
Exploiting Independent Instruments: Identification and Distribution Generalization

Sorawit Saengkyongam¹ Leonard Henckel¹ Niklas Pfister¹ Jonas Peters¹

Abstract

Instrumental variable models allow us to identify a causal function between covariates X and a response Y , even in the presence of unobserved confounding. Most of the existing estimators assume that the error term in the response Y and the hidden confounders are uncorrelated with the instruments Z . This is often motivated by a graphical separation, an argument that also justifies independence. Positing an independence restriction, however, leads to strictly stronger identifiability results. We connect to the existing literature in econometrics and provide a practical method called HSIC-X for exploiting independence that can be combined with any gradient-based learning procedure. We see that even in identifiable settings, taking into account higher moments may yield better finite sample results. Furthermore, we exploit the independence for distribution generalization. We prove that the proposed estimator is invariant to distributional shifts on the instruments and worst-case optimal whenever these shifts are sufficiently strong. These results hold even in the under-identified case where the instruments are not sufficiently rich to identify the causal function.

1. Introduction

When estimating the causal function between a vector of covariates X and a response Y in the presence of unobserved confounding, standard regression procedures such as ordinary least squares (OLS) are even asymptotically biased. Instrumental variable approaches (Wright, 1928; Imbens & Angrist, 1994; Newey, 2013) exploit the existence of exogenous heterogeneity in the form of an instrumental

variable (IV) Z and estimate, under suitable conditions, the causal function consistently. Importantly, the errors in Y and the hidden confounders U should be uncorrelated with the instruments Z . Usually, this has to be argued for with background knowledge. Often this is done by assuming that the data generating process follows a structural causal model (SCM) (Pearl, 2009; Bongers et al., 2021) (so that the distribution is Markov with respect to the induced graph), and that Y and U are d -separated from Z in the graph obtained by removing the edge from X to Y (this is the case for the DAG at the beginning of Section 2, for example). In particular, this requires an argument that Z is not causing Y directly but only via X . But this argument does not only imply that the errors in Y and U are uncorrelated but also that they are independent. This independence comes with several benefits.

For example, even in settings where the causal function can be identified by classical approaches based on uncorrelatedness, it has been observed that the independence can be exploited to construct estimators that achieve the semiparametric efficiency bound, at least when the error distribution comes from a known, parametric family (Hansen et al., 2010). Furthermore, the independence restriction is stronger than uncorrelatedness and therefore yields stronger identifiability results, which has been reported in the field of econometrics (e.g., Imbens & Newey, 2009; Chesher, 2003). For example, even binary instruments may be able to identify nonlinear effects (Dunker et al., 2014; Torgovitsky, 2015; Loh, 2019).

In this work, we investigate the independence restriction in more detail: we add to existing identifiability results, provide methods for exploiting the restriction for finite data and analyse implications for distribution generalisation.

More precisely, in Section 2 we discuss the identifiability conditions in a general and simple form, list some of the existing identifiability results that they imply, add novel results to this list, and extend the framework to conditional IV.

We also provide a practical method that exploits the above independence for estimating causal effects from data. It relies on the Hilbert-Schmidt independence criterion (HSIC) (Gretton et al., 2008) which has become a widely used tool

¹Department of Mathematical Sciences, University of Copenhagen, Denmark. Correspondence to: Sorawit Saengkyongam <ss@math.ku.dk>.

for testing independence in a joint distribution. Equipped with a characteristic kernel (such as the Gaussian kernel), HSIC is positive and equals zero if and only if the considered distribution factorizes. We propose an easy-to-use method, called HSIC-X (‘X’ for ‘exogenous’), that minimizes HSIC directly. The underlying problem is non-convex but can be tackled using widely used methods from stochastic optimization. While theoretical guarantees are hard to obtain, it has recently been shown empirically that reliable optimization of HSIC seems possible, at least when considering the different problem of minimizing independence of residuals and predictors in a regression problem (Greenfeld & Shalit, 2020; Mooij et al., 2009). Furthermore, the optimization can be initialized at informative starting points, such as the two-stage least squares (2SLS) solution, or even the OLS solution. HSIC-X can be combined with any machine learning method for nonparametric regression that is optimized by (stochastic) gradient descent. The details of our method are described in Section 3.

Furthermore, the independence restriction can be exploited for distribution generalization. In Section 4, we construct an estimator HSIC-X-pen (‘pen’ for ‘penalization’) that is worst-case optimal in nonlinear settings under distributional shifts corresponding to interventions on Z . This is particularly interesting in underidentified settings, where the causal function cannot be identified from the data. Our work thereby adds to an increasing literature connecting distributional robustness and causal inference (e.g., Schölkopf et al., 2012; Rojas-Carulla et al., 2018; Magliacane et al., 2018; Rothenhäusler et al., 2021; Arjovsky et al., 2019; Christiansen et al., 2021; Pfister et al., 2021; Yuan et al., 2021; Krueger et al., 2021; Creager et al., 2021). As for HSIC-X, the estimator HSIC-X-pen is modular in that it can be used with any machine learning method for nonparametric regression.

Experiments on both simulated and real world data in Section 5 confirm that HSIC-X can exploit the improved identifiability guarantees and can be more efficient in finite samples if wrong solutions yield both second and higher order dependencies between the residuals and the instruments. The code for all the experiments is available at <https://github.com/sorawitj/HSIC-X>. All proofs are provided in Appendix A.

1.1. Further Related Work

Independence between Z and $Y - f(X)$ can also be characterized differently; e.g., it is equivalent to

$$\mathbb{E}[\eta(Z)\psi(Y - f(X))] = \mathbb{E}[\eta(Z)]\mathbb{E}[\psi(Y - f(X))]$$

for all η and ψ in a sufficiently rich class (such as bounded continuous functions). This suggests estimating the causal function by a generalized methods of moments (GMM) ap-


proach. Indeed, Poirier (2017) focuses on the derivation of optimal weights for such estimating equations and the asymptotic behaviour of the estimator. Dunker et al. (2014); Dunker (2021) phrase the problem as an inverse problem and derive convergence rates for the iteratively regularized Gauss-Newton method. The above works focus on theoretical advancements (e.g., the experiments are restricted to univariate settings). Considering the independence of residuals has also been suggested, by Peters et al. (2016, Section 5) but no practical method was provided. Powerful regression techniques and machine learning methods to measure dependence in IV settings exist (Hartford et al., 2017; Singh et al., 2019; Bennett et al., 2019; Muandet et al., 2020) but to the best of our knowledge, none of these methods exploit the independence restriction.

In summary, exploiting independence does not seem to have played a major role in practice. Arguably, one of the reasons is that independence restrictions are difficult to work with in theory and practice. E.g., choosing the class of functions ψ is non-trivial.

2. Identifiability from Higher Order Moments

Unless stated otherwise, we consider the following SCM

$$\begin{aligned} Z &:= \epsilon_Z \\ U &:= \epsilon_U \\ X &:= g(Z, U, \epsilon_X) \\ Y &:= f(X) + h(U, \epsilon_Y) \end{aligned}$$



where $(\epsilon_Z, \epsilon_U, \epsilon_X, \epsilon_Y) \sim Q$ are jointly independent noise variables, $Z \in \mathbb{R}^r$ are *instruments*, $U \in \mathbb{R}^q$ are unobserved variables, $X \in \mathbb{R}^d$ are *predictors* and $Y \in \mathbb{R}$ is a *response*. For simplicity, we additionally assume that $\mathbb{E}[\epsilon_Z] = \mathbb{E}[X] = \mathbb{E}[h(U, \epsilon_Y)] = 0$. We call a collection $M = (f, g, h, Q)$ an *IV model*, with $f \in \mathcal{F} \subseteq \{f_\diamond \mid f_\diamond : \mathbb{R}^d \rightarrow \mathbb{R}\}$ and \mathcal{F} a pre-specified function class. The collection of all IV models of this form is denoted by \mathcal{M} . Any IV model $M \in \mathcal{M}$ induces a distribution \mathbb{P}_M over the observed variables (X, Y, Z) . We denote the data generating IV model by $M^0 = (f^0, g^0, h^0, Q^0)$ (parts of the data are unobserved). We refer to the function f^0 as the *causal function*. We can use it to compute the causal effect of any treatment contrast, such as the average treatment effect $f(1) - f(0)$ in case of a binary X . We assume that the causal function $f(\cdot)$ only depends on X , i.e., is homogeneous, and thus we do not need to distinguish between local and global treatment effects (e.g. Imbens & Angrist, 1994). Unless stated otherwise, we also assume that f^0 can be parameterized such that $f^0(\cdot) = \phi(\cdot)^\top \theta^0$ for some $\theta^0 \in \Theta \subseteq \mathbb{R}^p$, with a (known) function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ and, for simplicity,

we assume that $\theta \neq \theta^0$ implies $\phi(\cdot)^\top \theta \neq \phi(\cdot)^\top \theta^0$. In practice the true basis ϕ does not need to be known and can be approximated by a sufficiently flexible function approximator (e.g., neural networks), see Section 3.

Let us assume that we observe n i.i.d. observations from $(X, Y, Z) \sim \mathbb{P}_{M^0}$. In this paper, we consider two problems: estimating the causal function f^0 and predicting the response Y under distributional shifts on the variables Z (see Section 4). In order to solve either of these tasks, we make use of (parts of) the causal function f^0 . We therefore require that it is uniquely determined by the observed distribution \mathbb{P}_{M^0} (we relax this condition in Section 4). This is often termed *identifiability* of the causal function.

In a classical IV approach, identification of f^0 (or, equivalently, θ^0) is based on the *moment restriction*

$$\mathbb{E}[\eta(Z)(Y - \phi(X)^\top \theta)] = 0, \quad (1)$$

where $\eta : \mathbb{R}^r \rightarrow \mathbb{R}^k$ is a known function. Under the IV model M^0 this condition is satisfied for θ^0 . A sufficient and necessary condition for identifiability based on (1) is given by the (*moment*) *identifiability condition*

$$\mathbb{E}[\eta(Z)\phi(X)^\top] \tau = 0 \implies \tau = 0. \quad (2)$$

This condition is sometimes called the *rank condition* because it is equivalent to the matrix $\mathbb{E}[\eta(Z)\phi(X)^\top]$ having rank p , which, in particular, implies that $k \geq p$ (e.g., Wooldridge, 2010).

The moment restriction (1) aims to detect mean shifts of the residuals when varying the value of Z . For example, (1) can only identify θ^0 if we do not have for all $i \in \{1, \dots, p\}$ $\mathbb{E}[\phi_i(X)|Z] = 0$ almost surely (otherwise the function $x \mapsto f^0(x) + \phi_i(x)$ solves (1), too). But even if there are no mean shifts and (1) is not powerful enough to identify f^0 , we may still be able to identify f^0 from a stronger condition. To this end, we consider the *independence restriction*

$$Z \perp\!\!\!\perp Y - \phi(X)^\top \theta. \quad (3)$$

In this paper, we consider identifying f^0 based on (3), rather than (1). Under the IV model M^0 described above, this condition is satisfied for the true parameter θ^0 . A necessary and sufficient condition for identifiability using (3) is

$$Z \perp\!\!\!\perp h(U, \epsilon_Y) + \phi(X)^\top \tau \implies \tau = 0. \quad (4)$$

This, in particular, implies that for all $i \in \{1, \dots, p\}$ we have $Z \not\perp\!\!\!\perp h(U, \epsilon_Y) + \phi_i(X)$ (since, otherwise $\tau = e_i$ would violate the above condition).

If condition (4) is violated, condition (2) is violated, too, but the implication does not hold in the other direction. As a result, the independence restriction yields strictly stronger identifiability results (see also Table 1).

Proposition 2.1 (Identifiability based on independence). *Consider an IV model M^0 and assume $f^0(\cdot) = \phi(\cdot)^\top \theta^0$ for some $\theta^0 \in \mathbb{R}^p$. Then, the following statements hold.*

- (i) *If θ^0 is identifiable from the moment restriction (1) it is also identifiable from the independence restriction (3).*
- (ii) *There exist IV models such that θ^0 is identifiable from the independence restriction (3) but not from (1).*
- (iii) *In particular, there are examples of the following type that satisfy the conditions of (ii). (a) ‘Nonadditive Z ’:*

$$\begin{aligned} Z &:= \epsilon_Z & U &:= \epsilon_U & X &:= ZU + \epsilon_X \\ Y &:= X + U + \epsilon_Y, \end{aligned} \quad (5)$$

with $\mathcal{F} = \{f \mid f(x) = \theta x\}$, where $(\epsilon_Z, \epsilon_U, \epsilon_X, \epsilon_Y)$ are jointly independent standard Gaussian variables. (b) ‘Binary Z ’: $Z \in \{0, 1\}$, $X \in \mathbb{R}$, $p > 2$ (i.e., \mathcal{F} contains nonlinear functions). (c) ‘Independent Z ’: $Z \perp\!\!\!\perp X$.

Statements (i) and (ii) are known (e.g., Imbens & Newey, 2009; Chesher, 2003) but for completeness we nevertheless include their proofs in Appendix A. Intuitively, the causal function from the SCM (5) in (a) is not identifiable from the moment restriction because the instrument (or any transformation $\eta(Z)$) does not correlate with the mean of X , i.e., $\mathbb{E}[\eta(Z)X] = \mathbb{E}[\eta(Z) \mathbb{E}[X|Z]] = 0$. Therefore, for any $\gamma \in \mathbb{R}$, the shifted causal function $f(x) := f^0(x) + \gamma x$, for $x \in \mathbb{R}$, also satisfies the moment restriction. In the proof of Proposition 2.1, we argue that identifiability using (2) is impossible for an example of type (b). It may come as a surprise that it is indeed possible to identify nonlinear functions even if the instrument Z has a discrete support with small cardinality. This observation has been reported, e.g., by Dunker et al. (2014); Torgovitsky (2015); Loh (2019).

Even in cases where the causal function is identifiable from both (4) and (2), it may be beneficial to consider the independence restriction. This is, for example, the case when the effect of the instrument can be seen in both the conditional mean of $\phi(X)$, given Z , and in higher moments. In our simulation experiments in Section 5, we see that taking into account dependencies in higher moments may yield better statistical performance.

Condition (4) depends on the unknown $h(U, \epsilon_Y)$; at first glance, one might believe that the hidden term can be dropped from condition (4) but this is not the case.

Proposition 2.2. *Consider the IV model M^0 and assume that $\phi(\cdot) = (\phi_1(\cdot), \dots, \phi_p(\cdot))$ is a collection of basis functions such that $f^0(\cdot) = \phi(\cdot)^\top \theta^0$ for some $\theta^0 \in \mathbb{R}^p$. Then, in general, condition (4) does not imply*

$$Z \perp\!\!\!\perp \phi(X)^\top \tau \implies \tau = 0. \quad (6)$$

Furthermore, (6) does not imply (4), either.

Table 1. Overview of some of the identifiability results, described in Section 2.

$ \text{supp}(Z) $	Z acts on mean of X	Z acts on higher orders of X	identif. with (1) possible	identif. with (3) possible	Comments
∞	yes	yes/no	✓	✓	gain efficiency with (3), cf Sec. 5.1
∞	no	yes	✗	✓	cf Prop 2.1-iii-a
$m < \infty$	yes	yes/no	✓	✓	identif. $\leq m$ par. with (1), gain efficiency & identif. with (3), cf Prop 2.1-iii-b, Sec 5.1
$m < \infty$	no	yes	✗	✓	cf example in proof of Prop 2.2

Identifiability condition (2) depends only on the joint distribution of (Z, X) (which can be estimated from data), but we do not consider this an advantage: in practice, it is usually desirable to consider empirical relaxations of the identifiability conditions and output the set of θ 's that (approximately) satisfy the empirical version of (1) or (3), respectively. For example, one can invert statistical tests to construct confidence sets, see Section 3.4.

2.1. Conditional Instrumental Variables

In some applications, such as the one we consider in Section 5.3, restrictions (1) and (3) may be violated for the causal parameter, e.g., because there are confounding variables W between the instruments Z and the response Y . Under certain assumptions, however, the framework of conditional IV (CIV) (Frölich, 2007; Newey, 2013) still allows us to identify the causal function f^0 from the observed distribution. Both the identifiability point of view and the methodology we develop in Section 3 can be extended to CIV. More details are provided in Appendix B. Consider the following SCM.


$$W := m(\epsilon_W, V, U)$$

$$Z := q(W, V, \epsilon_Z)$$

$$X := g(Z, W, U, \epsilon_X)$$

$$V := \epsilon_V, U := \epsilon_U$$

$$Y := f(X) + h(W, U, \epsilon_Y)$$



where $(\epsilon_W, \epsilon_Z, \epsilon_V, \epsilon_U, \epsilon_X, \epsilon_Y)$ are jointly independent and $W \in \mathbb{R}^t$ are observed covariates. We assume that there is no edge from V to W or no edge from U to W , i.e., either Z and W or W and Y are not confounded.

Due to the unobserved confounding, (3) may not hold but as we assume that there is either no confounding between Z and W or between W and Y , we can instead use the *conditional independence restriction*

$$Z \perp\!\!\!\perp Y - \phi(X)^\top \theta \mid W \quad (7)$$

to identify f^0 . A corresponding sufficient and necessary

identifiability condition for f^0 is

$$Z \perp\!\!\!\perp h(W, U, \epsilon_Y) + \phi(X)^\top \tau \mid W \implies \tau = 0. \quad (8)$$

We can estimate f^0 using a loss that is minimized if restriction (7) is satisfied, e.g., using a conditional independence measure (e.g. Fukumizu et al., 2008; Zhang et al., 2011; Berrett & Samworth, 2019; Shah & Peters, 2020). In Appendix B we discuss cases, in which we can avoid using a conditional independence restriction.

3. Independence-based IV with HSIC

Starting from the independence restriction (3), our goal is to find a function \hat{f} such that the residuals $R^{\hat{f}} := Y - \hat{f}(X)$ are independent of the instruments Z . In the identifiable case, that is, if condition (4) is satisfied, only the causal function f^0 achieves independence. Thus, given an i.i.d. sample $(x_i, y_i, z_i)_{i=1}^n$ of the variables (X, Y, Z) , our method aims to find a function \hat{f} that minimizes the dependency between the residuals $(r_i^{\hat{f}})_{i=1}^n$, with $r_i^{\hat{f}} := y_i - \hat{f}(x_i)$, and the instruments $(z_i)_{i=1}^n$. In this work, we measure dependency using HSIC (Gretton et al., 2008). When using characteristic kernels (Fukumizu et al., 2008), this measure equals zero if and only if the considered joint distribution factorizes. HSIC has been used for optimization problems before (e.g., Greenfeld & Shalit, 2020; Mooij et al., 2009) and satisfies the conditions used to prove consistency (see Section 3.3) but other choices of independence measures are possible, too. Specifically, we consider the following learning problem:

$$\hat{f} := \arg \min_{f \in \mathcal{F}} \widehat{\text{HSIC}}((r_i^f, z_i)_{i=1}^n; k_{R^f}, k_Z), \quad (9)$$

where $\widehat{\text{HSIC}}((r_i^f, z_i)_{i=1}^n; k_{R^f}, k_Z) := \text{tr}(KHLH)$ is a consistent estimator of $\text{HSIC}((R^f, Z); k_{R^f}, k_Z)$ (Gretton et al., 2008); here, $K_{ij} = k_{R^f}(r_i^f, r_j^f)$ and $L_{ij} = k_Z(z_i, z_j)$ are the kernel matrices for the residuals R^f and the instruments Z with positive definite kernels k_{R^f} and k_Z , respectively, and $H_{ij} := \delta_{ij} - \frac{1}{n}$ is the centering matrix. We call the estimator in (9) HSIC-X.

The independence restriction (3) allows us to learn the causal function f^0 up to a bias term (for any $\alpha \in \mathbb{R}$, $\text{HSIC}((Y - f^0(X), Z); k_{R^{f^0}}, k_Z)$ and $\text{HSIC}((Y - \alpha -$

$f^0(X, Z); k_{Rf^0}, k_Z$) are identical). Nonetheless, we can correct for the bias by using the zero mean assumption of the noise ϵ_Y . The final estimate is then obtained as $\tilde{f}(\cdot) := \hat{f}(\cdot) - \frac{1}{n} \sum_{i=1}^n (y_i - \hat{f}(x_i))$.

3.1. Regularizing towards Predictive Functions

In many practical applications the identifiability condition (4) (or (2) for classical IV) is satisfied but many parameters approximately solve the empirical version of (3). This is the case, for example, if the influence of the instruments is weak, which usually yields subpar finite sample properties. Furthermore, classical estimators like 2SLS are known to only have moments up to the degree of over-identification (e.g., Mariano, 2001). To stabilize the estimation it has been proposed to regularize towards a predictive function, such as the OLS in linear settings, see, e.g., the K-class estimators (Theil, 1958; Jakobsen & Peters, 2022), which contain the OLS, 2SLS, the FULLER (Fuller, 1977) and LIML estimators (Anderson & Rubin, 1949) as special cases.

We propose an analogous regularization for our estimator and call this variant HSIC-X-pen. More specifically, for a convex loss function $\ell : \mathbb{R} \rightarrow \mathbb{R}$ we modify the optimization problem (9) as follows,

$$\hat{f}^\lambda = \arg \min_{f \in \mathcal{F}} \widehat{\text{HSIC}}((r_i^f, z_i)_{i=1}^n; k_{Rf}, k_Z) + \lambda \sum_{i=1}^n \ell(y_i - f(x_i)), \quad (10)$$

where $\lambda \in [0, \infty)$ is a tuning parameter. Unlike in the linear settings described above (e.g., Fuller, 1977), deriving a data-driven choice of the tuning parameter λ is non-trivial. We propose to select λ following a procedure analogue to the one described in Jakobsen & Peters (2022): we select the largest possible value of λ for which an HSIC-based independence test (e.g., Gretton et al., 2008; Pfister et al., 2017) between $R^{\hat{f}^\lambda}$ and Z is not rejected.

As discussed in Section 4, HSIC-X-pen can be understood in relation to distribution generalization, too. There, one starts from the objective of optimizing the predictive loss and adds the HSIC term as a penalty that regularizes the predictor to guard against distributional shifts.

3.2. Algorithm and Implementation Details

We now specify the details of HSIC-X. To solve (9), we fix any parametric function class $\mathcal{F} := \{f_\theta(\cdot) \mid \theta \in \Theta \subseteq \mathbb{R}^p\}$ (e.g., a linear combination of some basis functions or a neural network) and optimize the parameters θ by a gradient-based optimization method. We choose the Gaussian kernel k_R (e.g., Schölkopf & Smola, 2002) for the residuals, and the discrete or Gaussian kernel k_Z for the instruments depending on whether Z is discrete or continuous, respectively. The bandwidth parameter σ of the Gaussian kernel is chosen

by the median heuristic (e.g., Sriperumbudur et al., 2009) and is recomputed during the optimization process (as the residuals change at each iteration).

Since the optimization problem (9) is generally non-convex, the resulting parameter estimates may not be the global optimal solution. We alleviate this problem by introducing the following restarting heuristic. Let $\hat{\theta} \in \Theta$ be a solution to the optimization problem. We conduct an independence test between the resulting residuals $(r_i^{\hat{\theta}} := y_i - f_{\hat{\theta}}(x_i))_{i=1}^n$ and the instruments $(z_i)_{i=1}^n$ using HSIC with Gamma approximation (Gretton et al., 2008). We accept the parameters $\hat{\theta}$ if the test is not rejected, otherwise we randomly re-initialize the parameters and restart the optimization. In the spirit of (10), we initialize the parameters in the first trial at the OLS solution. Algorithm 1 in Appendix C illustrates the whole optimization procedure with the standard gradient descent update. The gradient step can be replaced with other gradient-based optimization algorithms such as Adam (Kingma & Ba, 2015) or Adagrad (Duchi et al., 2011); in all of our experiments, we used Adam. The Algorithm for HSIC-X-pen is also provided in Appendix C.

3.3. Consistency

We now prove consistency of the proposed approach in that the minimizer of (9) converges (in probability) against the causal function, as sample size increases.¹

Theorem 3.1. *Consider the IV model M^0 , assume that $f^0(\cdot) = \phi(\cdot)^\top \theta^0$ for some bounded function $\phi(\cdot)$ and some $\theta^0 \in \Theta$, with $\Theta \subseteq \mathbb{R}^p$ being compact, and assume that the identifiability condition (4) holds. Consider the function class $\mathcal{F} := \{f(\cdot) = \phi(\cdot)^\top \theta \mid \theta \in \Theta\}$ and fixed bounded, continuously differentiable, characteristic kernels k_Z and k_R with bounded derivatives. Then, the estimator \hat{f} defined in (9) is consistent, i.e., $\|\hat{f} - f^0\|_\infty \xrightarrow{\mathbb{P}_{M^0}} 0$.*

The same statement holds if in (9) we replace HSIC and $\widehat{\text{HSIC}}$ by any independence measure $H(\mathbb{P}_{M^0}, \theta)$ and its estimate $\hat{H}_n(\mathcal{D}_n, \theta)$ (based on data \mathcal{D}_n), such that (i) $H(\mathbb{P}_{M^0}, \theta) = 0$ if and only if $(Y - \phi(X)^\top \theta) \perp\!\!\!\perp Z$, (ii) for all $\theta \in \Theta$, we have $\hat{H}_n(\mathcal{D}_n, \theta) \rightarrow H(\mathbb{P}_{M^0}, \theta)$ in probability and, (iii), both $H(\mathbb{P}_{M^0}, \theta)$ and $\hat{H}_n(\mathcal{D}_n, \theta)$ for all n are Lipschitz continuous in θ with the same constant L .

3.4. Confidence Regions

Suppose $T_n : \Theta \times \mathbb{R}^{r \times n} \rightarrow \{0, 1\}$ tests, for a given parameter $\theta \in \Theta$, for independence between the residuals $R^\theta := Y - f_\theta(X)$ and the instruments Z , based on the n i.i.d. observations $(X_i, Y_i, Z_i)_{i=1}^n$. Denote by $\Theta_0 :=$

¹There is a slight mismatch between Theorem 3.1 and the described algorithm: in practice, the kernel bandwidth is not fixed but is chosen according to the median heuristic. This difference could be accounted for by sample splitting.

$\{\theta \mid R^\theta \perp\!\!\!\perp Z\}$ the set of θ 's satisfying the null hypothesis of independence. We say T_n has *pointwise asymptotic level* α if $\sup_{\theta \in \Theta_0} \lim_{n \rightarrow \infty} \mathbb{P}(T_n(\theta, (Z_i)_{i=1}^n) = 1) \leq \alpha$ (here, $T_n = 1$ corresponds to rejecting independence). The test has *uniform asymptotic level* if ‘lim’ and ‘sup’ in the definition can be exchanged and the statement still holds; it has *finite sample level* if it holds for all n when removing ‘lim’. If T_n has finite sample level, then, by construction,

$$\hat{C}_n := \{\theta \in \Theta \mid T_n(\theta, (Z_i)_{i=1}^n) = 0\}$$

is a $(1 - \alpha)$ -confidence region for θ^0 : $\mathbb{P}(\theta^0 \in C_n) = \mathbb{P}(T_n(\theta^0, (Z_i)_{i=1}^n) = 0) \geq 1 - \alpha$. Similarly for pointwise (uniform) asymptotic level.

Because the independence restriction is stronger than the moment restriction (see Proposition 2.1), we can use any existing test for (1), such as the Anderson-Rubin test (Anderson & Rubin, 1949), to construct such confidence regions, too. Alternatively, we can use HSIC to construct an independence test, which we then invert. For a fixed kernel, tests can be constructed either by permutation-based procedures or by approximating the distributon of $\widehat{\text{HSIC}}$ under the null hypothesis, e.g., by a gamma distribution (Gretton et al., 2008). For many tests, \hat{C}_n has to be approximated.

4. Distribution Generalization

Here, we follow a line of work, which connects distribution generalization with causality (see Section 1). In this framework, distributional shifts are modeled as interventions. We propose using HSIC-X-pen, introduced in Section 3.1, for this task and prove that it is worst-case optimal under interventions on the exogenous variables, even if the model is nonlinear and the causal function is not identifiable.

To formalize the result, we again assume that the data is generated by the IV model M^0 . Furthermore, in this section, we consider the function classes $\mathcal{F} := \mathcal{L}^2(\mathbb{R}^d, \mathbb{P}_{M^0}^X)$ consisting of all functions $f_\diamond : \mathbb{R}^d \rightarrow \mathbb{R}$ that satisfy $\mathbb{E}_{M^0}[f_\diamond(X)^2] < \infty$ and

$$\mathcal{F}_{\text{inv}} := \{f_\diamond \in \mathcal{F} \mid Z \perp\!\!\!\perp Y - f_\diamond(X) \text{ under } \mathbb{P}_{M^0}\}.$$

We model a distributional shift as an intervention on the exogenous Z , denoted by i ; it consists of replacing the distribution of Z with a new distribution. The intervened model is denoted by $M^0(i)$ and is again an IV model. In this setting, we now prove that for any predictor $f \in \mathcal{F}$ satisfying the independence restriction $Z \perp\!\!\!\perp Y - f(X)$, that is, $f \in \mathcal{F}_{\text{inv}}$, the expected loss is invariant to interventions on Z in the following sense.

Theorem 4.1 (Invariance with respect to interventions on Z). *Let $\ell : \mathbb{R} \rightarrow \mathbb{R}$ be a convex loss function and \mathcal{I} be a set of interventions on Z satisfying for all $i \in \mathcal{I}$ that $\mathbb{P}_{M^0(i)}$ is*

dominated by \mathbb{P}_{M^0} . Then, for all $f \in \mathcal{F}_{\text{inv}}$ it holds that

$$\mathbb{E}_{M^0}[\ell(Y - f(X))] = \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)}[\ell(Y - f(X))].$$

We consider this setting relevant as identifiability (that is, $|\mathcal{F}_{\text{inv}}| = 1$) is not achievable in most modern machine learning applications. The assumption that the intervened distributions are dominated by the observed distribution ensures that none of the interventions on Z extend the support of Z . Generalizing to distributions that extend the support is only possible under additional extrapolation assumptions that require the functions g^0 and f^0 in the IV model M^0 to be partially identifiable (Christiansen et al., 2021). In this sense it is not possible to strengthen our result without making additional assumptions on the identifiability of the IV model.

Theorem 4.1 motivates estimating a predictor based on

$$\arg \min_{f \in \mathcal{F}_{\text{inv}}} \mathbb{E}_{M^0}[\ell(Y - f(X))], \quad (11)$$

where $\ell : \mathbb{R} \rightarrow \mathbb{R}$ is a convex loss function. Our proposed HSIC-X-pen estimator from (10) for λ approaching 0 provides a flexible way of estimating the minimizer (11). By Theorem 4.1, it is guaranteed to control the test error under any distributional shift generated by an intervention on Z which does not extend the support. Moreover, we now prove that for a sufficiently rich class of distribution shifts, a predictor solving (11) on the training data is worst-case optimal. To formalize our result, define the *intervened subset* $S \subseteq \{1, \dots, d\}$ consisting of all X^j that are descendants of Z in the causal graph. When $|S| < d$, this may contain cases in which identifiability according to (3) is impossible.²

Theorem 4.2 (Generalization to interventions on Z). *Let $\ell : \mathbb{R} \rightarrow \mathbb{R}$ be a convex loss function and \mathcal{I} be a set of interventions on Z satisfying for all $i \in \mathcal{I}$ that $\mathbb{P}_{M^0(i)}$ is dominated by \mathbb{P}_{M^0} , which itself is absolutely continuous with respect to a product measure. If there exists $i_* \in \mathcal{I}$ such that $X^S \perp\!\!\!\perp U \mid X^{S^c}$ under $\mathbb{P}_{M^0(i_*)}$ and $\text{supp}(\mathbb{P}_{M^0(i_*)}^X) = \text{supp}(\mathbb{P}_{M^0}^X)$, then*

$$\inf_{f \in \mathcal{F}_{\text{inv}}} \mathbb{E}_{M^0}[\ell(Y - f(X))] = \inf_{f \in \mathcal{F}} \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)}[\ell(Y - f(X))].$$

The intervention i_* can be called partially confounding-removing (see also Christiansen et al., 2021) in the sense that in the generated distribution $\mathbb{P}_{M^0(i_*)}$ the variables X^S affected by the exogenous Z are, conditioned on X^{S^c} , no longer confounded with Y via U (if Z acts additively, this implies that, in general, U does not act on both Y and X^S). As shown in the proof, this intervention results in the worst-case loss. An example of this type of intervention is given in

²For example, if $Y = f^0(X_1, X_2) + U + \epsilon_Y$, $X_1 = U$ and $X_2 = Z + U$, then the causal function is not identifiable from (3).

Section 5.2. If $\ell(\cdot) = (\cdot)^2$ is the squared loss, the minimizer is attained at the conditional mean of Y given X under $\mathbb{P}_{M^0(i_*)}$, that is, $f^0(x) + \mathbb{E}_{M^0(i_*)} [h(U, \epsilon^Y) \mid X = x]$ (see proof of Theorem 4.2).

5. Experiments

5.1. Simulation: Instrumental Variable Estimation

We first evaluate the empirical performance of HSIC-X and HSIC-X-pen for estimating causal functions. To this end, we use the following IV models in our experiments:

$$M(\alpha, \mathbb{P}_{\epsilon_Z}, f^0) : \begin{cases} Z := \epsilon_Z & U := \epsilon_U \\ X := Z\epsilon_X + \alpha Z + U \\ Y := f^0(X) - 4U + \epsilon_Y, \end{cases} \quad (12)$$

where $\epsilon_U, \epsilon_X, \epsilon_Y \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$, $\epsilon_Z \sim \mathbb{P}_{\epsilon_Z}$ and $\epsilon_Z \perp\!\!\!\perp \epsilon_U, \epsilon_X, \epsilon_Y$. We consider different experiment settings by varying three parameters $\alpha, \mathbb{P}_{\epsilon_Z}$ and f^0 . The parameter α adjusts how the instruments Z influence the predictor X : the instruments Z only change the variance of X when $\alpha = 0$, while both the mean and the variance of X are changed by Z when $\alpha > 0$. For \mathbb{P}_{ϵ_Z} , we consider binary and Gaussian random variables. Lastly, for f^0 , we consider a linear function $f_{\text{lin}}^0(X) := -2X$ and a nonlinear function $f_{\text{nonlin}}^0(X) := 1.5X - 0.2X^2 + \sum_{j=1}^{10} w_j e^{-(X-c_j)^2}$, where c_j represents a partition of $[-7, 7]$ in 10 equally spaced intervals and $w_1, \dots, w_{10} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 2)$. We generate 1000 observations from the IV model (12) and evaluate the performance of our methods under different settings. In all the experiments, we use Adam as the optimizer with the learning rate set to 0.01 and a batch-size of 256.

Known basis functions. We first assume that the underlying function class containing the causal function is known; that is, we consider $\mathcal{F} := \{f(\cdot) = \phi(\cdot)^\top \theta \mid \theta \in \mathbb{R}^p\}$ with $p = 1$ and $\phi(x) = x$ in the linear case, and $p = 12$ and $\phi(x) = [x, x^2, e^{-(x-c_1)^2}, \dots, e^{-(x-c_{10})^2}]$ in the nonlinear case (c_1, \dots, c_{10} are the same as described above). We compare the performance of HSIC-X with the following baseline methods: 2SLS: this solves the moment restriction (1); we use the feature map ϕ defined above as the function η in the moment restriction. OLS: least square regression of Y on X (using the basis functions ϕ or neural networks). Oracle: least square regression of Y on X using non-confounded data (the confounder U is removed from the assignment of X). HSIC-Oracle: same as HSIC-X(-pen) but initialized at the true causal parameters; this is to investigate how much our method suffers from the non-convexity of the objective function. The last two methods serve as oracle benchmarks.

The performance of each method is measured by the integrated mean squared error (MSE) $\mathbb{E}[(\hat{f}(X) - f^0(X))^2]$




Figure 1. MSEs of different estimators when the correct basis functions are used. Each point represents an average over 10 simulations and the error bar indicates its 95% confident interval. 2SLS is inconsistent and underperforms OLS when $\alpha = 0$, and for all α in the bottom right setting, while HSIC-X yields a substantial improvement over OLS in such settings. Under weak instruments (small α), we observe an efficiency gain of HSIC-X over 2SLS.

between the estimate \hat{f} and the causal function f^0 , approximated using a test sample of size 10000. In all experiments, we report an average of the MSE values over 10 simulations.

Figure 1 reports the MSE as we increase the parameter α . Our method shows a significant improvement over 2SLS in almost all settings. The improvement is especially prominent in the nonlinear-binary setting (see bottom right Figure 1), where the moment identifiability condition (2) is not satisfied (see Proposition 2.1). We still see an improvement gain even in the identifiable cases (when $\alpha > 0$) which suggests a finite sample efficiency gain from using the independence restriction. Lastly, the performance of HSIC-X is on par with that of HSIC-Oracle, indicating that the optimization objective is reasonably well-behaved despite its non-convexity (only in the nonlinear/binary case, there is a slight deviation). We illustrate the estimated functions against the true causal function in Appendix D.2.

Approximate functions. We now consider a more flexible function class to approximate the causal function f_{nonlin}^0 by using a neural network (NN). For a fixed width and depth, f_{nonlin}^0 may not lie in the function class represented by the NN. Nonetheless, we expect our method to produce a reasonable estimate of the causal function.

We generate 1000 observations from the IV model (12) with f_{nonlin}^0 . In addition to the OLS and Oracle baselines, we compare our method to DeepGMM (Bennett et al., 2019) and DeepIV (Hartford et al., 2017) using their publicly available implementations. To investigate the effect of the MSE regularization (see Section 3.1), we add a variant of our method (HSIC-X-pen) where the MSE regularization is




Figure 2. MSEs of different estimators when the causal function is approximated by a neural network. HSIC-X shows a substantial gain over the baselines in all settings. HSIC-X-pen does not show a major improvement over HSIC-X.

employed. A neural network with one hidden layer of size 64 is used in our methods and all the baselines.

Figure 2 outlines the simulation results. In short, both HSIC-X and HSIC-X-pen outperform DeepGMM, DeepIV and OLS baselines and approach the Oracle’s performance as the instrument strength (α) increases. We speculate that the inferior performance of DeepGMM and DeepIV may be explained by their sole reliance on the moment restriction (1) as opposed to the full independence restriction (3). Lastly, HSIC-X-pen does not show a major improvement over HSIC-X.

Multi-dimensional setting. We also investigate the effect of X ’s and Z ’s dimensionality on our estimators. The exact experiment setup and detailed results are provided in Appendix D.1. The results suggest that for linear models with higher order effects of Z , HSIC-X outperforms 2SLS, in particular when the dimensions of Z is strictly smaller than that of X .

5.2. Simulation: Distribution Generalization

To empirically verify the theoretical generalization guarantees from Section 4, we consider for $i \in \{0.5, 1, \dots, 3.5, 3.99\}$, the collection of IV models

$$M(i) : \begin{cases} Z := \epsilon_i, U_1 := \epsilon_{U_1}, U_2 := \epsilon_{U_2}, X_2 := U_2 + \epsilon_{X_2} \\ X_1 := U_1 \mathbf{1}(Z \leq 3.5) + 0.1Z + 2Z\epsilon_{X_1}, \\ Y := f^0(X_1, X_2) + U_1 + U_2, \end{cases}$$

where $\epsilon_{U_1}, \epsilon_{U_2}, \epsilon_{X_1}, \epsilon_{X_2}$ are i.i.d. $\mathcal{N}(0, 1)$, $\epsilon_i = W_1 \mathbf{1}(K = 0) + W_2 \mathbf{1}(K = 1)$ with $K \sim \text{Ber}(i/4)$, $W_1 \sim \text{Unif}(0, i)$ and $W_2 \sim \text{Unif}(i, 4)$ and $\epsilon_i \perp\!\!\!\perp (\epsilon_{U_1}, \epsilon_{U_2}, \epsilon_{X_1}, \epsilon_{X_2})$. The parameter i determines how much the distribution of Z with support $(0, 4)$ is skewed towards the right. In particular, for $i \rightarrow 4$ the intervention removes the confounding effect of U_1 , see Theorem 4.2. We use $M(0.5)$ to generate 3000 observations from the training distribution and use it to fit HSIC-X-pen, OLS and Anchor regression (AR) (Rothenhäusler et al., 2021), which comes with generalization guarantees, too, but does not cover the higher-moments




Figure 3. Predictive performance of different predictors on shifted test distributions as the intervention strength increases. Causal is the most stable predictor but conservative, while OLS is markedly affected by the interventions. Due to the higher-moments influence from Z , HSIC-X outperforms AR in most cases and establishes the best trade-off between invariance and predictiveness.

influence of Z as in the SCM above. As in Section 5.1, we consider both linear and nonlinear functions for f^0 , and employ the correct basis (Known Basis) and neural network (NN) function classes in our method and the baselines, see Appendix D.3 for details. We evaluate the fitted estimators on shifted test distributions and add a causal baseline (Causal) in which the true causal function is used as a predictor. This serves as an oracle baseline as, in this setting, the causal function is not identifiable.

The results are shown in Figure 3. Our method (HSIC-X-pen) and Causal are significantly more robust to the interventions compared to OLS and AR, with the causal oracle being the most invariant predictor as described in Theorem 4.1. However, the causal oracle is conservative (it ignores the information from the hidden confounder U_2 that is unaffected by the interventions and is helpful in predicting Y) and yields subpar performance when the interventions i are small. Our method utilizes the invariant information from U_2 and yields the best trade-off among all candidates.

5.3. Real-world Application

We apply HSIC-X to estimate the causal effect of education on earnings using 3010 observations from the 1979 National Longitudinal Survey of Young Men (Card, 1995). The response variable Y is the logarithm of wage, X is years of education, and the (discrete) instrument Z is geographic proximity to colleges (whether an individual grew up near a four-year college). Card (1995) also considered several conditioning variables W including years of experience, race and geographic information and studied a linear causal effect of education on earnings, so we use $\mathcal{F} = \{f \mid f(x) = \theta x\}$ and apply the conditional IV estimator described in Appendix B. In addition, we obtain a confidence region as described in Section 3.4. We then compare our results with those by 2SLS, which was used in the original study of Card (1995), and confidence intervals constructed by the Anderson-Rubin test (Anderson & Rubin, 1949).

Table 2. Estimation results of the effect of education on earnings (Card, 1995), including 95% confidence sets. The independence restriction puts stronger constraints on the parameters and yields a smaller confidence set; e.g., it does not include the OLS solution, which is not rejected by an Anderson-Rubin test.

METHOD	POINT ESTIMATE	LOWER	UPPER
OLS	0.072	0.065	0.079
2SLS	0.142	0.050	0.273
HSIC-X	0.160	0.097	0.208

Table 2 reports the point estimates and the confidence intervals at the 95% confidence level. The point estimates from 2SLS and HSIC-X differ from the estimate from OLS suggesting a potential sizable unobserved confounding effect. Nonetheless, the difference between the OLS and 2SLS estimates is not statistically significant which was also observed in the original study (see Card (1995) and Card (2001)). We speculate that the linear effect of college proximity Z on education X may not be strong enough which leads to the imprecision of the 2SLS. On the other hand, the confidence region around HSIC-X does not contain the OLS. This suggests that taking into account full independence gains finite sample efficiency in this study.

6. Conclusion and Future Work

Exploiting independence between exogenous variables and the residual terms of a response variable can be beneficial for identifiability of a causal function, the empirical performance of corresponding estimates, and for constructing estimators that perform well even under interventions on such exogenous variables. We have proposed two estimators, HSIC-X and HSIC-X-pen that can be equipped with any machine learning regression method based on gradient descent. Empirical results on simulated and real data indicate that one may indeed benefit from considering independence restrictions in practice.

We believe it could be fruitful to construct fast approximate methods for inverting an independence test and analyze distribution generalization when the support of Z is extended.

Acknowledgments

SS, LH, JP were supported by a research grant (18968) from VILLUM FONDEN. NP was supported by a research grant (0069071) from Novo Nordisk Fonden. We thank Nicola Gnecco and Nikolaj Thams for helpful discussions.

References

Anderson, T. W. and Rubin, H. Estimation of the parameters of a single equation in a complete system of stochastic

equations. *Annals of Mathematical Statistics*, 20(1):46–63, 1949.

Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *ArXiv e-prints (1907.02893)*, 2019.

Bennett, A., Kallus, N., and Schnabel, T. Deep generalized method of moments for instrumental variable analysis. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*. Curran Associates, Inc., 2019.

Berrett, T. B. and Samworth, R. J. Nonparametric independence testing via mutual information. *Biometrika*, 106(3):547–566, 2019.

Bongers, S., Forre, P., Peters, J., and Mooij, J. M. Foundations of structural causal models with cycles and latent variables. *Annals of Statistics*, 49(5):2885–2915, 2021.

Card, D. Using geographic variation in college proximity to estimate the return to schooling. *Aspects of Labour Market Behaviour: Essays in Honour of John Vanderkamp*, pp. 201–222, 1995.

Card, D. Estimating the return to schooling: Progress on some persistent econometric problems. *Econometrica*, 69(5):1127–1160, 2001.

Chesher, A. Identification in nonseparable models. *Econometrica*, 71(5):1405–1441, 2003.

Christiansen, R., Pfister, N., Jakobsen, M. E., Gnecco, N., and Peters, J. A causal framework for distribution generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021.

Creager, E., Jacobsen, J.-H., and Zemel, R. Environment inference for invariant learning. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pp. 2189–2200. PMLR, 2021.

Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7), 2011.

Dunker, F. Adaptive estimation for some nonparametric instrumental variable models with full independence. *Electronic Journal of Statistics*, 15:6151–6190, 2021.

Dunker, F., Florens, J.-P., Hohage, T., Johannes, J., and Mammen, E. Iterative estimation of solutions to noisy nonlinear operator equations in nonparametric instrumental regression. *Journal of Econometrics*, 178:444–455, 2014.

Frölich, M. Nonparametric IV estimation of local average treatment effects with covariates. *Journal of Econometrics*, 139(1):35–75, 2007.

- Fukumizu, K., Gretton, A., Sun, X., and Schölkopf, B. Kernel measures of conditional dependence. In *Advances in Neural Information Processing Systems 20 (NeurIPS)*. Curran Associates, Inc., 2008.
- Fuller, W. A. Some properties of a modification of the limited information estimator. *Econometrica*, 45:939–53, 1977.
- Greenfeld, D. and Shalit, U. Robust learning with the Hilbert-schmidt independence criterion. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 3759–3768. PMLR, 2020.
- Gretton, A., Fukumizu, K., Teo, C. H., Song, L., Schölkopf, B., and Smola, A. A kernel statistical test of independence. In *Advances in Neural Information Processing Systems 20 (NeurIPS)*. Curran Associates, Inc., 2008.
- Hansen, C., McDonald, J. B., and Newey, W. K. Instrumental variables estimation with flexible distributions. *Journal of Business & Economic Statistics*, 28(1):13–25, 2010.
- Hartford, J., Lewis, G., Leyton-Brown, K., and Taddy, M. Deep IV: A flexible approach for counterfactual prediction. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pp. 1414–1423. PMLR, 2017.
- Imbens, G. W. and Angrist, J. D. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–75, 1994.
- Imbens, G. W. and Newey, W. K. Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica*, 77(5):1481–1512, 2009.
- Jakobsen, M. and Peters, J. Distributional robustness of K-class estimators and the PULSE. *The Econometrics Journal*, 25(2):404–432, 2022.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*; *arXiv:1412.6980*, 2015.
- Krueger, D., Caballero, E., Jacobsen, J.-H., Zhang, A., Binas, J., Zhang, D., Priol, R., and Courville, A. Out-of-distribution generalization via risk extrapolation (REx). In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pp. 5815–5826. PMLR, 2021.
- Loh, I. Nonparametric identification and estimation with independent, discrete instruments. *ArXiv e-prints (1906.05231)*, 2019.
- Magliacane, S., van Ommen, T., Claassen, T., Bongers, S., Versteeg, P., and Mooij, J. M. Domain adaptation by using causal inference to predict invariant conditional distributions. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pp. 10846–10856. Curran Associates, Inc., 2018.
- Mariano, R. S. Simultaneous equation model estimators: Statistical properties and practical implications. In Baltagi, B. H. (ed.), *A companion to theoretical econometrics*, pp. 122–43. Blackwell, Malden, MA, 2001.
- Mooij, J. M., Janzing, D., Peters, J., and Schölkopf, B. Regression by dependence minimization and its application to causal inference. In *Proceedings of the 26th International Conference on Machine Learning (ICML)*, pp. 745–752. ACM Press, 2009.
- Mooij, J. M., Peters, J., Janzing, D., Zscheischler, J., and Schölkopf, B. Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research*, 17(32):1–102, 2016.
- Muandet, K., Mehrjou, A., Lee, S. K., and Raj, A. Dual instrumental variable regression. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*. Curran Associates, Inc., 2020.
- Newey, W. K. Uniform convergence in probability and stochastic equicontinuity. *Econometrica: Journal of the Econometric Society*, pp. 1161–1167, 1991.
- Newey, W. K. Nonparametric instrumental variables estimation. *American Economic Review*, 103(3):550–556, 2013.
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, USA, 2nd edition, 2009.
- Peters, J., Bühlmann, P., and Meinshausen, N. Causal inference using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (with discussion)*, 78(5):947–1012, 2016.
- Pfister, N., Bühlmann, P., Schölkopf, B., and Peters, J. Kernel-based tests for joint independence. *Journal of the Royal Statistical Society: Series B*, 80:5–31, 2017.
- Pfister, N., William, E. G., Peters, J., Aebbersold, R., and Bühlmann, P. Stabilizing variable selection and regression. *Annals of Applied Statistics*, 15(3):1220–1246, 2021.
- Poirier, A. Efficient estimation of models with independence restrictions. *Journal of Econometrics*, 196(1):1–22, 2017.

- Rojas-Carulla, M., Schölkopf, B., Turner, R., and Peters, J. Causal transfer in machine learning. *Journal of Machine Learning Research*, 19(36):1–34, 2018.
- Rothenhäusler, D., Bühlmann, P., Meinshausen, N., and Peters, J. Anchor regression: heterogeneous data meets causality. *Journal of Royal Statistical Society, Series B*, 83(2):215–246, 2021.
- Schölkopf, B. and Smola, A. *Learning with Kernels*. MIT Press, Massachusetts, 2002.
- Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., and Mooij, J. M. On causal and anticausal learning. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*. Omnipress, 2012.
- Shah, R. and Peters, J. The hardness of conditional independence testing and the generalised covariance measure. *Annals of Statistics*, 48(3):1514–1538, 2020.
- Singh, R., Sahani, M., and Gretton, A. Kernel instrumental variable regression. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32 (NeurIPS)*. Curran Associates, Inc., 2019.
- Sriperumbudur, B. K., Fukumizu, K., Gretton, A., Lanckriet, G., and Schölkopf, B. Kernel choice and classifiability for rkhs embeddings of probability distributions. In *Advances in Neural Information Processing Systems 22 (NeurIPS)*. Curran Associates, Inc., 2009.
- Theil, H. *Economic forecasts and policy*. North-Holland, Amsterdam, NL, 1958.
- Torgovitsky, A. Identification of nonseparable models using instruments with small support. *Econometrica*, 83(3): 1185–1197, 2015.
- Wooldridge, J. M. *Econometric analysis of cross section and panel data*. MIT press, Cambridge, MA, 2010.
- Wright, P. G. *The Tariff on Animal and Vegetable Oils*. Investigations in International Commercial Policies. Macmillan, New York, NY, 1928.
- Yuan, J., Ma, X., Kuang, K., Xiong, R., Gong, M., and Lin, L. Learning domain-invariant relationship with instrumental variable for domain generalization. *ArXiv e-prints (2110.01438)*, 2021.
- Zhang, K., Peters, J., Janzing, D., and Schölkopf, B. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the 27th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 804–813. AUAI Press, 2011.

A. Proofs

A.1. Proof of Proposition 2.1

Proof. (i) We first show that if there exists k and η , such that condition (2) holds, then condition (4) holds. We do so by contraposition, so assume that there exists a $\tau \neq 0$ violating condition (4). The following then hold for all $k \in \mathbb{N}$ and all $\eta : \mathbb{R}^r \rightarrow \mathbb{R}^k$. First, $\text{Cov}[\eta(Z), h(U, \epsilon_Y) + \phi(X)^\top \tau] = 0$. Second, as $Z \perp\!\!\!\perp h(U, \epsilon_Y)$ it follows that $\text{Cov}[\eta(Z), h(U, \epsilon_Y)] = 0$. Third, by combining the previous two results it follows that $\text{Cov}[\eta(Z), \phi(X)^\top \tau] = 0$, which is a violation of condition (2).

(ii) This statement is proven by (iii) (a), for example.

(iii) We now consider the two types of examples (a), (b) and (c).

(a) In the example SCM, we have $\mathbb{E}[X|Z] = Z \mathbb{E}[U] + \mathbb{E}[\epsilon_X] = 0$ and therefore for all $k \in \mathbb{N}$ and functions $\eta : \mathbb{R}^r \rightarrow \mathbb{R}^k$, $\mathbb{E}[\eta(Z)X] = 0$. As a result, the identifiability condition (2) does not hold here. Specifically, any function of the form θX , $\theta \in \mathbb{R}$ satisfies the moment restriction, so we cannot identify the causal function $f^0(x) = x$. Consider now the independence based identifiability condition (4):

$$Z \perp\!\!\!\perp U + \epsilon_Y + \tau X \implies \tau = 0.$$

Plugging in the structural equation for X this is equivalent to $K_\tau := U + \tau ZU + \tau \epsilon_X + \epsilon_Y \not\perp\!\!\!\perp Z$ for all $\tau \neq 0$. We now show that this is true by showing that for all $\tau \neq 0$, $\mathbb{E}[Z^2 K_\tau^2] \neq \mathbb{E}[Z^2] \mathbb{E}[K_\tau^2]$. First,

$$\begin{aligned} \mathbb{E}[Z^2 K_\tau^2] &= \mathbb{E}[Z^2 (U + \tau ZU + \tau \epsilon_X + \epsilon_Y)^2] \\ &= \mathbb{E}[Z^2 U^2] + \tau^2 \mathbb{E}[Z^4 U^2] + \tau^2 \mathbb{E}[Z^2 \epsilon_X^2] + \mathbb{E}[Z^2 \epsilon_Y^2] \\ &= 2 + 4\tau^2. \end{aligned}$$

Second,

$$\begin{aligned} \mathbb{E}[Z^2] \mathbb{E}[K_\tau^2] &= \mathbb{E}[Z^2] \mathbb{E}[(U + \tau ZU + \tau \epsilon_X + \epsilon_Y)^2] \\ &= \mathbb{E}[Z^2] (\mathbb{E}[U^2] + \tau^2 \mathbb{E}[Z^2 U^2] + \tau^2 \mathbb{E}[\epsilon_X^2] + \mathbb{E}[\epsilon_Y^2]) \\ &= 2 + 2\tau^2. \end{aligned}$$

Therefore, condition (4) is fulfilled, meaning we can identify the causal function $f^0(x) = x$.

(b) We first give an argument why identifiability using condition (2) is impossible. For a discrete instrument Z , the number of almost surely linearly independent functions³ of Z is bounded by the cardinality of the support of Z ; therefore, if the number of basis functions is larger than the support of Z , there is no $k \in \mathbb{N}$ and $\eta : \mathbb{R}^r \rightarrow \mathbb{R}^k$ such that condition (2) is satisfied.

We now construct an explicit example as follows: First, we show for a class of SCMs that under certain assumptions condition (2) does not hold while (4) does. Second, we give an explicit SCM with a binomially distributed instrument that fulfills these conditions. Fix $k, p \in \mathbb{N}$ such that $p > k$ and consider the SCM

$$\begin{aligned} Z &:= \epsilon_Z \\ U &:= \epsilon_U \\ X &:= Z + U + \epsilon_X \\ Y &:= \sum_{i=1}^p \alpha_i X^i + U + \epsilon_Y, \end{aligned}$$

where $\epsilon_Z, \epsilon_U, \epsilon_X$ and ϵ_Y are jointly independent, real-valued errors $\epsilon_Z \sim \text{Unif}(\{0, \dots, k-1\})$ and $\alpha_i \in \mathbb{R} \setminus \{0\}$ for all $i \in \{1, \dots, p\}$. Assume that the random variables $\{\epsilon_X^i U^j | 1 \leq i+j \leq p\}$ are almost surely linearly independent

³Here, we say that a collection of functions l_1, \dots, l_k is almost surely linearly independent if for all vectors $\alpha \neq 0$, $\mathbb{P}\left(\sum_{i=1}^k \alpha_i l_i(Z) = 0\right) = 0$.

(see above). We are interested in estimating the causal function $f^0(x) = \sum_{i=1}^p \alpha_i x^i$ with the basis functions $\phi(x) = (x^1, \dots, x^p)$.

Consider first the moment identifiability condition (2). Because Z is discrete, the set of functions $\{\mathbf{1}_{\{Z=i\}}, i \in \{0, \dots, k-1\}\}$ form a linear basis for all functions from $\text{supp}(Z)$ to \mathbb{R} . As a result, for all functions $\eta = (\eta_1, \dots, \eta_q) : \mathbb{R} \rightarrow \mathbb{R}^q$, with $q > k$, condition (2) cannot hold. This implies that while condition (2) may hold if $p \leq k$, with for example $\eta(Z) = (\mathbf{1}_{\{Z=0\}}, \dots, \mathbf{1}_{\{Z=k-1\}})$, it cannot hold if $p > k$, irrespective of the choice of η .

Consider now the independence identifiability condition (4):

$$\sum_{i=1}^p \tau_i \alpha_i X^i + h(U, \epsilon_Y) \perp\!\!\!\perp Z \implies \tau_1 = \dots = \tau_p = 0.$$

Because the support of Z is finite, this is equivalent to the statement that the distributions of the random variables

$$\sum_{i=1}^p \tau_i \alpha_i (j + U + \epsilon_X)^i + U + \epsilon_Y,$$

are the same for all $j \in \{0, \dots, k-1\}$. Consider the cases $j = 0$ and $j = 1$. For $\sum_{i=1}^p \tau_i \alpha_i (U + \epsilon_X)^i + U + \epsilon_Y$ and $\sum_{i=1}^p \tau_i \alpha_i (1 + U + \epsilon_X)^i + U + \epsilon_Y$ to have the same distribution, the coefficients in front of the ϵ_X^i (when collecting all terms) for all $i \in \{1, \dots, p\}$ must be equal in both random variables. Consider first the ϵ_X^{p-1} term. The corresponding two coefficients are $\tau_{p-1} \alpha_{p-1}$ and $\tau_{p-1} \alpha_{p-1} + p \tau_p \alpha_p$. These two coefficients are only equal if $\tau_p = 0$. By iterating this argument we obtain that $\tau_2 = \dots = \tau_p = 0$. This implies that $\tau_1 \alpha_1 (U + \epsilon_X) + U + \epsilon_Y$ and $\tau_1 \alpha_1 (1 + U + \epsilon_X) + U + \epsilon_Y$ have the same distribution. Thus, $\tau_1 = 0$ and therefore condition (4) holds.

We now give an explicit example, where $k < p$ and where the random variables $\{\epsilon_X^i U^j | 1 \leq i + j \leq p\}$ are almost surely linearly independent. Let $\epsilon_U, \epsilon_Z, \epsilon_X$ and ϵ_Y standard normal, $Z \sim \text{Bernoulli}(0.5)$ and $p = 3$. Here, $k = 2 < p$ and we will now show that the set of random variables $\{\epsilon_X^i U^j | 1 \leq i + j \leq 3\}$ is almost sure linearly independent. Let $\tau_{ij}, 1 \leq i + j \leq 3$ be coefficients such that $K_\tau := \sum_{1 \leq i+j \leq 3} \tau_{ij} \epsilon_X^i U^j = 0$ almost surely. In particular, this implies that $\mathbb{E}[K_\tau \epsilon_X^i U^j] = 0$ for all $i, j \leq 0$. Using $\mathbb{E}[\epsilon_X^i U^j \epsilon_X^k U^l] = 0$ if either $i + k$ or $j + l$ are odd (which follows as ϵ_X and U are mean zero Gaussian and independent), we obtain the following nine equations.

$$\begin{aligned} \mathbb{E}[K_\tau U] &= \tau_{01} + 3\tau_{03} + \tau_{21} = 0 \\ \mathbb{E}[K_\tau U^2] &= 3\tau_{02} + \tau_{20} = 0 \\ \mathbb{E}[K_\tau U^3] &= 3\tau_{01} + 15\tau_{03} + 3\tau_{21} = 0 \\ \mathbb{E}[K_\tau \epsilon_X] &= \tau_{10} + \tau_{12} + 3\tau_{30} = 0 \\ \mathbb{E}[K_\tau \epsilon_X U] &= \tau_{11} = 0 \\ \mathbb{E}[K_\tau \epsilon_X U^2] &= \tau_{10} + 3\tau_{12} + 3\tau_{30} = 0 \\ \mathbb{E}[K_\tau \epsilon_X^2] &= \tau_{02} + 3\tau_{20} = 0 \\ \mathbb{E}[K_\tau \epsilon_X^2 U] &= \tau_{01} + 3\tau_{03} + 3\tau_{21} = 0 \\ \mathbb{E}[K_\tau \epsilon_X^3] &= 3\tau_{10} + 3\tau_{12} + 15\tau_{30} = 0. \end{aligned}$$

We can write this as a linear system $A\tau = 0$ with corresponding 9×9 matrix A . As A is invertible, $\tau = 0$ follows.

(c) Consider the SCM

$$\begin{aligned} Z &:= \epsilon_Z \\ U &:= \epsilon_U \\ X &:= 2ZU - U + \epsilon_X \\ Y &:= X + U + \epsilon_Y. \end{aligned}$$

where $\epsilon_Z, \epsilon_U, \epsilon_X$ and ϵ_Y are jointly independent, $\epsilon_Z \sim \text{Bernoulli}(0.5)$ and the remaining errors are standard normal. Consider the basis function $\phi(x) = x$. Here,

$$K_\tau := h(U, \epsilon_Y) + \tau X = \begin{cases} (1 + \tau)U + \tau \epsilon_X + \epsilon_Y & \text{if } Z = 1 \\ (1 - \tau)U + \tau \epsilon_X + \epsilon_Y & \text{if } Z = 0 \end{cases}$$

and therefore $\mathbb{E}[K_\tau^2|Z = 1] = 2\tau^2 + 2\tau + 2$ and $\mathbb{E}[K_\tau^2|Z = 0] = 2\tau^2 - \tau + 2$. It follows that condition (4) holds. On the other hand, $\tau X|(Z = 1) = \tau(U + \epsilon_X)$ and $\tau X|(Z = 0) = \tau(-U + \epsilon_X)$ and second, the distribution of U is symmetric around 0. Therefore, $Z \perp\!\!\!\perp \tau X$ for all $\tau \in \mathbb{R}$. We can therefore identify the causal function with the independence restriction, even though Z is independent of X . \square

A.2. Proof of Proposition 2.2

Proof. This is a proof by example. We have already constructed an SCM, such that condition (4) holds but condition (6) does not in the proof of statement (iii) (c) of Proposition 2.1.

We now construct a SCM, such that condition (6) holds but condition (4) does not. Consider the SCM

$$\begin{aligned} Z &:= \epsilon_Z \\ U &:= \epsilon_U \\ X &:= 2ZU + \epsilon_X \\ Y &:= X - U + \epsilon_Y. \end{aligned}$$

where $\epsilon_Z, \epsilon_U, \epsilon_X$ and ϵ_Y are jointly independent, $\epsilon_Z \sim \text{Bernoulli}(0.5)$ and the remaining errors are standard normal. Consider the basis functions $\phi(x) = (x)$. Here,

$$h(U, \epsilon_Y) + X = \begin{cases} U + \epsilon_X + \epsilon_Y & \text{if } Z = 1 \\ -U + \epsilon_X + \epsilon_Y & \text{if } Z = 0 \end{cases}$$

and therefore condition (4) does not hold for $\tau = 1$. On the other hand,

$$\tau X = \begin{cases} \tau(2U + \epsilon_X) & \text{if } Z = 1 \\ \tau\epsilon_X & \text{if } Z = 0 \end{cases}$$

and therefore $\mathbb{E}[(\tau X)^2|Z = 1] = 5\tau^2$ while $\mathbb{E}[(\tau X)^2|Z = 0] = \tau^2$. It follows that condition (6) holds. \square

A.3. Proof of Theorem 3.1

Proof. Let $H(\mathbb{P}_{M^0}, \theta) := \text{HSIC}((Y - \phi(X)^\top \theta, Z); k_R, k_Z)$ with non-negative bounded and Lipschitz continuous kernels k_R and k_Z with bounded and continuous derivatives.⁴ Consider an i.i.d. data set $\mathcal{D}_n = (X_i, Y_i, Z_i)_{i=1}^n$, such that that each triple (X_i, Y_i, Z_i) follows the distribution of M^0 . Let $\hat{H}_n(\mathcal{D}_n, \theta) := \widehat{\text{HSIC}}((R_i^\theta, Z_i)_{i=1}^n)$, with $R_i^\theta = Y_i - \phi(X_i)^\top \theta$.

Then the following three results hold. First, as the kernels k_R and k_L are by assumption non-negative and bounded it follows by Corollary 15 in Mooij et al. (2016) that for all θ and all $\epsilon > 0$, $\lim_{n \rightarrow \infty} \mathbb{P}_{M^0}(|\hat{H}_n(\mathcal{D}_n, \theta) - H(\theta)| > \epsilon) = 0$.

Second, it follows by Lemma 16 in Mooij et al. (2016) that for all $\theta_1, \theta_2 \in \Theta$ and all $n \geq 2$

$$\begin{aligned} \hat{H}_n(\mathcal{D}_n, \theta_1) - \hat{H}_n(\mathcal{D}_n, \theta_2) &= \widehat{\text{HSIC}}((R_i^{\theta_1}, Z_i)_{i=1}^n) - \widehat{\text{HSIC}}((R_i^{\theta_2}, Z_i)_{i=1}^n) \\ &\leq \frac{32\lambda C}{\sqrt{n}} \|R^{\theta_1} - R^{\theta_2}\| \\ &= \frac{32\lambda C}{\sqrt{n}} \|\phi(X)^\top (\theta_1 - \theta_2)\| \\ &\leq \frac{32\lambda C}{\sqrt{n}} \|\phi(X)^\top\| \|\theta_1 - \theta_2\| \\ &= \frac{32\lambda C}{\sqrt{n}} \left(\sum_{i=1}^n \|\phi(X_i)\|^2 \right)^{0.5} \|\theta_1 - \theta_2\| \\ &\leq 32\lambda C M \|\theta_1 - \theta_2\|, \end{aligned}$$

⁴From now on, we do not explicitly write down the kernels and write $H(\theta)$ rather than $H(\mathbb{P}_{M^0}, \theta)$ to ease notation

where $R^\theta = (Y_i - \phi(X_i)^\top \theta)_{i=1}^n$, $\phi(X) = (\phi(X_i))_{i=1}^n$, λ is the Lipschitz constant for k_R , C is an upper bound for k_Z and M is an upper bound for $\|\phi(\cdot)\|^2$ which exists by the assumption that $\phi(\cdot)$ is bounded.

Third, for example by Pfister et al. (2017) Proposition 2.5, the population HSIC can be written as

$$\begin{aligned} H(\theta) &= \text{HSIC}(Y - \phi(X)^\top \theta, Z) \\ &= \mathbb{E}[k_R(R_1^{\theta^1}, R_2^{\theta^2})k_Z(Z_1, Z_2)] + \mathbb{E}[k_R(R_1^{\theta^1}, R_2^{\theta^2})]\mathbb{E}[k_Z(Z_1, Z_2)] - 2\mathbb{E}[k_R(R_1^{\theta^1}, R_2^{\theta^2})k_Z(Z_1, Z_3)], \end{aligned}$$

with $R_1^{\theta^1}$ and $R_2^{\theta^2}$ as well as Z_1 , Z_2 and Z_3 being i.i.d. copies of R^{θ^1} and Z , respectively. The function $R^\theta : \theta \mapsto Y - \phi(X)^\top \theta$ is (surely) continuously differentiable, as is k_R by assumption. Further, k_R and k_Z have bounded derivative. We can therefore conclude by dominated convergence and the chain rule that $H(\theta)$ is continuously differentiable. Therefore, $H(\theta)$ is Lipschitz on Θ , since Θ is compact.

Jointly, these three results (which correspond to (i), (ii), and (iii) from the statement of the theorem) allow us to apply Corollary 2.2 from Newey (1991) and conclude that for all $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P}_{M^0}(\max_{\theta \in \Theta} |\hat{H}_n(\mathcal{D}_n, \theta) - H(\theta)| > \varepsilon) = 0. \quad (13)$$

Here and in the remainder of the proof we use that $\theta \mapsto \hat{H}_n(\mathcal{D}_n, \theta)$ and $\theta \mapsto H(\theta)$ are (surely) continuous functions that map to \mathbb{R} . Restricted to any compact set, they therefore (surely) attain a maximum and a minimum.

Consider now any sequence of estimators $\hat{\theta}_n^0 \in \arg \min_{\theta \in \Theta} \hat{H}_n(\mathcal{D}_n, \theta)$. By the assumption that condition (4) holds, θ^0 is the unique minimizer of the continuous function $H(\cdot)$ on the compact set Θ . Therefore, for all $\varepsilon > 0$, there exists a $\zeta(\varepsilon)$ such that for all $\theta \notin B_\varepsilon(\theta^0) = \{\theta : \|\theta - \theta^0\| < \varepsilon\}$ it holds that $H(\theta) - H(\theta^0) > \zeta(\varepsilon)$. Note also that $\Theta \setminus B_\varepsilon$ is compact. Fix $\varepsilon > 0$ and $\delta > 0$, then by (13) there exists an $n' \in \mathbb{N}$, such that for all $n \geq n'$

$$\mathbb{P}_{M^0}(\max_{\theta \in \Theta} |\hat{H}_n(\mathcal{D}_n, \theta) - H(\theta)| > \frac{1}{2}\zeta(\varepsilon)) < \frac{1}{2}\delta. \quad (14)$$

Then, for all $n \geq n'$ it holds that

$$\begin{aligned} \mathbb{P}_{M^0}(\|\hat{\theta}_n^0 - \theta^0\| > \varepsilon) &\leq \mathbb{P}_{M^0}(\min_{\theta \in \Theta \setminus B_\varepsilon(\theta^0)} (\hat{H}_n(\mathcal{D}_n, \theta) - \hat{H}_n(\mathcal{D}_n, \theta^0)) \leq 0) \\ &\leq \mathbb{P}_{M^0}(\max_{\theta \in \Theta \setminus B_\varepsilon(\theta^0)} |\hat{H}_n(\mathcal{D}_n, \theta) - H(\theta)| + |\hat{H}_n(\mathcal{D}_n, \theta^0) - H(\theta^0)| > \zeta(\varepsilon)) \\ &\leq \mathbb{P}_{M^0}(\{\max_{\theta \in \Theta \setminus B_\varepsilon(\theta^0)} |\hat{H}_n(\mathcal{D}_n, \theta) - H(\theta)| > \frac{1}{2}\zeta(\varepsilon)\} \cup \{|\hat{H}_n(\mathcal{D}_n, \theta^0) - H(\theta^0)| > \frac{1}{2}\zeta(\varepsilon)\}) \\ &< \delta. \end{aligned}$$

Here, we used the following four arguments:

- Firstly, the event $\{\|\hat{\theta}_n^0 - \theta^0\| > \varepsilon\}$ can only occur if there exists a $\theta \in \Theta \setminus B_\varepsilon(\theta^0)$ such that the event $\{\hat{H}_n(\mathcal{D}_n, \theta) - \hat{H}_n(\mathcal{D}_n, \theta^0) \leq 0\}$ occurs.
- Secondly, $\zeta(\varepsilon)$ is defined such that for all $\theta \in \Theta \setminus B_\varepsilon(\theta^0)$ it holds that $H(\theta) - H(\theta^0) > \zeta(\varepsilon)$. Therefore, if there exist a θ such that the event $\{\hat{H}_n(\mathcal{D}_n, \theta) - \hat{H}_n(\mathcal{D}_n, \theta^0) \leq 0\}$ occurs, then there must exist a θ such that the event $\{|\hat{H}_n(\mathcal{D}_n, \theta) - H(\theta)| + |\hat{H}_n(\mathcal{D}_n, \theta^0) - H(\theta^0)| > \zeta(\varepsilon)\}$ occurs.
- Thirdly, for $\{\max_{\theta \in \Theta \setminus B_\varepsilon(\theta^0)} |\hat{H}_n(\mathcal{D}_n, \theta) - H(\theta)| + |\hat{H}_n(\mathcal{D}_n, \theta^0) - H(\theta^0)| > \zeta(\varepsilon)\}$ to occur, either $\{\max_{\theta \in \Theta \setminus B_\varepsilon(\theta^0)} |\hat{H}_n(\mathcal{D}_n, \theta) - H(\theta)| > \frac{1}{2}\zeta(\varepsilon)\}$ or $\{|\hat{H}_n(\mathcal{D}_n, \theta^0) - H(\theta^0)| > \frac{1}{2}\zeta(\varepsilon)\}$ must occur.
- Fourthly, we use the union bound together with (14) for the last inequality.

We can therefore conclude that $\lim_{n \rightarrow \infty} \mathbb{P}_{M^0}(\|\hat{\theta}_n^0 - \theta^0\| > \varepsilon) = 0$.

Finally, consider our estimator $\hat{f}(\cdot) = \phi(\cdot)^\top \hat{\theta}_n^0$ for the causal function f^0 . By the assumption that ϕ is bounded, we get that

$$\|\hat{f} - f^0\|_\infty = \sup_{x \in \mathbb{R}^d} |\phi(x)^\top (\hat{\theta}_n^0 - \theta^0)| \leq \sup_{x \in \mathbb{R}^d} \|\phi(x)\| \cdot \|\hat{\theta}_n^0 - \theta^0\| \leq M \|\hat{\theta}_n^0 - \theta^0\|.$$

Hence we conclude that $\|\hat{f} - f^0\|_\infty \xrightarrow{\mathbb{P}_{M^0}} 0$.

Our second claim follows by the fact that under the conditions laid out, we can directly apply Corollary 2.2 from Newey (1991) to obtain uniform convergence in probability and then argue as in the case for $\widehat{\text{HSIC}}$.

□

A.4. Proof of Theorem 4.1

Proof. Let $f_\diamond \in \mathcal{F}_{\text{inv}}$ and $i \in \mathcal{I}$. Let \mathcal{G} be the directed graph induced by the SCM M^0 . Then, since Z is by construction a source node in \mathcal{G} and i an intervention on Z satisfying that $\mathbb{P}_{M^0(i)}$ is dominated by \mathbb{P}_{M^0} , it holds that

$$\begin{aligned} \mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X))] &= \mathbb{E}_{M^0(i)} \left[\mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X)) \mid Z] \right] \\ &= \mathbb{E}_{M^0(i)} \left[\mathbb{E}_{M^0} [\ell(Y - f_\diamond(X)) \mid Z] \right]. \end{aligned} \quad (15)$$

Now, since $f_\diamond \in \mathcal{F}_{\text{inv}}$ it holds that $Z \perp\!\!\!\perp Y - f_\diamond(X)$ under \mathbb{P}_{M^0} , which implies that

$$\mathbb{E}_{M^0(i)} \left[\mathbb{E}_{M^0} [\ell(Y - f_\diamond(X)) \mid Z] \right] = \mathbb{E}_{M^0(i)} \left[\mathbb{E}_{M^0} [\ell(Y - f_\diamond(X))] \right] = \mathbb{E}_{M^0} [\ell(Y - f_\diamond(X))]. \quad (16)$$

Since $i \in \mathcal{I}$ was arbitrary, combining (15) and (16) directly implies

$$\mathbb{E}_{M^0} [\ell(Y - f(X))] = \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f(X))],$$

which completes the proof of Theorem 4.1.

□

A.5. Proof of Theorem 4.2

Proof. We consider two optimization problems: (A) Minimize $\mathbb{E}_{M^0} [\ell(Y - f_\diamond(X))]$ over all $f_\diamond \in \mathcal{F}_{\text{inv}}$ and (B) minimize $\sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X))]$ over all $f_\diamond \in \mathcal{F}$. To prove the result, we fix an arbitrary $\epsilon > 0$ and find a function $\bar{\phi} \in \mathcal{F}_{\text{inv}}$ that satisfies

$$\mathbb{E}_{M^0} [\ell(Y - \bar{\phi}(X))] \leq \inf_{f_\diamond \in \mathcal{F}_{\text{inv}}} \mathbb{E}_{M^0} [\ell(Y - f_\diamond(X))] + \epsilon \quad (17)$$

and

$$\sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - \bar{\phi}(X))] \leq \inf_{f_\diamond \in \mathcal{F}} \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X))] + \epsilon. \quad (18)$$

Since by (15) and (16) in the proof of Theorem 4.1 it holds that $\mathbb{E}_{M^0} [\ell(Y - \bar{\phi}(X))] = \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - \bar{\phi}(X))]$, the result we wish to prove follows immediately. The proof now proceeds in two steps:

- (1) We construct the function $\bar{\phi}$ using the intervention i_* from the statement.
- (2) We use that $\bar{\phi} \in \mathcal{F}_{\text{inv}}$ to conclude the proof.

Step (1): Let $i_* \in \mathcal{I}$ be the intervention from the statement of the theorem. For all $f_\diamond \in \mathcal{F}$, it holds that

$$\sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X))] \geq \mathbb{E}_{M^0(i_*)} [\ell(Y - f_\diamond(X))].$$

As $f_\diamond \in \mathcal{F}$ was arbitrary, we can take the infimum on both sides and get

$$\inf_{f_\diamond \in \mathcal{F}} \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X))] \geq \inf_{f_\diamond \in \mathcal{F}} \mathbb{E}_{M^0(i_*)} [\ell(Y - f_\diamond(X))]. \quad (19)$$

By the definition of the infimum there exists⁵ $\psi \in \mathcal{F}$ satisfying

$$\mathbb{E}_{M^0(i_*)} [\ell(Y - \psi(X))] \leq \inf_{f_\diamond \in \mathcal{F}} \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X))] + \epsilon/2. \quad (20)$$

⁵In the case of squared loss $\ell(\cdot) = (\cdot)^2$, the minimum is attained at the conditional mean $\mathbb{E}_{M^0(i_*)}[Y \mid X] = f^0(X) + \mathbb{E}_{M^0(i_*)}[h(U, \epsilon_Y) \mid X]$.

Moreover, setting $\bar{\psi} := \psi - f^0$ and expanding Y from the structural assignment we get that

$$\mathbb{E}_{M^0(i_*)} [\ell(Y - \psi(X))] = \mathbb{E}_{M^0(i_*)} [\ell(f^0(X) + h(U, \epsilon_Y) - \psi(X))] = \mathbb{E}_{M^0(i_*)} [\ell(h(U, \epsilon_Y) - \bar{\psi}(X))].$$

Next, recall that $X^S \perp\!\!\!\perp U \mid X^{S^c}$. Since $\epsilon_Y \perp\!\!\!\perp (X, U)$, it also holds (using the properties of conditional independence) that $X^S \perp\!\!\!\perp (U, \epsilon_Y) \mid X^{S^c}$ under $\mathbb{P}_{M^0(i_*)}$. Let μ denote the common dominating product measure and let p_* denote the density corresponding to $\mathbb{P}_{M^0(i_*)}$ with respect to μ . Then, by conditional independence it holds that

$$p_*(x^S, x^{S^c}, u, e) = p_*(x^S \mid x^{S^c})p_*(u, e \mid x^{S^c})p_*(x^{S^c}).$$

Expressing the expectation as an integral and applying Jensen's inequality once (ℓ is convex) we get that

$$\begin{aligned} & \mathbb{E}_{M^0(i_*)} [\ell(h(U, \epsilon_Y) - \bar{\psi}(X))] \\ &= \int \ell(h(u, e) - \bar{\psi}(x^S, x^{S^c}))p_*(x^S \mid x^{S^c})p_*(u, e \mid x^{S^c})p_*(x^{S^c})d\mu(u, e, x^S, x^{S^c}) \\ &\geq \int \ell\left(\int h(u, e) - \bar{\psi}(x^S, x^{S^c})p_*(x^S \mid x^{S^c})d\mu(x^S)\right)p_*(u, e \mid x^{S^c})p_*(x^{S^c})d\mu(u, e, x^{S^c}) \\ &= \int \ell\left(h(u, e) - \int \bar{\psi}(x^S, x^{S^c})p_*(x^S \mid x^{S^c})d\mu(x^S)\right)p_*(u, e \mid x^{S^c})p_*(x^{S^c})d\mu(u, e, x^{S^c}). \end{aligned}$$

Let $\phi \in \mathcal{F}$ be a function that satisfies for all $x^{S^c} \in \text{supp}(\mathbb{P}_{M^0(i_*)}^{X^{S^c}})$ that

$$\phi(x^{S^c}) = \int \bar{\psi}(x^S, x^{S^c})p_*(x^S \mid x^{S^c})d\mu(x^S).$$

Such a $\phi \in \mathcal{F}$ exists because \mathcal{F} consists of all square-integrable functions. We then get that

$$\begin{aligned} \mathbb{E}_{M^0(i_*)} [\ell(Y - \psi(X))] &\geq \int \ell\left(h(u, e) - \phi(x^{S^c})\right)p_*(u, e \mid x^{S^c})p_*(x^{S^c})d\mu(u, e, x^{S^c}) \\ &= \mathbb{E}_{M^0(i_*)} [\ell(h(U, \epsilon_Y) - \phi(X^{S^c}))] \\ &= \mathbb{E}_{M^0(i_*)} [\ell(Y - (f^0(X) + \phi(X^{S^c})))]. \end{aligned}$$

Combining this with (20) and defining $\bar{\phi} \in \mathcal{F}$ for all $x \in \mathbb{R}^d$ by $\bar{\phi}(x) := f^0(x) + \phi(x^{S^c})$ leads to

$$\mathbb{E}_{M^0(i_*)} [\ell(Y - \bar{\phi}(X))] \leq \inf_{f_\diamond \in \mathcal{F}} \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X))] + \varepsilon/2. \quad (21)$$

Moreover, by construction it holds $\mathbb{P}_{M^0(i_*)}$ -almost surely that

$$Y - \bar{\phi}(X) = h(U, \epsilon_Y) - \phi(X^{S^c}). \quad (22)$$

Since $\mathbb{P}_{M^0(i_*)}$ is dominated by \mathbb{P}_{M^0} and $\text{supp}(\mathbb{P}_{M^0(i_*)}^X) = \text{supp}(\mathbb{P}_{M^0}^X)$, (22) also holds \mathbb{P}_{M^0} -almost surely. Finally, using that X^{S^c} are not descendants of Z and Z is a source node, it holds that $Z \perp\!\!\!\perp (X^{S^c}, U, \epsilon_Y)$ under \mathbb{P}_{M^0} , which implies

$$Z \perp\!\!\!\perp Y - \bar{\phi}(X) \quad \text{under } \mathbb{P}_{M^0}.$$

Hence $\bar{\phi} \in \mathcal{F}_{\text{inv}}$.

Step (2): Since $\bar{\phi} \in \mathcal{F}_{\text{inv}}$ we can use (15) and (16) from the proof of Theorem 4.1 together with (21) to get that

$$\sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - \bar{\phi}(X))] = \mathbb{E}_{M^0(i_*)} [\ell(Y - \bar{\phi}(X))] \leq \inf_{f_\diamond \in \mathcal{F}} \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_\diamond(X))] + \varepsilon/2.$$

This proves (18). Next, by the definition of the infimum there exists $f_* \in \mathcal{F}_{\text{inv}}$ satisfying

$$\mathbb{E}_{M^0} [\ell(Y - f_*(X))] \leq \inf_{f_\diamond \in \mathcal{F}_{\text{inv}}} \mathbb{E}_{M^0} [\ell(Y - f_\diamond(X))] + \varepsilon/2. \quad (23)$$

Now, assume that

$$\mathbb{E}_{M^0} [\ell(Y - f_*(X))] \leq \mathbb{E}_{M^0} [\ell(Y - \bar{\phi}(X))].$$

If this was not the case then (17) would be true and the proof would be complete. Again, applying Theorem 4.1 and defining $c_{\min} := \inf_{f_\circ \in \mathcal{F}} \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_\circ(X))]$ this implies

$$c_{\min} \leq \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_*(X))] \leq \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - \bar{\phi}(X))] \leq c_{\min} + \epsilon/2.$$

This implies that


$$\sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - \bar{\phi}(X))] \leq \sup_{i \in \mathcal{I}} \mathbb{E}_{M^0(i)} [\ell(Y - f_*(X))] + \epsilon/2 = \mathbb{E}_{M^0} [\ell(Y - f_*(X))] + \epsilon/2.$$

Combined with (23) this proves (17), which concludes the proof of Theorem 4.2. \square

B. Additional details - Conditional IV

We consider the following SCM.⁶

$$\begin{aligned} W &:= m(\epsilon_W, V, U) \\ Z &:= q(W, V, \epsilon_Z) \\ X &:= g(Z, W, U, \epsilon_X) \\ V &:= \epsilon_V \\ U &:= \epsilon_U \\ Y &:= f(X) + h(W, U, \epsilon_Y), \end{aligned}$$



where $(\epsilon_W, \epsilon_Z, \epsilon_V, \epsilon_U, \epsilon_X, \epsilon_Y) \sim Q$ are jointly independent noise variables, $Z \in \mathbb{R}^r$ are *instruments*, $V \in \mathbb{R}^{s_1}$ and $U \in \mathbb{R}^{s_2}$ are unobserved variables, $W \in \mathbb{R}^t$ are observed *covariates*, $X \in \mathbb{R}^d$ are *predictors* and $Y \in \mathbb{R}$ is a *response*. For simplicity, we additionally assume that $\mathbb{E}[W] = \mathbb{E}[Z] = \mathbb{E}[X] = \mathbb{E}[h(W, U, \epsilon_Y)] = 0$. We consider models of the form $M = (m, q, g, f, h)$ and assume that either there is no edge from V to W , i.e., $W = m(\epsilon_W, U)$ or no edge between W and U , i.e., $W = m(\epsilon_W, V)$. As before, $M^0 = (m^0, q^0, g^0, f^0, h^0)$ denotes the data generating model and $f^0(\cdot) = \phi(\cdot)^\top \theta^0$ for some $\theta^0 \in \Theta \subseteq \mathbb{R}^p$, with a (known) function $\phi: \mathbb{R}^d \rightarrow \mathbb{R}^p$.

Due to the unobserved confounding, it may be the case that $Y - f^0(X)$ and Z are dependent and, as a result, restriction (3) may not hold for the true model M^0 . But as we assume that there is either no confounding between Z and W or between W and Y , we can instead use the *conditional independence restriction*

$$Z \perp\!\!\!\perp Y - \phi(X)^\top \theta \mid W \quad (24)$$

to identify f^0 . A corresponding sufficient and necessary identifiability condition for f^0 is

$$Z \perp\!\!\!\perp h(W, U, \epsilon_Y) + \phi(X)^\top \tau \mid W \implies \tau = 0.$$

Under this condition we could proceed as for the unconditional independence restriction and estimate f^0 using a loss that is minimized if restriction (7) is satisfied, e.g., using a conditional independence measure. We can avoid using a conditional independence restriction, in the following two special cases.

First, consider the case in which $W \perp\!\!\!\perp V$ and, in addition, assume $q(W, \epsilon_Z) = q_1(W) + \epsilon_Z$. Then we can use the following independence restriction

$$Z - q_1(W) \perp\!\!\!\perp Y - \phi(X)^\top \theta. \quad (25)$$

Under our assumptions, $q_1^0(W) = E[Z|W]$ and therefore q_1^0 is identifiable. The corresponding identifiability condition becomes

$$Z - q_1(W) \perp\!\!\!\perp h(W, U, \epsilon_Y) + \phi(X)^\top \tau \implies \tau = 0. \quad (26)$$

⁶The results can be extended to other SCMs, too; we only show one for simplicity. E.g., we could also allow for the edge between Z and W to point towards W , that is, for Z to be a cause of W .

Second, consider the case in which we instead assume that $W \perp\!\!\!\perp U$ and that $h(W, U, \epsilon_Y) = h_1(W) + h_2(U, \epsilon_Y)$. Then we can use the independence restriction:

$$Y - \phi(X)^\top \theta - \phi(W)^\top \gamma \perp\!\!\!\perp (Z, W), \quad (27)$$

where we assume that⁷ $h_1^0(\cdot) = \phi(\cdot)^\top \gamma^0$ for $\gamma^0 \in \mathbb{R}^t$. The identifiability condition becomes

$$(Z, W) \perp\!\!\!\perp h_2(U, \epsilon_Y) + \phi(X)^\top \tau_1 + \phi(W)^\top \tau_2 \implies \tau_1 = 0. \quad (28)$$

This identifiability condition translates to a corresponding moment restriction and as in the case of unconditional IV, we can again show that the independence version leads to strictly stronger identifiability.

Proposition B.1 (CIV - Identifiability of the independence restriction is strictly stronger). *Consider a conditional IV model M^0 and assume that $\phi(\cdot) = (\phi_1(\cdot), \dots, \phi_p(\cdot))$ is a collection of basis functions such that $f^0(\cdot) = \phi(\cdot)^\top \theta^0$ for some $\theta^0 \in \mathbb{R}^p$ and $h_1^0(\cdot) = \phi(\cdot)^\top \gamma^0$ for some $\gamma^0 \in \mathbb{R}^t$. Then, if there exists $k \in \mathbb{N}$ and $\eta : \mathbb{R}^{(r+t)} \rightarrow \mathbb{R}^k$ such that θ^0 is identifiable from the moment restriction*

$$\mathbb{E}[\eta(Z, W)(Y - \phi(X)^\top \theta - \phi(W)^\top \gamma)] = 0,$$

then θ^0 is also identifiable from the conditional independence restriction (27). Furthermore, there exist a conditional IV model such that θ^0 is identifiable from the independence restriction (27) but not from the corresponding moment restriction. An example for that is the model

$$\begin{aligned} V &:= \epsilon_V, U := \epsilon_U, W := V + \epsilon_W, Z := W + V + \epsilon_Z, X := ZU + \epsilon_X \\ Y &:= X + W + U + \epsilon_Y, \end{aligned}$$

where $\epsilon_Z, \epsilon_V, \epsilon_U, \epsilon_W, \epsilon_X$ and ϵ_Y are jointly independent standard Gaussian variables and we consider the base functions $\phi(X) = (X)$ and $\phi(W) = (W)$.

Proof. For the first claim, consider a θ^0 that is not identifiable from restriction (27). Then there exists a $\tau_1 \neq 0$ for which condition (28) is violated. By the independence statement for τ_1 , it holds for all $k \in \mathbb{N}$, functions $\eta : \mathbb{R}^{(r+t)} \rightarrow \mathbb{R}^k$ and some $\tau_2 \in \mathbb{R}$ that

$$\text{Cov}[\eta(Z, W), h(U, \epsilon_Y) + \phi(X)^\top \tau_1 + \phi(W)^\top \tau_2] = 0.$$

But as $(Z, W) \perp\!\!\!\perp h(U, \epsilon_Y)$ it also holds that $\text{Cov}[\eta(Z, W), h(U, \epsilon_Y)] = 0$. By the additivity of the covariance it follows that

$$\text{Cov}[\eta(Z, W), \phi(X)^\top \tau_1 + \phi(W)^\top \tau_2] = 0,$$

which is the identifiability condition corresponding to the moment restriction.

For the second claim, note that in our example SCM, $\mathbb{E}[X|(Z, W)] = 0$. It follows that for all $k \in \mathbb{N}$ and functions $\eta : \mathbb{R}^{(r+t)} \rightarrow \mathbb{R}^k$, $\text{Cov}[\eta(Z, W), X] = 0$. We can therefore not uniquely identify the causal function $f^0(x) = x$ with the moment restriction.

We now consider condition (28). Consider $K_\tau := U + \epsilon_Y + \tau_1(ZU + \epsilon_X) + \tau_2 W$. It holds that

$$\begin{aligned} \mathbb{E}[Z^2 K_\tau^2] &= \mathbb{E}[Z^2 (U + \epsilon_Y + \tau_1(ZU + \epsilon_X) + \tau_2 W)^2] \\ &= \mathbb{E}[Z^2 U^2] + \mathbb{E}[Z^2 \epsilon_Y^2] + \tau_1^2 \mathbb{E}[Z^4 U^2] + \tau_1^2 \mathbb{E}[Z^2 \epsilon_X^2] + \tau_2^2 \mathbb{E}[Z^2 W^2] \\ &= 12 + 114\tau_1^2 + 30\tau_2^2, \end{aligned}$$

while

$$\begin{aligned} \mathbb{E}[Z^2] \mathbb{E}[K_\tau^2] &= \mathbb{E}[Z^2] \mathbb{E}[(U + \epsilon_Y + \tau_1(ZU + \epsilon_X) + \tau_2 W)^2] \\ &= \mathbb{E}[Z^2] (\mathbb{E}[U^2] + \mathbb{E}[\epsilon_Y^2] + \tau_1^2 \mathbb{E}[Z^2 U^2] + \tau_1^2 \mathbb{E}[\epsilon_X^2] + \tau_2^2 \mathbb{E}[W^2]) \\ &= 6(2 + 7\tau_1^2 + 2\tau_2^2) \\ &= 12 + 42\tau_1^2 + 12\tau_2^2. \end{aligned}$$

These two terms are only equal if $72\tau_1^2 + 18\tau_2^2 = 0$, which requires that $\tau_1 = 0$. Therefore, condition (28) holds, meaning we can identify the causal function $f^0(x) = x$. \square

⁷For simplicity, we use the same basis ϕ .

We now propose two procedures; one for each of the identifiability conditions we discussed. We conjecture that combining both leads to a doubly robust estimator.

For restriction (25), we propose the following estimation procedure. Given an i.i.d. sample $(x_i, y_i, z_i, w_i)_{i=1}^n$ of the variables (X, Y, Z, W) , we can first use any flexible non-parametric estimator to estimate q_1^0 . Based on our estimate \hat{q}_1 we then learn \hat{f} as

$$\arg \min_{f \in \mathcal{F}} \widehat{\text{HSIC}}((r_i^f, z_i^{\hat{q}_1})_{i=1}^n; k_R, k_{Z^{\hat{q}_1}}),$$

as our estimate for f^0 , where $r_i^f = y_i - \hat{f}(x_i)$ and $z_i^{\hat{q}_1} = z_i - \hat{q}_1(w_i)$ and \mathcal{F} is a function class. Here, $\widehat{\text{HSIC}}$ is the same empirical HSIC estimator as in Section 3.

For restriction (27), we propose the following estimation procedure. Given an i.i.d. sample $(x_i, y_i, z_i, w_i)_{i=1}^n$ of the variables (X, Y, Z, W) , we estimate f^0 as the first entry of

$$\arg \min_{f \in \mathcal{F}, k \in \mathcal{K}} \widehat{\text{HSIC}}((r_i^{f,k}, (z_i, w_i))_{i=1}^n; k_{R_Y}, k_{Z,W}),$$

where $r_i^{f,k} = y_i - f(x_i) - k(w_i)$ and, \mathcal{F} and \mathcal{K} are function classes. Here, $\widehat{\text{HSIC}}$ is again the same empirical HSIC estimator as in Section 3.

C. Algorithms: HSIC-X and HSIC-X-pen

We provide details for HSIC-X in Algorithm 1 and details for HSIC-X-pen in Algorithm 2. We propose to choose the tuning parameter λ for HSIC-X-pen as the largest possible value for which an HSIC-based independence test between the estimated residuals and the instruments is not rejected (see Section 3.1).

Algorithm 1 HSIC-X

Input: observations $(x_i, y_i, z_i)_{i=1}^n$, kernels k_R and k_Z , batch size m , learning rate γ , number of gradient steps per cycle k , significance level α and maximum restarting times t

Initialize a restart counter $\ell \leftarrow 0$

Compute the bandwidth σ_Z of k_Z with median heuristic

repeat

if $\ell = 0$ **then**

 | Initialize parameters θ at the OLS solution

else

 | Randomly initialize parameters θ

end

repeat

 | Compute the residuals $r_i^\theta := y_i - f_\theta(x_i)$

 | Compute the bandwidth σ_{R^θ} of k_{R^θ} with median heuristic

repeat k **times**

 | Sample a mini-batch $(x_j, y_j, z_j)_{j=1}^m$

 | Compute the residuals $r_j^\theta := y_j - f_\theta(x_j)$

 | Compute the HSIC loss $\mathcal{L}(\theta) := \text{tr}(KHLH)$

 | Update parameters $\theta \leftarrow \theta - \gamma \nabla \mathcal{L}(\theta)$

end

until *convergence*

 | Compute the residuals $r_i^\theta := y_i - f_\theta(x_i)$

 | Compute the p-value p of the independence test between $(r_i^\theta)_{i=1}^n$ and $(z_i)_{i=1}^n$

 | Update the counter $\ell \leftarrow \ell + 1$

until $p \geq \alpha$ or $\ell \geq t$

Output: Final estimate $\tilde{f}_\theta(\cdot) := f_\theta(\cdot) - \frac{1}{n} \sum_{i=1}^n y_i - f_\theta(x_i)$

Algorithm 2 HSIC-X-pen

Input: observations $(x_i, y_i, z_i)_{i=1}^n$, kernels k_R and k_Z , batch size m , learning rate γ , number of gradient steps per cycle k , significance level α and maximum restarting times t , penalization parameter λ , prediction loss ℓ

Initialize a restart counter $l \leftarrow 0$

Compute kernel bandwidth σ_Z of k_Z with median heuristic

repeat

if $l = 0$ **then**

 | Initialize parameters θ at the OLS solution

else

 | Randomly initialize parameters θ

end

repeat

 | Compute the residuals $r_i^\theta := y_i - f_\theta(x_i)$

 | Compute kernel bandwidth σ_R of k_R with median heuristic

repeat k **times**

 | Sample a mini-batch $(x_j, y_j, z_j)_{j=1}^m$

 | Compute the residuals $r_j^\theta := y_j - f_\theta(x_j)$

 | Compute loss $\mathcal{L}(\theta) := \lambda \sum_j \ell(r_j^\theta) + (1 - \lambda) \text{tr}(KHLH)$

 | Update parameter $\theta \leftarrow \theta - \gamma \nabla \mathcal{L}(\theta)$

end

until *convergence*

 | Compute the residuals $r_i^\theta := y_i - f_\theta(x_i)$

 | Compute the p-value p of the independence test between $(r_i^\theta)_{i=1}^n$ and $(z_i)_{i=1}^n$

 | Update the counter $l \leftarrow l + 1$

until $p \geq \alpha$ or $l \geq t$

Output: Final estimate $\tilde{f}_\theta(\cdot) := f_\theta(\cdot) - \frac{1}{n} \sum_{i=1}^n y_i - f_\theta(x_i)$

D. Additional Experiment Details and Results

D.1. Multi-dimensional Setting

We consider the following IV models in our experiments:

$$M : \begin{cases} Z := \epsilon_Z \\ U := \epsilon_U \\ X := JZ^2\epsilon_X + BZ + U \\ Y := \beta^\top X - 2U + \epsilon_Y, \end{cases}$$

where $\epsilon_U, \epsilon_Y \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$, $\epsilon_Z \sim \mathcal{N}(0, I_{d_Z})$, $\epsilon_X \sim \mathcal{N}(0, I_{d_X})$ are independent noise variables, d_X and d_Z represent the dimensions of X and Z , J and B are $d_X \times d_Z$ matrices controlling the influence of the instruments Z and β is the causal parameters. In the experiment, J is an $d_X \times d_Z$ all-ones matrix, $B_{i,j} \sim \text{Uniform}(-4, 4)$ and $\beta \sim \mathcal{N}(0, I_{d_X})$.

Figure 4 reports the MSE of each method as dimension (d_Z) of the instruments varies for some fixed predictor's dimensionality (d_X). When $d_Z < d_X$, the identifiability result suggests that 2SLS cannot consistently estimate the causal function (which is reflected by the high values of the MSE); while HSIC-X can still gain some improvement over the OLS solution. When $d_Z \geq d_X$, the performance of our method is on par with that of 2SLS.

D.2. Known Basis Functions

Figure 5 shows some of the functions estimated by HSIC-X, OLS, 2SLS, along with the underlying causal function under various settings. In short, when the instrument has no effect on the mean of X ($\alpha = 0$), 2SLS fails to produce sensible estimates because of the non-identifiability under the moment restriction. On the other hand, the proposed method (HSIC-X) can still identify the causal function thanks to the independence restriction and yields reasonable estimates in all of the settings.




Figure 4. MSEs of different estimators under varying size d_Z of the instruments when (a) the predictors’ dimension d_X is 3, and (b) the predictors’ dimension is 5. 2SLS is inconsistent and underperforms OLS when $d_Z < d_X$, while HSIC-X shows a substantial improvement over OLS in such settings.

D.3. Distribution Generalization

In Section 5.2, we consider both linear and nonlinear underlying causal functions. The linear function is defined as $f_{\text{lin}}^0(X_1, X_2) := X_1 + X_2$, and the nonlinear function is defined as $f_{\text{nonlin}}^0(X) := \sum_{j=1}^{10} w_j e^{-\frac{1}{3}\|X - c_j\|_2^2}$, where c_j is drawn from a uniform distribution over a two-dimensional grid $[-5, 5] \times [-5, 5]$ and $w_1, \dots, w_{10} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 4)$. We employ the correct basis and neural network function classes in our method and the baselines. For the correct basis, we consider $\mathcal{F} := \{f(\cdot) = \phi(\cdot)^\top \theta \mid \theta \in \mathbb{R}^p\}$ with $p = 3$ and $\phi(x) = [1, x_1, x_2]$ in the linear case, and $p = 11$ and $\phi(x) = [1, e^{-\frac{1}{3}\|x - c_1\|_2^2}, \dots, e^{-\frac{1}{3}\|x - c_{10}\|_2^2}]$ in the nonlinear case. For the neural network function class, we use a neural network with one hidden layer of size 64. We optimize our model using Adam optimizer with the learning rate of 0.005 and batch-size of 256, and use the R package ‘AnchorRegression’ (<https://github.com/simzim96/AnchorRegression>) for the Anchor Regression baseline. Lastly, the tuning parameter γ of Anchor Regression is set to 100, and the tuning parameter λ for HSIC-X-pen is chosen as the largest possible value for which an HSIC-based independence test between the estimated residuals and the instruments is not rejected (see Section 3.1).

Exploiting Independent Instruments

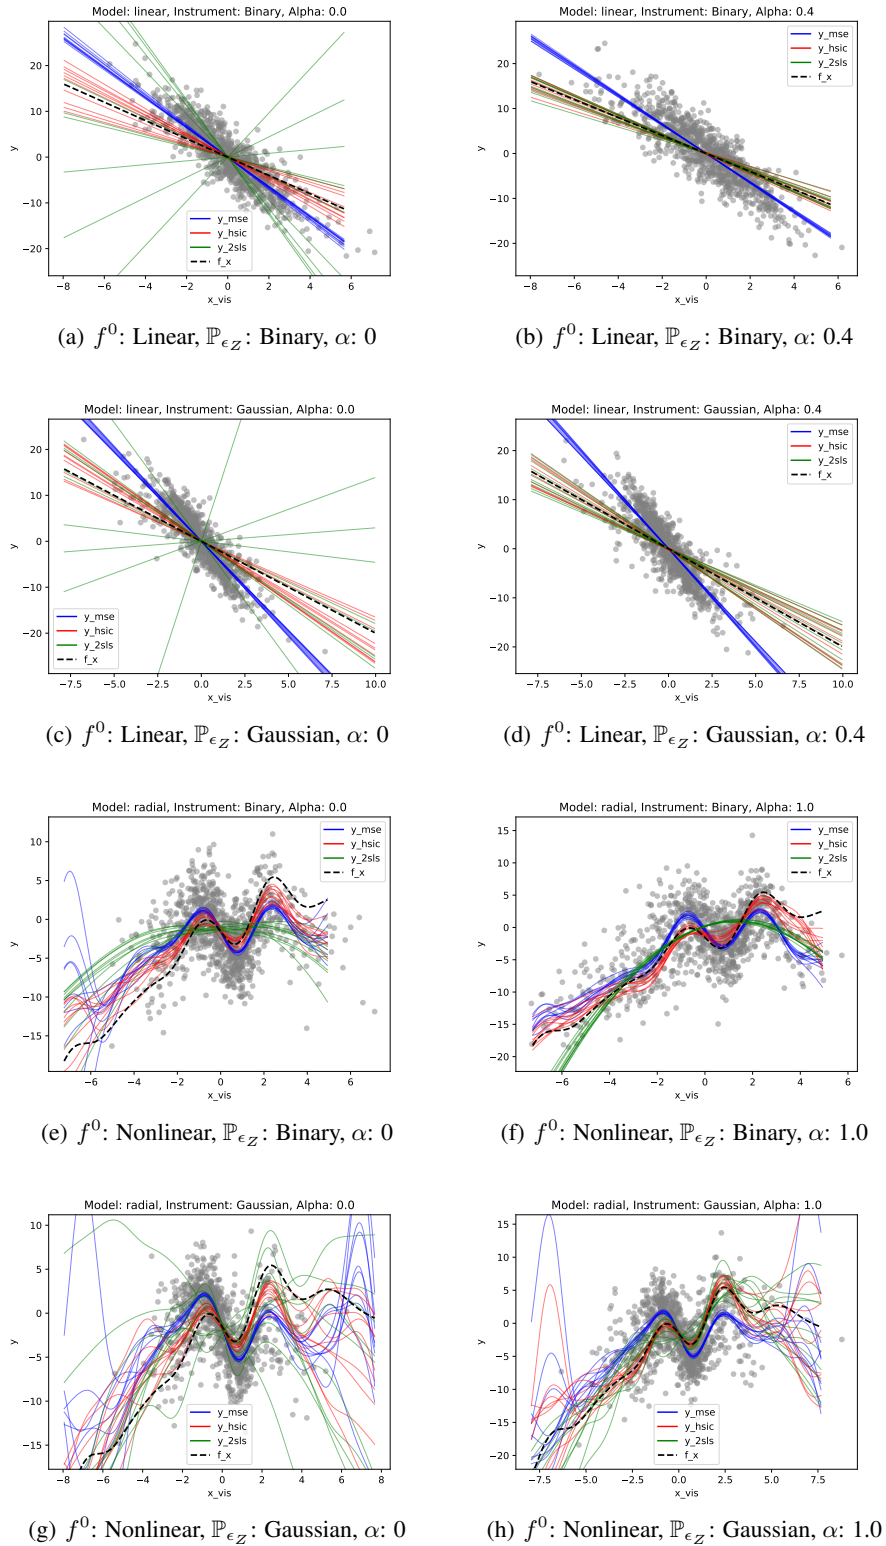


Figure 5. Estimated causal functions with different estimators under varying α values from the experiment in Section 5.1 (the known basis functions setting). The legend reads as follows: "y_mse" represents OLS, "y_hsic" represents HSIC-X, "y_2sls" represents 2SLS and "f_x" represents the underlying causal function. Each line represents an estimated function from one simulation run. The grey dots are the observations drawn from the corresponding IV model.