

Detecting Unobserved Confounders: A Kernelized Regression Approach

Yikai Chen^{1*}, Yunxin Mao^{1*}, Chunyuan Zheng², Hao Zou³, Shanzhi Gu¹,
Shixuan Liu¹, Yang Shi⁴, Wenjing Yang¹, Kun Kuang⁵, Haotian Wang^{1†}

¹College of Computer Science, National University of Defense Technology, Changsha, China

²School of Mathematical Sciences, Peking University, Beijing, China

³Department of Computer Science, Tsinghua University, Beijing, China

⁴School of Computer Science, Peking University, Beijing, China

⁵College of Computer Science and Technology, Zhejiang University, Hangzhou, China
chenyikai1998@nudt.edu.cn

Abstract

Detecting unobserved confounders is crucial for reliable causal inference in observational studies. Existing methods require either linearity assumptions or multiple heterogeneous environments, limiting applicability to nonlinear single-environment settings. To bridge this gap, we propose Kernel Regression Confounder Detection (KRCD), a novel method for detecting unobserved confounding in nonlinear observational data under single-environment conditions. KRCD leverages reproducing kernel Hilbert spaces to model complex dependencies. By comparing standard and higher-order kernel regressions, we derive a test statistic whose significant deviation from zero indicates unobserved confounding. Theoretically, we prove two key results: First, in infinite samples, regression coefficients coincide if and only if no unobserved confounders exist. Second, finite-sample differences converge to zero-mean Gaussian distributions with tractable variance. Extensive experiments on synthetic benchmarks and the Twins dataset demonstrate that KRCD not only outperforms existing baselines but also achieves superior computational efficiency.

1 Introduction

Causal effect estimation plays a central role in diverse high-stakes domains, such as healthcare (Prosperi et al. 2020; Dahabreh and Bibbins-Domingo 2024; Wang et al. 2024), public policy (Matthay and Glymour 2022), and financial decision-making (Kühne et al. 2022; Vanderschueren 2024). While randomized controlled trials (RCTs) remain the gold standard for causal inference, their practical deployment is often hindered by ethical concerns or limited scalability. Consequently, observational studies have emerged as a primary alternative, valued for their broader external validity and reduced cost (Pearl 2009; Listl, Jürges, and Watt 2016).

However, the reliability of observational studies critically depends on the *unconfoundedness assumption*—the requirement that all variables confounding the treatment-outcome relationship are fully measured and accounted for. When

the unconfoundedness assumption does not hold, spurious correlations will undermine causal conclusions drawn from observational data (Angrist, Imbens, and Rubin 1996). Intuitively, before mitigating the effect of unobserved confounders (Lee and Lemieux 2010), we need first to determine whether such confounders exist, which is the main focus of this paper. Thus, detecting the presence of unobserved confounders plays a foundational role in guiding subsequent methodological decisions.

Existing methods for detecting unobserved confounding can be broadly divided into two categories. The first one requires linear model assumptions, which imposes linear functional relationships among all variables (Schultheiss, Bühlmann, and Yuan 2024; Janzing and Schölkopf 2018; Liu and Chan 2018). While straightforward, these assumptions are not testable and severely hinder detection performance on nonlinear data. The second approach requires multiple environments with mechanism-varying Structural Causal Models (SCMs), leveraging causal mechanism shifts to detect confounders via propagation signatures (Karlsson and Krijthe 2023; Mameche, Vreeken, and Kaltenpoth 2024; Reddy and N Balasubramanian 2024). However, in practice, it is challenging to obtain sufficiently heterogeneous environments satisfying the required causal invariance and mechanism independence assumptions. Note that both these two types of methods make assumptions on the data structure, which raises an important open question:

Can unobserved confounders be detected in nonlinear settings within a single data environment?

To this end, we propose an effective method named “**Kernel Regression Confounder Detection**” (KRCD), which detects unobserved confounders on the nonlinear observational data drawn from a single environment. Specifically, the proposed KRCD method consists of three main steps: nonlinear structure modeling, kernelized regression, and hypothesis construction. To characterize the non-linear data structure, our KRCD first maps the treatment variable (T), the outcome variable (Y) and observed covariates (X) to reproducing kernel Hilbert space (RKHS), enabling a linear inner-product representation of the under-

*These authors contributed equally.

†Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

lying nonlinear structural functions. Subsequently, the proposed KRCD regresses Y on (T, X) in the RKHS twice: the first is a standard kernelized regression, while the second is a higher-order regression weighted by $\|T, X\|^2$, designed to capture nonlinear dependencies indicative of unobserved confounding. Finally, we construct the hypothesis test, which rejects the presence of unobserved confounders if the normalized difference between the two regression coefficients becomes sufficiently small. The main contributions of this paper can be summarized as follows:

- We investigate the problem of detecting unobserved confounders from nonlinear data in a single-environment setting. To address this, we propose KRCD, a novel method that compares standard kernel regression and higher-order kernel regression in RKHS, enabling efficient detection of unobserved confounding.
- Theoretically, we show that standard kernel regression and higher-order kernel regression coefficients coincide if and only if no unobserved confounders exist in infinite samples. For finite samples, their difference asymptotically converges to a zero-mean Gaussian distribution with tractable asymptotic variance.
- Through extensive experiments on synthetic benchmark datasets and the Twins dataset, our method stably outperforms the baselines in detecting unobserved confounders.

2 Related Work

2.1 Detection of Unobserved Confounders in Single-Environment Settings

Traditional approaches to detect unobserved confounding—including instrumental variable (IV) regression (Angrist, Imbens, and Rubin 1996) and the Durbin-Wu-Hausman test (Hausman 1978)—rely on external instruments or estimator comparisons (Wang et al. 2022; Kou et al. 2025). Subsequent innovations like ROCD (Schultheiss, Bühlmann, and Yuan 2024) derived from reweighting strategies (Buja et al. 2019a,b) bypass instrument requirements by leveraging non-Gaussian higher-order moments, offering a resampling-free OLS/HOLS comparison framework (Wang et al. 2023a). For high-dimensional linear models, spectral methods such as (Liu and Chan 2018; Janzing and Schölkopf 2018) detect confounding through spectral analysis of regression-induced associations or deviations in covariance structures, with (Janzing and Schölkopf 2018) explicitly linking confounding to multivariate linear model perturbations. Beyond linearity, early efforts to address nonlinear confounding relied on strong structural assumptions (e.g., nonparametric IV (Newey and Powell 2003)) or specific non-Gaussian distributions (e.g., LiNGAM (Shimizu et al. 2006)), yet lacked general nonparametric tests for unobserved confounders. While these methods advance instrument-free detection, they remain fundamentally constrained to linear models—a limitation overcome by our KRCD method, which extends confounder detection to nonlinear relationships.

2.2 Detection of Unobserved Confounders in Multi-Environment Settings

Beyond linear constraints, recent methods exploit causal mechanism shifts across environments to detect unobserved confounding without causal sufficiency assumptions. Theoretical work established testable conditional independence violations indicating confounding under shifts (Karlsson and Krijthe 2023). Subsequent studies developed operational frameworks: information-theoretic detection measures (Mameche, Vreeken, and Kaltenpoth 2024) and nonparametric quantification of confounding strength via mechanism shifts (Reddy and N Balasubramanian 2024). These approaches treat environmental variations as natural experiments to derive confounding signals (Yin et al. 2023; Wang 2025; Wang et al. 2023b). Joint causal inference (Mooij, Magliacane, and Claassen 2020) and causal discovery under heterogeneity (Perry, Von Kügelgen, and Schölkopf 2022) also facilitate confounder detection. However, such methods typically require explicit distribution shifts or auxiliary experimental data (Kallus, Puli, and Shalit 2018; Hu et al. 2024), limiting their use in purely observational, single-context settings. In contrast, KRCD achieves robust confounding detection in single-environment observational data while maintaining nonparametric flexibility, eliminating the need for multi-source data.

2.3 Kernelized Causal Inference

The theoretical foundation of Reproducing Kernel Hilbert Spaces (RKHS) was established by Aronszajn’s seminal work on the reproducing property (Aronszajn 1950), with the Representer Theorem (Kimeldorf and Wahba 1971) enabling efficient regularization via kernel expansions. This framework supports three key causal strategies: (1) Kernel mean embeddings (KME) (Hazlett 2016; Muandet et al. 2017) for covariate balancing assume no unobserved confounders; (2) RKHS-based IV regression (Singh, Sahani, and Gretton 2019; Mastouri et al. 2021) handles observed confounding but requires external instruments; (3) Latent state identification (Zhang, Gong, and Schölkopf 2015; Pfister et al. 2018) infers unmeasured variables without explicit confounder detection. Despite enhancing nonparametric flexibility, these methods remain limited by causal sufficiency assumptions or auxiliary data dependencies.

Critically, no existing RKHS approach directly detects nonlinear unobserved confounding in standard observational data. Current techniques either ignore latent biases (KME), depend on untestable instruments (kernel IV), or lack bias quantification (independence tests) (Li et al. 2023; Wang et al. 2025a,b). To address this gap, we unify RKHS theory with the ROCD framework (Schultheiss, Bühlmann, and Yuan 2024) through a novel Representer Theorem variant. This unification enables data-adaptive kernel bases for nonlinear confounding detection, forming our KRCD method. KRCD pioneers direct detection of nonlinear unobserved confounders by embedding ROCD’s higher-order moment structures into RKHS—requiring neither instruments nor multi-environment data.

3 Preliminaries and Problem Setup

In this paper, we denote the treatment, observed covariates, outcome, and the unobserved confounder as T , X , Y and U , respectively. We follow the potential outcome framework and characterize the potential outcome of Y under the assignment $T = t$ as $Y(t)$. We denote the random variables by uppercase letters (e.g., T and Y) and their realizations by lowercase letters (e.g., t and y).

3.1 Previous Problem Setup

Existing approaches for detecting unobserved confounders bifurcate into two constrained paradigms:

- **Linear constraints:** Single-environment methods universally require parametric functional assumptions to achieve identifiability. This precludes their application to general nonlinear systems with complex dependencies.
- **Multi-environment dependency:** Multi-environment methods universally require observable causal mechanism shifts in SCMs. This fundamentally precludes application to single-environment observational studies.

3.2 Our Problem: Detecting Unobserved Confounders in Nonlinear Single-Environment (NSE)

We aim to *detect the existence of unobserved confounders U in nonlinear single-environment settings*. Under the null hypothesis H_0 , no unobserved confounders exist, specified by the structural equations:

$$H_0 : \begin{cases} T = f_T(X) + \epsilon_T, \\ Y = g(T, X) + \epsilon_Y. \end{cases} \quad (1)$$

Conversely, under the alternative hypothesis H_1 , unobserved confounders are present, characterized by:

$$H_1 : \begin{cases} T = f_T(X) + h_T(U) + \epsilon_T, \\ Y = g(T, X) + h_Y(U) + \epsilon_Y, \end{cases} \quad (2)$$

where f_T , g , h_T , and h_Y are nonlinear functions, ϵ_T and ϵ_Y are exogenous noise terms. For subsequent analysis, we define $Z = (T, X)^\top$ with $\mathbb{E}[\epsilon_T] = \mathbb{E}[\epsilon_Y] = 0$.

Remark 1 (Distinction from Previous Problem Setups). *Our proposed method differs fundamentally from prior works in two key aspects:*

(1) **Nonparametric formulation:** *We do not assume parametric functional forms (e.g., linearity) for f_T , g , h_T , h_Y , or distributions of ϵ_T , ϵ_Y , making our approach applicable to general nonlinear systems.*

(2) **Single-environment setting:** *Unlike methods relying on multiple environments (e.g., stochastic mechanism shifts, or explicit interventions), we address the more challenging scenario where only observational data from a single environment is available.*

This positions our work as the first nonparametric detection for unobserved confounders using single-environment observational data.

4 Kernel Regression for Confounder Detection

In this section, we propose **Kernel Regression for Confounder Detection (KRCD)**, a novel and powerful method that leverages Reproducing Kernel Hilbert Space (RKHS) theory to extend the Residual Orthogonality Condition for Diagnostics (ROCD) framework for nonparametric detection of unobserved nonlinear confounding in high-dimensional settings. This approach significantly enhances the diagnostic power of moment-based methods beyond linear relationships.

4.1 Motivation

In the NSE setting with potential nonlinear relationships, conventional linear detection methods face fundamental limitations. This motivates our reexamination of the ROCD framework's core premise:

Regression-Oriented Confounder Detection (ROCD). ROCD exploits the asymptotic equivalence of OLS and Higher-order least squares (HOLS) estimators under correct linear specification. Systematic differences indicate misspecification (e.g., unobserved confounding). The method:

- Computes OLS and HOLs coefficients (Y on X, T).
- Tests coefficient equality via standardized differences; significance implies unobserved confounders.

However, ROCD is *limited to linear regression*, failing for nonlinear data. To extend this paradigm to nonlinear settings, we define two types of least squares estimators for the nonlinear function $g(Z)$:

$$g^{OLS} = \operatorname{argmin}_{g \in \mathcal{H}} \mathbb{E} \left[\|Y - g(Z)\|^2 \right], \quad (3)$$

$$g^{HOLS} = \operatorname{argmin}_{g \in \mathcal{H}} \mathbb{E} \left[\|Y - g(Z)\|^2 \|Z\|^2 \right], \quad (4)$$

where \mathcal{H} is an appropriate function space.

Challenges. In linear settings where g resides in finite-dimensional Euclidean space, g^{OLS} and g^{HOLS} admit explicit vector representations. However, for nonlinear g in infinite-dimensional function spaces, representing g for comparative analysis and downstream detection remains a fundamental theoretical and computational challenge.

4.2 Representer for Functional Regression

To address this challenge, we leverage the Representer Theorem in Reproducing Kernel Hilbert Spaces. This approach transforms the infinite-dimensional optimization problem into a finite-dimensional convex optimization over coefficients, enabling the capture of complex nonlinear dependencies indicative of unobserved confounding.

Reproducing Kernel Hilbert Space (RKHS). (Aronszajn 1950) A RKHS \mathcal{H}_k is a function space with reproducing kernel k , satisfying $\langle f, k(\cdot, x) \rangle_{\mathcal{H}_k} = f(x)$. Its feature map $\phi(x) = k(\cdot, x)$ embeds inputs into \mathcal{H}_k .

Representer Theorem. (Kimeldorf and Wahba 1971) For g in RKHS \mathcal{H}_k induced by kernel k , the minimizer of a regularized empirical risk admits the form: $\hat{g}(\cdot) = \sum_{i=1}^N \alpha_i \phi_i(\cdot)$,

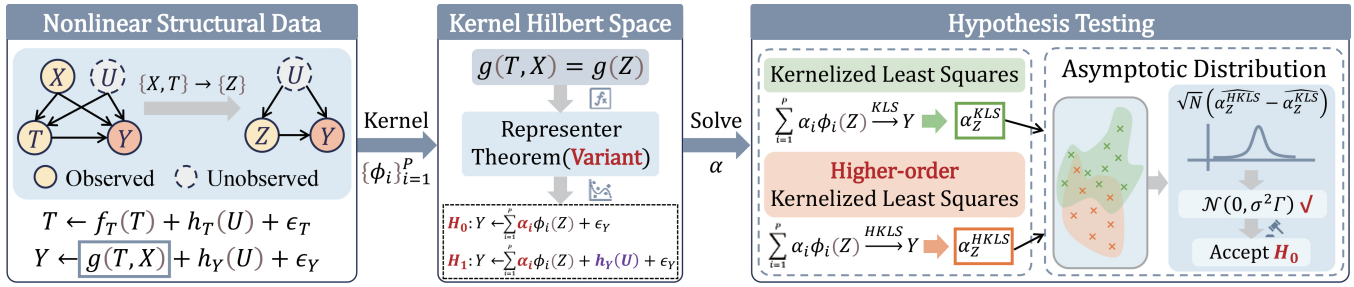


Figure 1: Workflow for Proposed Method (KRCD).

where $\phi_i(\cdot) = k(x_i, \cdot)$, $\{x_i\}_{i=1}^N$ are training sample points and $\alpha_i \in \mathbb{R}$ are coefficients.

Unfortunately, we observe that *such representer theorem cannot be directly adapt to our detection problem*, as the dimensionality of $\alpha \in \mathbb{R}^N$ scales with sample size N . Consequently, with significantly enlarging N (very often in realistic analysis), the computational tractability of the designed detection method (i.e., the complexity) by checking two representers of the regression coefficients $\hat{g}(\cdot) \in \mathcal{R}^N$ becomes prohibitively infeasible.

To fill this gap, we establish a novel variant of the Representer Theorem operating in a subspace with fixed and finite dimensional $P < N$ (P is the feature size):

Theorem 1 (Variant of Representer Theorem). *Let $X \in \mathbb{R}^{N \times d}$ be the design matrix with rank $r < N$, and let $\{\phi_i\}_{i=1}^P$ ($r \leq P < N$) form a basis for a P -dimensional subspace $\mathcal{S}_P \subseteq \mathcal{H}_k$ (RKHS). Consider the regularized empirical risk minimization problem:*

$$\hat{g}(\cdot) = \operatorname{argmin}_{g \in \mathcal{S}_P} \left(\frac{1}{N} \sum_{i=1}^N \ell(y_i, g(\mathbf{x}_i)) + \xi \Omega(\|g\|_{\mathcal{H}_k}) \right),$$

with convex loss ℓ , convex non-decreasing regularizer Ω , and $\xi > 0$. Then the minimizer \hat{g} admits the representation:

$$\hat{g}(\cdot) = \sum_{i=1}^P \alpha_i \phi_i(\cdot), \quad (5)$$

and the coefficient vector $\alpha \in \mathbb{R}^P$ fully encodes the solution based on the sample.

Remark 2. *This theorem (see detailed proof in Appendix A.1) allows us to recast the infinite-dimensional functional estimation problems for g^{OLS} and g^{HOLS} into fixed, finite-dimensional optimization over $\alpha \in \mathbb{R}^P$.*

4.3 Theoretical Foundation

Leveraging Theorem 1, we formulate the Kernelized Least Squares (KLS) and Higher-order Kernelized Least Squares (HKLS) estimators within a P -dimensional subspace \mathcal{S}_P spanned by basis $\{\phi_i\}_{i=1}^P$ selected from the eigenfunctions of a symmetric positive definite kernel. Consequently, the functional regression results in (3) and (4) can be reduced to the P -dimensional α_Z^{KLS} and α_Z^{HKLS} solving the following minimization problems:

1. **Kernelized Least Squares (KLS)** minimizes the expected squared error loss:

$$\alpha_Z^{KLS} = \operatorname{argmin}_{\alpha \in \mathbb{R}^P} \mathbb{E} \left[\left\| Y - \sum_{i=1}^P \alpha_i \phi_i(Z) \right\|^2 \right]. \quad (6)$$

2. **Higher-order Kernelized Least Squares (HKLS)** minimizes the expected $\|Z\|^2$ -weighted squared error loss:

$$\alpha_Z^{HKLS} = \operatorname{argmin}_{\alpha \in \mathbb{R}^P} \mathbb{E} \left[\left\| Y - \sum_{i=1}^P \alpha_i \phi_i(Z) \right\|^2 \|Z\|^2 \right]. \quad (7)$$

Core Principle of our KRCD Method. Based on such reformulations, we are now ready to detect the non-linear, unobserved confounding behind the observations, by checking and comparing α_Z^{KLS} and α_Z^{HKLS} in each coordinate. To be specific, denoting the i -th coordinate of α_Z^{KLS} as $\alpha_Z^{KLS}[i]$, we prove that the coordinate-wise equality between $\alpha_Z^{KLS}[i]$ and $\alpha_Z^{HKLS}[i]$ is sufficient to inform the existence of unobserved confounding:

Theorem 2. *Under the causal model defined by Equations (1) and (2), the following holds regarding the equivalence of the population estimators:*

$$H_0 : \forall i \in [P], \quad \alpha_Z^{KLS}[i] = \alpha_Z^{HKLS}[i]$$

(No unobserved confounder U)

$$H_1 : \exists i \in [P], \quad \alpha_Z^{KLS}[i] \neq \alpha_Z^{HKLS}[i]$$

(Unobserved confounder U present)

Proofs of establishing equality under H_0 and inequality under H_1 are provided in Appendix A.2. In summary, this core theorem establishes the rigorous population-level theoretical foundation for KRCD: the systematic difference between α_Z^{KLS} and α_Z^{HKLS} serves as a valid statistical indicator of unobserved confounding presence.

4.4 Hypothesis Testing with Finite Samples

Although Theorem 2 provides theoretical guidance for unobserved confounding detection by comparing regression representers α_Z^{HKLS} and α_Z^{KLS} , a significant obstacle remains for finite-sample algorithmic implementation.

Specifically, the minimization problems defined in (6) and (7) assume infinite samples (i.e., population expectations \mathbb{E}), while finite-sample estimates of α_Z^{HKLS} and

α_Z^{KLS} may deviate from their population counterparts. Consequently, direct comparison of these coefficients with finite samples risks incorrect detection.

To overcome this issue, we first derive the closed-form solutions of empirical estimators $\widehat{\alpha}_Z^{KLS}$ and $\widehat{\alpha}_Z^{HKLS}$ based on empirical samples $\{(Z_j, Y_j)\}_{j=1}^N$. With the kernel basis $k(Z_i, \cdot)$ (setting $\phi_i(\cdot) = k(Z_i, \cdot)$, $\Omega(\|g\|_{\mathcal{H}_k}) = \sum_{i=1}^P \alpha_i^2$, and $P < N$), then we have:

Theorem 3. *The empirical KLS and HKLS estimators are given by:*

$$\begin{aligned}\widehat{\alpha}_Z^{KLS} &= (K_{ZZ}K_{ZZ}^\top + N\xi I_P)^{-1}K_{ZZ}\mathbf{Y}, \\ \widehat{\alpha}_Z^{HKLS} &= (K_{ZZ}\Psi K_{ZZ}^\top + N\xi I_P)^{-1}K_{ZZ}\Psi\mathbf{Y},\end{aligned}\quad (8)$$

where $K_{ZZ} \in \mathbb{R}^{P \times N}$ with $[K_{ZZ}]_{ij} = k(Z_i, Z_j)$, $\mathbf{Y} = [Y_1, \dots, Y_N]^\top$, and $\Psi = \text{diag}(\|Z_1\|^2, \dots, \|Z_N\|^2)$.

As mentioned before, diverse risk factors, e.g., sampling variability and outcome noise, will lead to perturbation of the difference vector $\hat{\delta} = \widehat{\alpha}_Z^{HKLS} - \widehat{\alpha}_Z^{KLS}$, served as the diagnostic marker for unobserved confounding. We therefore present the last theorem in below to indicate the asymptotic normality of $\hat{\delta}$, which further guides the construction of the hypothesis testing with finite samples:

Theorem 4 (Asymptotic Distribution of $\hat{\delta}$). *Assume standard regularity conditions (specified in Appendix A.3) hold, and the causal model (1), (2) is correct. Under H_0 ($\alpha^{KLS} = \alpha^{HKLS}$), as $N \rightarrow \infty$:*

$$\sqrt{N}\hat{\delta} = \sqrt{N}(\widehat{\alpha}_Z^{HKLS} - \widehat{\alpha}_Z^{KLS}) \xrightarrow{D} \mathcal{N}(\mathbf{0}, \sigma^2\Gamma), \quad (9)$$

where the asymptotic covariance matrix $\sigma^2\Gamma$ is consistently estimable.

This asymptotic characterization enables practical hypothesis testing for H_0 . Specifically, the established asymptotic normality allows constructing component-wise standardized tests whose null distributions are approximately standard normal. Accounting for multiple testing via Bonferroni correction, the null hypothesis of no unobserved confounding is rejected if any standardized component difference shows statistically significant deviation at the adjusted significance level.

4.5 Complexity Analysis

We derive the overall complexity of our proposed KRCD method as $O(N^2d + P^3 + P^2N)$, in terms of the sample size N , parameter dimension P , and feature dimension d (detailed derivation provided in Appendix A.4).

Algorithm 1: Kernel Regression Confounder Detection

Input: Observed covariates $X \in \mathbb{R}^{N \times d}$, treatment $T \in \mathbb{R}^N$, outcome $Y \in \mathbb{R}^N$, parameter dimension P , regularization parameter $\lambda = N\xi$.

Output: Detect conclusion

```

1:  $Z \leftarrow [T, X]$ 
2:  $K \leftarrow \text{rows}_{1:P}(\text{Kernel Function}(Z))$ 
3:  $K_\psi \leftarrow K \odot (\text{diag}(ZZ^\top))^\top$ 
4:  $\alpha_{KLS} \leftarrow (KK^\top + \lambda I)^{-1}KY$ 
5:  $\alpha_{HKLS} \leftarrow (K_\psi K^\top + \lambda I)^{-1}(K_\psi Y)$ 
6:  $V_0 \leftarrow (K_\psi K^\top + \lambda I)^{-1}K_\psi - (KK^\top + \lambda I)^{-1}K$ 
7:  $V \leftarrow N \cdot (V_0 V_0^\top)$ 
8:  $\sigma^2 \leftarrow \frac{1}{N}(Y - K^\top \alpha_{KLS})^\top (Y - K^\top \alpha_{KLS})$ 
9: for  $j = 1$  to  $P$  do
10:    $p_j \leftarrow 2(1 - \Phi(|\frac{\sqrt{N}(\alpha_{HKLS,j} - \alpha_{KLS,j})}{\sqrt{\sigma^2 V_{jj}}}|))$ 
11: end for
12: Apply Bonferroni correction:  $\delta_j \leftarrow \mathbb{1}[p_j < 0.05/P]$ 
13: if  $\sum \delta_j = 0$  then
14:   return "Support null hypothesis"
15: else
16:   return "Reject null hypothesis"
17: end if

```

5 Experiments

Benchmarks. To comprehensively evaluate KRCD's capability for unobserved confounding detection, we conducted simulation studies using both synthetic datasets and the Twins dataset (Almond, Chay, and Lee 2005). Unless explicitly stated otherwise, all experiments employed non-Gaussian independent noise components and a default regularization parameter of $\lambda = 10^{-8}$. Each parameter combination was repeated 30 times with sample sizes of 1000 observations, and statistical significance was assessed at the conventional $\alpha = 0.05$ threshold.

Baselines. This paper focuses on state-of-the-art (SOTA) confounding detection methods for comparative examination, including: ROCD (Schultheiss, Bühlmann, and Yuan 2024), ME-ICM (Karlsson and Krijthe 2023), and CNF (Reddy and N Balasubramanian 2024). (See details of baselines in Appendix B)

Evaluation Protocols. We evaluate the proposed KRCD in a variety of scenarios, including (see details of scenarios in Appendix C):

- *Nonlinear Single-Environment.* All data is generated by identical nonlinear structural causal mechanisms under a fixed causal model.
- *Nonlinear Multi-Environment.* Distinct environments exhibit different nonlinear causal mechanisms within a shared causal framework.

The evaluation metrics including:

- *ROC Curve & AUC.* Sensitivity and False Positive Rate over thresholds; AUC quantifies overall performance of this trade-off.

- *Detection Rate.* The method’s detection rate for unobserved confounders, quantifying sensitivity (with confounders) and false positive rate (without confounders).
- *Runtime.* Total time required to complete all detection runs at a fixed confounding strength, benchmarking computational efficiency.

Questions. The empirical evaluation addresses the following three research questions:

- Whether the proposed KRCD method outperforms detection baselines in our NSE problem?
- Whether the proposed KRCD method outperforms detection baselines in multi-environment settings?
- Can KRCD achieve superior computational efficiency in detecting hidden confounding, outperforming baselines?
- Can KRCD exhibit robustness when the grounding structural function, fail to fall inside the kernel space?
- How sensitive is KRCD’s performance to changes in the regularization parameter?

5.1 Performance in Our NSE Setup

To validate KRCD’s performance in our NSE setup, we performed systematic experiments on the Twins dataset where all environments share identical causal mechanisms.

As shown in Figure 2a, KRCD remains robust under null confounding and shows high sensitivity once confounding strength reaches a threshold (≥ 0.5). Figure 2c indicates that KRCD detects unobserved confounding with only 200 samples, highlighting its sample efficiency.

Figure 2b compares single-environment nonlinear confounding detection. ME-ICM, requiring multiple environments, fails completely here. ROCD, based on linearity assumptions, misinterprets nonlinearity as confounding, yielding high false positives under null confounding and declining detection under strong confounding as true confounders mask nonlinear features. In contrast, KRCD maintains 0% detection under null confounding (controlling Type I error) and exceeds 95% detection at confounding strengths ≥ 1 .

These results confirm KRCD’s strong single-environment performance: robustness to null confounding, high sensitivity beyond a strength threshold, low sample requirements, and clear advantages over methods needing linear assumptions or multi-environment data.

5.2 Performance in Multi-Environment

To validate KRCD’s performance in multi-environment settings, we conducted systematic experiments where environments exhibit distribution shifts in causal mechanisms while maintaining consistent causal structures. Due to divergent environmental definitions across baselines, evaluations were performed separately on Twins and synthetic datasets.

Comparison with ME-ICM. On the Twins dataset, KRCD fundamentally differs by integrating environmental variables as covariates rather than using ME-ICM’s explicit partitioning. Figure 3a shows both methods achieve comparable performance, demonstrating KRCD’s competitive multi-environment adaptability.

Dataset	Method	Runtime(s/it)
Twins	KRCD	68.748
	ME-ICM	2657.838
Synthetic	KRCD	0.184
	CNF-1	0.072
	CNF-2	0.048
	CNF-3	0.046

Table 1: Comparison of Runtime for KRCD and Baselines

Comparison with CNF. On the synthetic dataset, KRCD outperforms CNF in detecting unobserved confounders, particularly under weak confounding, as evidenced by Figure 3b. This performance gap occurs because CNF’s blocking approach, while computationally efficient, loses information and reduces sensitivity to weaker confounders.

5.3 Computational Efficiency Analysis

To evaluate KRCD’s computational efficiency, we compared runtime of all baselines in experiments for Figure 3a and 3b, defined as the time for 100 trials. Table 1 shows that on the Twins dataset, KRCD is substantially faster than ME-ICM, markedly reducing computational time. On the Synthetic dataset, KRCD has efficiency similar to CNF.

Notably, KRCD’s significant efficiency variation across datasets stems directly from its parameter dimension P : Setting $P = 999$ on the Twins dataset prioritizes modeling precision, while $P = 40$ on the Synthetic dataset maximizes computational efficiency. This tunability enables KRCD to flexibly balance predictive performance against computational demands through P adjustment, ensuring adaptability to diverse resource constraints.

5.4 Robustness under Kernel Misspecification

To evaluate the computational efficiency of our KRCD method, we compared runtime across all baselines in the experiments for Figure 3a and 3b, defined as the time for 100 trials. Table 1 shows that on the Twins dataset, KRCD is substantially faster than ME-ICM, markedly reducing computational time. On the Synthetic dataset, KRCD has efficiency similar to CNF baselines.

As demonstrated in Figure 3c, both kernel configurations of KRCD maintain strict Type I error control (no false positives). Polynomial-KRCD achieves high sensitivity even with minimal confounding conditions, maintaining full detection as confounding strength increases. Gaussian-KRCD needs stronger signals but achieves full detection at high confounding strength. KRCD thus provides robust, reliably detects unobserved confounding without false alarms and excels in high-confounding scenarios.

5.5 Parameter Sensitivity Analysis

The regularization parameter λ in Algorithm 1 prevents overfitting and ensures numerical stability during kernel matrix inversion. As Table 2 shows, KRCD exhibits minimal sensitivity to λ across wide parameter ranges (10^{-12} to




Figure 2: KRCD Performance in Nonlinear Single-Environment—(a) ROC Curves across Confounding Strengths. (b) Detection Rate of KRCD, ROCD, and ME-ICM. (c) Detection Rate across Confounding Strengths and Sample Sizes.




Figure 3: (a) Detection Rate of KRCD and ME-ICM in Nonlinear Multi-Environment. (b) Detection Rate of KRCD and CNF in Nonlinear Multi-Environment. (c) Performance of Polynomial-KRCD and Gaussian-KRCD.

Conf Strength	Parameter λ			
	10^{-12}	10^{-8}	10^{-4}	1
0.25	0.468	0.507	0.473	0.404
0.5	0.959	0.980	0.981	0.257
1	0.998	1.000	1.000	0.720
2	1.000	1.000	1.000	1.000

Table 2: KRCD AUC Performance Sensitivity to λ

10^{-4}), maintaining stable AUC performance. Within this broad operational window, the method consistently achieves high accuracy (≥ 0.95 AUC for confounding strengths ≥ 0.5), demonstrating significant robustness to parameter selection. Performance degradation occurs only under extreme regularization ($\lambda = 1$), where over-smoothing obscures essential confounding signals—particularly detrimental at medium confounding strength with clinically significant effects. This resilience to λ variation confirms KRCD’s practical usability in diverse settings.

6 Conclusion

This paper proposes the Kernel Regression Confounder Detection (KRCD) method, which can detect unobserved confounders in nonlinear single-environment observational data, addressing the limitations of existing methods. The core principle of KRCD involves using kernel regression within reproducing kernel Hilbert spaces to model complex relationships; it compares the results of standard kernel regression and higher-order kernel regression to construct a test statistic—a significant deviation of this statistic from zero indicates the presence of unobserved confounding. Experiments on synthetic benchmark datasets and the Twins dataset demonstrate that KRCD not only significantly outperforms baseline methods in detection performance but also achieves superior computational efficiency.

While KRCD effectively detects unobserved confounding in nonlinear single-environment settings, it cannot quantify confounding strength. A promising direction is developing a novel framework for quantifying confounding strength based on KRCD methodology.

Ethics Statement

This work does not raise any ethical concerns. All experiments are conducted on publicly available datasets, and no human subjects or sensitive attributes are involved.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China under Grants No. 62372459, 62525213, and 62376243, in part by the NUDT Youth Independent Innovation Science Fund under Grant No. ZK25-20, in part by the "Pioneer" and "Leading Goose" R&D Program of Zhejiang (2025C02037).

References

- Almond, D.; Chay, K. Y.; and Lee, D. S. 2005. The costs of low birth weight. *The Quarterly Journal of Economics*, 120(3): 1031–1083.
- Angrist, J. D.; Imbens, G. W.; and Rubin, D. B. 1996. Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434): 444–455.
- Aronszajn, N. 1950. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3): 337–404.
- Buja, A.; Brown, L.; Berk, R.; George, E.; Pitkin, E.; Traskin, M.; Zhang, K.; and Zhao, L. 2019a. Models as approximations I. *Statistical Science*, 34(4): 523–544.
- Buja, A.; Brown, L.; Kuchibhotla, A. K.; Berk, R.; George, E.; and Zhao, L. 2019b. Models as approximations II. *Statistical Science*, 34(4): 545–565.
- Dahabreh, I. J.; and Bibbins-Domingo, K. 2024. Causal inference about the effects of interventions from observational studies in medical journals. *Jama*, 331(21): 1845–1853.
- Hausman, J. A. 1978. Specification tests in econometrics. *Econometrica: Journal of the econometric society*, 1251–1271.
- Hazlett, C. 2016. Kernel Balancing: A flexible non-parametric weighting procedure for estimating causal effects. arXiv:1605.00155.
- Hu, K.; Yang, W.; Huang, W.; Zhou, X.; Cao, M.; Ren, J.; and Tan, H. 2024. Sequential fusion based multi-granularity consistency for space-time transformer tracking. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 12519–12527.
- Janzing, D.; and Schölkopf, B. 2018. Detecting confounding in multivariate linear models via spectral analysis. *Journal of Causal Inference*, 6(1): 20170013.
- Kallus, N.; Puli, A. M.; and Shalit, U. 2018. Removing hidden confounding by experimental grounding. *Advances in neural information processing systems*, 31.
- Karlsson, R.; and Krijthe, J. 2023. Detecting hidden confounding in observational data using multiple environments. *Advances in Neural Information Processing Systems*, 36: 44280–44309.
- Kimeldorf, G.; and Wahba, G. 1971. Some results on Tchebycheffian spline functions. *Journal of mathematical analysis and applications*, 33(1): 82–95.
- Kou, Z.; Qin, S.; Wang, H.; Xie, M.; Chen, S.; Jia, Y.; Liu, T.; Sugiyama, M.; and Geng, X. 2025. Label Distribution Learning with Biased Annotations by Learning Multi-Label Representation. *arXiv preprint arXiv:2502.01170*.
- Kühne, F.; Schomaker, M.; Stojkov, I.; Jahn, B.; Conrads-Frank, A.; Siebert, S.; Sroczynski, G.; Puntischer, S.; Schmid, D.; Schnell-Inderst, P.; et al. 2022. Causal evidence in health decision making: methodological approaches of causal inference and health decision science. *GMS German Medical Science*, 20: Doc12.
- Lee, D. S.; and Lemieux, T. 2010. Regression discontinuity designs in economics. *Journal of economic literature*, 48(2): 281–355.
- Li, H.; Dai, Q.; Li, Y.; Lyu, Y.; Dong, Z.; Zhou, X.-H.; and Wu, P. 2023. Multiple robust learning for recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 4417–4425.
- Listl, S.; Jürges, H.; and Watt, R. G. 2016. Causal inference from observational data. *Community dentistry and oral epidemiology*, 44(5): 409–415.
- Liu, F.; and Chan, L. 2018. Confounder detection in high-dimensional linear models using first moments of spectral measures. *Neural Computation*, 30(8): 2284–2318.
- Mameche, S.; Vreeken, J.; and Kaltenpoth, D. 2024. Identifying confounding from causal mechanism shifts. In *International Conference on Artificial Intelligence and Statistics*, 4897–4905. PMLR.
- Mastouri, A.; Zhu, Y.; Gultchin, L.; Korba, A.; Silva, R.; Kusner, M.; Gretton, A.; and Muandet, K. 2021. Proximal causal learning with kernels: Two-stage estimation and moment restriction. In *International conference on machine learning*, 7512–7523. PMLR.
- Matthay, E. C.; and Glymour, M. M. 2022. Causal inference challenges and new directions for epidemiologic research on the health effects of social policies. *Current Epidemiology Reports*, 9(1): 22–37.
- Mooij, J. M.; Magliacane, S.; and Claassen, T. 2020. Joint causal inference from multiple contexts. *Journal of machine learning research*, 21(99): 1–108.
- Muandet, K.; Fukumizu, K.; Sriperumbudur, B.; Schölkopf, B.; et al. 2017. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1-2): 1–141.
- Newey, W. K.; and Powell, J. L. 2003. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5): 1565–1578.
- Pearl, J. 2009. *Causality*. Cambridge university press.
- Perry, R.; Von Kügelgen, J.; and Schölkopf, B. 2022. Causal discovery in heterogeneous environments under the sparse mechanism shift hypothesis. *Advances in Neural Information Processing Systems*, 35: 10904–10917.
- Pfister, N.; Bühlmann, P.; Schölkopf, B.; and Peters, J. 2018. Kernel-based tests for joint independence. *Journal of the*

- Royal Statistical Society Series B: Statistical Methodology, 80(1): 5–31.
- Prosperi, M.; Guo, Y.; Sperrin, M.; Koopman, J. S.; Min, J. S.; He, X.; Rich, S.; Wang, M.; Buchan, I. E.; and Bian, J. 2020. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nature Machine Intelligence*, 2(7): 369–375.
- Reddy, A. G.; and N Balasubramanian, V. 2024. Detecting and measuring confounding using causal mechanism shifts. *Advances in Neural Information Processing Systems*, 37: 61677–61699.
- Schultheiss, C.; Bühlmann, P.; and Yuan, M. 2024. Higher-order least squares: assessing partial goodness of fit of linear causal models. *Journal of the American Statistical Association*, 119(546): 1019–1031.
- Shimizu, S.; Hoyer, P. O.; Hyvärinen, A.; Kerminen, A.; and Jordan, M. 2006. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10).
- Singh, R.; Sahani, M.; and Gretton, A. 2019. Kernel instrumental variable regression. *Advances in Neural Information Processing Systems*, 32.
- Vanderschueren, T. 2024. *Operational decision-making with machine learning and causal inference*. Ph.D. thesis, University of Antwerp.
- Wang, H.; Chen, Z.; Zhang, H.; Li, Z.; Pan, L.; Li, H.; and Gong, M. 2025a. Debiased Recommendation via Wasserstein Causal Balancing. *ACM Transactions on Information Systems*.
- Wang, H.; Kuang, K.; Chi, H.; Yang, L.; Geng, M.; Huang, W.; and Yang, W. 2023a. Treatment effect estimation with adjustment feature selection. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2290–2301.
- Wang, H.; Kuang, K.; Lan, L.; Wang, Z.; Huang, W.; Wu, F.; and Yang, W. 2023b. Out-of-distribution generalization with causal feature separation. *IEEE Transactions on Knowledge and Data Engineering*, 36(4): 1758–1772.
- Wang, H.; Yang, W.; Yang, L.; Wu, A.; Xu, L.; Ren, J.; Wu, F.; and Kuang, K. 2022. Estimating individualized causal effect with confounded instruments. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1857–1867.
- Wang, M.; Li, J.; Su, H.; Yin, N.; Yang, L.; and Li, S. 2024. GraphCL: Graph-based Clustering for Semi-Supervised Medical Image Segmentation. *arXiv preprint arXiv:2411.13147*.
- Wang, M.; Ren, W.; Zhang, Y.; Fan, Y.; Shi, D.; Jing, L.; and Yin, N. 2025b. Gaussian Mixture Model for Graph Domain Adaptation. In Kwok, J., ed., *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, 1963–1972. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Wang, M.-z. 2025. SimProF: A Simple Probabilistic Framework for Unsupervised Domain Adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 21153–21161.
- Yin, N.; Shen, L.; Wang, M.; Lan, L.; Ma, Z.; Chen, C.; Hua, X.-S.; and Luo, X. 2023. Coco: A coupled contrastive framework for unsupervised domain adaptive graph classification. In *International Conference on Machine Learning*, 40040–40053. PMLR.
- Zhang, K.; Gong, M.; and Schölkopf, B. 2015. Multi-source domain adaptation: A causal view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.