

Article

Not peer-reviewed version

Kernel-Based Independence Tests for Causal Structure Learning on Functional Data

[Felix Laumann](#)*, [Julius Von Kügelgen](#), Junhyung Park, Bernhard Schölkopf, [Mauricio Barahona](#)

Posted Date: 30 October 2023

doi: 10.20944/preprints202310.1816.v1

Keywords: causal discovery; independence tests; functional data analysis; kernel methods



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Article

Kernel-Based Independence Tests for Causal Structure Learning on Functional Data

Felix Laumann^{1,*}, Julius von Kügelgen^{2,3}, Junhyung Park², Bernhard Schölkopf² and Mauricio Barahona¹

¹ Department of Mathematics, Imperial College London, London, UK

² MPI for Intelligent Systems, Tübingen, Germany

³ Department of Engineering, University of Cambridge, Cambridge, UK

* Correspondence: f.laumann18@imperial.ac.uk

Abstract: Measurements of systems taken along a continuous functional dimension, such as time or space, are ubiquitous in many fields, from the physical and biological sciences to economics and engineering. Such measurements can be viewed as realisations of an underlying smooth process sampled over the continuum. However, traditional methods for independence testing and causal learning are not directly applicable to such data, as they do not take into account the dependence along the functional dimension. By using specifically designed kernels, we introduce statistical tests for bivariate, joint, and conditional independence for functional variables. Our method not only extends the applicability to functional data of the Hilbert-Schmidt independence criterion (HSIC) and its d -variate version (d -HSIC), but also allows us to introduce a test for conditional independence by defining a novel statistic for the conditional permutation test (CPT) based on the Hilbert-Schmidt conditional independence criterion (HSCIC), with optimised regularisation strength estimated through an evaluation rejection rate. Our empirical results of the size and power of these tests on synthetic functional data show good performance, and we then exemplify their application to several constraint- and regression-based causal structure learning problems, including both synthetic examples and real socio-economic data.

Keywords: causal discovery; independence tests; functional data analysis; kernel methods

1. Introduction

Uncovering the causal relationships between measured variables, a discipline known as *causal structure learning* or *causal discovery*, is of great importance across various scientific fields, such as climatology [1], economics [2], or biology [3]. Doing so from passively collected ('observational') data enables the inference of causal interactions between variables without performing experiments or randomised control trials, which are often expensive, unethical, or impossible to conduct [4]. Causal structure learning is the inference, under a given set of assumptions, of directed and undirected edges in graphs representing the data generating process, where the nodes represent variables and the inferred edges capture causal (directed) or non-causal (undirected) relationships between them.

Much research in various areas collates *functional* data consisting of multiple series of measurements that are observed conjointly over a given continuum (e.g., time, space, or frequency), where each series is assumed to be a realisation of an underlying *smooth* process [5, §3]. By viewing the series of measurements as discretisations of functions, the observations are not required to be collected over regular meshes of points along the continuum. If the variables are measured over time as the underlying continuum, there is a long history of methods that have been developed to infer (time-based) causality between variables. Among those, the classic Granger-causality [6] declares that variable 'X causes Y' ($X \rightarrow Y$) if predicting the future of Y becomes more accurate with, as compared to without, access to the past of X, conditional on all other relevant variables. However, these methods

assume that the observed time-series are stationary, the causal dependency of X on Y is linear, and no confounding variables are considered. More recently, Sugihara et al. [7] developed convergent cross mappings (CCMs), a method that relaxes the assumption of linearity and finds causal relationships based on time-embeddings of the (stationary) time-series at each point, while also assuming no confounding variables are present. While useful in many situations, as we will show, Granger-causality and CCM can perform weakly when the time-series for X and Y are non-stationary and related through non-linear relationships, meaning that the influence of X on Y varies over time. Here we present a method that uses kernel-based independence tests to detect statistically significant causal relationships by extending constraint- and regression-based causal structure learning to functional data. The key advantage over Granger-causality and CCM is the consideration of confounders and the relaxation of their assumptions around linear relationships or stationarity in the data, which can lead to different causal relationships between variables. As a motivating example, consider the relationship between two variables: ‘corruption’ and ‘income inequality’, as measured by the World Governance Indicators (WGIs) [8] and the World Bank [9], respectively. Using data for 48 African countries from 1996 to 2016, Sulemana and Kpienbaareh [2] investigated their cause-effect relationship and found that corruption ‘Granger-causes’ income inequality. Independently, we have also confirmed that applying CCM to the same data leads to the same conclusion. However, by considering the time-series data as realisations of functions over time, and thus avoiding linearity and stationarity assumptions, our proposed kernel-based approach suggests the reverse result, i.e., causal influence of income inequality on corruption appears as the more statistically likely direction. This is in agreement with other quantitative findings, which draw on different data sources [10–12]. We will return to this example in Section 4.2.2 where we analyse causal dependencies between all six WGIs.

Methodologically, we extend the applicability of two popular paradigms in causal structure learning—*constraint-based* [13, § 5] and *regression-based* methods [14]—to functional data. Independence tests play a crucial role in uncovering causal relationships in both paradigms, and kernels provide a powerful framework for such tests by embedding probability distributions in reproducing kernel Hilbert spaces (RKHS) [15, § 2.2]. Until now, however, related methods for causal learning had only been applicable to univariate and multivariate data, but not to functional data. Here, we address this limitation by using recently derived kernels over functions [16] to embed probability distributions over functions in RKHS, thus widening the applicability of kernel-based independence tests to functional data settings. To test for conditional independence, we compute HSCIC [17] in a conditional permutation test (CPT) [18], and we propose a straightforward search to determine the optimised regularisation rate in HSCIC.

We structure our paper as follows. Section 2 provides a brief overview of prior literature on functional data analysis, kernel-based independence tests, and causal structure learning methods. Section 3 presents our methodological contributions, from the definition of a conditional independence test for functional data to causal structure learning on such data. We then empirically analyse the performance of our independence tests and causal structure learning algorithms on synthetic and real-world data in Section 4. We conclude with a discussion in Section 5.

We highlight the following contributions:

- In Section 3 we propose a conditional independence test for functional data that combines a novel test statistic based on HSCIC with CPT to generate samples under the null hypothesis. The algorithm also searches for the optimised regularisation strength λ^* required to compute HSCIC, by pre-test permutations and calculating an *evaluation rejection rate*.
- In Section 4.1.2, we extend the historical functional linear model [19] to the multivariate case ($\{X_1, \dots, X_i, \dots, X_d\} \rightarrow Y$) for regression-based causal structure learning, and we show how a joint independence test can be used to verify candidate directed acyclic graphs (DAGs) [20,

§ 5.2] that embed the causal structure of function-valued random variables. This model has been contributed to the Python package `scikit-fda` [21].

- On synthetic data, we show empirically that our bivariate, joint and conditional independence tests achieve high test power, and that our causal structure learning algorithms outperform previously proposed methods.
- Using a real-world data set (World Governance Indicators) we demonstrate how our method can yield insights into cause-effect relationships amongst socio-economic variables measured in every country worldwide.
- Implementations of our algorithms are made available at <https://github.com/felix-laumann/causal-fda/> in an easily usable format that builds on top of `scikit-fda` and `causal DAG` [22].

2. Background and Related Work

2.1. Functional Data Analysis

In functional data analysis [5], a variable X is described by a set of n samples (or realisations), $\mathbf{X} = \{x_i(s)\}_{i=1}^n$, where each functional sample $x_i(s)$ corresponds to a series of observations over the continuum s , also called the *functional dimension*. Typical functional dimensions are time or space. In practical settings, the observations are taken at a set of S discrete values s_1, \dots, s_S of the continuum variable s . Examples of functional data sets include the vertical position of the lower lip over time when speaking out a given word [19], the muscle soreness over the duration of a tennis match [23], or the ultrafine particle concentration in air measured over the distance to the nearest motorway [24].

In applications, the functional samples are usually represented as linear combinations of a finite set of M basis functions $\{\phi_m(s)\}_{m=1}^M$ (e.g., Fourier or monomial basis functions):

$$x_i(s) = \sum_{m=1}^M c_{i,m} \phi_m(s), \quad (1)$$

where the coefficients $C_i = (c_{i,1}, c_{i,2}, \dots, c_{i,M})$ characterise each sample. If the number of basis functions is equal to the number of observations, $M = S$, each observed value $x_i(s_k)$ can be fitted exactly by obtaining the coefficients C_i using standard non-linear least squares fitting techniques (provided the $\phi_m(s)$ are valid basis functions), and Equation (1) allows us to interpolate between any two observations. When the number of basis functions is smaller than the number of observations, $M < S$, as it is commonly the case in practice, the basis function expansion (1) provides a smoothed approximation to the set of observations, $\hat{x}_i(s_k)$.

For the many applications where the continuum is time, historical functional linear models (HFLMs) [19] provide a comprehensive framework to map the relationship between two sets of functional samples. Let 0 and T be the initial and final time points for a set of samples $y_i(t)$. HFLMs describe dependencies that can vary over time using the function $\beta(s, t)$, which encapsulates the influence of $x(s), s \in [0, S]$ on another variable $y(t), t \in [0, T]$ at any two points in time, s_k and t_k :

$$y_i(t) = \int_{s_0(t)}^t x_i(s) \beta(s, t) ds, \quad t \in [0, T] \quad (2)$$

where $s_0(t)$ is the maximum allowable lag for any influence of X on Y and $s_0(t) \leq s \leq t$. Typical choices for $\beta(s, t)$ are exponential decay and hyperbolic paraboloid (or “saddle”) functions. The continuum is not required to be time but can also be space, frequency or others (see Ramsay and Hooker [25] for an extensive collection of function-to-function models and applications).

2.2. Kernel Independence Tests

Let \mathcal{H}_{k_1} and \mathcal{H}_{k_2} be separable RKHSs with kernels $k_1 : \mathbb{R}^X \times \mathbb{R}^X \rightarrow \mathbb{R}$ and $k_2 : \mathbb{R}^Y \times \mathbb{R}^Y \rightarrow \mathbb{R}$ such that the tensor product kernel $(k_1 \otimes k_2)(x, y, x', y') = k_1(x, x') \cdot k_2(y, y')$ implies $(\mathbb{R}^X \times \mathbb{R}^Y) \times (\mathbb{R}^X \times \mathbb{R}^Y) \rightarrow \mathbb{R}$. If the kernel k_1 uniquely embeds a distribution \mathbb{P}_X in an RKHS by a mean embedding,

$$\mu_X := \int_{\mathcal{X}} k_1(x, \cdot) d\mathbb{P}_X \quad (3)$$

which captures any information about \mathbb{P}_X , we call k_1 a characteristic kernel on $\mathbb{R}^X \times \mathbb{R}^Y$ [26]. Characteristic kernels have thus been extensively used in bivariate ($\mathbb{P}_{XY} = \mathbb{P}_X \mathbb{P}_Y$), joint ($\mathbb{P}_{XYZ} = \mathbb{P}_X \mathbb{P}_Y \mathbb{P}_Z$) and conditional ($\mathbb{P}_{XY|Z} = \mathbb{P}_{X|Z} \mathbb{P}_{Y|Z}$) independence tests [e.g., 20,27,28].

For the bivariate independence test, let \mathbb{P}_{XY} denote the joint distribution of X and Y . Then HSIC is defined as

$$\text{HSIC}(\mathcal{H}_{k_1}, \mathcal{H}_{k_2}, \mathbb{P}_{XY}) := \|\mu_{XY} - \mu_X \otimes \mu_Y\|_{\mathcal{H}_{k_1} \otimes \mathcal{H}_{k_2}}^2 \geq 0, \quad (4)$$

with equality iff $\mathbb{P}_{XY} = \mathbb{P}_X \mathbb{P}_Y$.

We refer to Gretton et al. [27] for the definition of an estimator for finite samples that constitutes the test statistic in the bivariate independence test with null hypothesis $H_0 : X \perp Y$. The test statistic is then computed on the original data (X, Y) and statistically compared to random permutations $\{(X, Y_{p'})\}_{p'=1}^P$ under the null hypothesis.

For distributions with more than two variables, let $\mathbb{P}_{X_1, X_2, \dots, X_d}$ denote the joint distribution on X_1, X_2, \dots, X_d . To test for joint independence we compute

$$\begin{aligned} d\text{-HSIC}(\mathcal{H}_{k_1}, \mathcal{H}_{k_2}, \dots, \mathcal{H}_{k_d}, \mathbb{P}_{X_1, X_2, \dots, X_d}) := \\ \|\mu_{X_1, X_2, \dots, X_d} - \mu_{X_1} \otimes \mu_{X_2} \otimes \dots \otimes \mu_{X_d}\|_{\mathcal{H}_{k_1} \otimes \mathcal{H}_{k_2} \otimes \dots \otimes \mathcal{H}_{k_d}}^2 \geq 0 \\ \text{with equality iff } \mathbb{P}_{X_1, X_2, \dots, X_d} = \mathbb{P}_{X_1} \mathbb{P}_{X_2} \dots \mathbb{P}_{X_d}. \end{aligned} \quad (5)$$

Pfister et al. [20] derive a numerical estimator, which serves as the basis for a joint independence test on finite samples. Here, the distribution under the null hypothesis of joint independence is generated by randomly permuting all sample sets in the same way as Y is in the bivariate independence test.

Lastly, the conditional independence test relies on accurately sampling from the distribution under the null hypothesis $H_0 : X \perp Y | Z$. At the core of the conditional permutation test [18] lies a randomisation procedure that generates permutations of X , denoted $\{X_{p'}\}_{p'=1}^P$, which are generated without altering the conditional distribution $\mathbb{P}_{X|Z}$, so

$$\mathbb{P}_{X_{p'}, Y, Z} = \mathbb{P}_{X, Y, Z} \quad (6)$$

under H_0 while breaking any dependence between X and Y . The null distribution can therefore be generated by repeating this procedure multiple times, and we can decide whether H_0 should be rejected by comparing a test statistic on the original data against its results on the generated null distribution.

The existing literature on kernel-based independence tests is extensive, see e.g., Berrett et al. [18] for a relevant review, but only a small part of those tests investigates independence among functional variables [29,30]. There have been particularly strong efforts in developing conditional independence tests and understanding their applicable settings. The kernel conditional independence test (KCIT) [28], for example, gave promising results in use with univariate data but increasingly suffered when the number of conditional variables was large. In contrast, the kernel conditional

independence permutation test (KCIPT) [31] repurposed the well-established kernel two-sample test [32] to a conditional independence setting which delivered stable results for multiple conditional variables. However, Lee and Honavar [33] pointed out that as the number of permutations increases, while its power increases, its calibratedness decreases. This issue was overcome by their proposed self-discrepancy conditional independence test (SDCIT), which is based on a modified unbiased estimate of the maximum mean discrepancy (MMD).

2.3. Causal Structure Learning

The aim of causal structure learning, or causal discovery, is to infer the qualitative causal relationships among a set of observed variables, typically in the form of a causal diagram or DAG. Once learnt, such a causal structure can then be used to construct a causal model such as a causal Bayesian network or a *structural causal model* (SCM) [34]. Causal models are endowed with a notion of manipulation and, unlike a statistical model, do not just describe a single distribution, but many distributions indexed by different interventions and counterfactuals. They can be used for causal reasoning, that is, to answer causal questions such as computing the average causal effect of a treatment on a given outcome variable. Such questions are of interest across many disciplines, and causal discovery is thus a highly topical area. We refer to Glymour et al. [4], Mooij et al. [35], Peters et al. [36], Schölkopf and von Kügelgen [37], Squires and Uhler [38], Vowels et al. [39] for comprehensive surveys and accounts of the main research concepts. In particular, we focus here on causal discovery methods for causally sufficient systems, for which there are no unobserved confounders influencing two or more of the observed variables. Existing causal discovery methods can roughly be categorised into three families:

- Score-based approaches assign a score, such as a penalised likelihood, to each candidate graph and then pick the highest scoring graph(s). A common drawback of score-based approaches is the need for a combinatorial enumeration of all DAGs in the optimisation, although greedy approaches have been proposed to alleviate such issues [40].
- Constraint-based methods start by characterising the set of *conditional independences* in the observed data [13]. They then determine the graph(s) consistent with the detected conditional independences by using a graphical criterion called d-separation, as well as the causal Markov and faithfulness assumptions, which establish a one-to-one connection between d-separation and conditional independence (see Appendix D for definitions). When only observational i.i.d. data are available, this yields a so-called Markov equivalence class, possibly containing multiple candidate graphs. For example, the graphs $X \rightarrow Y \rightarrow Z$, $X \leftarrow Y \leftarrow Z$, and $X \leftarrow Y \rightarrow Z$ are Markov equivalent, as they all imply $X \perp\!\!\!\perp Z | Y$ and no other conditional independence relations.
- Regression-based approaches directly fit the structural equations $X_i := f_i(\mathbf{PA}_i, U_i)$ of an underlying SCM for each X_i , where \mathbf{PA}_i denote the parents of X_i in the causal DAG and $U = (U_1, \dots, U_n)$ are *jointly independent* exogenous noise variables. Provided that the function class of the f_i is sufficiently restricted, e.g., by considering only linear relationships [41] or additive noise models [42], the true causal graph is identified as the unique choice of parents for each i such that the resulting residuals $\tilde{U}_i = X_i - \hat{f}_i(\mathbf{PA}_i)$ are jointly independent.

As can be seen from these definitions, conditional, bivariate, and joint independence tests are an integral part of constraint- and regression-based causal discovery methods. Our main focus in the present work is therefore to extend the applicability of these causal discovery frameworks to functional data by generalising the underlying independence tests to such domains.

3. Methods

In all three of our independence tests (bivariate, joint, conditional), we employ kernels over functions, also known as squared-exponential T (SE-T) kernels [16]. Let \mathcal{X} and \mathcal{Y} be real, separable Hilbert spaces with norms $\|\cdot\|_{\mathcal{X}}$ and $\|\cdot\|_{\mathcal{Y}}$, respectively. Then, for $T : \mathcal{X} \rightarrow \mathcal{Y}$, SE-T kernels are defined as

$$k_T(x, y) = e^{-\frac{1}{2\sigma^2} \|T(x) - T(y)\|_{\mathcal{Y}}^2}. \quad (7)$$

where σ^2 is commonly defined as the median heuristic, $\sigma^2 = \text{Median}\{\|T(a) - T(b)\|_{\mathcal{Y}}^2 : a, b \in \{x_i\}_{i=1}^n \cup \{y_i\}_{i=1}^n, a \neq b\}$. Replacing any characteristic kernel by the SE-T kernels for bivariate independence tests (based on HSIC) and for joint independence tests (based on d -variable Hilbert-Schmidt independence criterion (d -HSIC)) is straightforward and does not require further theoretical investigation besides evaluating numerically the validity and power of the tests (see Section 4). However, the application of SE-T kernels to conditional independence tests needs further theoretical results, as we discuss now.

3.1. Conditional independence test on functional data

We consider the conditional independence test, which generates samples under the null hypothesis based on the CPT, and uses the sum of HSCICs over all samples $z \in \mathbf{Z}$ as its test statistic. The CPT defines a permutation procedure that preserves the dependence of Z on both X and Y while resampling data for X that eliminates any potential dependence between X and Y . This procedure results in samples according to the null hypothesis, $X \perp\!\!\!\perp Y|Z$. We use this procedure whenever permutations are required as part of the conditional independence test. Given the computation of the HSCIC is based on a kernel ridge regression, it requires to set a regularisation strength λ , which has been manually chosen previously [17]. Generally, our aim is to have type-I error rates close to the allowable false positive rate α . However, choosing λ inappropriately may result in an invalid test (type-I error rates exceed α if λ is chosen too large), or in a deflated test power (type-I error rates are well below α and type-II error rates are high if λ is chosen too small). Thus, we must define an algorithm that conducts a search over a range of potentially suitable values for λ and assesses each candidate value by, what we will call, an *evaluation rejection rate*.

The search proceeds by iterating over a range of values $\{\lambda_l\}_{l=1}^L$ to find the optimised value λ^* , as follows. For each λ_l , we start by producing one permutation of the samples \mathbf{X} , which we denote as \mathbf{X}_b , and we compute its corresponding *evaluation test statistic* given by the sum of HSCICs over all samples $z \in \mathbf{Z}$. We then apply the usual strategy for permutation-based statistical tests: we produce an additional set of P permuted sample sets of \mathbf{X} , which we denote $\{\mathbf{X}_\pi\}_{\pi=1}^P$, and for each \mathbf{X}_π , we determine the sum of HSCICs over $z \in \mathbf{Z}$ to generate a distribution over statistics under the null hypothesis, which we call the *evaluation null statistics*. We then compute the percentile where the evaluation test statistic on \mathbf{X}_b falls within the distribution of evaluation null statistics on the permutation set $\{\mathbf{X}_\pi\}_{\pi=1}^P$. This results in an evaluation p-value which is compared to the allowable false positive rate α to establish if the hypothesis $H_0 : \mathbf{X}_b \perp\!\!\!\perp \mathbf{Y}|\mathbf{Z}$ can be rejected. Given that both the evaluation test statistic and the evaluation null statistics are computed on conditionally independent samples, we repeat this procedure for $b = 1, \dots, B \geq 100$ times to estimate an *evaluation rejection rate* for each value of λ_l . Having completed this procedure over all values $\{\lambda_l\}_{l=1}^L$, we select the λ_l that produces an evaluation rejection rate closest to α as the optimised regularisation strength, λ^* . Finally, we apply a CPT-based conditional independence test using the optimised λ^* to test the null hypothesis $H_0 : X \perp\!\!\!\perp Y|Z$. This entire procedure is summarised in Algorithm 1.

Algorithm 1 Search for λ^* with subsequent conditional independence test

Require: Samples $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, Range $\{\lambda_l\}_{l=1}^L$, Significance level α , Permutation iterations P , Rejection iterations B

Initialise: evaluation rejection rate $[\lambda^*] = 0$

```

1: for  $1 \leq l \leq L$  do ▷ Start: Search for  $\lambda^*$ 
2:   for  $1 \leq b \leq B$  do
3:     Permute  $\mathbf{X}$  by CPT, call them  $\mathbf{X}_b$ 
4:     evaluation test statistic $[b] \leftarrow \sum_{z \in \mathbf{Z}} \text{HSCIC}(\mathbf{X}_b, \mathbf{Y}, \mathbf{Z}, \lambda_l)$ 
5:     for  $1 \leq \pi \leq P$  do
6:       Permute  $\mathbf{X}_b$  by CPT, call them  $\mathbf{X}_{b,\pi}$ 
7:       evaluation null statistics $[b, \pi] \leftarrow \sum_{z \in \mathbf{Z}} \text{HSCIC}(\mathbf{X}_{b,\pi}, \mathbf{Y}, \mathbf{Z}, \lambda_l)$ 
8:       evaluation p-value  $\leftarrow 1 -$ 
9:         percentile(evaluation test statistic $[b]$ , evaluation null statistics $[b, :]$ )
10:      if evaluation p-value  $\geq \alpha$  then
11:        Fail to reject  $H_0 : \mathbf{X}_b \perp\!\!\!\perp \mathbf{Y} | \mathbf{Z}$ 
12:        rejects $[b] = 0$ 
13:      else
14:        Reject  $H_0 : \mathbf{X}_b \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{Z}$ 
15:        rejects $[b] = 1$ 
16:      evaluation rejection rate $[\lambda_l] \leftarrow \text{mean}_{1 \leq b \leq B}(\text{rejects})$ 
17:      if  $|\text{evaluation rejection rate}[\lambda^*] - \alpha| \geq |\text{evaluation rejection rate}[\lambda_l] - \alpha|$  then
18: return  $\lambda^* \leftarrow \lambda_l$  ▷ End: Search for  $\lambda^*$ 

```

```

20: test statistic  $\leftarrow \sum_{z \in \mathbf{Z}} \text{HSCIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \lambda^*)$  ▷ Start: Conditional independence test
21: for  $1 \leq p' \leq P$  do
22:   Permute  $\mathbf{X}$  by CPT, call them  $\mathbf{X}_{p'}$ 
23:   null statistics $[p'] \leftarrow \sum_{z \in \mathbf{Z}} \text{HSCIC}(\mathbf{X}_{p'}, \mathbf{Y}, \mathbf{Z}, \lambda^*)$ 
24: p-value  $\leftarrow 1 - \text{percentile}(\text{test statistic}, \text{null statistics})$ 
25: if p-value  $\geq \alpha$  then
26:   Fail to reject  $H_0 : X \perp\!\!\!\perp Y | Z$ 
27: else
28:   Reject  $H_0 : X \not\perp\!\!\!\perp Y | Z$  ▷ End: Conditional independence test

```

The following Theorem 1 guarantees the consistency of the conditional independence test in Algorithm 1 with respect to the regularisation parameter λ .

Theorem 1. Let \mathcal{H}_{k_1} and \mathcal{H}_{k_2} be separable RKHSs with kernels $k_1 : \mathbb{R}^X \times \mathbb{R}^X \rightarrow \mathbb{R}$ and $k_2 : \mathbb{R}^Y \times \mathbb{R}^Y \rightarrow \mathbb{R}$ such that the tensor product kernel $k_1 \otimes k_2 : (\mathbb{R}^X \times \mathbb{R}^Y) \times (\mathbb{R}^X \times \mathbb{R}^Y) \rightarrow \mathbb{R}$ is a characteristic kernel on $\mathbb{R}^X \times \mathbb{R}^Y$ [43]. If the regularisation parameter $\lambda = \lambda_n$ decays as $n \rightarrow \infty$ at a slower rate than $n^{-1/2}$, then the test based on the test statistic in Algorithm 1 (lines 19–27) is consistent.

Proof. See Appendix B. \square

3.2. Causal structure learning on functional data

To infer the existence of (directed) edges in an underlying causal graph G based on the joint data distribution \mathbb{P} over a set of observed variables, we must assume that G and \mathbb{P} are intrinsically linked. Spirtes et al. [13, § 2.3.3] define the *faithfulness* assumption, from which it follows that if two random variables are (conditionally) independent in the observed distribution \mathbb{P} , then they are d-separated in the underlying causal graph G [34, § 1.2.3]. Using this fact, constraint-based causal structure learning methods [13, § 5] take advantage of bivariate and conditional independence tests to infer whether two nodes are d-separated in the underlying graph. These methods yield completed partially directed acyclic graphs (CPDAGs), which are graphs with undirected and/or directed edges.

In contrast, regression-based methods [14], utilise the joint independence test to discover DAGs that have all edges oriented. We now describe both of these approaches in more detail.

3.2.1. Constraint-based causal structure learning

Constraint-based causal structure learning relies on performing conditional independence tests for each pair of variables, X and Y , conditioned on any possible subset of the remaining variables, that is, any subset within the power set of the $d \setminus \{X, Y\}$ remaining variables Z_1, Z_2, \dots, Z_* . Therefore, for d variables we need to carry out 2^{d-2} conditional independence tests for every pair of variables, which results in $\binom{d}{2}2^{d-2}$ tests in total.

Conditional dependences found in the data can then be used to delete and orient edges in the graph G , as follows. We start with a complete graph with d nodes, where each node corresponds to one of the variables. The edge connecting X and Y is deleted if there exists a subset \mathcal{Z}_1 of the $d - 2$ remaining variables such that $Z_* \notin \mathcal{Z}_1$ and $X \perp\!\!\!\perp Y | \mathcal{Z}_1$. If we then find another subset \mathcal{Z}_2 , such that $Z_* \in \mathcal{Z}_2$, $\mathcal{Z}_1 = \mathcal{Z}_2 \setminus \{Z_*\}$, and for which $X \not\perp\!\!\!\perp Y | \mathcal{Z}_2$, we can orient the edges $X - Z$ and $Y - Z$ to form colliders (or v-structures), $X \rightarrow Z_* \leftarrow Y$.

Based on these oriented edges, ‘Meek’s orientation rules’ defined by Meek [44] [see e.g., 36, § 7.2.1] can be applied to direct additional edges based on certain graphical compositions, as shown in Appendix F. Briefly, these rules follow from the avoidance of cycles, which would violate the acyclicity of a DAG, and colliders, which violate the found conditional independence. Algorithms that implement these criteria are the SGS algorithm and the more efficient PC algorithm [13], which we summarise in Appendix F.

3.2.2. Regression-based causal structure learning

To carry out our regression-based causal learning, we choose additive noise models (ANMs), a special class of SCMs where the noise terms are assumed to be independent and additive. This assumption guarantees that the causal graph can be identified if the function f_i is non-linear with Gaussian additive noise [42,45], a typical setting in functional data. Our model over the set of random variables $\{X\}_{i=1}^d$ is given by:

$$\mathbf{X}_i := f_i(\mathbf{PA}_i) + U_i, \quad (8)$$

where the additive noise terms U_i are jointly independent, i.e., any noise term U_i is independent of any intersection of the other noise terms,

$$\mathbb{P}(\cap_{j=1}^k U_{i_j}) = \prod_{j=1}^k \mathbb{P}(U_{i_j}), \quad \forall k \leq d.$$

Based on these assumptions, we follow RESIT [Regression with Subsequent Independence Testing; 14], an approach to causal discovery that can be briefly described as follows. If we only have two variables X and Y , we start by assuming that the samples in \mathbf{Y} are the effect and the samples in \mathbf{X} are the cause. Therefore we write \mathbf{Y} as a function \hat{f}_Y of \mathbf{X} plus some noise, and we test whether the residual $r_Y = \mathbf{Y} - \hat{f}_Y(\mathbf{X})$ is independent of \mathbf{X} . We then exchange the roles of assumed effect and assumed cause to obtain the residual $r_X = \mathbf{X} - \hat{f}_X(\mathbf{Y})$, which is tested for independence from \mathbf{Y} . To overcome issues with finite samples and test power, Peters et al. [14] used the p-value of the independence tests as a measure of strength of independence, which we follow in our experiments in Section 4. Alternatively, one may determine the causal direction for this bivariate case by comparing both cause-effect directions with respect to a score defined by Bühlmann et al. [46].

For d variables, the joint independence test replaces the bivariate independence test. Firstly, we consider the set of candidate causal graphs, $\Gamma = \{G_1, \dots, G_\gamma, \dots, G_\Gamma\}$, which contains every potential DAG with d nodes. For each candidate DAG, G_γ , we regress each descendant \mathbf{X}_i on its parents \mathbf{PA}_i :

$$\mathbf{X}_i := \sum_{k \in \mathbf{PA}_i} \hat{f}_{i,k}(\mathbf{X}_k) + U_i, \quad i \in 1, \dots, d, \quad (9)$$

and we compute the residuals $r_{\mathbf{X}_i} = \mathbf{X}_i - \sum_{k \in \mathbf{PA}_i} \hat{f}_{i,k}(\mathbf{X}_k)$. We then apply the joint independence test to all d residuals. The candidate DAG G_γ is accepted as the ‘true’ causal graph if the null hypothesis of joint independence amongst the residuals is not rejected. This process is repeated over all candidate causal graphs in the set Γ . Again, because in finite samples this procedure may not always lead to only one candidate DAG being accepted, the one with the highest p-value is chosen [14].

4. Experiments

In practice, functional data analysis requires samples of continuous functions evaluated over a mesh of discrete observations. In our synthetic data, we consider the interval $s = [0, 1]$ for our functional dimension and draw 100 equally spaced points over s to evaluate the functions. For real-world data, we map the space in which the data live (e.g., the years 1996 to 2020) to the interval $s = [0, 1]$ and interpolate over the discrete measurement points. We then use this interpolation to evaluate the functional samples on 100 equally spaced points. Unless otherwise mentioned, henceforth we use SE-T kernels with $T = I$ where I is the identity matrix and σ^2 as the median heuristic, as previously defined in Section 3.

4.1. Evaluation of independence tests for functional data

Before applying our proposed independence tests to causal structure learning problems, we use appropriate synthetic data to evaluate their type-I error rate (test size) and type-II error rate (test power). Type-I errors are made when the samples are independent (achieved by setting $a = 0$ in our experiments below) but the test concludes that they are dependent, i.e., the test falsely rejects the null hypothesis H_0 . In contrast, type-II errors appear when the samples are dependent but the test fails to reject H_0 (achieved by setting $a > 0$). Specifically, we compute the error rates on 200 independent trials, which correspond to different random realisations of the data sets, i.e., synthetically generated data sets with random Fourier basis coefficients, random coefficients of the β -functions and random additive noise. We set the test significance level at $\alpha = 0.05$ and approximate the null hypothesis of independence by 1000 permutations using the CPT scheme described above. Note that although our synthetic data are designed to replicate typical behaviour in time-series, the independence tests are not limited to time as the functional dimension, and can be used more generally.

4.1.1. Bivariate independence test

We consider n functional samples $\{x_i(s)\}_{i=1}^n$ of a random variable X defined over the interval $s = [0, 1]$. To generate our samples, we sum $M = 3$ Fourier basis functions $\phi_{m,T}(s)$ (definitions in Appendix A.1) with period $T = 0.1$ and coefficients $c_{i,m}$ randomly drawn from a standard normal distribution:

$$x_i(s) = \sum_{m=1}^M c_{i,m} \phi_{m,T}(s). \quad (10)$$

To mimic real-world measurements, we draw each sample over a distinct, irregular mesh of evaluation points. We then interpolate them with splines, and generate realisations over a regular mesh of points. Multivariate random noise $\epsilon_i^X(s) \sim \mathcal{N}(0, I)$, where I is the identity matrix, is then added. The samples $\{y_i(t)\}_{i=1}^n$ of random variable Y are defined as a HFLM [19] over $t = [0, 1]$ by:

$$y_i(t) = a \int_0^t x_i(s) \beta(s, t) ds, \quad (11)$$

where $a \in [0, 1]$. The function $\beta(s, t)$ maps the dependence from x_i at any point s to y_i at any point t , and is defined here as a hyperbolic paraboloid function:

$$\beta(s, t) = 8(s - c_1)^2 - 8(t - c_2)^2, \quad (12)$$

with coefficients c_1 and c_2 drawn from uniform distributions (see Appendix A for lower and upper bounds). Afterwards, the samples $y_i(t)$ are evaluated over a regular mesh of 100 points $t \in [0, 1]$ and random noise $\epsilon_i^Y(t) \sim \mathcal{N}(0, I)$ is added. Clearly, for $a = 0$ our samples are independent, as the samples $y_i(t) = \epsilon_i^Y(t)$ are just random noise. As a is increased, the dependence between the samples $x_i(s)$ and $y_i(t)$ becomes easier to detect. Figure 1a shows that our test has no difficulties in staying within the allowable false-positive rate α and quickly detects the dependence as soon as $a > 0$ even with very low number of samples.

4.1.2. Joint independence test

To produce the synthetic data for joint independence tests, we first generate random DAGs with d nodes using an Erdős-Rényi model with density 0.5 [20, § 5.2]. For each DAG, we first use Equation (10) to produce function-valued random samples for the variables X^i without parents. We then generate samples for other variables using a historical functional linear model:

$$x_i^j(t) = a \sum_{p \in \text{PA}^j} \int_0^t x_i^p(s) \beta^p(s, t) ds, \quad (13)$$

where PA^j are the parents of node X^j , and the function $\beta^p(s, t)$ is given by Equation (12) with random coefficients c_1 and c_2 independently generated for each descendant-parent pair indexed by p . After being evaluated at a regular mesh of 100 points within $t \in [0, 1]$, random noise $\epsilon_i^j(t) \sim \mathcal{N}(0, I)$ is added. Note that, again, an increase in the factor $a \in [0, 1]$ should make the dependence structure of the DAG easier to detect. Figure 1b shows the test size where $a = 0$, resulting in independent variables, and test power where $a > 0$. We evaluate the joint independence test for $d = 4$ variables over various values of a and a range of sample sizes.

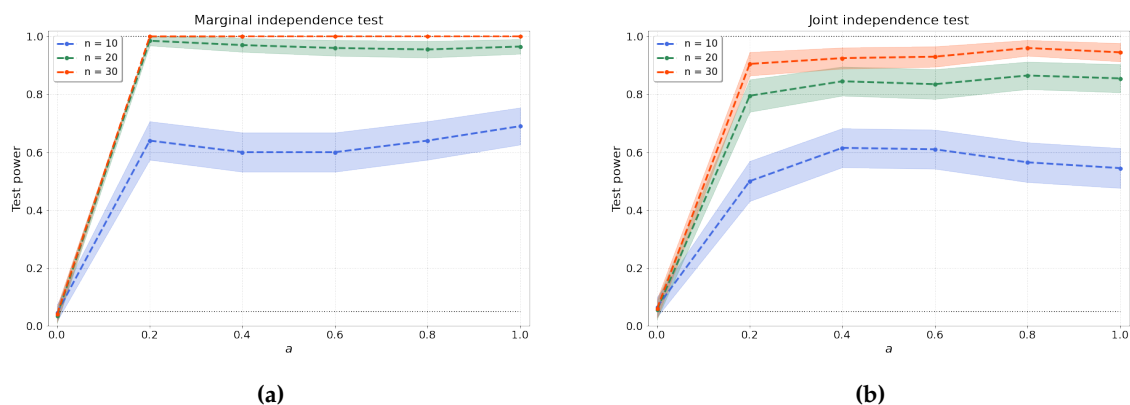


Figure 1. (a) Bivariate and (b) joint independence tests over various values of the dependence factor a and sample size n . The joint independence test is conducted with $d = 4$ variables.

4.1.3. Conditional independence test

To evaluate whether X is independent of Y given \mathcal{Z} , where \mathcal{Z} may be any subset within the power set of the $d \setminus \{X, Y\}$ remaining variables, we generate data samples according to:

$$\begin{aligned} z_i^j(u) &= \sum_{m=1}^M c_{i,m}^j \phi_{m,T}(u), \quad j = 1, \dots, |\mathcal{Z}| \\ x_i(s) &= \sum_{p \in \mathbf{PA}^X} \int_0^s z_i^p(u) \beta^p(u, s) du \\ y_i(t) &= \sum_{p \in \mathbf{PA}^Y} \int_0^t z_i^p(u) \beta^p(u, t) du + a' \int_0^t x_i(s) \beta^Y(s, t) ds, \end{aligned} \quad (14)$$

where $|\mathcal{Z}|$ is the cardinality of the set \mathcal{Z} , β 's are given by Equation (12), and noise terms $\epsilon_i^{\mathcal{Z}^j}(u)$, $\epsilon_i^X(s)$ and $\epsilon_i^Y(t)$, are added to the discrete observation values.

We then apply Algorithm 1 to compute the statistics for our conditional independence test. Firstly, we find that the optimised regularisation strength λ^* is robust and reproducible for different random realisations with the same sample size and model parameters. Consequently, we search for one optimised λ^* (line 1–18 in Algorithm 1) for each sample size, and fix this value for all 200 independent trials. Figure 2 summarises the results for λ^* for $n \in \{100, 200, 300\}$ samples after conducting a grid search over a range of 18 possible values $10^{-4} \leq \lambda \leq 10^{-1}$, with $P = 1000$ permutations, and $B = 100$ rejection iterations. We find that the optimised λ^* exhibits a saturation as the number of conditional variables (the dimension of the conditional set) is increased. Given that we perform a kernel ridge regression, a larger number of samples n should result in lower requirements for regularisation—which aligns with our observations of a decreasing λ^* over the increase in number of samples n . Therefore, Algorithm 1 optimises the regularisation parameter λ^* such that the evaluation rejection rate of the test is closest to the allowable false positive rate $\alpha = 0.05$.

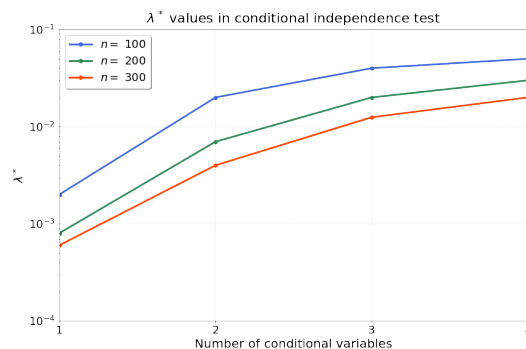


Figure 2. The optimised λ^* for increasing dimension d of the evaluated size of the conditional set and different sample sizes $n \in \{100, 200, 300\}$.

Figures 3a–3d show the results of the conditional independence test for increasing dimension d of the conditional set (i.e., number of conditional variables). We find that the test power is well preserved as d increases through a concomitant increase of λ^* that partly ameliorates the “curse of dimensionality” [47]. Furthermore, the values of the test power for $a' = 0$ correspond to the type-I error rates.

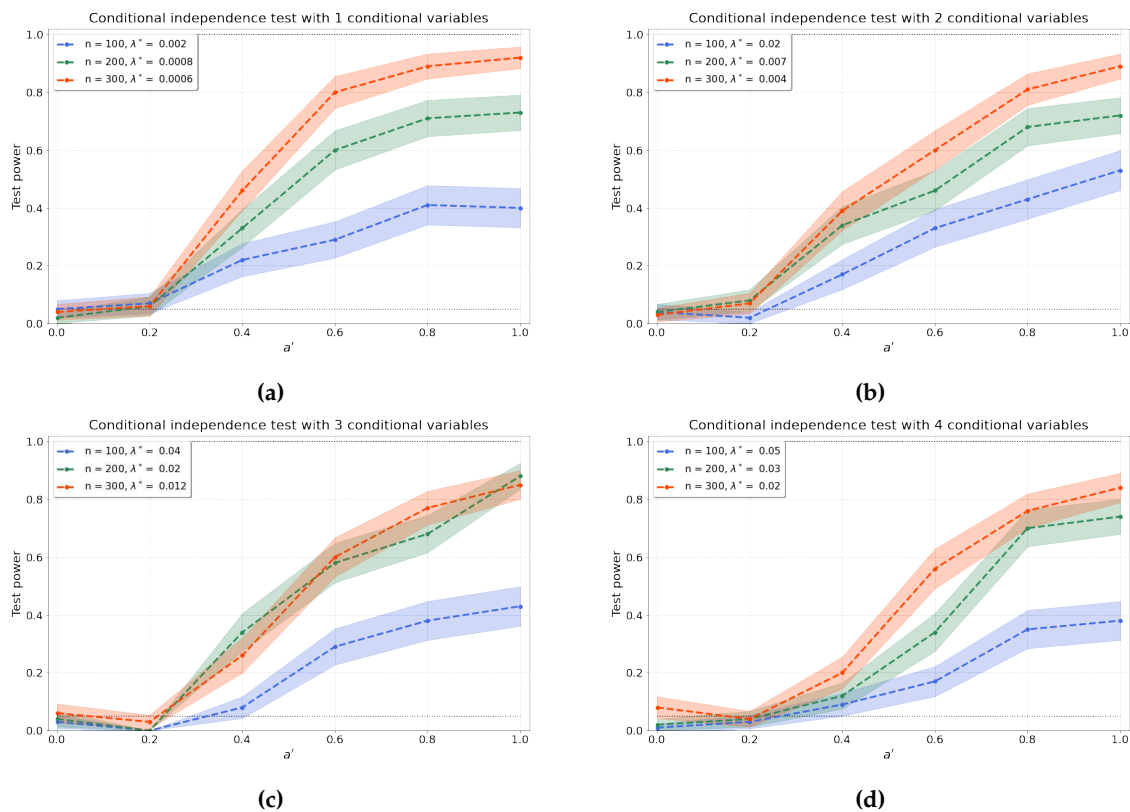


Figure 3. The test power of the conditional independence tests over various values for a' and sample sizes n , from (a) 1 to (d) 4 conditional variables. The regularisation parameter λ^* is given in the legend.

4.2. Causal structure learning

We use the three independence tests evaluated numerically in Section 4.1 to learn different causal structures among function-valued random variables. We start with regression-based methods for the bivariate case $X \rightarrow Y$ and extend to the multivariate case through regression-, constraint-based, and a combination of constraint- and regression-based causal structure learning approaches.

4.2.1. Synthetic data

Regression-based causal discovery. The synthetic data for a bivariate system $X \rightarrow Y$ is generated according to Equations (10)–(11) with $a = 1$, and we measure the performance of the method by the structural Hemming distance (SHD), whose definition can be found in Appendix E. As seen by the low value of the SHD over 200 independent trials, Figure 4a shows that for this case our regression-based algorithm using RESIT outperforms two well-known algorithms for causal discovery: Granger-causality, which assumes linearity in the relationships [6], and CCM, which allows for non-linearity through a time-series embedding [7].

We also evaluate how the regression-based approach performs in a system of three variables where data are generated according to random DAGs with three nodes that entail historical non-linear dependence. The above methods (Granger, CCM) are commonly applied to sets with two variables only, hence we compare our method on a system with three variables to PC algorithm with momentary conditional independence (PCMCI), an algorithm where each time point is represented as a node in a graph [48] from which we extract a directed graph for the variables. Figure 4b also shows improved performance of our regression-based algorithm, with a reduced SHD. Details of the comparing methods are described in Appendix G.

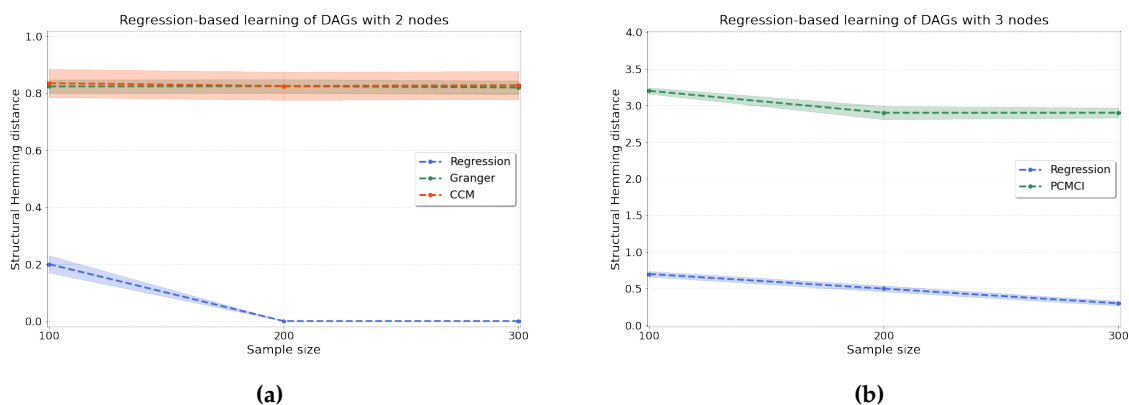


Figure 4. Regression-based causal discovery using kernel-based joint independence tests among residuals with (a) two and (b) three variables. In the case of two variables (a), we compare to Granger-causality and CCM. In the case of three variables (b), we used PCMCI to produce comparable results. Our method significantly outperforms both Granger-causality and CCM in the bivariate setting, with just a few mistakes made for low sample numbers ($n = 100$) and no mistakes for higher sample sizes. For three variables, our method is substantially more accurate than PCMCI, with an average SHD of 0.3 at $n = 300$ versus 2.9 for PCMCI. The accuracy is measured in terms of the structural Hemming distance (SHD).

Constraint-based causal discovery. The accuracy of our constraint-based approach is evaluated by computing three metrics: normalised SHD (SHD_{norm}), precision, and recall (see Appendix E for definitions). Figure 5 shows the accuracy of our algorithm when applied to $d \in \{3, 4, 5, 6\}$ variables.

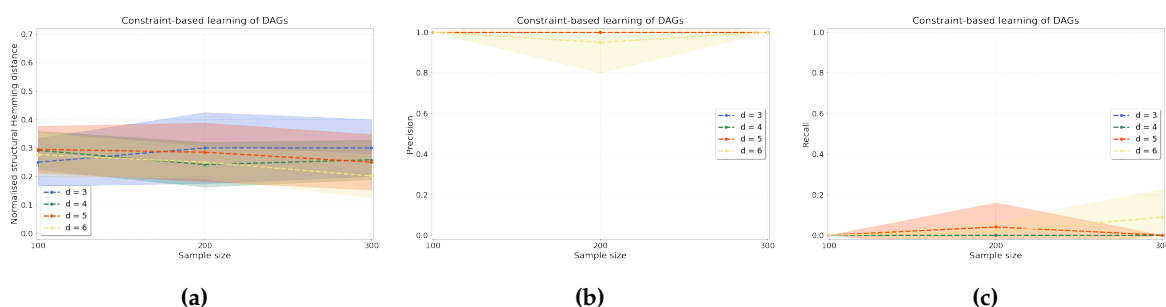


Figure 5. Constraint-based causal structure learning experiments. From left to right, we compute (a) the normalised SHD, (b) precision and (c) recall over $d \in \{3, 4, 5, 6\}$ variables and $n \in \{100, 200, 300\}$ samples with $a = 1$ as the factor of the dependence between any two variables in the data that are connected by an edge in the true DAG.

Combined approach. To be able to scale up our causal discovery techniques more efficiently to larger graphs, it is also possible to combine constraint- and regression-based causal learning, yielding DAGs. In this “combined” approach we start with a constraint-based step through which we learn CPDAGs by applying conditional independence tests to any two variables conditioned on any subset of the remaining $d - 2$ variables and orient edges according to the Algorithm in Appendix F, which finds v-structures, applies Meek’s orientation rules and returns the Markov equivalence class. Often, the Markov equivalence class entails undirected edges, but we can then take a second step using a regression-based approach, under a different set of assumptions, to orient edges by applying RESIT. This two-step process finally yields a DAG, i.e., every edge in the graph is oriented, in a more scalable manner than applying regression-based causal discovery directly. We then measure the accuracy of our approach computing the normalised SHD (A4) as above. Figure 6 shows that our results for $d = 3, 4, 5, 6$ variables compare favourably to PCMCI.

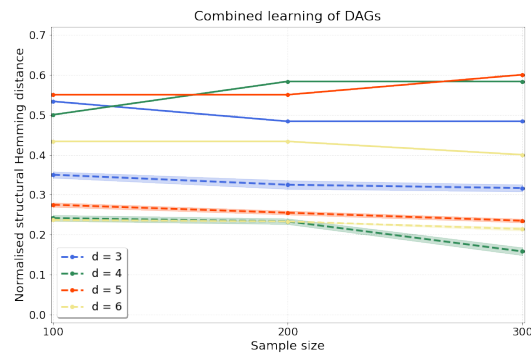


Figure 6. Causal structure learning with the “combined” approach, where we first apply the constraint-based method to find the Markov equivalence class, followed by the regression-based method to orient the undirected edges of the Markov equivalence class. Our results (*dashed lines*) are compared to results obtained from PCMCI (*solid lines*), as described in Appendix G. We compute the normalised SHD over $d \in \{3, 4, 5, 6\}$ variables and $n \in \{100, 200, 300\}$ samples with $a = 1$.

4.2.2. Real-world data

To showcase the application of our methods to real data, we return now to the motivating example on socio-economic indicators from the Introduction (Section 1). Sulemana and Kpienbaareh [2] tested for the causal direction between corruption and income inequality, as measured by the Control of Corruption index [8] and Gini coefficient [49], respectively. They applied Granger-causality with Wald- χ^2 tests, which statistically test against the null hypothesis that an assumed explanatory variable does not have an influence on other considered variables. Both variables, corruption and income inequality, were tested against being the explanatory variable (i.e., the cause) in their bivariate relationship. Their findings showed that income inequality as the explanatory variable results in a p-value of 0.079, whereas corruption is more likely to be the cause with a p-value of 0.276. Analysing the same data using CCM, which does not assume linearity, we also find the same cause-effect direction: 70.9% of the sample pairs (x_i, y_i) validate corruption as the cause, against 29.7% of pairs which detect the opposite direction. However, using our proposed regression-based kernel approach, which does not assume linearity or stationarity, we find the opposite result, i.e., the causal dependence of corruption on income inequality is statistically more likely ($p = 0.0859$) than the reverse direction ($p = 0.0130$).

Going beyond pairwise interactions, we have also applied our causal structure learning method to the set of World Governance Indicators (WGIs) [8]. The WGIs are a collection of six variables, which have been measured in 182 countries over 25 years from 1996 to 2020 (see Appendix H). We view the time-series of the countries as independent functional samples of each variable (so that $n = 182$ and the functional dimension is time), and we apply our methods as described above to establish the causal structure amongst the six WGIs. In this case, we use the “combined” approach by successively applying our constraint- and regression-based causal structure learning methods. We first learn an undirected causal skeleton from the WGIs (Figure 7a), and we find a triangle between Government Effectiveness (GE), Rule of Law (RL) and Control of Corruption (CC), and a separate link between Voice and Accountability (VA) and Regulatory Quality (RQ). We then orient these edges (Figure 7b) and find that Government Effectiveness (GE) causes both Rule of Law (RL) and Control of Corruption (CC), and Rule of Law (RL) causes Control of Corruption (CC). We also find that Voice and Accountability (VA) causes Regulatory Quality (RQ).




Figure 7. Undirected (a) and directed (b) causal networks on the World Governance Indicators data set to evaluate the results of the constraint-based approach alone (a) and the subsequent regression-based approach (b). The labels are abbreviations of the official names: Voice and Accountability (VA), Political Stability (PS), Government Effectiveness (GE), Regulatory Quality (RQ), Rule of Law (RL), and Control of Corruption (CC) (see Appendix H).

5. Discussion and Conclusion

We present a causal structure learning framework on functional data that utilises kernel-based independence tests to extend the applicability of the widely used regression- and constraint-based approaches. The foundation of the framework originates from the functional data analysis literature by interpreting any discrete-measured observations of a random variable as finite-dimensional realisations of functions.

Using synthetic data, we demonstrate that our regression-based approach significantly outperforms existing methods such as Granger-causality, CCM and PCMCI when learning causal relationships between two and three variables. Further, we show that our conditional independence test, which is the cornerstone of the constraint-based causal discovery approach, achieves type-I error rates close to the acceptable false-positive rate α and high type-II error rates, even when the number of variables in the conditional set increases. Shah et al. [47] rightly state that any conditional independence test can suffer from an undesirable test size or low test power in finite samples, and our method is obviously no exception. However, Figure 3 demonstrates the counterbalance of the “curse of dimensionality” through the optimised regularisation strength λ^* . Indeed, while with larger numbers of conditional variables, we would generally expect the test power to diminish, the optimisation of λ^* offsets this reduction, resulting in no significant decrease in test power with a growing number of conditional variables. Although our proposed method is computationally expensive and significantly benefits from high sample sizes, a suitable regularisation strength could be chosen for large conditional sets in principle.

Furthermore, we demonstrate that constraint- and regression-based causal discovery methods can be combined to learn DAGs with large number of nodes, which would otherwise be computationally very expensive when relying on regression-based methods to yield DAGs. By comparing Figure 5a and 6 we see however that it does not necessarily result in a lower SHD. After learning the Markov equivalence class through the constraint-based approach, we apply regression with subsequent independence test (RESIT) to orient undirected edges in the “combined” approach. Here, mistakes in the orientation of undirected edges add 2 to the SHD whereas an undirected edge only adds 1 to the SHD.

The presented work contributes to opening the field of causal discovery to functional data. Beyond the results presented here, we believe that more research needs to be conducted to (i) increase the efficiency of the independence tests, meaning smaller sample sets can achieve higher test power; (ii)

learn about upper bounds of the regularisation strength λ^* with respect to the size of the sample and conditional sets; and (iii) make the computations in the conditional independence test less expensive.

Appendix A. Data

Appendix A.1. Fourier basis functions

Throughout our paper, we use the following Fourier basis functions:

$$\begin{aligned}\phi_0 &= \frac{1}{\sqrt{2}} \\ \phi_{2m-1,T}(s) &= \frac{\sin(\frac{2\pi m}{T}s)}{\sqrt{\frac{T}{2}}} \\ \phi_{2m,T}(s) &= \frac{\cos(\frac{2\pi m}{T}s)}{\sqrt{\frac{T}{2}}}\end{aligned}$$

Appendix A.2. Lower and upper bounds of β -functions

	Lower bound	Upper bound
c_1	0.25	0.75
c_2	0.25	0.75

Table A1. Lower and upper bounds of uniform distribution that define c_1 and c_2 in β .




Figure A1. Function $\beta(s, t)$ generated with random coefficients c_1 and c_2 as specified in Table A1.

Appendix B. Proof of Theorem 1

Proof. All norms $\|\cdot\|$ and inner products $\langle \cdot, \cdot \rangle$ in this proof are taken with respect to the RKHS in which the normand and the arguments of the inner product reside in.

Denote by $\text{HSCIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ the *true* Hilbert-Schmidt conditional independence criterion between the random variables \mathbf{X} and \mathbf{Y} given \mathbf{Z} , i.e.

$$\text{HSCIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = \|\mu_{\mathbf{XY}|\mathbf{Z}} - \mu_{\mathbf{X}|\mathbf{Z}} \otimes \mu_{\mathbf{Y}|\mathbf{Z}}\|,$$

and denote by $\text{H}\hat{\text{S}}\text{CIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \lambda_n)$ the empirical estimate of $\text{HSCIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ based on samples, i.e.

$$\text{H}\hat{\text{S}}\text{CIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \lambda_n) = \|\hat{\mu}_{\mathbf{XY}|\mathbf{Z}} - \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \otimes \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\|,$$

where $\hat{\mu}_{\mathbf{XY}|\mathbf{Z}}$, $\hat{\mu}_{\mathbf{X}|\mathbf{Z}}$ and $\hat{\mu}_{\mathbf{Y}|\mathbf{Z}}$ are the empirical estimates of the conditional mean embeddings $\mu_{\mathbf{XY}|\mathbf{Z}}$, $\mu_{\mathbf{X}|\mathbf{Z}}$ and $\mu_{\mathbf{Y}|\mathbf{Z}}$ obtained with the regularisation parameter λ_n .

Note that by the reverse triangle inequality, followed by the ordinary triangle inequality, we have

$$\begin{aligned} |\text{HSCIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) - \text{H}\hat{\text{S}}\text{CIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \lambda_n)| &= \left| \|\mu_{\mathbf{XY}|\mathbf{Z}} - \mu_{\mathbf{X}|\mathbf{Z}} \otimes \mu_{\mathbf{Y}|\mathbf{Z}}\| - \|\hat{\mu}_{\mathbf{XY}|\mathbf{Z}} - \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \otimes \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\| \right| \\ &\leq \|\mu_{\mathbf{XY}|\mathbf{Z}} - \mu_{\mathbf{X}|\mathbf{Z}} \otimes \mu_{\mathbf{Y}|\mathbf{Z}} - \hat{\mu}_{\mathbf{XY}|\mathbf{Z}} + \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \otimes \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\| \\ &\leq \|\mu_{\mathbf{XY}|\mathbf{Z}} - \hat{\mu}_{\mathbf{XY}|\mathbf{Z}}\| + \|\mu_{\mathbf{X}|\mathbf{Z}} \otimes \mu_{\mathbf{Y}|\mathbf{Z}} - \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \otimes \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\|. \end{aligned}$$

Here,

$$\begin{aligned} \|\mu_{\mathbf{X}|\mathbf{Z}} \otimes \mu_{\mathbf{Y}|\mathbf{Z}} - \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \otimes \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\|^2 &= \langle \mu_{\mathbf{X}|\mathbf{Z}} \otimes \mu_{\mathbf{Y}|\mathbf{Z}}, \mu_{\mathbf{X}|\mathbf{Z}} \otimes \mu_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\quad - 2\langle \mu_{\mathbf{X}|\mathbf{Z}} \otimes \mu_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \otimes \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\quad + \langle \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \otimes \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \otimes \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &= \langle \mu_{\mathbf{X}|\mathbf{Z}}, \mu_{\mathbf{X}|\mathbf{Z}} \rangle \langle \mu_{\mathbf{Y}|\mathbf{Z}}, \mu_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\quad - 2\langle \mu_{\mathbf{X}|\mathbf{Z}}, \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \rangle \langle \mu_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\quad + \langle \hat{\mu}_{\mathbf{X}|\mathbf{Z}}, \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \rangle \langle \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &= \langle \mu_{\mathbf{X}|\mathbf{Z}}, \mu_{\mathbf{X}|\mathbf{Z}} \rangle \langle \mu_{\mathbf{Y}|\mathbf{Z}}, \mu_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\quad - 2\langle \mu_{\mathbf{X}|\mathbf{Z}}, \mu_{\mathbf{X}|\mathbf{Z}} \rangle \langle \mu_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle + 2\langle \mu_{\mathbf{X}|\mathbf{Z}}, \mu_{\mathbf{X}|\mathbf{Z}} - \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \rangle \langle \mu_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\quad + \langle \mu_{\mathbf{X}|\mathbf{Z}}, \mu_{\mathbf{X}|\mathbf{Z}} \rangle \langle \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\quad + \langle \hat{\mu}_{\mathbf{X}|\mathbf{Z}} - \mu_{\mathbf{X}|\mathbf{Z}}, \hat{\mu}_{\mathbf{X}|\mathbf{Z}} + \mu_{\mathbf{X}|\mathbf{Z}} \rangle \langle \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &= \langle \mu_{\mathbf{X}|\mathbf{Z}}, \mu_{\mathbf{X}|\mathbf{Z}} \rangle \|\mu_{\mathbf{Y}|\mathbf{Z}} - \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\|^2 + 2\langle \mu_{\mathbf{X}|\mathbf{Z}}, \mu_{\mathbf{X}|\mathbf{Z}} - \hat{\mu}_{\mathbf{X}|\mathbf{Z}} \rangle \langle \mu_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\quad + \langle \hat{\mu}_{\mathbf{X}|\mathbf{Z}} - \mu_{\mathbf{X}|\mathbf{Z}}, \hat{\mu}_{\mathbf{X}|\mathbf{Z}} + \mu_{\mathbf{X}|\mathbf{Z}} \rangle \langle \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle \\ &\leq \|\mu_{\mathbf{X}|\mathbf{Z}}\|^2 \|\mu_{\mathbf{Y}|\mathbf{Z}} - \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\|^2 \\ &\quad + \|\mu_{\mathbf{X}|\mathbf{Z}} - \hat{\mu}_{\mathbf{X}|\mathbf{Z}}\| \left(2\|\mu_{\mathbf{X}|\mathbf{Z}}\| \langle \mu_{\mathbf{Y}|\mathbf{Z}}, \hat{\mu}_{\mathbf{Y}|\mathbf{Z}} \rangle + \|\hat{\mu}_{\mathbf{X}|\mathbf{Z}} + \mu_{\mathbf{X}|\mathbf{Z}}\| \|\hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\| \right), \end{aligned}$$

where we used the Cauchy-Schwarz inequality in the last inequality. We note that Park and Muandet [17, Theorem 4.4] ensures that each of $\|\mu_{\mathbf{XY}|\mathbf{Z}} - \hat{\mu}_{\mathbf{XY}|\mathbf{Z}}\|$, $\|\mu_{\mathbf{Y}|\mathbf{Z}} - \hat{\mu}_{\mathbf{Y}|\mathbf{Z}}\|$ and $\|\mu_{\mathbf{X}|\mathbf{Z}} - \hat{\mu}_{\mathbf{X}|\mathbf{Z}}\|$ converge to 0 in probability, hence $\text{H}\hat{\text{S}}\text{CIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \lambda_n) \rightarrow \text{HSCIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ in probability with respect to the L_2 -norm.

Now we show that the empirical average of the empirical HSCIC converges to the population integral of the population HSCIC. Writing $\text{H}\hat{\text{S}}\text{CIC}$ as a shorthand for $\text{H}\hat{\text{S}}\text{CIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \lambda_n)$, we see that

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=1}^n \text{H}\hat{\text{S}}\text{CIC}(z_i) - \int \text{HSCIC}(z) d\mathbb{P}(z) \right| &\leq \left| \frac{1}{n} \sum_{i=1}^n \text{H}\hat{\text{S}}\text{CIC}(z_i) - \int \text{H}\hat{\text{S}}\text{CIC}(z) d\mathbb{P}(z) \right| \\ &\quad + \left| \int \text{H}\hat{\text{S}}\text{CIC}(z) d\mathbb{P}(z) - \int \text{HSCIC}(z) d\mathbb{P}(z) \right| \\ &\leq \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(z_i) - \int f(z) d\mathbb{P}(z) \right| \\ &\quad + \left| \int \text{H}\hat{\text{S}}\text{CIC}(z) d\mathbb{P}(z) - \int \text{HSCIC}(z) d\mathbb{P}(z) \right|. \end{aligned}$$

Here, the first term converges to 0 thanks to the uniform law of large numbers over the function class \mathcal{F} in which the empirical estimation of $\text{H}\hat{\text{S}}\text{CIC}$ is carried out, and the second term converges to 0 as shown above.

Now, by Park and Muandet [17, Theorem 5.4], since our kernel $k_1 \otimes k_2$ is characteristic, we have that $\text{HSCIC}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is the identically zero function if and only if $\mathbf{X} \perp \mathbf{Y} \mid \mathbf{Z}$. Hence, if the alternative hypothesis H_1 holds, we have

$$\int \text{HSCIC}(z) d\mathbb{P}(z) > 0.$$

See that, for any $r > 0$, we have

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{P} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \text{H}\hat{\text{S}}\text{CIC}(z_i) \geq r \right) &= 1 - \lim_{n \rightarrow \infty} \mathbb{P} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \text{H}\hat{\text{S}}\text{CIC}(z_i) < r \right) \\ &= 1 - \lim_{n \rightarrow \infty} \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n \text{H}\hat{\text{S}}\text{CIC}(z_i) - \frac{r}{\sqrt{n}} < 0 \right) \\ &= 1 - \lim_{n \rightarrow \infty} \mathbb{P} \left(\int \text{HSCIC}(z) d\mathbb{P}(z) < 0 \right) \\ &= 1, \end{aligned}$$

where we used Portmanteau's theorem [50, p.6, Lemma 2.2] to get from the second line to the third, using the convergence of $\frac{1}{n} \sum_{i=1}^n \text{H}\hat{\text{S}}\text{CIC}(z_i)$ to $\int \text{HSCIC}(z) d\mathbb{P}(z)$ shown above. Hence, our test is consistent. \square

Appendix C. Supplementary information about computational complexity

When applying the regression-based causal discovery approach, one needs to iterate over all possible candidate DAGs. Each candidate DAG is one potential configuration of the causal graph and needs to be evaluated against the possibility that it generated the observed data. The total number of candidate DAGs κ_d of a graph with d nodes is given by

$$\kappa_d = \sum_{i=1}^d (-1)^{i+1} \binom{d}{i} 2^{i(d-i)} \kappa_{d-i}. \quad (\text{A1})$$

When utilising the constraint-based approach, one needs to exhaustively test for conditional independence between any two nodes of the graph given any possible subset of the remaining nodes. Including the empty conditional set, this results in 2^{d-2} conditional independence tests for every pair

of nodes in the graph. Then, the total number of conditional independence tests η_d , that is, the pairs of nodes in a graph of d nodes multiplied by the number of conditional independence tests is given by

$$\eta_d = \binom{d}{2} 2^{d-2}. \quad (\text{A2})$$

Appendix D. Supplementary definitions

Appendix D.1. D-separation

If X, Y and Z are three pairwise disjoint (sets of) nodes in a DAG G , then Z d-separates X from Y in G if Z blocks every possible path between (any node in) X and (any node in) Y in G [34, § 1.2.3]. We then write $X \perp_G Y|Z$. The symbol \perp_G is commonly used to denote d-separation.

Appendix D.2. Faithfulness assumption

If two random variables are (conditionally) independent in the observed distribution \mathbb{P} , then they are d-separated in the underlying causal graph G [36, Definition 6.33].

Appendix D.3. Causal Markov assumption

The distribution \mathbb{P}_{XYZ} is Markov with respect to a DAG G if $X \perp_G Y|Z \implies X \perp Y|Z$ for all disjoint (sets of) nodes X, Y, Z , where \perp_G denotes d-separation [36, Definition 6.21].

Appendix E. Definitions of metrics

Appendix E.1. SHD

For two graphs, G_1 and G_2 , SHD is the sum of the element-wise absolute differences between the adjacency matrices A_1 and A_2 of G_1 and G_2 , respectively.

$$\text{SHD} = \sum_i |A_1^i - A_2^i|, \quad (\text{A3})$$

where we iterate over each element i in the adjacency matrices. Thus, when a learnt causal graph G_1 including an edge, e.g., $X \rightarrow Y$ is oriented opposite to the corresponding edge in the true graph G_2 , meaning $X \leftarrow Y$ is true, we add 2 to the SHD score (“double penalty”). Note that others [e.g., 51] only penalise by 1 for a learnt edge being directed opposite to the true edge.

Appendix E.2. Normalised SHD

We normalise the SHD by the number of possible directed edges in a graph of size d , thus allowing comparison across graphs with different number of nodes.

$$\text{SHD}_{\text{norm}} = \frac{\text{SHD}}{d(d-1)}. \quad (\text{A4})$$

Appendix E.3. Precision

We define the total number of edges in a learnt graph G_1 as the sum of truly and falsely learnt edges, $e = e_t + e_f$. The truly learnt edges are the edges that correspond to the zero elements in the matrix containing the element-wise differences between the learnt A_1 and the true A_2 . In contrast, the falsely learnt edges are the non-zero elements. The total number of truly and falsely learnt edges are e_t and e_f , respectively. Precision is the proportion of correctly oriented edges e_t in a learnt graph G_1 in comparison to a true graph G_2 out of all learnt edges e in G_1 .

$$\text{Precision} = \frac{e_t}{e} \quad (\text{A5})$$

Appendix E.4. Recall

We separate between oriented and unoriented edges, e_o and e_u , respectively, which can again be summed up to the total number of edges, $e = e_o + e_u$. Recall is the fraction of edges in the CPDAG that are oriented:

$$\text{Recall} = \frac{e_o}{e}. \quad (\text{A6})$$

Appendix F. Meek's orientation rules, the SGS and the PC algorithms

In Figure A2, a node is placed at each end of the edges. Rule 1 states that if one directed edge (here the vertical edge) points towards another undirected edge (the horizontal edge) that is not further connected to any other edges, the undirected edge must point in the same direction as the adjacent edge, because the three vertices would otherwise form a v-structure (which are only formed with conditional dependence). Rule 2, in contrast, points the undirected edge in both adjacent vertices' direction because it would violate the acyclicity of a DAG otherwise. Rule 3 and 4 avoid new v-structures which would come into existence by applying Rule 2 twice if the oriented edges pointed in the opposite direction. These new v-structures would be in the left top corners of the rectangles.




Figure A2. Meek rules to orient edges that remain in the graph after conditional independence tests and edges are oriented based on detected colliders. Proofs that edges cannot be oriented in the opposite direction without violating the acyclicity of the graph and found conditional independencies are given in Meek [44].

Based on Meek's orientation rules, the SGS algorithm tests each pair of variables as follows [13, § 5.4.1]:

1. Form the complete undirected graph G from the set of vertices (or nodes) \mathbf{V} .
2. For each pair of vertices X and Y , if there exists a subset S of $\mathbf{V} \setminus \{X, Y\}$ such that X and Y are d-separated, i.e., conditionally independent (causal Markov condition), given Z , remove the edge between X and Y from G .
3. Let K be the undirected graph resulting from the previous step 2. For each triple of vertices X , Y , and Z such that the pair X and Y and the pair Y and Z are each adjacent in K (written as $X - Y - Z$) but the pair X and Z are not adjacent in K , orient $X - Y - Z$ as $X \rightarrow Y \leftarrow Z$ if and only if there is no subset S of Y that d-separates X and Z .

4. Repeat the following steps until no more edges can be oriented:
 - (a) If $X \rightarrow Y$, Y and Z are adjacent, X and Z are not adjacent, and there is no arrowhead of other vertices at Y , then orient $Y - Z$ as $Y \rightarrow Z$ (Rule 1 of Meek's orientation rules).
 - (b) If there is a directed path over some other vertices from X to Y , and an edge between X and Y , then orient $X - Y$ as $X \rightarrow Y$ (Rule 2 of Meek's orientation rules).

The large computational cost of applying the SGS algorithm (especially step 2) to data, due to the large number of combinations of variables as the potential conditional sets, opened the door for computationally more efficient methods. The PC algorithm [13, § 5.4.2] is amongst the most popular alternatives and minimises the computational cost by searching for conditional independence in a structured manner, as opposed to step 2 of the SGS algorithm that iterates over all possible conditional subsets to any two variables A and B . Starting again with a fully connected undirected graph, the PC algorithm begins by testing for marginal independence between any two variables A and B and deletes the edge connecting A and B , $A - B$, if $A \perp B$. Then, it proceeds with testing for $A \perp B|C$ and erases $A - B$ if one conditional variable C is found that makes A and B independent given C . Afterwards, the conditional set is extended to two variables, $\{C, D\}$, and the edge $A - B$ is deleted if this conditional set makes A and B independent. The conditional set is extended, round after round, until no more conditional independencies are found, resulting in the sparsified graph K . Step 3 and 4 are then pursued as in the SGS algorithm.

Appendix G. Comparison methods

Appendix G.1. Granger-causality

We implement Granger-causality through software provided by Seabold and Perktold [52] and define the optimal lag as the first local minimum in the mutual information function over measurement points. Given that Granger-causality is applied to a pair of samples, (x_i, y_i) , we iterate over every pair and determine the percentage of pairs that result in $X \rightarrow Y$ against $X \leftarrow Y$.

Appendix G.2. CCM

We implement CCM through software provided by [?] and define the optimal lag as for Granger-causality. The embedding dimensionality is determined by a false nearest neighbour method [53]. As with Granger-causality, CCM is also applied to a pair of samples, (x_i, y_i) . We take the same summarising approach and iterate over every pair and determine the percentage of pairs that result in $X \rightarrow Y$ against $X \leftarrow Y$.

Appendix G.3. PCMCI

Runge et al. [48] proposed PCMCI, a method that produces *time-series graphs* where each node represents a variable at a certain state in time. When applying the method in our experiments, we estimate the maximum possible lag as the average of the first local minimum of mutual information of all considered data series. Then, we test for the presence of an edge by applying distance correlation-based independence tests [54] between the residuals of X and Y that remain after regressing out the influence of the remaining nodes in the graph through Gaussian processes. The distance correlation coefficients of the significant edges between two variables are summed up to measure the prevalence of a causal direction of edges connecting these two variables. The causal direction with the greater sum is then considered to be the directed edge, and assessed against the edge between the same two variables in the true DAG.

Appendix H. World Governance Indicators

In Figure 7, we abbreviate the variables according to the official names that are mentioned in Table A2. The Descriptions are taken from the Documentation on <https://info.worldbank.org/governance/wgi/>. The data set consists of six variables (normalised to a range from 0 to 100), of which each is a summary of multiple indicators measured over 25 years from 1996 to 2020 in 182 countries. The methodology about how the summarising measure is calculated can be found in Kaufmann et al. [8, §4].

Table A2. Abbreviations together with official names and descriptions of the variables in the data set of the World Governance Indicators.

Abbreviation	Official Name	Description
VA	Voice and Accountability	Voice and accountability captures perceptions of the extent to which a country's citizens are able to participate in selecting their government, as well as freedom of expression, freedom of association, and a free media.
PS	Political Stability and No Violence	Political Stability and Absence of Violence/Terrorism measures perceptions of the likelihood of political instability and/or politically-motivated violence, including terrorism.
GE	Government Effectiveness	Government effectiveness captures perceptions of the quality of public services, the quality of the civil service and the degree of its independence from political pressures, the quality of policy formulation and implementation, and the credibility of the government's commitment to such policies.
RQ	Regulatory Quality	Regulatory quality captures perceptions of the ability of the government to formulate and implement sound policies and regulations that permit and promote private sector development.
RL	Rule of Law	Rule of law captures perceptions of the extent to which agents have confidence in and abide by the rules of society, and in particular the quality of contract enforcement, property rights, the police, and the courts, as well as the likelihood of crime and violence.
CC	Control of Corruption	Control of corruption captures perceptions of the extent to which public power is exercised for private gain, including both petty and grand forms of corruption, as well as "capture" of the state by elites and private interests.

References

1. Runge, J.; Bathiany, S.; Bollt, E.; Camps-Valls, G.; Coumou, D.; Deyle, E.; Glymour, C.; Kretschmer, M.; Mahecha, M.D.; Muñoz-Marí, J.; et al. Inferring causation from time series in Earth system sciences. *Nature communications* **2019**, *10*, 1–13.
2. Sulemana, I.; Kpienbaareh, D. An empirical examination of the relationship between income inequality and corruption in Africa. *Economic Analysis and Policy* **2018**, *60*, 27–42.
3. Finkle, J.D.; Wu, J.J.; Bagheri, N. Windowed granger causal inference strategy improves discovery of gene regulatory networks. *Proceedings of the National Academy of Sciences* **2018**, *115*, 2252–2257.
4. Glymour, C.; Zhang, K.; Spirtes, P. Review of causal discovery methods based on graphical models. *Frontiers in genetics* **2019**, *10*, 524.
5. Ramsay, J.O.J.O. *Functional data analysis*, 2nd ed. ed.; Springer series in statistics, Springer: New York, 2005.
6. Granger, C.W. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society* **1969**, pp. 424–438.
7. Sugihara, G.; May, R.; Ye, H.; Hsieh, C.h.; Deyle, E.; Fogarty, M.; Munch, S. Detecting causality in complex ecosystems. *science* **2012**, *338*, 496–500.
8. Kaufmann, D.; Kraay, A.; Mastruzzi, M. The worldwide governance indicators: Methodology and analytical issues1. *Hague journal on the rule of law* **2011**, *3*, 220–246.
9. World Bank. Gini index, 2022.
10. Jong-Sung, Y.; Khagram, S. A comparative study of inequality and corruption. *American sociological review* **2005**, *70*, 136–157.
11. Alesina, A.; Angeletos, G.M. Corruption, inequality, and fairness. *Journal of Monetary Economics* **2005**, *52*, 1227–1244.
12. Dobson, S.; Ramlogan-Dobson, C. Is there a trade-off between income inequality and corruption? Evidence from Latin America. *Economics letters* **2010**, *107*, 102–104.
13. Spirtes, P.; Glymour, C.N.; Scheines, R.; Heckerman, D. *Causation, prediction, and search*; MIT press, 2000.
14. Peters, J.; Mooij, J.M.; Janzing, D.; Schölkopf, B. Causal discovery with continuous additive noise models. *The Journal of Machine Learning Research* **2014**, *15*, 2009–2053.
15. Muandet, K.; Fukumizu, K.; Sriperumbudur, B.; Schölkopf, B.; et al. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends in Machine Learning* **2017**, *10*, 1–141.
16. Wynne, G.; Duncan, A.B. A kernel two-sample test for functional data. *Journal of Machine Learning Research* **2022**, *23*, 1–51.
17. Park, J.; Muandet, K. A measure-theoretic approach to kernel conditional mean embeddings. *arXiv preprint arXiv:2002.03689* **2020**.
18. Berrett, T.B.; Wang, Y.; Barber, R.F.; Samworth, R.J. The conditional permutation test for independence while controlling for confounders. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **2020**, *82*, 175–197.
19. Malfait, N.; Ramsay, J.O. The historical functional linear model. *Canadian Journal of Statistics* **2003**, *31*, 115–128.
20. Pfister, N.; Bühlmann, P.; Schölkopf, B.; Peters, J. Kernel-based tests for joint independence. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **2018**, *80*, 5–31.
21. Ramos-Carreño, C.; Suárez, A.; Torrecilla, J.L.; Carbajo Berrocal, M.; Marcos Manchón, P.; Pérez Manso, P.; Hernando Bernabé, A.; García Fernández, D.; Hong, Y.; Rodríguez-Ponga Eyriès, P.M.; et al. GAA-UAM/scikit-fda: Version 0.7.1, 2022. <https://doi.org/10.5281/zenodo.5903557>.
22. Squires, C. *causal dag: creation, manipulation, and learning of causal models*, 2018.
23. Girard, O.; Lattier, G.; Micallef, J.P.; Millet, G.P. Changes in exercise characteristics, maximal voluntary contraction, and explosive strength during prolonged tennis playing. *British journal of sports medicine* **2006**, *40*, 521–526.
24. Zhu, Y.; Kuhn, T.; Mayo, P.; Hinds, W.C. Comparison of daytime and nighttime concentration profiles and size distributions of ultrafine particles near a major highway. *Environmental science & technology* **2006**, *40*, 2531–2536.
25. Ramsay, J.; Hooker, G. *Dynamic data analysis*; Springer, 2017.

26. Fukumizu, K.; Gretton, A.; Sun, X.; Schölkopf, B. Kernel Measures of Conditional Dependence. In Proceedings of the Advances in Neural Information Processing Systems; Platt, J.C.; Koller, D.; Singer, Y.; Roweis, S., Eds., Red Hook, NY, USA, 2008; Vol. 20, pp. 489–496.
27. Gretton, A.; Fukumizu, K.; Teo, C.H.; Song, L.; Schölkopf, B.; Smola, A.J. A kernel statistical test of independence. In Proceedings of the Advances in neural information processing systems, 2008, pp. 585–592.
28. Zhang, K.; Peters, J.; Janzing, D.; Schölkopf, B. Kernel-based conditional independence test and application in causal discovery. *arXiv preprint arXiv:1202.3775* **2012**.
29. Lai, T.; Zhang, Z.; Wang, Y.; Kong, L. Testing independence of functional variables by angle covariance. *Journal of Multivariate Analysis* **2021**, *182*, 104711.
30. Górecki, T.; Krzyśko, M.; Wołyński, W. Independence test and canonical correlation analysis based on the alignment between kernel matrices for multivariate functional data. *Artificial Intelligence Review* **2020**, *53*, 475–499.
31. Doran, G.; Muandet, K.; Zhang, K.; Schölkopf, B. A Permutation-Based Kernel Conditional Independence Test. In Proceedings of the UAI. Citeseer, 2014, pp. 132–141.
32. Gretton, A.; Borgwardt, K.M.; Rasch, M.J.; Schölkopf, B.; Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research* **2012**, *13*, 723–773.
33. Lee, S.; Honavar, V.G. Self-discrepancy conditional independence test. In Proceedings of the Uncertainty in artificial intelligence, 2017, Vol. 33.
34. Pearl, J. *Causality*; Cambridge University Press, 2009.
35. Mooij, J.M.; Peters, J.; Janzing, D.; Zscheischler, J.; Schölkopf, B. Distinguishing cause from effect using observational data: methods and benchmarks. *The Journal of Machine Learning Research* **2016**, *17*, 1103–1204.
36. Peters, J.; Janzing, D.; Schölkopf, B. *Elements of causal inference*; The MIT Press, 2017.
37. Schölkopf, B.; von Kügelgen, J. From statistical to causal learning. *arXiv preprint arXiv:2204.00607* **2022**.
38. Squires, C.; Uhler, C. Causal structure learning: a combinatorial perspective. *Foundations of Computational Mathematics* **2022**, pp. 1–35.
39. Vowels, M.J.; Camgoz, N.C.; Bowden, R. D’ya like dags? a survey on structure learning and causal discovery. *ACM Computing Surveys* **2022**, *55*, 1–36.
40. Chickering, D.M. Optimal structure identification with greedy search. *Journal of machine learning research* **2002**, *3*, 507–554.
41. Shimizu, S.; Hoyer, P.O.; Hyvärinen, A.; Kerminen, A.; Jordan, M. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research* **2006**, *7*.
42. Hoyer, P.; Janzing, D.; Mooij, J.M.; Peters, J.; Schölkopf, B. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems* **2008**, *21*.
43. Szabó, Z.; Sriperumbudur, B.K. Characteristic and Universal Tensor Product Kernels. *J. Mach. Learn. Res.* **2017**, *18*, 1–29.
44. Meek, C. *Complete orientation rules for patterns*; Carnegie Mellon [Department of Philosophy], 1995.
45. Peters, J.; Mooij, J.; Janzing, D.; Schölkopf, B. Identifiability of causal graphs using functional models. *arXiv preprint arXiv:1202.3757* **2012**.
46. Bühlmann, P.; Peters, J.; Ernest, J. CAM: Causal additive models, high-dimensional order search and penalized regression **2014**.
47. Shah, R.D.; Peters, J.; et al. The hardness of conditional independence testing and the generalised covariance measure. *Annals of Statistics* **2020**, *48*, 1514–1538.
48. Runge, J.; Nowack, P.; Kretschmer, M.; Flaxman, S.; Sejdinovic, D. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science advances* **2019**, *5*.
49. Gini, C. On the measure of concentration with special reference to income and statistics. *Colorado College Publication, General Series* **1936**, *208*, 73–79.
50. Van der Vaart, A.W. *Asymptotic statistics*; Vol. 3, Cambridge university press, 2000.
51. Peters, J.; Bühlmann, P. Structural intervention distance for evaluating causal graphs. *Neural computation* **2015**, *27*, 771–799.
52. Seabold, S.; Perktold, J. statsmodels: Econometric and statistical modeling with python. In Proceedings of the 9th Python in Science Conference, 2010.
53. Munch, E.; Khasawneh, F.; Myers, A.; Yesilli, M.; Tymochko, S.; Barnes, D.; Guzel, I.; Chumley, M. teaspoon: Topological Signal Processing in Python, 2022.

54. Székely, G.J.; Rizzo, M.L.; Bakirov, N.K. Measuring and testing dependence by correlation of distances **2007**.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.