate \hat{f} via Proposition 6 and $\hat{u}_\lambda := (\hat{\mathcal{C}}_W + \lambda \mathcal{I})^{-1} (\hat{\mathcal{C}}_{WX} \hat{f} - \hat{\mathbf{b}})$ on the first half of the data $(x_i, y_i, z_i)_{i=1}^n$. Next, the out-of-sample loss of \hat{f} is evaluated on the second half $(x_i, y_i, z_i)_{i=n+1}^{2n}$ by

$$\hat{R}(\hat{f}) = \frac{1}{n} \sum_{i=n+1}^{2n} (\mathbb{E}_{X|z_i}[\hat{f}(X)] - y_i)^2 \approx \frac{1}{n} \sum_{i=n+1}^{2n} \hat{u}_\lambda(y_i, z_i)^2. \quad (11)$$

Note that $\hat{u}_\lambda = \Upsilon(\mathbf{L} + n\lambda I)^{-1}(\mathbf{K}\beta - \mathbf{y}) = \Upsilon\alpha$ where $\alpha := (\mathbf{L} + n\lambda I)^{-1}(\mathbf{K}\beta - \mathbf{y})$ and $\mathbf{K}, \mathbf{L} \in \mathbb{R}^{n \times n}$ are kernel matrices evaluated on $(x_i, y_i, z_i)_{i=1}^n$. Hence, $\hat{R}(\hat{f}) \approx \alpha^\top \tilde{\mathbf{L}} \mathbf{1} / n$ where $\tilde{\mathbf{L}}_{ij} = l((y_i, z_i), (y_j, z_j))$ for $i = 1, \dots, n$ and $j = n+1, \dots, 2n$ and $\mathbf{1}$ is the all-ones column vector. In practice, we fix λ in \hat{u}_λ to a small constant to stabilize the loss (11) and only optimize (λ_1, λ_2) that appear in β . Note that this procedure differs from the two-stage causal validation procedures used in [7, 8]. Alternatively, one may choose the hyperparameters by cross-validation with respect to (11).

Algorithm 1 outlines the kernelized DualIV method whose consistency is studied in Appendix D. We note that above algorithm involves matrix inversions ($\mathcal{O}(n^3)$) which become the primary computational bottlenecks when scaling to large datasets. To improve the scalability of our algorithm,

Table 1: Comparisons of IV regression methods in small (top) and medium (bottom) sample size regimes. We report the \log_{10} mean squared error (MSE) and its standard deviations over 20 trials.

$n = 50$	Log ₁₀ Mean Squared Error (MSE)				
	$\rho = 0.1$	$\rho = 0.25$	$\rho = 0.5$	$\rho = 0.75$	$\rho = 0.9$
2SLS	5.814 ± 1.214	6.013 ± 0.827	5.895 ± 0.718	5.625 ± 1.182	5.308 ± 1.031
DeepIV	5.127 ± 0.043	5.131 ± 0.031	5.133 ± 0.072	5.130 ± 0.124	5.127 ± 0.061
KernelIV	4.481 ± 0.134	4.460 ± 0.095	4.438 ± 0.132	4.433 ± 0.100	4.462 ± 0.114
DeepGMM	3.848 ± 1.096	2.899 ± 1.638	3.952 ± 0.900	4.148 ± 0.556	3.738 ± 0.587
DualIV	4.257 ± 0.108	4.210 ± 0.126	4.285 ± 0.170	4.286 ± 0.126	4.232 ± 0.152
<hr/>					
$n = 1000$	Log ₁₀ Mean Squared Error (MSE)				
2SLS	8.236 ± 0.117	7.242 ± 1.232	8.290 ± 1.132	8.371 ± 0.865	8.544 ± 1.109
DeepIV	4.613 ± 0.052	4.618 ± 0.048	4.614 ± 0.068	4.701 ± 0.040	4.731 ± 0.032
KernelIV	4.189 ± 0.046	4.209 ± 0.040	4.199 ± 0.043	4.195 ± 0.045	4.194 ± 0.055
DeepGMM	4.090 ± 0.691	3.953 ± 1.076	4.392 ± 0.561	4.272 ± 0.595	4.415 ± 0.522
DualIV	4.143 ± 0.117	4.221 ± 0.185	4.104 ± 0.102	4.142 ± 0.105	4.127 ± 0.106

we can leverage the rich literature on large-scale kernel machines such as random Fourier features and Nyström method; see, e.g., Yang et al. [50] and references therein. Alternatively, we can employ stochastic gradient descent-based (SGD) algorithms similar to those proposed in Dai et al. [11, Algorithm 1] to solve the dual formulation (7) directly. This would also allow us to employ flexible models such as neural networks to parameterize the function classes \mathcal{F} and \mathcal{U} . Recently, Liao et al. [44] has taken an important step in this direction and provided convergence analysis for neural networks under a similar formulation.

5 Experiments

In this section, we compare kernelized DualIV² with: (i) vanilla two-stage least squares (2SLS) [23], (ii) DeepIV [8], (iii) KernelIV [7] and (iv) DeepGMM [39]. To provide a fair comparison, we adhered to the provided hyperparameter settings. Given the low-dimensional nature of our experiments, we used DeepGMM’s settings for low-dimensional scenarios in [39, Appendix B.2.1]. We ran 20 simulations of each algorithm for sample sizes of 50 and 1000 and calculated the \log_{10} mean squared error and its standard deviations w.r.t. the true function f for 2800 out-of-sample test points.

Demand design. We consider the same simulation as in [7, 8]: $Y = f(X) + \varepsilon$ where Y is outcome, $X = (P, T, S)$ are inputs, and $Z = (C, T, S)$ are instruments. Specifically, Y is sales, P is price, which is endogeneous, C is a supply cost shifter (instrument), and (T, S) are time of year and customer sentiment acting as exogeneous variables. The aim is to estimate the demand function $f(p, t, s) = 100 + (10 + p)s\psi(t) - 2p$ where $\psi(t) = 2[(t - 5)^4/600 + \exp(-4(t - 5)^2) + t/10 - 2]$. Training data is sampled according to $S \sim \text{Unif}\{1, \dots, 7\}$, $T \sim \text{Unif}[0, 10]$, $(C, V) \sim \mathcal{N}(\mathbf{0}, I_2)$, $\varepsilon \sim \mathcal{N}(\rho V, 1 - \rho^2)$, and $P = 25 + (C + 3)\psi(T) + V$. The parameter $\rho \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$ controls the extent to which price P is confounded with outcome Y by supply-side market forces. In our notation, $X = (P, T, S)$, $Z = (C, T, S)$ and $W = (Y, Z) = (Y, C, T, S)$.

For DualIV, we used the Gaussian RBF kernel for both k and l . In the experiments, the kernels on X and Z are product kernels, i.e., $k(x_i, x_j) = k_p(p_i, p_j)k_t(t_i, t_j)k_s(s_i, s_j)$ and $k(z_i, z_j) = k_c(c_i, c_j)k_t(t_i, t_j)k_s(s_i, s_j)$, and $l([y_i, z_i], [y_j, z_j]) = \exp([y_i - y_j, z_i - z_j]^\top V_{yz}^{-1} [y_i - y_j, z_i - z_j])$ where V_{yz} is a symmetric bandwidth matrix. The values of all bandwidth parameters are determined via the median heuristic. We choose (λ_1, λ_2) from $\{10^{-10}, 10^{-9}, \dots, 10^{-1}\}$ via cross-validation. Once (λ_1, λ_2) is chosen, we refit \hat{f} on the entire training set.

Table 1 reports the results of different methods evaluated on the test data. First, we observe that 2SLS achieves the largest MSE in both regimes as expected because the linearity assumption is violated here. Second, in the small sample size regime, DeepIV achieves relatively larger MSE than the other non-linear methods. KernelIV, DeepGMM, and DualIV, on the other hand, have comparable performance, with DeepGMM having the lowest MSE. However, we note that the results attained

²Our implementation is available at <https://github.com/krikamol/DualIV-NeurIPS2020>.

by DeepGMM were unstable out of the box and we had to reduce the variance of the initialization of the neural networks to 0.1 to obtain some degree of stability which is reflected in the standard deviations. We can fully attribute this variability to initialization as DeepGMM’s default batch size of 1024 is larger than that of both training datasets so there is no sampling variability in the optimization process. This suggests that DeepGMM, like DeepIV, is relatively brittle compared to kernel-based methods in the small sample size regime. Furthermore, DeepGMM comes with an extensive hyperparameter selection process, which highlights its need for fine-tuning. Last but not least, DualIV is competitive to KernelIV across ρ with slightly smaller MSE, which lends weight to our hypothesis that estimating the real-valued dual function is easier than vector-valued regression.

In the medium sample size regime, we observe that performance of DeepIV is in the same ballpark as the rest of the non-linear IV regression methods and the variance of DeepGMM is reduced, albeit still highest among the non-linear methods. The results of DualIV, KernelIV and DeepGMM are almost indistinguishable with DualIV having an edge as ρ increases. This could mean accounting for both Y and Z is perhaps slightly more effective than Z alone in the presence of greater confounding.

6 Conclusion

This paper proposes a general framework for non-linear IV regression called DualIV. Unlike previous work, DualIV does not require the first-stage regression which is the critical bottleneck of modern two-stage procedures. By exploiting tools in stochastic programming, we were able to reformulate the two-stage problem as the convex-concave saddle-point problem which is relatively simpler to solve. Instead of first-stage regression, DualIV requires the dual function $u(y, z)$ to be estimated, which is arguably easier than first-stage regression, especially when the instruments and treatments are high-dimensional. We demonstrate the validity of our framework with a kernel-based algorithm. Results show the competitiveness of our algorithm with respect to existing ones. Finally, potential directions for future work include (i) a minimax convergence analysis which could provide additional insight into the benefits of our framework, (ii) more flexible and scalable models such as deep neural networks as dual functions with stochastic gradient descent (SGD) [11], and (iii) applications to other two-stage problems in causal inference such as double ML [51].

Broader impact

This work provides a new framework for non-linear instrumental variable regression which allows one to perform causal analysis under the presence of unobserved confounders. This could have a profound impact in other fields such as economics, social science, and epidemiology, among others. Understanding the role of instruments in the context of learning theory may also pave the way towards creating more robust and trustworthy machine learning algorithms that are capable of surviving in the world full of hidden biases.

Acknowledgments and Disclosure of Funding

We are indebted to Rahul Singh and Arthur Gretton for their help with the KernelIV code which was used in our experiments. We thank Victor Chernozhukov, Elias Bareinboim, Sorawit Saengkyongam, Uri Shalit, Konrad Kording, Rahul Singh, Arthur Gretton, and You-Lin Chen for fruitful discussions as well as anonymous reviewers for the helpful feedback on our initial submission.

This work is funded by the federal and state governments of Germany through the Max Planck Society (MPG).

References

- [1] Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press, 2008.
- [2] Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, New York, NY, USA, 2015.
- [3] David Card. The causal effect of education on earnings. In *Handbook of Labor Economics*, volume 3 of *Handbook of Labor Economics*, chapter 30, pages 1801–1863. Elsevier, 1999.

- [4] Joshua D. Angrist and Alan B. Krueger. Does Compulsory School Attendance Affect Schooling and Earnings? *The Quarterly Journal of Economics*, 106(4):979–1014, 1991.
- [5] Stephen Burgess, Dylan S Small, and Simon G Thompson. A review of instrumental variable estimators for Mendelian randomization. *Statistical Methods in Medical Research*, 26(5): 2333–2355, 2017.
- [6] Halbert White. Instrumental variables regression with independent observations. *Econometrica*, 50(2):483–499, 1982.
- [7] Rahul Singh, Maneesh Sahani, and Arthur Gretton. Kernel instrumental variable regression. In *Advances in Neural Information Processing Systems 32*, pages 4593–4605. 2019.
- [8] Jason Hartford, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy. Deep IV: A flexible approach for counterfactual prediction. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1414–1423. PMLR, 2017.
- [9] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. 2017.
- [10] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory, Second Edition*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2014.
- [11] Bo Dai, Niao He, Yunpeng Pan, Byron Boots, and Le Song. Learning from Conditional Distributions via Dual Embeddings. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54, pages 1458–1467. PMLR, 2017.
- [12] Kelvin Hsu and Fabio Ramos. Bayesian deconditional kernel mean embeddings. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 2830–2838. PMLR, 2019.
- [13] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.
- [14] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009.
- [15] Whitney K. Newey and James L. Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003.
- [16] Joel L Horowitz. Applied nonparametric instrumental variables estimation. *Econometrica*, 79(2):347–394, 2011.
- [17] Reiner Kress. *Linear Integral Equations*, volume 3. Springer, 1989.
- [18] M. Nashed and G. Wahba. Generalized inverses in reproducing kernel spaces: An approach to regularization of linear operator equations. *SIAM Journal on Mathematical Analysis*, 5(6): 974–987, 1974.
- [19] D.B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974.
- [20] Donald Rubin. Causal inference using potential outcomes. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- [21] John Bound, David A. Jaeger, and Regina M. Baker. Problems with instrumental variables estimation when the correlation between the instruments and the endogeneous explanatory variable is weak. *Journal of the American Statistical Association*, 90(430):443–450, 1995.
- [22] Douglas Staiger and James H. Stock. Instrumental variables regression with weak instruments. *Econometrica*, 65(3):557–586, 1997.
- [23] Joshua D. Angrist, Guido W. Imbens, and Donald B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.
- [24] Jeffrey M. Wooldridge. *Econometric Analysis of Cross Section and Panel Data*, volume 1 of *MIT Press Books*. The MIT Press, March 2001.
- [25] Richard Blundell, Xiaohong Chen, and Dennis Kristensen. Semi-nonparametric IV estimation of shape-invariant Engel curves. *Econometrica*, 75(6):1613–1669, 2007.
- [26] Xiaohong Chen and Demian Pouzo. Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. *Econometrica*, 80(1):277–321, 2012.

- [27] Xiaohong Chen and Timothy M. Christensen. Optimal sup-norm rates and uniform inference on nonlinear functionals of nonparametric IV regression. *Quantitative Economics*, 9(1):39–84, 2018.
- [28] Peter Hall and Joel L. Horowitz. Nonparametric methods for inference in the presence of instrumental variables. *The Annals of Statistics*, 33(6):2904–2929, 12 2005.
- [29] S. Darolles, Y. Fan, J. P. Florens, and E. Renault. Nonparametric instrumental regression. *Econometrica*, 79(5):1541–1565, 2011.
- [30] Alexandre Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008.
- [31] Le Song, Jonathan Huang, Alex Smola, and Kenji Fukumizu. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *Proceedings of the 26th International Conference on Machine Learning (ICML)*, June 2009.
- [32] Le Song, Kenji Fukumizu, and Arthur Gretton. Kernel embeddings of conditional distributions: A unified kernel framework for nonparametric inference in graphical models. *IEEE Signal Processing Magazine*, 30(4):98–111, 2013.
- [33] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends in Machine Learning*, 10(1-2):1–141, 2017.
- [34] Vladimir N. Vapnik. *Statistical Learning Theory*. Wiley-Interscience, 1998.
- [35] Joshua D. Angrist and Alan B. Krueger. Split sample instrumental variables. Technical Report 699, Princeton University, Department of Economics, Industrial Relations Section., October 1993.
- [36] Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054, 1982.
- [37] A.R. Hall. *Generalized Method of Moments*. Advanced texts in econometrics. Oxford University Press, 2005.
- [38] Greg Lewis and Vasilis Syrgkanis. Adversarial generalized method of moments. 03 2018.
- [39] Andrew Bennett, Nathan Kallus, and Tobias Schnabel. Deep generalized method of moments for instrumental variable analysis. In *Advances in Neural Information Processing Systems 32*, pages 3564–3574. 2019.
- [40] Krikamol Muandet, Wittawat Jitkrittum, and Jonas Kübler. Kernel conditional moment test via maximum moment restriction. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence*, volume 124, pages 41–50. PMLR, 2020.
- [41] R. Tyrrell Rockafellar. *Convex analysis*. Princeton Mathematical Series. Princeton University Press, 1970.
- [42] Le Song, Kenji Fukumizu, and Arthur Gretton. Kernel embeddings of conditional distributions: A unified kernel framework for nonparametric inference in graphical models. *IEEE Signal Processing Magazine*, 30:98–111, 07 2013.
- [43] R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational Analysis*. Springer Verlag, Heidelberg, Berlin, New York, 1998.
- [44] Luofeng Liao, You-Lin Chen, Zhuoran Yang, Bo Dai, Zhaoran Wang, and Mladen Kolar. Provably efficient neural estimation of structural equation model: An adversarial approach. In *Advances in Neural Information Processing Systems 33*. 2020.
- [45] Bernhard Schölkopf and Alexander Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2002.
- [46] I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, 2008.
- [47] Charles R. Baker. Joint measures and cross-covariance operators. *Transactions of the American Mathematical Society*, 186:pp. 273–289, 1973.
- [48] Kenji Fukumizu, Francis Bach, and Michael Jordan. Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. *Journal of Machine Learning Research*, 5: 73–99, 2004.

- [49] B. Schölkopf, R. Herbrich, and A.J. Smola. A generalized representer theorem. In *Lecture Notes in Computer Science, Vol. 2111*, number 2111 in LNCS, pages 416–426, 2001.
- [50] Tianbao Yang, Yu-feng Li, Mehrdad Mahdavi, Rong Jin, and Zhi-Hua Zhou. Nyström method vs random Fourier features: A theoretical and empirical comparison. In *Advances in Neural Information Processing Systems 25*, pages 476–484, 2012.
- [51] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.
- [52] K. Fukumizu, L. Song, and A. Gretton. Kernel Bayes’ rule: Bayesian inference with positive definite kernels. *Journal of Machine Learning Research*, 14:3753–3783, 2013.

A Conjugate loss function

Let $\ell_y(v) := \frac{1}{2}(y - v)^2$ be a proper, convex, and lower semi-continuous function for all $y \in \mathbb{R}$. It follows from the definition of Fenchel conjugate (see, *e.g.*, Definition 2 or [43, Ch. 14] and [10, Ch. 7]) that for any $y \in \mathbb{R}$,

$$\ell_y^*(u) := \sup \{uv - \ell_y(v) : v \in \mathbb{R}\} = \sup \left\{ uv - \frac{1}{2}(y - v)^2 : v \in \mathbb{R} \right\}. \quad (12)$$

Hence, $\ell_y^*(u)$ is also a proper, concave, and upper semi-continuous function. Taking a derivative of $uv - \frac{1}{2}(y - v)^2$ w.r.t. v and setting it to zero yield a critical point $v^* = u + y$ for any $u, y \in \mathbb{R}$. Since $uv - \frac{1}{2}(y - v)^2$ is a concave function in v , we can substituting v^* back into (12) to obtain

$$\ell_y^*(u) = u(u + y) - \frac{1}{2}(y - (u + y))^2 = u^2 + uy - \frac{1}{2}u^2 = uy + \frac{1}{2}u^2,$$

as required.

B Connection to generalized method of moments (GMM)

To understand the connection between DualIV and GMM, let us consider $\mathcal{U} := \text{span}(g_1, \dots, g_m)$ where g_1, \dots, g_m are arbitrary real-valued functions on $\mathcal{Y} \times \mathcal{Z}$. That is, for any $u \in \mathcal{U}$, we have $u = \sum_{j=1}^m \alpha_j g_j$ for some $(\alpha_1, \dots, \alpha_m)^\top \in \mathbb{R}^m$. Then, we have

$$\begin{aligned} J(f) &:= \max_{u \in \mathcal{U}} \Psi(f, u) \\ &= \max_{\alpha \in \mathbb{R}^m} \mathbb{E}_{XYZ} \left[(f(X) - Y) \left(\sum_{j=1}^m \alpha_j g_j(Y, Z) \right) \right] - \frac{1}{2} \mathbb{E}_{YZ} \left[\left(\sum_{j=1}^m \alpha_j g_j(Y, Z) \right)^2 \right] \\ &= \max_{\alpha \in \mathbb{R}^m} \sum_{j=1}^m \alpha_j \mathbb{E}_{XYZ} [(f(X) - Y) g_j(Y, Z)] - \frac{1}{2} \mathbb{E}_{YZ} \left[\left(\sum_{j=1}^m \alpha_j g_j(Y, Z) \right)^2 \right] \\ &= \max_{\alpha \in \mathbb{R}^m} \alpha^\top \boldsymbol{\psi} - \frac{1}{2} \alpha^\top \Lambda \alpha, \end{aligned}$$

where we define $\boldsymbol{\alpha} := (\alpha_1, \dots, \alpha_m)^\top$, $\boldsymbol{\psi} := (\psi(f, g_1), \dots, \psi(f, g_m))^\top$, and $\Lambda := \mathbb{E}_{YZ}[\mathbf{g}(Y, Z) \otimes \mathbf{g}(Y, Z)]$ with $\mathbf{g}(Y, Z) := (g_1(Y, Z), \dots, g_m(Y, Z))^\top$. Taking the derivative w.r.t. $\boldsymbol{\alpha}$ and setting it to zero yield

$$J(f) = \frac{1}{2} \boldsymbol{\psi}^\top \Lambda^{-1} \boldsymbol{\psi}. \quad (13)$$

In this case, the DualIV objective can be expressed in a quadratic form. In the language of GMM, $\boldsymbol{\psi}$ acts as a vector of moment conditions and Λ acts as a weighting matrix [36, 37]. However, we have to reiterate that there is a fundamental difference here: the dual function $u \in \mathcal{U}$ cannot act as an instrument since it depends on Y and thereby violates the exclusion restriction assumption. Recently, Liao et al. [44, Appendix F] provided a clarification on this connection by resorting to an alternative formulation similar to (4) and (7).

C Dual formulation in RKHS

In this section, we provide a detailed derivation of the dual formulation (9) when \mathcal{F} and \mathcal{U} are both reproducing kernel Hilbert spaces (RKHSs) associated with positive definite kernels $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $l : \mathcal{W} \times \mathcal{W} \rightarrow \mathbb{R}$. Let $\phi : x \mapsto k(x, \cdot)$ and $\varphi : w \mapsto l(w, \cdot)$ be the canonical feature maps of k and l , respectively [45]. We assume throughout that both \mathcal{F} and \mathcal{U} are universal such that they are both dense in the space of bounded continuous functions (see, *e.g.*, [46, Ch. 4]). Furthermore, let $\text{HS}(\mathcal{F}, \mathcal{U})$ be a Hilbert space of Hilbert-Schmidt operators mapping from \mathcal{F} to \mathcal{U} with an inner product $\langle \cdot, \cdot \rangle_{\text{HS}}$ (see, *e.g.*, [33, Sec. 2.3]).

Then, for any $f \in \mathcal{F}$ and $u \in \mathcal{U}$, we can rewrite the objective in (7) as a functional

$$\begin{aligned}
\Psi(f, u) &= \mathbb{E}_{XW}[f(X)u(W)] - \mathbb{E}_{YZ}[Yu(Y, Z)] - \frac{1}{2}\mathbb{E}_W[u(W)^2] \\
&= \mathbb{E}_{XW}[\langle f, \phi(X) \rangle_{\mathcal{F}} \langle u, \varphi(W) \rangle_{\mathcal{U}}] - \mathbb{E}_{YZ}[Y \langle u, \varphi(Y, Z) \rangle_{\mathcal{U}}] - \frac{1}{2}\mathbb{E}_W[\langle u, \varphi(W) \rangle_{\mathcal{U}}^2] \\
&= \mathbb{E}_{XW}[\langle f \otimes u, \phi(X) \otimes \varphi(W) \rangle_{\text{HS}}] - \langle u, \mathbb{E}_{YZ}[Y\varphi(Y, Z)] \rangle_{\mathcal{U}} \\
&\quad - \frac{1}{2}\mathbb{E}_W[\langle u \otimes u, \varphi(W) \otimes \varphi(W) \rangle_{\text{HS}}] \\
&= \langle \mathcal{C}_{XW}, f \otimes u \rangle_{\text{HS}} - \langle u, \mathbf{b} \rangle_{\mathcal{U}} - \frac{1}{2}\langle \mathcal{C}_W, u \otimes u \rangle_{\text{HS}} \\
&= \langle f, \mathcal{C}_{XW}u \rangle_{\mathcal{F}} - \langle u, \mathbf{b} \rangle_{\mathcal{U}} - \frac{1}{2}\langle u, \mathcal{C}_W u \rangle_{\mathcal{U}},
\end{aligned}$$

where $\mathbf{b} := \mathbb{E}_{YZ}[Y\varphi(Y, Z)] \in \mathcal{U}$, $\mathcal{C}_W := \mathbb{E}_W[\varphi(W) \otimes \varphi(W)] \in \mathcal{U} \otimes \mathcal{U}$ is a covariance operator, and $\mathcal{C}_{XW} := \mathbb{E}_{XW}[\phi(X) \otimes \varphi(W)] \in \mathcal{F} \otimes \mathcal{U}$ is a cross-covariance operator [47, 48]. We used the reproducing property of \mathcal{F} and \mathcal{U} in the second equality. The third equality follows from the property of the rank-one operator, *i.e.*,

$$\langle f, g \rangle_{\mathcal{F}} \langle v, u \rangle_{\mathcal{U}} = \langle f \otimes v, u \otimes g \rangle_{\text{HS}(\mathcal{F}, \mathcal{U})}.$$

We then used the definition of \mathcal{C}_W , \mathcal{C}_{XW} , and \mathbf{b} to get the fourth equation. The last equation follows from the fact that $\langle L, f \otimes u \rangle_{\text{HS}(\mathcal{U}, \mathcal{F})} = \langle f, Lu \rangle_{\mathcal{F}}$ for any Hilbert-Schmidt operator $L \in \text{HS}(\mathcal{U}, \mathcal{F})$, $f \in \mathcal{F}$, and $u \in \mathcal{U}$.

D Consistency

In this section, we show that the kernelized DualIV estimator is asymptotically consistent under the assumption that \mathcal{C}_W^{-1} and $(\mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathcal{C}_{WX})^{-1}$ exist. Under these assumptions, we show in (10) that the solution f^* of the saddle-point problem (9) can be expressed as $f^* = (\mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathcal{C}_{WX})^{-1}\mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathbf{b}$. For simplicity, we assume further that the operator norm of the inverse covariance functions are bounded from below. The following theorem shows the consistency for the estimator \hat{f} obtained via Proposition 6.

Theorem 7. *Let \hat{f}_λ be an empirical estimator of f^* obtained from Proposition 6 with the regularization parameters $\lambda := (\lambda_1, \lambda_2)$. Assume that \mathcal{C}_W^{-1} and $(\mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathcal{C}_{WX})^{-1}$ exist and the operator norm of the inverse are bounded. Then, for sufficiently slow decay of regularization parameters λ_1 and λ_2 , \hat{f}_λ is a consistent estimator of f^* in RKHS norm, *i.e.*, $\|\hat{f}_\lambda - f^*\|_{\mathcal{F}} \rightarrow 0$ as $n \rightarrow \infty$.*

The proof of this theorem can be found in Appendix E.4.

The critical drawback of Theorem 7 is that it assumes the existence of \mathcal{C}_W^{-1} and $(\mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathcal{C}_{WX})^{-1}$ which may not hold in general. Similar assumption was also made in Fukumizu et al. [52] who provided counterexamples of cases in which such an assumption does not hold; see, also, [31, 33, 52] for the detailed discussion. One potential direction for future work is thus to provide the consistency result of DualIV under relatively weaker assumptions.

E Proofs

This section contains the detailed proofs.

E.1 Proof of Proposition 3

Proposition 3. *Let $\ell(y, y') = \frac{1}{2}(y - y')^2$. Then, for any fixed f , we have $R(f) = \max_u \Psi(f, u)$.*

Proof. Taking ℓ in (4) to be $\frac{1}{2}(y - y')^2$, plugging $u^*(y, z) = \mathbb{E}_{X|z}[f(X)] - y$ into (7) yields

$$\Psi(f, u^*) = \mathbb{E}_{XYZ}[(f(X) - Y)u^*(Y, Z)] - \frac{1}{2}\mathbb{E}_{YZ}[u^*(Y, Z)^2]$$

$$\begin{aligned}
&= \mathbb{E}_{XYZ}[(f(X) - Y)(\mathbb{E}_{X|Z}[f(X)] - Y)] - \frac{1}{2}\mathbb{E}_{YZ}[(\mathbb{E}_{X|Z}[f(X)] - Y)^2] \\
&= \mathbb{E}_{YZ}[(\mathbb{E}_{X|Z}[f(X)] - Y)(\mathbb{E}_{X|Z}[f(X)] - Y)] - \frac{1}{2}\mathbb{E}_{YZ}[(\mathbb{E}_{X|Z}[f(X)] - Y)^2] \\
&= \mathbb{E}_{YZ}[(\mathbb{E}_{X|Z}[f(X)] - Y)^2] - \frac{1}{2}\mathbb{E}_{YZ}[(\mathbb{E}_{X|Z}[f(X)] - Y)^2] \\
&= \frac{1}{2}\mathbb{E}_{YZ}[(\mathbb{E}_{X|Z}[f(X)] - Y)^2] = R(f),
\end{aligned}$$

as required. \square

E.2 Proof of Lemma 5

Lemma 5. For any $f \in \mathcal{F}$ and $u \in \mathcal{U}$, there exist $f_\beta = \sum_{i=1}^n \beta_i k(x_i, \cdot)$ and $u_\alpha = \sum_{i=1}^n \alpha_i l(w_i, \cdot)$ for some $\alpha, \beta \in \mathbb{R}^n$ such that $\widehat{\Psi}(f, u) = \widehat{\Psi}(f_\beta, u_\alpha)$.

Proof. Given a fixed sample $(x_i, y_i, z_i)_{i=1}^n$ of size n , any RKHSes \mathcal{F} and \mathcal{U} can be decomposed as $\mathcal{F} = \mathcal{F}_\beta \oplus \mathcal{F}_\perp$ and $\mathcal{U} = \mathcal{U}_\alpha \oplus \mathcal{U}_\perp$ where \mathcal{F}_β and \mathcal{U}_α are respectively subspaces consisting of functions of the following forms:

$$f_\beta = \sum_{i=1}^n \beta_i k(x_i, \cdot), \quad u_\alpha = \sum_{i=1}^n \alpha_i l(w_i, \cdot),$$

for some $\beta \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}^n$. The orthogonal subspaces \mathcal{F}_\perp and \mathcal{U}_\perp consist of elements which are orthogonal to \mathcal{F}_β and \mathcal{U}_α , respectively, i.e., for any $f_\beta \in \mathcal{F}_\beta$, $f_\perp \in \mathcal{F}_\perp$, $u_\alpha \in \mathcal{U}_\alpha$, $u_\perp \in \mathcal{U}_\perp$, we have $\langle f_\beta, f_\perp \rangle_{\mathcal{F}} = 0$ and $\langle u_\alpha, u_\perp \rangle_{\mathcal{U}} = 0$. Any elements $f \in \mathcal{F}$ and $u \in \mathcal{U}$ can thus be expressed as $f = f_\beta + f_\perp$ and $u = u_\alpha + u_\perp$ where $f_\beta \in \mathcal{F}_\beta$, $f_\perp \in \mathcal{F}_\perp$, $u_\alpha \in \mathcal{U}_\alpha$, and $u_\perp \in \mathcal{U}_\perp$.

Next, recall that

$$\widehat{\Psi}(f, u) = \langle f, \widehat{\mathcal{C}}_{XW} u \rangle_{\mathcal{F}} - \langle u, \widehat{\mathbf{b}} \rangle_{\mathcal{U}} - \frac{1}{2} \langle u, \widehat{\mathcal{C}}_W u \rangle_{\mathcal{U}}$$

where $\widehat{\mathcal{C}}_{XW} = n^{-1} \Phi \Upsilon^\top$, $\widehat{\mathcal{C}}_W = n^{-1} \Upsilon \Upsilon^\top$, $\widehat{\mathbf{b}} = n^{-1} \Upsilon \mathbf{y}$, $\Phi = [k(x_1, \cdot), \dots, k(x_n, \cdot)]$, $\Upsilon = [l(w_1, \cdot), \dots, l(w_n, \cdot)]$, and $\mathbf{y} = [y_1, \dots, y_n]^\top$. Using the above decomposition, we have

$$\begin{aligned}
\widehat{\Psi}(f, u) &= \langle f, \widehat{\mathcal{C}}_{XW} u \rangle_{\mathcal{F}} - \langle u, \widehat{\mathbf{b}} \rangle_{\mathcal{U}} - \frac{1}{2} \langle u, \widehat{\mathcal{C}}_W u \rangle_{\mathcal{U}} \\
&= \langle f_\beta + f_\perp, \sum_{i=1}^n \beta'_i k(x_i, \cdot) \rangle_{\mathcal{F}} - \langle u, \widehat{\mathbf{b}} \rangle_{\mathcal{U}} - \frac{1}{2} \langle u, \widehat{\mathcal{C}}_W u \rangle_{\mathcal{U}} \\
&= \langle f_\beta, \sum_{i=1}^n \beta'_i k(x_i, \cdot) \rangle_{\mathcal{F}} - \langle u, \widehat{\mathbf{b}} \rangle_{\mathcal{U}} - \frac{1}{2} \langle u, \widehat{\mathcal{C}}_W u \rangle_{\mathcal{U}},
\end{aligned}$$

where $\beta'_i := n^{-1} \langle l(w_i, \cdot), u \rangle_{\mathcal{U}}$. Since the choice of u is arbitrary, the minimizer of $\widehat{\Psi}(f, u)$ with respect to f lives in the subspace \mathcal{F}_β .

Similarly, we can write $\widehat{\Psi}(f, u)$ for any $f \in \mathcal{F}$ as a function of $u \in \mathcal{U}$ as

$$\begin{aligned}
\widehat{\Psi}(f, u) &= \langle \widehat{\mathcal{C}}_{WX} f, u \rangle_{\mathcal{U}} - \langle u, \widehat{\mathbf{b}} \rangle_{\mathcal{U}} - \frac{1}{2} \langle u, \widehat{\mathcal{C}}_W u \rangle_{\mathcal{U}} \\
&= \left\langle \sum_{i=1}^n \alpha'_i l(w_i, \cdot), u_\alpha + u_\perp \right\rangle_{\mathcal{U}} - \langle u_\alpha + u_\perp, \sum_{i=1}^n \alpha''_i l(w_i, \cdot) \rangle_{\mathcal{U}} \\
&\quad - \frac{1}{2} \langle u_\alpha + u_\perp, \widehat{\mathcal{C}}_W (u_\alpha + u_\perp) \rangle_{\mathcal{U}} \\
&= \left\langle \sum_{i=1}^n \alpha'_i l(w_i, \cdot), u_\alpha \right\rangle_{\mathcal{U}} - \langle u_\alpha, \sum_{i=1}^n \alpha''_i l(w_i, \cdot) \rangle_{\mathcal{U}} - \frac{1}{2} \langle u_\alpha, \widehat{\mathcal{C}}_W u_\alpha \rangle_{\mathcal{U}}.
\end{aligned}$$

The first equality follows from $\langle f, \widehat{\mathcal{C}}_{XW} u \rangle_{\mathcal{F}} = \langle \widehat{\mathcal{C}}_{WX} f, u \rangle_{\mathcal{U}}$ as $\widehat{\mathcal{C}}_{XW}$ is an adjoint operator of $\widehat{\mathcal{C}}_{WX}$. Since the choice of f is arbitrary, the maximizer of $\widehat{\Psi}(f, u)$ with respect to u also lives in the subspace \mathcal{U}_α .

Consequently, $\widehat{\Psi}(f, u) = \widehat{\Psi}(f_\beta, u_\alpha)$ for some $\beta \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}^n$. This completes the proof. \square

E.3 Proof of Proposition 6

Proposition 6. *Given an i.i.d. sample $(x_i, y_i, z_i)_{i=1}^n$ from $\mathbb{P}(X, Y, Z)$, let $\mathbf{K} := \Phi^\top \Phi$ and $\mathbf{L} := \Upsilon^\top \Upsilon$ be the Gram matrices such that $\mathbf{K}_{ij} = k(x_i, x_j)$ and $\mathbf{L}_{ij} = l(w_i, w_j)$ where $w_i := (y_i, z_i)$. Then, $\hat{f} = \Phi\beta$ where $\beta = (\mathbf{MK} + n\lambda_2\mathbf{K})^{-1}\mathbf{M}\mathbf{y}$ and $\mathbf{M} := \mathbf{K}(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}$.*

Proof. It follows from (10) that the structural function $f^* \in \mathcal{F}$ satisfies

$$(\mathcal{C}_{XW}(\mathcal{C}_W + \lambda_1\mathcal{I})^{-1}\mathcal{C}_{WX} + \lambda_2\mathcal{I})f^* = \mathcal{C}_{XW}(\mathcal{C}_W + \lambda_1\mathcal{I})^{-1}\mathbf{b}.$$

Replacing the population quantities with the empirical counterparts $\widehat{\mathcal{C}}_{XW} = \frac{1}{n}\Phi\Upsilon^\top$, $\widehat{\mathcal{C}}_W = \frac{1}{n}\Upsilon\Upsilon^\top$, and $\widehat{\mathbf{b}} = \frac{1}{n}\Upsilon\mathbf{y}$ yields

$$(\Phi\Upsilon^\top(\Upsilon\Upsilon^\top + n\lambda_1\mathcal{I})^{-1}\Upsilon\Phi^\top + n\lambda_2\mathcal{I})f^* = \Phi\Upsilon^\top(\Upsilon\Upsilon^\top + n\lambda_1\mathcal{I})^{-1}\Upsilon\mathbf{y}.$$

Using the identity $\Upsilon^\top(\Upsilon\Upsilon^\top + n\lambda_1\mathcal{I})^{-1} = (\Upsilon^\top\Upsilon + n\lambda_1\mathbf{I})^{-1}\Upsilon^\top$, the above equation can be rewritten as

$$\begin{aligned} (\Phi(\Upsilon^\top\Upsilon + n\lambda_1\mathbf{I})^{-1}\Upsilon^\top\Upsilon\Phi^\top + n\lambda_2\mathcal{I})f^* &= \Phi(\Upsilon^\top\Upsilon + n\lambda_1\mathbf{I})^{-1}\Upsilon^\top\Upsilon\mathbf{y} \\ (\Phi(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\Phi^\top + n\lambda_2\mathcal{I})f^* &= \Phi(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{y}. \end{aligned}$$

By Lemma 5, $f^* = \Phi\beta$ for some $\beta \in \mathbb{R}^n$. Substituting this back into the equation above yields

$$\begin{aligned} \Phi(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\Phi^\top\Phi\beta + n\lambda_2\Phi\beta &= \Phi(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{y} \\ \Phi(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{K}\beta + n\lambda_2\Phi\beta &= \Phi(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{y}. \end{aligned}$$

Multiplying both sides of the equation by Φ^\top gives

$$\begin{aligned} \Phi^\top\Phi(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{K}\beta + n\lambda_2\Phi^\top\Phi\beta &= \Phi^\top\Phi(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{y} \\ \mathbf{K}(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{K}\beta + n\lambda_2\mathbf{K}\beta &= \mathbf{K}(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{y} \\ (\mathbf{K}(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{K} + n\lambda_2\mathbf{K})\beta &= \mathbf{K}(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}\mathbf{y}. \end{aligned}$$

Setting $\mathbf{M} = \mathbf{K}(\mathbf{L} + n\lambda_1\mathbf{I})^{-1}\mathbf{L}$ yields the result. \square

E.4 Proof of Theorem 7

Theorem 7. *Let \hat{f}_λ be an empirical estimator of f^* obtained from Proposition 6 with the regularization parameters $\lambda := (\lambda_1, \lambda_2)$. Assume that \mathcal{C}_W^{-1} and $(\mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathcal{C}_{WX})^{-1}$ exist and the operator norm of the inverse are bounded. Then, for sufficiently slow decay of regularization parameters λ_1 and λ_2 , \hat{f}_λ is a consistent estimator of f^* in RKHS norm, i.e., $\|\hat{f}_\lambda - f^*\|_{\mathcal{F}} \rightarrow 0$ as $n \rightarrow \infty$.*

For ease of understanding, we will use the following notation throughout the proof:

$$\begin{aligned} \mathcal{R} &:= \mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathcal{C}_{WX} & \widehat{\mathcal{R}} &:= \widehat{\mathcal{C}}_{XW}\widehat{\mathcal{C}}_W^{-1}\widehat{\mathcal{C}}_{WX} \\ \mathcal{R}_{\lambda_1} &:= \mathcal{C}_{XW}(\mathcal{C}_W + \lambda_1\mathcal{I})^{-1}\mathcal{C}_{WX} & \widehat{\mathcal{R}}_{\lambda_1} &:= \widehat{\mathcal{C}}_{XW}(\widehat{\mathcal{C}}_W + \lambda_1\mathcal{I})^{-1}\widehat{\mathcal{C}}_{WX} \\ \mathcal{C}_{\lambda_1} &:= \mathcal{C}_W + \lambda_1\mathcal{I} & \widehat{\mathcal{C}}_{\lambda_1} &:= \widehat{\mathcal{C}}_W + \lambda_1\mathcal{I}. \end{aligned}$$

The following identity will be used heavily in our proof:

$$(B^{-1} - A^{-1}) = B^{-1}(A - B)A^{-1}. \quad (14)$$

Proof. First, it follows from the assumption that $\|\mathcal{C}_W^{-1}\|_{op} \leq \delta_1^{-1}$ and $\|(\mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathcal{C}_{WX}^\top)^{-1}\|_{op} \leq \delta_2^{-1}$ for some $\delta_1, \delta_2 > 0$. Moreover, we can write the empirical estimate as

$$\hat{f}_\lambda = (\widehat{\mathcal{R}}_{\lambda_1} + \lambda_2\mathcal{I})^{-1}\widehat{\mathcal{C}}_{XW}\widehat{\mathcal{C}}_{\lambda_1}^{-1}\widehat{\mathbf{b}}.$$

Similarly, under our assumption, the true population function can be expressed as

$$f^* = \mathcal{R}^{-1}\mathcal{C}_{XW}\mathcal{C}_W^{-1}\mathbf{b}.$$

The goal is then to bound the difference of \hat{f}_λ and f^* in RKHS norm, *i.e.*,

$$\begin{aligned}
\|\hat{f}_\lambda - f^*\|_{\mathcal{F}} &= \left\| (\widehat{\mathcal{R}}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \widehat{\mathcal{C}}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} - \mathcal{R}^{-1} \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \left\| (\widehat{\mathcal{R}}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \widehat{\mathcal{C}}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} - (\mathcal{R} + \lambda_2 \mathcal{I})^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\quad + \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{R}^{-1} \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&=: T_1 + T_2.
\end{aligned} \tag{15}$$

Bounding T_2 : Let us first consider the second term T_2 in (15):

$$\begin{aligned}
T_2 &:= \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{R}^{-1} \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{R}_{\lambda_1}^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\quad + \left\| \mathcal{R}_{\lambda_1}^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{R}^{-1} \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&=: T_{21} + T_{22}.
\end{aligned} \tag{16}$$

Bounding T_{21} : Let us consider T_{21} in (16) first.

$$\begin{aligned}
T_{21} &= \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{R}_{\lambda_1}^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&= \lambda_2 \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \mathcal{R}_{\lambda_1}^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \lambda_2 \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \right\|_{op} \left\| \mathcal{R}_{\lambda_1}^{-1} \right\|_{op} \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \lambda_2 \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} (\mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathcal{C}_{WX}) \right\|_{op} \left\| \mathcal{R}_{\lambda_1}^{-1} \right\|_{op}^2 \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \lambda_2 \left\| \mathcal{R}_{\lambda_1}^{-1} \right\|_{op}^2 \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \frac{\lambda_2}{\delta_2^2} \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \leq \frac{\lambda_2}{\delta_1 \delta_2^2} \left\| \mathcal{C}_{XW} \right\|_{op} \left\| \mathbf{b} \right\|_{\mathcal{F}}.
\end{aligned} \tag{17}$$

The second equality in (17) follows from the identity in (14). Hence, we have $T_{21} = \mathcal{O}(\lambda_2)$. From the above argument, it is also clear that there exists a positive constant C such that $\max(\|\mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b}\|_{\mathcal{F}}, \|\mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b}\|_{\mathcal{F}}) \leq C$.

Bounding T_{22} : Let us now consider the term T_{22} in (16).

$$\begin{aligned}
T_{22} &= \left\| \mathcal{R}_{\lambda_1}^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{R}^{-1} \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \left\| \mathcal{R}_{\lambda_1}^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{R}^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} + \left\| \mathcal{R}^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{R}^{-1} \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&= \left\| (\mathcal{R}_{\lambda_1}^{-1} - \mathcal{R}^{-1}) \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} + \left\| \mathcal{R}^{-1} (\mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} - \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b}) \right\|_{\mathcal{F}} \\
&\leq C \left\| \mathcal{R}_{\lambda_1}^{-1} \right\|_{op} \left\| \mathcal{R}^{-1} \right\|_{op} \left\| \mathcal{R}_{\lambda_1} - \mathcal{R} \right\|_{op} + \left\| \mathcal{R}^{-1} \right\|_{op} \left\| \mathcal{C}_{XW} (\mathcal{C}_{\lambda_1}^{-1} - \mathcal{C}_W^{-1}) \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq C \left\| \mathcal{R}_{\lambda_1}^{-1} \right\|_{op} \left\| \mathcal{R}^{-1} \right\|_{op} \left\| \mathcal{C}_{XW} (\mathcal{C}_{\lambda_1}^{-1} - \mathcal{C}_W^{-1}) \mathcal{C}_{WX} \right\|_{op} \\
&\quad + \left\| \mathcal{R}^{-1} \right\|_{op} \left\| \mathcal{C}_{XW} (\mathcal{C}_{\lambda_1}^{-1} - \mathcal{C}_W^{-1}) \mathbf{b} \right\|_{\mathcal{F}}.
\end{aligned} \tag{18}$$

Further, we have

$$T_{22} \leq \frac{\lambda C}{\delta_2} \left\| \mathcal{R}_{\lambda_1}^{-1} \right\|_{op} \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathcal{C}_W^{-1} \mathcal{C}_{WX} \right\|_{op} + \left\| \mathcal{R}^{-1} \right\|_{op} \left\| \mathcal{C}_{XW} (\mathcal{C}_{\lambda_1}^{-1} - \mathcal{C}_W^{-1}) \mathbf{b} \right\|_{\mathcal{F}}. \tag{19}$$

Let us consider now the following term in the above inequality:

$$\begin{aligned}
\left\| \mathcal{R}_{\lambda_1}^{-1} \right\|_{op} &\leq \left\| \mathcal{R}^{-1} \right\|_{op} + \left\| \mathcal{R}_{\lambda_1}^{-1} - \mathcal{R}^{-1} \right\|_{op} \\
&\leq \frac{1}{\delta_2} + \left\| \mathcal{R}_{\lambda_1}^{-1} - \mathcal{R}^{-1} \right\|_{op} \\
&\leq \frac{1}{\delta_2} + \left\| \mathcal{R}_{\lambda_1}^{-1} \right\|_{op} \left\| \mathcal{R}^{-1} \right\|_{op} \left\| \mathcal{R}_{\lambda_1} - \mathcal{R} \right\|_{op}
\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{\delta_2} + \|\mathcal{R}_{\lambda_1}^{-1}\|_{op} \|\mathcal{R}^{-1}\|_{op} \|\mathcal{C}_{XW}(\mathcal{C}_{\lambda_1}^{-1} - \mathcal{C}_W^{-1})\mathcal{C}_{WX}\|_{op} \\
&= \frac{1}{\delta_2} + \lambda_1 \|\mathcal{R}_{\lambda_1}^{-1}\|_{op} \|\mathcal{R}^{-1}\|_{op} \|\mathcal{C}_{XW}\mathcal{C}_{\lambda_1}^{-1}\mathcal{C}_W^{-1}\mathcal{C}_{WX}\|_{op}.
\end{aligned}$$

Now, $\|\mathcal{C}_{XW}\mathcal{C}_{\lambda_1}^{-1}\mathcal{C}_W^{-1}\mathcal{C}_{WX}\|_{op} \leq \tilde{c}/\delta_1^2$ and similarly $\|\mathcal{C}_{XW}(\mathcal{C}_{\lambda_1}^{-1} - \mathcal{C}_W^{-1})\mathbf{b}\|_{\mathcal{F}} \leq \hat{c}/\delta_1^2$ for some positive real numbers \hat{c} and \tilde{c} . Hence,

$$\|\mathcal{R}_{\lambda_1}^{-1}\|_{op} \leq \frac{1}{\delta_2} + \frac{\lambda_1 \tilde{c}}{\delta_1^2 \delta_2} \|\mathcal{R}_{\lambda_1}^{-1}\|_{op} \Rightarrow \|\mathcal{R}_{\lambda_1}^{-1}\|_{op} \leq \frac{1/\delta_2}{1 - \frac{\lambda_1 \tilde{c}}{\delta_1^2 \delta_2}}.$$

Hence, if $\lambda_1 \rightarrow 0$ sufficiently fast, then

$$T_{22} \leq \frac{\lambda_1 \tilde{C}}{\delta_1^2 \delta_2} \quad (20)$$

for a positive real number \tilde{C} . This implies $T_{22} = \mathcal{O}(\lambda_1)$.

Bounding T_1 : We now consider the first term in (15).

$$\begin{aligned}
T_1 &= \left\| (\hat{\mathcal{R}}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} - (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \left\| (\hat{\mathcal{R}}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} - (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} \right\|_{\mathcal{F}} \\
&\quad + \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} - (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&=: T_{11} + T_{12}.
\end{aligned} \quad (21)$$

Bounding T_{11} : Consider the first term in (21):

$$\begin{aligned}
T_{11} &= \left\| (\hat{\mathcal{R}}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} - (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} \right\|_{\mathcal{F}} \\
&\leq \left\| \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathbf{b}} \right\|_{\mathcal{F}} \left\| (\hat{\mathcal{R}}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \right\|_{op} \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \right\|_{op} \left\| \hat{\mathcal{R}}_{\lambda_1} - \mathcal{R}_{\lambda_1} \right\|_{op} \\
&\leq \frac{C}{\lambda_1 \lambda_2^2} \left\| \hat{\mathcal{R}}_{\lambda_1} - \mathcal{R}_{\lambda_1} \right\|_{op}
\end{aligned} \quad (22)$$

for some positive constant C . Next, we have

$$\begin{aligned}
&\left\| \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathcal{C}_{WX} \right\|_{op} \\
&\leq \left\| \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} - \hat{\mathcal{C}}_{XW} \mathcal{C}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} \right\|_{op} + \left\| \hat{\mathcal{C}}_{XW} \mathcal{C}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathcal{C}_{WX} \right\|_{op} \\
&\leq \left\| \hat{\mathcal{C}}_{XW} (\hat{\mathcal{C}}_{\lambda_1}^{-1} - \mathcal{C}_{\lambda_1}^{-1}) \hat{\mathcal{C}}_{WX} \right\|_{op} + \left\| \hat{\mathcal{C}}_{XW} \mathcal{C}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathcal{C}_{WX} \right\|_{op} \\
&\quad + \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathcal{C}_{WX} \right\|_{op} \\
&\leq \left\| \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} (\hat{\mathcal{C}}_W - \mathcal{C}_W) \mathcal{C}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} \right\|_{op} + \left\| (\hat{\mathcal{C}}_{XW} - \mathcal{C}_{XW}) \mathcal{C}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} \right\|_{op} \\
&\quad + \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} (\hat{\mathcal{C}}_{WX} - \mathcal{C}_{WX}) \right\|_{op}.
\end{aligned} \quad (23)$$

From the \sqrt{n} -consistency of covariance and cross-covariance operators [47, 48], we have

$$\left\| \hat{\mathcal{C}}_{XW} \hat{\mathcal{C}}_{\lambda_1}^{-1} \hat{\mathcal{C}}_{WX} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathcal{C}_{WX} \right\|_{op} = \mathcal{O}\left(\frac{1}{\lambda_1 \sqrt{n}}\right).$$

Hence, if $(\lambda_1 \lambda_2)^2$ converges to zero slower than $1/\sqrt{n}$, then T_{11} converges to zero asymptotically.

Bounding T_{12} : Let us now consider the second term in (21):

$$\begin{aligned}
T_{12} &= \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \widehat{\mathcal{C}}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathbf{b}} - (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \left\| (\mathcal{R}_{\lambda_1} + \lambda_2 \mathcal{I})^{-1} \right\|_{op} \left\| \widehat{\mathcal{C}}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathbf{b}} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \frac{1}{\lambda_2} \left\| \widehat{\mathcal{C}}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathbf{b}} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq \frac{1}{\lambda_2} \left[\left\| \widehat{\mathcal{C}}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathbf{b}} - \mathcal{C}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathbf{b}} \right\|_{\mathcal{F}} + \left\| \mathcal{C}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathbf{b}} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \widehat{\mathbf{b}} \right\|_{\mathcal{F}} \right. \\
&\quad \left. + \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \widehat{\mathbf{b}} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{F}} \right] \\
&\leq \frac{1}{\lambda_2} \left[\left\| (\widehat{\mathcal{C}}_{XW} - \mathcal{C}_{XW}) \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathbf{b}} \right\|_{\mathcal{F}} + \left\| \mathcal{C}_{XW} (\widehat{\mathcal{C}}_{\lambda_1}^{-1} - \mathcal{C}_{\lambda_1}^{-1}) \widehat{\mathbf{b}} \right\|_{\mathcal{F}} \right. \\
&\quad \left. + \left\| \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \widehat{\mathbf{b}} - \mathcal{C}_{XW} \mathcal{C}_{\lambda_1}^{-1} \mathbf{b} \right\|_{\mathcal{G}} \right].
\end{aligned} \tag{24}$$

By the \sqrt{n} -consistency of mean element and covariance operator in RKHS [47, 48], we have that $T_{12} = \mathcal{O}\left(\frac{1}{\lambda_1 \lambda_2 \sqrt{n}}\right)$. Moreover, it follows from what we have shown so far that

$$\begin{aligned}
&\left\| (\widehat{\mathcal{C}}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathcal{C}}_{WX} + \lambda_2 \mathcal{I})^{-1} \widehat{\mathcal{C}}_{XW} \widehat{\mathcal{C}}_{\lambda_1}^{-1} \widehat{\mathbf{b}} - (\mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathcal{C}_{WX})^{-1} \mathcal{C}_{XW} \mathcal{C}_W^{-1} \mathbf{b} \right\|_{\mathcal{F}} \\
&\leq T_1 + T_2 \leq T_{11} + T_{12} + T_{13} + T_{14}.
\end{aligned} \tag{25}$$

Hence, if λ_1 and λ_2 converge to zero with the sample size n such that $\frac{1}{\lambda_1^2 \lambda_2^2 \sqrt{n}}$ also converges to zero, then $\|\widehat{f}_\lambda - f^*\|_{\mathcal{F}} \rightarrow 0$. That is, our estimator is consistent in RKHS norm. \square