Christina Heinze-Deml*, Jonas Peters, and Nicolai Meinshausen

Invariant Causal Prediction for Nonlinear Models

<https://doi.org/10.1515/jci-2017-0016>

Received January 7, 2017; revised May 9, 2018; accepted August 24, 2018

Abstract: An important problem in many domains is to predict how a system will respond to interventions. This task is inherently linked to estimating the system's underlying causal structure. To this end, Invariant Causal Prediction (ICP) [1] has been proposed which learns a causal model exploiting the invariance of causal relations using data from different environments. When considering linear models, the implementation of ICP is relatively straightforward. However, the nonlinear case is more challenging due to the difficulty of performing nonparametric tests for conditional independence.

In this work, we present and evaluate an array of methods for nonlinear and nonparametric versions of ICP for learning the causal parents of given target variables. We find that an approach which first fits a nonlinear model with data pooled over all environments and then tests for differences between the residual distributions across environments is quite robust across a large variety of simulation settings. We call this procedure “invariant residual distribution test”. In general, we observe that the performance of all approaches is critically dependent on the true (unknown) causal structure and it becomes challenging to achieve high power if the parental set includes more than two variables.

As a real-world example, we consider fertility rate modeling which is central to world population projections. We explore predicting the effect of hypothetical interventions using the accepted models from nonlinear ICP. The results reaffirm the previously observed central causal role of child mortality rates.

Keywords: causal structure learning, structural equation models, invariance

1 Introduction

Invariance based causal discovery [1] relies on the observation that the conditional distribution of the target variable Y given its direct causes remains invariant if we intervene on variables other than Y . While the proposed methodology in [1] focuses on linear models, we extend Invariant Causal Prediction to nonlinear settings. We first introduce the considered structural causal models in Section 1.1 and review related approaches to causal discovery in Section 1.2. The invariance approach to causal discovery from [1] is briefly summarized in Section 1.3 and we outline our contribution in Section 1.4. In Section 1.5 we introduce the problem of fertility rate modeling which we consider as a real-world example throughout this work.

1.1 Structural causal models

Assume an underlying structural causal model (also called structural equation model) [e. g. 2]

$$\begin{aligned} Z_1 &\leftarrow g_1(Z_{\text{pa}_1}) + \eta_1, \\ Z_2 &\leftarrow g_2(Z_{\text{pa}_2}) + \eta_2, \\ &\vdots \\ Z_q &\leftarrow g_q(Z_{\text{pa}_q}) + \eta_q, \end{aligned}$$

*Corresponding author: Christina Heinze-Deml, Seminar für Statistik, ETH Zurich, Zurich, Switzerland, e-mail: heinzedeml@stat.math.ethz.ch

Jonas Peters, University of Copenhagen, Department of Mathematics, Copenhagen, Denmark, e-mail: jonas.peters@math.ku.dk

Nicolai Meinshausen, Seminar für Statistik, ETH Zurich, Zurich, Switzerland, e-mail: meinshausen@stat.math.ethz.ch

for which the functions g_k , $k = 1, \dots, q$, as well as the parents $\text{pa}_k \subseteq \{1, \dots, q\} \setminus \{k\}$ of each variable are unknown. Here, we have used the notation $Z_S = (Z_{i_1}, \dots, Z_{i_s})$ for any set $S = \{i_1, \dots, i_s\} \subseteq \{1, \dots, q\}$. We assume the corresponding directed graph to be acyclic. We further require the noise variables $\eta_1, \dots, \eta_q$ to be jointly independent and to have zero mean, i. e. we assume that there are no hidden variables.

Due to its acyclic structure, it is apparent that such a structural causal model induces a joint distribution P over the observed random variables. Interventions on the system are usually modeled by replacing some of the structural assignments [e. g. 2]. If one intervenes on variable Z_3 , for example, and sets it to the value 5, the system again induces a distribution over $Z_1, \dots, Z_q$, that we denote by $P(\cdot | \text{do}(Z_3 \leftarrow 5))$. It is usually different from the observational distribution P . We make no counterfactual assumptions here: we assume a new realization η is drawn from the noise distribution as soon as we make an intervention.¹

1.2 Causal discovery

In causal discovery (also called structure learning) one tries to reconstruct the structural causal model or its graphical representation from its joint distribution [e. g. 2, 3, 4, 5, 6, 7, 8].

Existing methods for causal structure learning can be categorized along a number of dimensions, such as (i) using purely observational data vs. using a combination of interventional and observational data; (ii) score-based vs. constraint-based vs. “other” methods; (iii) allowing vs. precluding the existence of hidden confounders; (iv) requiring vs. not requiring faithfulness;² (v) type of object that the method estimates. Moreover, different methods vary by additional assumptions they require. In the following, we give brief descriptions of the most common methods for causal structure learning.³

The PC algorithm [3] uses observational data only and estimates the Markov equivalence class of the underlying graph structure, based on (conditional) independence tests under a faithfulness assumption. The presence of hidden confounders is not allowed. Based on the PC algorithm, the IDA algorithm [12] computes bounds on the identifiable causal effects.

The FCI algorithm is a modification of the PC algorithm. It also relies on purely observational data while it allows for hidden confounders. The output of FCI is a partial ancestral graph (PAG), i. e. it estimates the Markov equivalence class of the underlying maximal ancestral graph (MAG). Faster versions, RFCI and FCI+, were proposed by Colombo et al. [13] and Claassen et al. [14], respectively.

The PC, FCI, RFCI and FCI+ algorithms are formulated such that they allow for an independence oracle that indicates whether a particular (conditional) independence holds in the distribution. These algorithms are typically applied in the linear Gaussian setting where testing for conditional independence reduces to testing for vanishing partial correlation.

One of the most commonly known score-based methods is greedy equivalence search (GES). Using observational data, it greedily searches over equivalence classes of directed acyclic graphs for the best scoring graph (all graphs within the equivalence class receive the same score) where the score is given by the Bayesian information criterion, for example. Thus, GES is based on an assumed parametric model such as linear Gaussian structural equations or multinomial distributions. The output of GES is the estimated Markov equivalence class of the underlying graph structure. Heckerman [7] describe a score-based method with a Bayesian score.

Greedy interventional equivalence search (GIES) extends GES to operate on a combination of interventional and observational data. The targets of the interventions need to be known and the output of GIES is

¹ The new realization of η under an intervention and the realization under observational data can be assumed to be independent. However, such an assumption is untestable since we can never observe realizations under different interventions simultaneously and we do not make statements or assumptions about the joint distribution of observational and interventional settings.

² A distribution satisfies faithfulness and the global Markov condition with respect to a graph G if the following statement holds for all disjoint sets A , B , and C of variables: A is independent of B , given C , if and only if A is d -separated (in G) from B , given C . The concept of d -separation [9, 10] is defined in Peters et al. [4, Def. 6.1], for example.

³ Also see [11] for a review and empirical comparison of recently proposed causal structure learning algorithms.

the estimated interventional Markov equivalence class. The latter is typically smaller than the Markov equivalence class obtained when using purely observational data.

Another group of methods makes restrictive assumptions which allows for obtaining full identifiability. Such assumptions include non-Gaussianity [15] or equal variances [16] of the errors or non-linearity of the structural equations in additive noise models [17, 6].

Instead of trying to infer the whole graph, we are here interested in settings, where there is a target variable Y of special interest. The goal is to infer both the parental set S^* for the target variable Y and confidence bands for the causal effects.

1.3 Invariance based causal discovery

This work builds on the method of Invariant Causal Prediction (ICP) [1] and extends it in several ways. The method's key observation is that the conditional distribution of the target variable Y given its direct causes remains invariant if we intervene on variables other than Y . This follows from an assumption sometimes called autonomy or modularity [18, 19, 20, 2, 21]. In a linear setting, this implies, for example, that regressing Y on its direct causes yields the same regression coefficients in each environment, provided we have an infinite amount of data. In a nonlinear setting, this can be generalized to a conditional independence between an index variable indicating the interventional setting and Y , given X ; see (3). The method of ICP assumes that we are given data from several environments. It searches for sets of covariates, for which the above property of invariance cannot be rejected. The method then outputs the intersection of all such sets, which can be shown to be a subset of the true set with high probability, see Section 2.1 and Algorithm 1 in Appendix II for more details. Such a coverage guarantee is highly desirable, especially in causal discovery, where information about ground truth is often sparse.

In many real life scenarios, however, relationships are not linear and the above procedure can fail: The true set does not necessarily yield an invariant model and the method may lose its coverage guarantee, see Example 2. Furthermore, environments may not come as a categorical variable but as a continuous variable instead. In this work, we extend the concept of ICP to nonlinear settings and continuous environments. The following paragraph summarizes our contributions.

1.4 Contribution

Our contributions are fivefold.

Conditional independence tests

We extend the method of ICP to nonlinear settings by considering conditional independence tests. We discuss in Section 3 and in more length in Appendix II several possible nonlinear and nonparametric tests for conditional independence of the type (3) and propose alternatives. There has been some progress towards nonparametric independence tests [22, 23, 24, 25, 26, 27]. However, in the general nonparametric case, no known non-trivial test of conditional independence has (even asymptotically) a type I error rate less than the pre-specified significance level. This stresses the importance of empirical evaluation of conditional independence tests.

Defining sets

We discuss in Section 2.2 cases of poor identifiability of the causal parents. If there are highly correlated variables in the dataset, we might get an empty estimator if we follow the approach proposed in [1]. We can, however, extract more information via defining sets. The results are to some extent comparable to similar issues arising in multiple testing [28]. For example, if we know that the parental set of a variable Y is either $S = \{1, 3\}$ or $S = \{2, 3\}$, we know that $\{3\}$ has to be a parent of Y . Yet we also want to explore the information

that one variable out of the set $\{1, 2\}$ also has to be causal for the target variable Y , even if we do not know which one out of the two.

Confidence bands for causal effects

Beyond identifying the causal parents, we can provide nonparametric or nonlinear confidence bands for the strength of the causal effects, as shown in Section 2.3.

Prediction under interventions

Using the accepted models from nonlinear ICP, we are able to forecast the average causal effect of external interventions. We will discuss this at hand of examples in Section 2.4.

Software

R [29] code for nonlinear ICP is provided in the package `nonlinearICP`. The proposed conditional independence tests are part of the package `CondIndTests`. Both packages are available from CRAN.

1.5 Fertility rate modeling

At the hand of the example of fertility rate modeling, we shall explore how to exploit the invariance of causal models for causal discovery in the nonlinear case.

Developing countries have a significantly higher fertility rate compared to Western countries. The fertility rate can be predicted well from covariates such as ‘infant mortality rate’ or ‘GDP per capita’. Classical prediction models, however, do not answer whether an active intervention on some of the covariates leads to a change in the fertility rate. This can only be answered by exploiting causal knowledge of the system.

Traditionally, in statistics the methods for establishing causal relations rely on carefully designed randomized studies. Often, however, such experiments cannot be performed. For instance, factors like ‘infant mortality rate’ are highly complex and cannot be changed in isolation. We may still be interested in the effect of a policy that aims at reducing the infant mortality rate but this policy cannot be randomly assigned to different groups of people within a country.

There is a large body of work that is trying to explain changes in fertility; for an interesting overview of different theories see Hirschman [30] and the more recent Huinink et al. [31]. There is not a single established theory for changes in fertility and we should clarify in the beginning that all models we will be using will have shortcomings, especially the shortcoming that we might not have observed all relevant variables. We would nevertheless like to take the fertility data as an example to establish a methodology that allows data-driven answers; discussing potential shortfalls of the model is encouraged and could be beneficial in further phrasing the right follow-up questions and collecting perhaps more suitable data.

An interesting starting point for us was the work of Raftery et al. [32] and very helpful discussions with co-author Adrian Raftery. That work tries to distinguish between two different explanatory models for a decline in fertility in Iran. One model argues that economic growth is mainly responsible; another argues that transmission of new ideas is the primary factor (ideation theory). What allows a distinction between these models is that massive economic growth started in 1955 whereas ideational changes occurred mostly 1967 and later. Since the fertility began to drop measurably already in 1959, the demand theory seems more plausible and the authors conclude that reduced child mortality is the key explanatory variable for the reduction in fertility (responsible for at least a quarter of the reduction).

Note the way we decide between two potential causal theories for a decline in fertility: if a causal model is valid, it has to be able to explain the decline consistently. In particular, the predictions of the model have to be valid for all time-periods, including the time of 1959 with the onset of the fertility decline. The ideation theory wrongly places the onset of fertility decline later and is thus dismissed as less plausible.

The invariance approach of Peters et al. [1] we follow here for linear models has a similar basic idea: a causal model has to work consistently. In our case, we choose geographic location instead of time for the

example and demand that a causal model has to work consistently across geographic locations or continents. We collect all potential models that show this invariance and know that *if* the underlying assumption of causal sufficiency is true and we have observed all important causal variables *then* the causal model will be in the set of retained models. Clearly, there is room for a healthy and interesting debate to what extent the causal sufficiency assumption is violated in the example. It has been argued, however, that missing variables do not allow for any invariant model, which renders the method to remain conservative [1, Prop. 5].

We establish a framework for causal discovery in nonlinear models. Incidentally, the approach also identifies reduced child mortality as one of key explanatory variables for a decline in fertility.

2 Nonlinear invariant causal prediction

We first extend the approach of [1] to nonlinear models, before discussing defining sets, nonparametric confidence bands and prediction under interventions.

2.1 Invariance approach for causal discovery

[1] proposed an invariance approach in the context of linear models. We describe the approach here in a notationally slightly different way that will simplify statements and results in the nonlinear case and allow for more general applications. Assume that we are given a structural causal model (SCM) over variables (Y, X, E) , where Y is the target variable, X the predictors and E so-called environmental variables.

Definition 1 (Environmental variables). *We know or assume that the variables E are neither descendants nor parents of Y in the causal DAG of (Y, X, E) . If this is the case, we call E environmental variables.*

In [1], the environmental variables were given and non-random. Note that the definition above treats the variables as random but we can in practice condition on the observed values of E . The definition above excludes the possibility that there is a direct causal connection between one of the variables in E and Y . We will talk in the following about the triple of random variables (Y, X, E) , where the variable X of predictor variables is indexed by $X_1, \dots, X_p$. With a slight abuse of notation, we let $S^* \subseteq \{1, \dots, p\}$ be the indices of X that are causal parents pa_Y of Y . Thus, the structural equation for Y can be written as

$$Y \leftarrow f(X_{S^*}) + \varepsilon, \quad (1)$$

where $f : \mathbb{R}^{|S^*|} \rightarrow \mathcal{Y}$. We let $\mathcal{F}$ be the function class of f and let $\mathcal{F}_S$ be the subclass of functions that depend only on the set $S \subseteq \{1, \dots, p\}$ of variables. With this notation we have $f \in \mathcal{F}_{S^*}$.

The assumption of no direct effect of E on Y is analogous to the typical assumptions about instrumental variables [33, 34]. See Section 5 in [1] for a longer discussion on the relation between environmental variables and instrumental variables. The two main distinctions between environmental and instrumental variables are as follows. First, we do not need to test for the “weakness” of instrumental/environmental variables since we do not assume that there is a causal effect from E on the variables in X . Second, the approaches are used in different contexts. With instrumental variables, we assume the graph structure to be known typically and want to estimate the strength of the causal connections, whereas the emphasis is here on both causal discovery (what are the parents of a target?) and then also inference for the strength of causal effects. With a single environmental variable, we can identify in some cases multiple causal effects whereas the number of instrumental variables needs to match or exceed the number of variables in instrumental variable regression. The instrumental variable approach, on the other hand, can correct for unobserved confounders between the parents and the target variable if their influence is linear, for example. In these cases, our approach could remain uninformative [1, Proposition 5].

Example (fertility data)

In this work, we analyze a data set provided by the [35]. Here, Y , X and E correspond to the following quantities:

- (a) $Y \in \mathbb{R}$ is the total fertility rate (TFR) in a country in a given year,
- (b) $X \in \mathbb{R}^9$ are potential causal predictor variables for TFR:
 - IMR – infant mortality rate
 - Q5 – under-five mortality rate
 - Education expenditure (% of GNI)
 - Exports of goods and services (% of GDP)
 - GDP per capita (constant 2005 US\$)
 - GDP per capita growth (annual %)
 - Imports of goods and services (% of GDP)
 - Primary education (% female)
 - Urban population (% of total)
- (c) $E \in \{C_1, C_2, C_3, C_4, C_5, C_6\}$ is the continent of the country, divided into the categories Africa, Asia, Europe, North and South America and Oceania. If viewed as a random variable (which one can argue about), the assumption is that the continent is not a descendant of the fertility rate, which seems plausible. For an environmental variable, the additional assumption is that the TFR in a country is only indirectly (that is via one of the other variables) influenced by which continent it is situated on (cf. Figure 1).

Clearly, the choices above are debatable. We might for example also want to include some ideation-based variables in X (which are harder to measure, though) and also take different environmental variables E such as time instead of geographic location. We could even allow for additive effects of the environmental variable on the outcome of interest (such as a constant offset for each continent) but we do not touch this debate much more here as we are primarily interested in the methodological development.

Figure 1: Three candidates for a causal DAG with target total fertility rate (TFR) and four potential causal predictor variables. We would like to infer the parents of TFR in the true causal graph. We use the continent as the environment variable E . If the true DAG was one of the two graphs on the left, the environmental variable would have no direct influence on the target variable TFR and ‘Continent’ would be a valid environmental variable, see Definition 1.

The basic yet central insight underlying the invariance approach is the fact that for the true causal parental set $S^* := \text{pa}_Y$ we have the following conditional independence relation under Definition 1 of environmental variables:

$$Y \perp\!\!\!\perp E \mid X_{S^*}. \quad (2)$$

This follows directly from the local Markov condition [e. g. 36]. The goal is to find S^* by exploiting the above relation (2). Suppose we have a test for the null hypothesis

$$H_{0,S}: \quad Y \perp\!\!\!\perp E \mid X_S. \quad (3)$$

It was then proposed in [1] to define an estimate $\hat{S}$ for the parental set S^* by setting

$$\hat{S} := \bigcap_{S: H_{0,S} \text{ not rejected}} S. \quad (4)$$

Here, the intersection runs over all sets S , s.t. $E \cap S = \emptyset$. If the index set is empty, i. e. $H_{0,S}$ is rejected for all sets S , we define $\hat{S}$ to be the empty set. If we can test (3) with the correct type I error rate in the sense that

$$P(H_{0,S^*} \text{ is rejected at level } \alpha) \leq \alpha, \quad (5)$$

then we have as immediate consequence the desired statement

$$P(\hat{S} \subseteq S^*) \geq P(H_{0,S^*} \text{ accepted}) \geq 1 - \alpha.$$

This follows directly from the fact that S^* is accepted with probability at least $1 - \alpha$ since H_{0,S^*} is true; see [1] for details.

In the case of linear models, the method proposed by Peters et al. [1, Eq. (16)] considers a set S as invariant if there exist linear regression coefficients β and error variance σ which are identical across all environments. We consider the conditional independence relation in (3) as a generalization, even for linear relations. In the following example the regression coefficients are the same in all environments, and the residuals have the same mean and variance, but differ in higher order moments [cf. 1, Eq. (3)]:

Example 1. Consider a discrete environmental variable E . If in $E = 1$ we have

$$Y = 2X + N, N \perp\!\!\!\perp X,$$

and in $E = 2$

$$Y = 2X + M, M \perp\!\!\!\perp X,$$

where M and N have the same mean and variance but differ in higher order moments. In this case, we would have $E \perp\!\!\!\perp Y | X$, but the hypothesis “same linear regression coefficients and error variance” cannot be rejected.

The question remains how to test (3). If we assume a linear function f in the structural equation (1), then tests that can guarantee the level as in (5) are available [1]. The following examples show what could go wrong if the data contain nonlinearities that are not properly taken into account.

Example 2 (Linear model and nonlinear data). Consider the following SCM, in which X_2 and X_3 are direct causes of Y .

$$\begin{aligned} X_1 &\leftarrow E + \eta_X \\ X_2 &\leftarrow \sqrt{3X_1 + \eta_{X_1}} \\ X_3 &\leftarrow \sqrt{2X_1 + \eta_{X_2}} \\ Y &\leftarrow X_2^2 - X_3^2 + \eta_Y \end{aligned}$$

Due to the nonlinear effect, a linear regression from Y on X_2 and X_3 does not yield an invariant model. If we regress Y on X_1 , however, we obtain invariant prediction and independent residuals. In this sense, the linear version of ICP fails but it still chooses a set of ancestors of Y (it can be argued that this failure is not too severe).

Example 3 (Linear model and nonlinear data). In this example, the model misspecification leads to a wrong set that includes a descendant of Y . Consider the following SCM

$$\begin{aligned} X_1 &\leftarrow E + \eta_1 \\ Y &\leftarrow f(X_1) + \eta_Y \end{aligned}$$

$$X_2 \leftarrow g(X_1) + \gamma Y + \eta_2$$

with independent Gaussian error terms. Furthermore, assume that

$$\begin{aligned}\forall x \in \mathbb{R} : \quad & f(x) = \alpha x + \beta h(x) \\ \forall x \in \mathbb{R} : \quad & g(x) = h(x) - \gamma f(x) \\ & \beta \gamma^2 - \gamma = -\beta \text{var}(\eta_2) / \text{var}(\eta_Y)\end{aligned}$$

for some α, β and $h : \mathbb{R} \rightarrow \mathbb{R}$. Then, in the limit of an infinite sample size, the set $\{X_1, X_2\}$ is the only set that, after a linear regression, yields residuals that are independent of E . (To see this write $Y = f(X_1) + \eta_Y$ as a linear function in X_1, X_2 and show that the covariance between the residuals and X_2 is zero.) Here, the functions have to be “fine-tuned” in order to make the conditional $Y|X_1, X_2$ linear in X_1 and X_2 .⁴ As an example, one may choose $Y \leftarrow X_1 + 0.5X_1^2 + \eta_Y$ and $X_2 \leftarrow 0.5X_1^2 - X_1 + Y + \eta_2$ and η_1, η_Y, η_2 i.i.d. with distribution $\mathcal{N}(0, \sigma^2 = 0.5)$.

The examples show that ICP loses its coverage guarantee if we assume linear relationships for testing (3) while the true data generating process is nonlinear.

In the general nonlinear and nonparametric case, however, it becomes more difficult to guarantee the type I error rate when testing the conditional independence (3) [38]. This in contrast to nonparametric tests for (unconditional) independence [22, 26]. In a nonlinear conditional independence test setting, where we know an appropriate parametric basis expansion for the causal effect of the variables we condition on, we can of course revert back to unconditional independence testing. Apart from such special circumstances, we have to find tests that guarantee the type I error rate in (5) as closely as possible under a wide range of scenarios. We describe some methods that test (3) in Section 3 but for now let us assume that we are given such a test. We can then apply the method of nonlinear ICP (4) to the example of fertility data.

Example (fertility data)

The following sets were accepted at the level $\alpha = 0.1$ when using nonlinear ICP with invariant conditional quantile prediction (see Appendix II for details) as a conditional independence test:

$$\begin{aligned}S_1 &= \{Q5\} \\ S_2 &= \{\text{IMR, Imports of goods and services, Urban pop. (\% of total)}\} \\ S_3 &= \{\text{IMR, Education expend. (\% of GNI), Exports of goods and services, GDP per capita}\}\end{aligned}$$

As the intersection of $S_1, \dots, S_3$ is empty, we have $\hat{S} = \emptyset$. This motivates the concept of defining sets.

2.2 Defining sets

It is often impossible to distinguish between highly correlated variables. For example, infant mortality IMR and under-five mortality Q5 are highly correlated in the data and can often be substituted for each other. We accept sets that contain either of these variables. When taking the intersection as in (4), this leads to exclusion of both variables in $\hat{S}$ and potentially to an altogether empty set $\hat{S}$. We can instead ask for the defining sets [28], where a defining set $\hat{D} \subseteq \{1, \dots, p\}$ has the properties

- (i) $S \cap \hat{D} \neq \emptyset$ for all S such that $H_{0,S}$ is accepted.
- (ii) there exists no strictly smaller set D' with $D' \subset \hat{D}$ for which property (i) is true.

In words, we are looking for subsets $\hat{D}$, such that each accepted set S has at least one element that also appears in $\hat{D}$. If the intersection $\hat{S}$ (4) is non-empty, any subset of $\hat{S}$ that contains only one variable is a defining set.

⁴ This example is motivated by theory that combines linear and nonlinear models with additive noise [37].

Defining set are especially useful, however, in cases where the intersection $\hat{S}$ is empty. We still know that, with high probability, at least one of the variables in the defining set $\hat{D}$ has to be a parent. Defining sets are not necessarily unique. Given a defining set $\hat{D}$, we thus know that

$$P(S^* \cap \hat{D} = \emptyset) \leq P(H_{0,S^*} \text{ rejected}) \leq \alpha.$$

That is, a) at least one of the variables in the defining set $\hat{D}$ is a parent of the target, and b) the data do not allow to resolve it on a finer scale.

Example (fertility data)

We obtain seven defining sets:

$$\begin{aligned}\hat{D}_1 &= \{\text{IMR}, \text{Q5}\} \\ \hat{D}_2 &= \{\text{Q5}, \text{Education expenditure (\% of GNI)}, \text{Imports of goods and services}\} \\ \hat{D}_3 &= \{\text{Q5}, \text{Education expenditure (\% of GNI)}, \text{Urban pop. (\% of total)}\} \\ \hat{D}_4 &= \{\text{Q5}, \text{Exports of goods and services}, \text{Imports of goods and services}\} \\ \hat{D}_5 &= \{\text{Q5}, \text{Exports of goods and services}, \text{Urban pop. (\% of total)}\} \\ \hat{D}_6 &= \{\text{Q5}, \text{GDP per capita}, \text{Imports of goods and services}\} \\ \hat{D}_7 &= \{\text{Q5}, \text{GDP per capita}, \text{Urban pop. (\% of total)}\}\end{aligned}$$

Thus the highly-correlated variables infant mortality IMR and under-five mortality Q5 indeed form one of the defining sets in this example in the sense that we know at least one of the two is a causal parent for fertility but we cannot resolve which one it is or whether both of them are parents.

2.3 Confidence bands

For a given set S , we can in general construct a $(1 - \alpha)$ -confidence band $\hat{\mathcal{F}}_S$ for the regression function when predicting Y with the variables X_S . Note that if f is the regression function when regressing Y on the true set of causal variables X_{S^*} and hence, then, with probability $1 - \alpha$, we have

$$P(f \in \hat{\mathcal{F}}_{S^*}) \geq 1 - \alpha.$$

Furthermore, from Section 2.1 we know that H_{0,S^*} is accepted with probability $1 - \alpha$. We can hence construct a confidence band for the causal effects as

$$\hat{\mathcal{F}} := \bigcup_{S: H_{0,S} \text{ not rejected}} \hat{\mathcal{F}}_S. \quad (6)$$

Using a Bonferroni correction, we have the guarantee that

$$P(f \in \hat{\mathcal{F}}) \geq 1 - 2\alpha,$$

where the coverage guarantee is point-wise or uniform, depending on the coverage guarantee of the underlying estimators $\hat{\mathcal{F}}_S$ for all given $S \subseteq \{1, \dots, p\}$.

2.4 Average causal effects

The confidence bands $\hat{\mathcal{F}}$ themselves can be difficult to interpret. Interpretability can be guided by looking at the average causal effect in the sense that we compare the expected response at $\tilde{x}$ and x :

$$ACE(\tilde{x}, x) := E(Y | \text{do}(X = \tilde{x})) - E(Y | \text{do}(X = x)). \quad (7)$$

For the fertility data, this would involve a hypothetical scenario where we fix the variables to be equal to x for a country in the second term and, for the first term, we set the variables to $\tilde{x}$, which might differ from x just in

one or a few coordinates. Eq. (7) then compares the average expected fertility between these two scenarios. Note that the expected response under a do-operation is just a function of the causal variables $S^* \subseteq \{1, \dots, p\}$. That is—in the absence of hidden variables—we have

$$E(Y|\text{do}(X = x)) = E(Y|\text{do}(X_{S^*} = x_{S^*})),$$

and the latter is then equal to

$$E(Y|\text{do}(X_{S^*} = x_{S^*})) = E(Y|X_{S^*} = x_{S^*}),$$

that is it does not matter whether we set the causal variables to a specific value x_{S^*} or whether they were observed in this state.

Once we have a confidence band as defined in (6), we can bound the average causal effect (7) by the interval

$$\widehat{ACE}(\tilde{x}, x) := \left[\inf_{g \in \hat{\mathcal{F}}} (g(\tilde{x}) - g(x)), \sup_{g \in \hat{\mathcal{F}}} (g(\tilde{x}) - g(x)) \right],$$

with the immediate guarantee that

$$P(ACE(\tilde{x}, x) \in \widehat{ACE}(\tilde{x}, x)) \geq 1 - 2\alpha, \quad (8)$$

where the factor 2α is guarding, by a Bonferroni correction, against both a probability α that S^* will not be accepted—and hence $\hat{S} \subseteq S^*$ is not necessarily true—and another probability α that the confidence bands will not provide coverage for the parental set S^* .

Example (fertility data)

The confidence bands $\hat{\mathcal{F}}$, required for the computation of $\widehat{ACE}(\tilde{x}, x)$, are obtained by a time series bootstrap [39] as the fertility data contain temporal dependencies. The time series bootstrap procedure is described in Appendix I. We use a level of $\alpha = 0.1$ which implies a coverage guarantee of 80 % as per (8). In the examples below, we set x to an observed data point and vary only $\tilde{x}$.

In the first example, we consider the observed covariates for Nigeria in 1993 as x . The point of comparison $\tilde{x}$ is set equal to x , except for the variables in the defining set $\hat{D}_1 = \{\text{IMR}, \text{Q5}\}$. In Figures 2(a) and (b), these are

Figure 2: Data for Nigeria in 1993: The union of the confidence bands $\hat{\mathcal{F}}_S$, denoted by $\hat{\mathcal{F}}$, bounds the average causal effect of varying the variables in the defining set $\hat{D}_1 = \{\text{IMR}, \text{Q5}\}$ on the target $\log(\text{TFR})$. IMR and Q5 have been varied individually, see panels (a) and (b), as well as jointly, see panel (c), over their respective quantiles. In panels (a) and (b), we do not observe consistent effects different from zero as some of the accepted models do not contain IMR and some do not contain Q5. However, when varying the variables $\hat{D}_1 = \{\text{IMR}, \text{Q5}\}$ jointly (see panel (c)), we see that all accepted models predict an increase in expected $\log(\text{TFR})$ as IMR and Q5 increase.

Figure 3: (a) Bounds for the average causal effect of setting the variables IMR and Q5 in the African countries in 2013 to European levels, that is $\bar{x}$ differs from the country-specific observed values x in that the child mortality rates IMR and Q5 have been set to their respective European average. The implied coverage guarantee is 80 % as we chose $\alpha = 0.1$. (b) Random Forest regression model using all covariates as input. The (non-causal) regression effect coverage is again set to 80 %. We will argue below that the confidence intervals obtained by random forest are too small, see Table 1 and Figure 5.

varied individually over their respective quantiles. The overall confidence interval $\hat{\mathcal{F}}$ consists of the union of the shown confidence intervals $\hat{\mathcal{F}}_S$. If $x = \bar{x}$ (shown by the vertical lines), the average causal effect is zero, of course. In neither of the two scenarios shown in Figures 2(a) and (b), we observe consistent effects different from zero as some of the accepted models do not contain IMR and some do not contain Q5. However, when varying the variables $\hat{D}_1 = \{\text{IMR}, \text{Q5}\}$ jointly (see Figure 2(c)), we see that all accepted models predict an increase in expected log(TFR) as IMR and Q5 increase.

In the second example, we compare the expected fertility rate between countries where all covariates are set to the value x , which is here chosen to be equal to the observed values of all African countries in 2013. The expected value of log-fertility under this value x of covariates is compared to the scenario where we take as $\bar{x}$ the same value but set the values of the child-mortality variables IMR and Q5 to their respective European averages. The union of intervals in Figure 3(a) (depicted by the horizontal line segments) correspond to $\widehat{ACE}(\bar{x}, x)$ for each country under nonlinear ICP with invariant conditional quantile prediction. The accepted models make largely coherent predictions for the effect associated with this comparison. For most countries, the difference is negative, meaning that the average expected fertility declines if the child mortality rate in a

Table 1: Coverage.

Coverage guarantee	0.95	0.90	0.8	0.5
Coverage with nonlinear ICP	0.99	0.95	0.88	0.58
Coverage with Random Forest	0.76	0.71	0.61	0.32
Coverage with mean change	0.95	0.88	0.68	0.36

Figure 4: The confidence intervals show the predicted change in log (TFR) from 1973–2008 for all African countries when not using their data in the nonlinear ICP estimation (using invariant conditional quantile prediction with $\alpha = 0.1$). In other words, only the data from the remaining five continents was used during training. The horizontal line segments mark the union over the accepted models' intervals for the predicted change; the blue squares show the true change in log (TFR) from 1973–2008. Only those countries are displayed for which the response was not missing in the data, i. e. where log (TFR) in 1973 and in 2008 were recorded. The coverage is $25/26 \approx 0.96$.

country decreases to European levels. The countries where $\widehat{ACE}_{IMR+Q5}(\tilde{x}, x)$ contains 0 typically have a child mortality rate that is close to European levels, meaning that there is no substantial difference between the two points $\tilde{x}, x$ of comparison.

For comparison, in Figure 3(b), we show the equivalent computation as in Figure 3(a) when *all* covariates are *assumed* to have a direct causal effect on the target and a Random Forest is used for estimation [40]. We observe that while the resulting regression bootstrap confidence intervals often overlap with $\widehat{ACE}_{IMR+Q5}(\tilde{x}, x)$, they are typically much smaller. This implies that if the regression model containing all covariates was—wrongly—used as a surrogate for the causal model, the uncertainty of the prediction would be underestimated. Furthermore, such an approach ignoring the causal structure can lead to a significant bias in the prediction of causal effects when we consider interventions on descendants of the target variable, for example.

Lastly, we consider a cross validation scheme over time to assess the coverage properties of nonlinear ICP. We leave out the data corresponding to one entire continent and run nonlinear ICP with invariant conditional quantile prediction using the data from the remaining five continents. We perform this leave-one-continent-out scheme for different values of α . For each value of α , we then compute the predicted change in the response log (TFR) from 1973–2008 for each country belonging to the continent that was left out during the estimation procedure. The predictions are obtained by using the respective accepted models.⁵ We then compare the union of the associated confidence intervals with the real, observed change in log (TFR). This allows us to compute the coverage statistics shown in Table 1. We observe that nonlinear ICP typically achieves more accurate coverage compared to (i) a Random Forest regression model including all variables and (ii) a baseline where the predicted change in log (TFR) for a country is the observed mean change in log (TFR) across all continents other than the country's own continent. Figures 4 and 5 show the confidence intervals and the

⁵ Their number differs according to α : for a smaller α , additional models can be accepted compared to using a larger value of α . In other words, the accepted models for α_2 where $\alpha_1 < \alpha_2$ are a subset of the accepted models for α_1 .

(a) Nonlinear ICP with invariant conditional quantile prediction. Coverage: 0.88.

(b) Random Forest regression model with bootstrap confidence intervals. Coverage: 0.61.

(c) Prediction based on mean change on continents other than country's own continent. Coverage: 0.68.

Figure 5: The confidence intervals show the predicted change in log (TFR) from 1973–2008 for all countries when not using the data of the country's continent in the estimation (with implied coverage guarantee 80 %). Only those countries are displayed for which log (TFR) in 1973 and 2008 were not missing in the data. For nonlinear ICP the shown intervals are the union over the accepted models' intervals.

observed values for all African countries (Figure 4) and all countries (Figure 5) with observed log (TFR) in 1973 and 2008.

Recall that one advantage of a causal model is that, in the absence of hidden variables, it does not matter whether certain variables have been intervened on or whether they were observed in this state – the resulting

prediction remains correct in any of these cases. On the contrary, the predictions of a non-causal model can become drastically incorrect under interventions. This may be one reason for the less accurate coverage statistics of the Random Forest regression model—in this example, it seems plausible that some of the predictors were subject to different external ‘interventions’ across continents and countries.

3 Conditional independence tests

We present and evaluate an array of methods for testing conditional independence in a nonlinear setting, many of which exploit the invariance of causal models across different environments. Here, we briefly sketch the main ideas of the considered tests, their respective assumptions and further details are provided in Appendix II. All methods (A)–(F) are available in the package `CondIndTests` for the R language. Table 2 in Appendix II.7 shows the supported methods and options. An experimental comparison of the corresponding power and type I error rates of these tests can be found in Section 4.

- (A) **Kernel conditional independence test.** Use a kernel conditional independence test for $Y \perp\!\!\!\perp E \mid X_S$ [41, 27]. See Appendix II.1 for further details.
- (B) **Residual prediction test.** Perform a nonlinear regression from Y on X_S , using an appropriate basis expansion, and apply a variant of a Residual Prediction (RP) test [42]. The main idea is to scale the residuals of the regression such that the resulting test statistic is not a function of the unknown noise variance. This allows for a straight-forward test for dependence between the residuals and (E, X_S) . In cases where a suitable basis expansion is unknown, random features [43, 44] can be used as an approximation. See Appendix II.2 for further details.
- (C) **Invariant environment prediction.** Predict the environment E , once with a model that uses X_S as predictors only and once with a model that uses (X_S, Y) as predictors. If the null is true and we find the optimal model in both cases, then the out-of-sample performance of both models is statistically indistinguishable. See Appendix II.3 for further details.
- (D) **Invariant target prediction.** Predict the target Y , once with a model that uses X_S as predictors only and once with a model that uses (X_S, E) as predictors. If the null is true and we find the optimal model in both cases, then the out-of-sample performance of both models is statistically indistinguishable. See Appendix II.4 for further details.
- (E) **Invariant residual distribution test.** Pool the data across all environments and predict the response Y with variables X_S . Then test whether the distribution of the residuals is identical in all environments E . See Appendix II.5 for further details.
- (F) **Invariant conditional quantile prediction.** Predict a $1 - \beta$ quantile of the conditional distribution of Y , given X_S , by pooling the data over all environments. Then test whether the exceedance of the conditional quantiles is independent of the environment variable. Repeat for a number of quantiles and aggregate the resulting individual p -values by Bonferroni correction. See Appendix II.6 for further details.

Another interesting possibility for future work would be to devise a conditional independence test based on model-based recursive partitioning [45, 46].

Non-trivial, assumption-free conditional independence tests with a valid level do not exist [38]. It is therefore not surprising that all of the above tests assume the dependence on the conditioning variable to be “simple” in one form or the other. Some of the above tests require the noise variable in (1) to be additive in the sense that we do not expect the respective test to have the correct level when the noise is not additive. As additive noise is also used in Sections 2.3 and 2.4, we have written the structural equations above in an additive form.

One of the inherent difficulties with these tests is that the estimation bias when conditioning on potential parents in (3) can potentially lead to a more frequent rejection of a true null hypothesis than the nominal level suggests. In approaches (C) and (D), we also need to test whether the predictive accuracy is identical under both models and in approaches (E) and (F) we need to test whether univariate distributions remain invariant across environments. While these additional tests are relatively straightforward, a choice has to be made.

Discussion of power

Conditional independence testing is a statistically challenging problem. For the setting where we condition on a continuous random variable, we are not aware of any conditional independence test that holds the correct level and still has (asymptotic) power against a wide range of alternatives. Here, we want to briefly mention some power properties of the tests we have discussed above.

Invariant target prediction (D), for example, has no power to detect if the noise variance is a function of E , as shown by the following example

Example 4. Assume that the distribution is entailed by the following model

$$\begin{aligned} E &\leftarrow 0.2\eta_E \\ X &\leftarrow \eta_X \\ Y &\leftarrow X^2 + E \cdot \eta_Y, \end{aligned}$$

where $\eta_E, \eta_X, \eta_Y \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$. Then, any regression from Y on X and E yields the same results as regressing Y on X only. That is,

$$\text{for all } x, e : \mathbf{E}[Y | X = x] = \mathbf{E}[Y | X = x, E = e]$$

although

$$Y \not\perp\!\!\!\perp E | X.$$

The invariant residual distribution test (E), in contrast, assumes homoscedasticity and might have wrong coverage if this assumption is violated. Furthermore, two different linear models do not necessarily yield different distributions of the residuals when performing a regression on the pooled data set.

Example 5. Consider the following data generating process

$$\begin{aligned} Y^{e=1} &\leftarrow 2X^{e=1} + N^{e=1} \\ Y^{e=2} &\leftarrow -X^{e=2} + 0.3N^{e=2}, \end{aligned}$$

where the input variables $X^{e=1}$ and $X^{e=2}$ and the noise variables $N^{e=1}$ and $N^{e=2}$ have the same distribution in each environment, respectively. Then, approach (E) will accept the null hypothesis of invariant prediction.

It is possible to reject the null hypothesis of invariant prediction in Example 5 by testing whether in each environment the residuals are uncorrelated from the input.

Invariant conditional quantile prediction (F) assumes neither homoscedasticity nor does it suffer from the same issue of (D), i. e. no power against an alternative where the noise variance σ is a function of E . However, it is possible to construct examples where (F) will have no power if the noise variance is a function of both E and the causal variables X_{S^*} . Even then, though, the noise level would have to be carefully balanced to reduce the power to 0 with approach (F) as the exceedance probabilities of various quantiles (a function of X_{S^*}) would have to remain constant if we condition on various values of E .

4 Simulation study

For the simulations, we generate data from different nonlinear additive noise causal models and compare the performance of the proposed conditional independence tests. The structural equations are of the form $Z_k \leftarrow g_k(Z_{\text{pa}_k}) + \eta_k$, where the structure of the DAG is shown in Figure 6 and kept constant throughout the simulations for ease of comparison. We vary the nonlinearities used, the target, the type and strength of interventions, the noise tail behavior and whether parental contributions are multiplicative or additive. The simulation settings are described in Appendix III in detail.

Figure 6: The structure of the causal graph used in the simulations. The causal order is unknown for the simulations. All edge weights are 1 in absolute value.

We apply all the conditional independence tests (CITs) that we have introduced in Section 3, implemented with the following methods and tests as subroutines:

CIT	Implementation
(A)	KCI without Gaussian process estimation
(B)(i)	RP w/ Fourier random features
(B)(ii)	RP w/ Nyström random features and RBF kernel
(B)(iii)	RP w/ Nyström random features and polynomial kernel (random degree)
(B)(iv)	RP w/ provided polynomial basis (random degree)
(C)	Random forest and χ^2 -test
(D)(i)	GAM with F-Test
(D)(ii)	GAM with Wilcoxon test
(D)(iii)	Random forest with F-Test
(D)(iv)	Random forest with Wilcoxon test
(E)(i)	GAM with Kolmogorov-Smirnov test
(E)(ii)	GAM with Levene's test + Wilcoxon test
(E)(iii)	Random forest with Kolmogorov-Smirnov test
(E)(iv)	Random forest with Levene's test + Wilcoxon test
(F)	Quantile regression forest with Fisher's exact test

As a disclaimer we have to note that KCI is implemented without Gaussian process estimation. The KCI results shown below might improve if the latter is added to the algorithm.

Baselines

We compare against a number of baselines. Importantly, most of these methods contain various model misspecifications when applied in the considered problem setting. Therefore, they would not be suitable in practice. However, it is instructive to study the effect of the model misspecifications on performance.

1. The method of Causal Additive Models (CAM) [47] identifies graph structure based on nonlinear additive noise models [6]. Here, we apply the method in the following way. We run CAM separately in each environment and output the intersection of the causal parents that were retrieved in each environment. Note that the method's underlying assumption of Gaussian noise is violated.
2. We run the PC algorithm [3] in two different variants. We consider a variable to be the parent of the target if a *directed* edge between them is retrieved; we discard undirected edges. In the first variant of PC we consider, the environment variable is part of the input; conditional independence testing within the PC algorithm is performed with KCI, for unconditional independence testing we use HSIC [48, 49], using the implementation from [50] (denoted with 'PC(i)' in the figures). In the second variant, we run the PC algorithm on the pooled data (ignoring the environment information), testing for zero partial correlations (denoted with 'PC(ii)' in the figures). Here, the model misspecification is the assumed linearity of the structural equations.
3. We compare against linear ICP [6] where the model misspecification is the assumed linearity of the structural equations.

4. We compare against LiNGAM [15], run on the pooled data without taking the environment information into account. Here, the model misspecifications are the assumed linearity of the structural equations and the i.i.d. assumption which does not hold.
5. We also show the outcome of a random selection of the parents that adheres to the FWER-limit by selecting the empty set ($\hat{S} = \emptyset$) with probability $1 - \alpha$ and setting $\hat{S} = \{k\}$ for k randomly and uniformly picked from $\{1, \dots, p\} \setminus k'$ with probability α , where k' is the index of the current target variable. The random selection is guaranteed to maintain FWER at or below $1 - \alpha$.

Thus, all considered baseline models in 1.-4. —except for ‘PC(i)’—contain at least slight model misspecifications.

Metrics

Error rates and power are measured in the following by

- (i) Type I errors are measured by the **family-wise error rate** (FWER), the probability of making one or more erroneous selections

$$P(\hat{S} \not\subseteq S^*).$$

- (ii) Power is measured by the **Jaccard similarity**, the ratio between the size of the intersection and the size of the union of the estimated set $\hat{S}$ and the true set S^* . It is defined as 1 if both $S^* = \hat{S} = \emptyset$ and otherwise as

$$\frac{|\hat{S} \cap S^*|}{|\hat{S} \cup S^*|}.$$

The Jaccard similarity is thus between 0 and 1 and the optimal value 1 is attained if and only if $\hat{S} = S^*$.

Type-I-error rate of conditional independence tests

Figure 7 shows the average FWER on the x -axis (and the average Jaccard similarity on the y -axis) for all methods. The FWER is close but below the nominal FWER rate of $\alpha = 0.05$ for all conditional independence tests, that is $P(\hat{S} \subseteq S^*) \geq 1 - \alpha$. The same holds for the baselines linear ICP and random selection. Notably, the average Jaccard similarity of the random selection baseline is on average not much lower than for the other methods. The reason is mostly a large variation in average Jaccard similarity across the different target variables, as discussed further below and as will be evident from Figure 8 (top right plot). In fact, as can be seen from Figure 9, random guessing is much worse than the optimal methods on each target variable. The FWER of the remaining baselines CAM, LiNGAM, PC(i) and PC(ii) lies well above α .

A caveat of the FWER control seen in Figure 7 is that while the FWER is maintained at the desired level, the test H_{0,S^*} might be rejected more often than with probability α . The error control rests on the fact that H_{0,S^*} is accepted with probability higher than $1 - \alpha$ (since the null is true for S^*). However, if a mistake is made and H_{0,S^*} is falsely rejected, then we might still have $\hat{S} \subseteq S^*$ because either all other sets are rejected, too, in which case $\hat{S} = \emptyset$, or because other sets (such as the empty set) are accepted and the intersection of all accepted sets is—by accident—again a subset of S^* . In other words: some mistakes might cancel each other out but overall the FWER is very close to the nominal level, even if we stratify according to sample size, target, type of nonlinearity and other parameters, as can be seen from Figure 8.

Power

Figure 7 shows on the y -axis the average Jaccard similarity for all methods. The optimal value is 1 and is attained if and only if $\hat{S} = S^*$. A value 0 corresponds to disjoint sets $\hat{S}$ and S^* . The average Jaccard similarity is around 0.4 for most methods and not clearly dependent on the type I errors of the methods. Figure 8 shows the average FWER and Jaccard similarities stratified according to various parameters.

Figure 7: Average Jaccard similarity (y-axis) against average FWER (x-axis), stratified according to which conditional independence test (A)–(F) or baseline method has been used. The nominal level is $\alpha = 0.05$, illustrated by the vertical dotted line. The shown results are averaged over all target variables. Since the empty set is the correct solution for target variable 1 and 5, methods that mostly return the empty set (such as random or linear ICP) perform still quite well in terms of average Jaccard similarity. Since all variables are highly predictive for the target variable Y , see Figure 6, classical variable selection techniques as LASSO have a FWER that lies far beyond α . Importantly, the considered baselines are not suitable for the considered problem setting due to various model misspecifications. We show their performance for comparison to illustrate the influence of these misspecifications.

One of the most important determinants of success (or the most important) is the target, that is the variable for which we would like to infer the causal parents; see top right panel in Figure 8. Variables 1 and 5 as targets have a relatively high average Jaccard similarity when trying to recover the parental set. These two variables have an empty parental set ($S^* = \emptyset$) and the average Jaccard similarity thus always exceeds $1 - \alpha$ if the level of the procedure is maintained as then $\hat{S} = \emptyset = S^*$ with probability at least $1 - \alpha$ and the Jaccard similarity is 1 if both $\hat{S}$ and S^* are empty. As testing for the true parental set corresponds to an unconditional independence test in this case, maintaining the level of the test procedure is much easier than for the other variables.

Figure 9 shows the same plot as Figure 7 for each of the more difficult target variables 2, 3, 4, and 6 separately. As can be seen from the graph in Figure 6 and the detailed description of the simulations in Appendix III, the parents of target variable 3 are difficult to estimate as the paths $1 \rightarrow 2 \rightarrow 3$ and $1 \rightarrow 3$ cancel each other exactly in the linear setting (and approximately for nonlinear data), thus creating a non-faithful distribution. The cancellation of effects holds true if interventions occur on variable 1 and not on variable 2. A local violation of faithfulness leaves type I error rate control intact but can hurt power as many other sets besides the true S^* can get accepted, especially the empty set, thus yielding $\hat{S} = \emptyset$ when taking the intersection across all accepted sets to compute the estimate $\hat{S}$ in (4). Variable 4, on the other hand, has only a single parent, namely $S^* = \{3\}$, and the recovery of the single parent is much easier, with average Jaccard similarity up to a third. Variable 6 finally again has average Jaccard similarity of up to around a tenth only. It does not suffer from a local violation of faithfulness as variable 3 but the size of the parental set is now three, which again hurts the power of the procedure, as often already a subset of the true parents will be accepted and hence $\hat{S}$ in (4) will not be equal to S^* any longer but just a subset. For instance, when variable 5 is not intervened on in any environment it cannot be identified as a causal parent of variable 6, as it is then indistinguishable from the noise term. Similarly, in the linear setting, merely variable 3 can be identified as a parent of variable 6 if the interventions act on variables 1 and/or 2 only.

Figure 8: Average Jaccard similarity over the conditional independence tests (A)–(F) (y-axis) against average FWER (x-axis), stratified according to various parameters (from top left to bottom right): sample size ‘n’, type of nonlinearity ‘id’, ‘target variable’, intervention location ‘interv’, multiplicative effects indicator ‘multiplic’, ‘strength’ of interventions, mean value of interventions ‘meanshift’, shift intervention indicator ‘shift’ and degrees of freedom for t-distributed noise ‘df’. For details, see the description in Appendix III. The FWER is within the nominal level in general for all conditional independence tests. The average Jaccard similarity is mostly determined by the target variable under consideration, see top right panel.

The baselines LiNGAM and PC show a larger Jaccard similarity for target variables 3, 4 (only LiNGAM), and 6 at the price of large FWER values.

In Appendix IV, Figures 10–13 show the equivalent to Figure 8, separately for target variables 2, 3, 4 and 6. For the sample size n , we observe that increasing it from 2000 to 5000 decreases power in case of target variables 4. This behavior can be explained by the fact that when testing S^* in Eq. (3), the null is rejected too often as the bias in the estimation performed as part of the conditional independence test yields deviations from the null that become significant with increasing sample size. For the nonlinearity, we find that the function $f_4(x) = \sin(2\pi x)$ is the most challenging one among the nonlinearities considered. It is associated with very low Jaccard similarity values for the target variables that do have parents. For the intervention type, it may seem surprising that ‘all’ does not yield the largest power. A possible explanation is that intervening on all variables except for the target yields more similar intervention settings—the intervention targets do not differ between environments 2 and 3, even though the strength of the interventions is different. So more heterogeneity between the intervention environments, i. e. also having different intervention targets, seems to improve performance in terms of Jaccard similarity. Lastly, we see that power is often higher for additive parental contributions than for multiplicative ones.

Figure 9: The identical plot to Figure 7 separately for target variables 2, 3, 4 and 6. For all target variables, method (E)(ii)—an invariant residual distribution test using GAM with Levene’s test + Wilcoxon test—performs constantly as good or nearly as good as the optimal method among the considered tests.

In summary, all tests (A)–(F) seem to maintain the desired type I error, chosen here as the family-wise error rate, while the power varies considerably. An invariant residual distribution test using GAM with Levene’s test and Wilcoxon test produces results here that are constantly as good or nearly as good as the optimal methods for a range of different settings. However, it is only applicable for categorical environmental variables. For continuous environmental variables, the results suggest that the residual prediction test with random features might be a good choice.

5 Discussion and future work

Causal structure learning with the invariance principle was proposed [1]. However, the assumption of linear models in [1] is unrealistic in many applications. In this work, we have shown how the framework can be ex-

tended to nonlinear and nonparametric models by using suitable nonlinear and nonparametric conditional independence tests. The properties of these conditional independence tests are critically important for the power of the resulting causal discovery procedure. We evaluated many different test empirically in the given context and highlighted approaches that seem to work robustly in different settings. In particular we find that fitting a nonlinear model with pooled data and then testing for differences between the residual distributions across environments results in desired coverage and high power if compared against a wide range of alternatives.

Our approach allowed us to model how several interventions may affect the total fertility rate of a country, using historical data about decline and rise of fertilities across different continents. In particular, we provided bounds on the average causal effect under certain (hypothetical) interventions such as a reduction in child mortality rates. We showed that the causal prediction intervals for hold-out data have better coverage than various baseline methods. The importance of infant mortality rate and under-five mortality rate on fertility rates is highlighted, reconfirming previous studies that have shown or hypothesized these factors to be important [30, 32]. We stress that the results rely on causal sufficiency of the used variables, an assumption that can and should be debated for this particular example.

We also introduced the notion of ‘defining sets’ in the causal discovery context that helps in situations where the signal is weak or variables are highly correlated by returning sets of variables of which we know that at least one variable (but not necessarily all) in this set are causal for the target variable in question.

Finally, we provide software in the R [29] package `nonlinearICP`. A collection of the discussed conditional independence tests are part of the package `CondIndTests` and are hopefully of independent interest.

In applications where it is unclear whether the underlying models are linear or not, we suggest the following. While our proposed methods also hold the significance level if the underlying models are linear, we expect the linear version of ICP to have more power. Therefore, it is advisable to use the linear version of ICP if one has strong reasons to believe that the underlying model is indeed linear. In practice, one might first apply ICP with linear models and apply a nonlinear version if, for example, all linear models are rejected. One would then need to correct for multiple testing by a factor of 2.

Acknowledgment: We thank Jan Ernest and Adrian Raftery for helpful discussions and an AE and two referees for very helpful comments on an earlier version of the manuscript.

Appendix I. Time series bootstrap procedure

In the time series bootstrap procedure used to obtain the confidence bands $\hat{\mathcal{F}}$ in Section 2.4, B bootstrap samples of the response Y are generated by first fitting the model on all data points. We then use the fitted values and residuals from this model. Each bootstrap sample is generated by resampling the residuals of this fit block- and country-wise. In more detail, we define the block-length l_b of residuals that should be sampled consecutively (we use $l_b = 3$) and we sample a number of time points $t_{s_1}, \dots, t_{s_k}$ from which the residuals are resampled. For a country a and the first time point t_1 , consider the fitted values at point t_1 and the fitted values for the $l_b - 1$ consecutive observations. We then sample a country b and add country b ’s residuals from time points $t_{s_1}, t_{s_1+1}, \dots, t_{s_1+l_b-1}$ to the fitted values of country a for the considered period $t_1, \dots, t_{l_b}$. We then proceed with the next l_b consecutive fitted values for country a and add country b ’s residuals from observations $t_{s_2}, t_{s_2+1}, \dots, t_{s_2+l_b-1}$, until all fitted values of country a are covered. This procedure is applied to each country. Finally, to obtain the confidence intervals, we fit the model on each of the B bootstrap samples (Y^b, X) , consisting of the response Y^b generated from the fitted values and the resampled residuals, and the observations X which have not been modified.

Appendix II. Conditional independence tests

For completeness, we first restate the generic method for Invariant Causal Prediction from [1]:

Algorithm 1 Generic method for Invariant Causal Prediction.

Input: i.i.d. sample of (Y, X, E) , α

- 1: **for each** $S \subseteq \{1, \dots, p\}$ **do**
- 2: Test whether $H_{0,S}$ holds at level α .
- 3: **end for**
- 4: Set $\hat{S} := \bigcap_{S: H_{0,S} \text{ not rejected}} S$

Output: $\hat{S}$

The conditional independence tests discussed in this work can be used to perform the test in Step 2 of Algorithm 1. Therefore, the inputs to these tests consist of an i.i.d. sample of (Y, X_S, E) and α where X_S contains the variables corresponding to $S \subseteq \{1, \dots, p\}$, i. e. the subset to be tested. Additionally, some test specific parameters might need to be specified. The return value of the tests is the respective test's decision about $H_{0,S}$.

For most tests, $E \in \mathbb{R}^d$ can be either discrete or continuous. As all empirical results in this work use an environment variable that is discrete and one-dimensional, the descriptions below focus on this setting. We then denote the index set of different environments with $\mathcal{E}$. We will comment on the required changes for the continuous and higher-dimensional case in the respective sections. Whenever applying the test for environmental variables $E \in \mathbb{R}^d$ with $d > 1$ is infeasible with the method, each test can be applied separately for each variable in E . The overall p -value is obtained by multiplying the minimum of the individual p -values by d , i. e. by applying a Bonferroni correction for the number of environmental variables. When applying the function `CondIndTest()` from the R package `CondIndTests` with a conditional independence test that does not support a multidimensional environment variable, the described Bonferroni correction is applied.

II.1 Kernel conditional independence test

Setting and assumptions

We use the kernel conditional independence test proposed in [27]. When E is discrete, we use a delta kernel for E , and otherwise an RBF kernel. The test is also applicable when E contains more than one environmental variable as the inputs can be sets of random variables.

II.2 Residual prediction test

Setting and assumptions

We do not expect this test to have the correct level when the noise in Eq. (1) is not additive. The described procedure does not need to be modified for higher-dimensional and/or continuous environmental variables E .

We consider a version of a Residual Prediction test as proposed in [42] to determine whether $H_{0,S}$ holds at level α for a particular set of variables S . The main idea is to find a suitable basis expansion of f that allows us to regress Y on X_S by reverting back to the linear case. Given an appropriate basis expansion, the *scaled* ordinary least squares residuals can then be tested for possible remaining nonlinear dependencies between the scaled residuals and (E, X_S) . The scaling ensures that the resulting test statistic is not a function of the noise variance. Under the null, the scaled residuals are expected to behave roughly like the noise term. In other words, there should be no dependence between the scaled residuals and the environmental variables and X_S , so there should be no signal left in the residuals that could be fitted by a nonparametric method like a

random forest using E and X_S as predictors. This necessitates to make an assumption on the noise distribution F_ε , e. g. $\varepsilon \sim \mathcal{N}(0, 1)$.

In order to generalize the method to settings where an appropriate basis expansion of f is unknown, we look at ways to find such a suitable basis expansion automatically by using random features [43, 44].

Algorithm 2 Residual Prediction tests applied to nonlinear ICP.

Input: i.i.d. sample of (Y, X_S, E) , α , F_ε , B , a subroutine to compute the basis functions $h_m(\cdot)$ for $m = 1, \dots, M$

- 1: Compute the non-linear transformations $h_m(X_S)$, $m = 1, \dots, M$ and create the design matrix $\mathbf{H}_{X_S} \in \mathbb{R}^{n \times M}$ comprising these M nonlinear features.
- 2: Regress Y on $\mathbf{H}_{X_S}$ with ordinary least squares.
- 3: Predict (a function of) the scaled residuals with the environment variable E and X_S .
- 4: Compute a statistic for the prediction accuracy to be used as test statistic.
- 5: **for** b from 1 to B **do**
- 6: Simulate one sample of size n from the assumed noise distribution F_ε .
- 7: Predict (a function of) these simulated values after rescaling with the environment variable E and X_S .
- 8: Compute a statistic for the prediction accuracy.
- 9: **end for**
- 10: The B simulated values for prediction accuracy yield the empirical null distribution from which the p -value is obtained.

Output: Decision about $H_{0,S}$

Step 1

The choice of $h_m(X_S)$, $m = 1, \dots, M$ can be based on domain knowledge, e. g. when the nonlinearity in Eq. (1) is known to be a polynomial of a given order. If such domain knowledge is not available, the linear basis expansion can be approximated by random features, e. g. using the Nyström method or by random Fourier features. For these methods, the kernel function needs to be chosen as well as the kernel parameters and the number of random features to be generated.

Step 3

For instance, a random forest can be used for the estimation. If the residuals only differ in the second moments, predicting the expectation of the residuals is not sufficient as the predictors E have no discriminative power for this task. In such a setting, the absolute value of the residuals can be predicted to exploit the heterogeneity in the second moments across environments.

Step 4

For instance, the mean squared error can be used here.

Step 5

If the error term is non-Gaussian, the appropriate distribution can be used at this stage to accommodate non-Gaussianity of the noise.

Parameter settings used in simulations

In the simulations, we use $B = 250$ and $\varepsilon \sim \mathcal{N}(0, 1)$. In step 1, we consider the following options: (a) Fourier random features (approach (B)(i) in Section 4), (b) Nyström random features and RBF kernel ((B)(ii)), (c) Nyström random features and polynomial kernel of random degree ((B)(iii)), (d) polynomial basis of random degree ((B)(iv)). The number of random features in (c) and (d) is chosen to be $\lceil n/4 \rceil$. In step 7, we predict the mean as well as the absolute value of the residuals and aggregate the results using a Bonferroni correction.

II.3 Invariant environment prediction

Setting and assumptions

The described procedure does not need to be modified for continuous environmental variables E . For higher-dimensional E the test would need to be applied for each variable separately and the resulting p -values would need to be aggregated with a Bonferroni correction.

Algorithm 3 Invariant environment prediction for nonlinear ICP.

Input: i.i.d. sample of (Y, X_S, E) , α , subroutine for test in step 5.

- 1: Split the sample into training and test set.
- 2: Use the training set to train a model to predict E with (Y, X_S) as predictors.
- 3: Use the training set to train a model to predict E with X_S as predictors.
- 4: For both fits, compute the prediction accuracy on the test set.
- 5: Use a one-sided test at the significance level α to assess whether the prediction accuracy of the fit using (Y, X_S) as predictors is larger than the prediction accuracy of the fit using only X_S as predictors.

Output: Decision about $H_{0,S}$

Step 3

When a random forest is used to predict the environment variable, one can also use X_S and a permutation of Y as predictors to ensure the random forest fits are based on the same number of predictor variables. As the number of variables considered for each split in the random forest estimation procedure is a function of the total number of predictor variables, this helps to mitigate differences between the prediction accuracies that are only due to artefacts of the estimation procedure. This is especially relevant for small sets S .

Step 5

For instance, a χ^2 test can be used here. If the null is true and we find the optimal model in both cases, then the out-of-sample performance of both models is statistically indistinguishable as Y is independent of E given X_S . If the null is not true, we expect the model containing the response to perform better as Y contains additional information in this case (since Y is not independent of E given X_S).

Parameter settings used in simulations

In step 1, we use 2/3 of the data points for training and 1/3 for testing. In step 3, we use a random forest to predict the environment variable and use X_S and a permutation of Y as predictors. In step 4, we use the χ^2 test implemented in `prop.test()` [51] from the `stats` package in R.

II.4 Invariant target prediction

Setting and assumptions

The described procedure does not need to be modified for continuous and/or higher-dimensional environmental variables E .

Step 3

When a random forest is used, one can also use X_S and a permutation of E as predictors to ensure the random forest fit is based on the same number of predictor variables. As the number of variables considered for each split in the random forest estimation procedure is a function of the total number of predictor variables, this helps to mitigate differences between the prediction accuracies that are only due to artefacts of the estimation procedure. This is especially relevant for small sets S . As an alternative to using a random forest, one could use GAMs as the estimation procedure, implying the implicit assumption that the components in $f(X)$ in Eq. (1) are additive.

Algorithm 4 Invariant target prediction for nonlinear ICP.**Input:** i.i.d. sample of (Y, X_S, E) , α , subroutine for test in step 5.

- 1: Split the sample into training and test set.
- 2: Use the training set to train a model to predict Y with (X_S, E) as predictors.
- 3: Use the training set to train a model to predict Y with X_S as predictors.
- 4: For both fits, compute the prediction accuracy on the test set.
- 5: Use a one-sided test at the significance level α to assess whether the prediction accuracy of the fit using (X_S, E) as predictors is larger than the prediction accuracy of the fit using only X_S as predictors.

Output: Decision about $H_{0,S}$ **Step 5**

For instance, an F-test can be used here. Another option is a Wilcoxon test using the difference between the absolute residuals. If the null is true and we find the optimal model in both cases, then the out-of-sample performance of both models is statistically indistinguishable as Y is independent of E given X_S . If the null is not true, we expect the model additionally containing E to perform better as E contains additional information in this case (since Y is not independent of E given X_S).

Parameter settings used in simulations

In step 1, we use 2/3 of the data points for training and 1/3 for testing. In step 3, to predict Y we use a GAM or a random forest. In step 5, we use an F-test or a Wilcoxon test (`wilcox.test()` from the `stats` package in R). These combinations yield approaches (D)(i)–(iv) in Section 4. When using a random forest in step 3, we use X_S and a permutation of E as predictors.

II.5 Invariant residual distribution test**Setting and assumptions**

We do not expect this test to have the correct level when the noise in Eq. (1) is not additive. It is only applicable to discrete environmental variables. For higher-dimensional E the test would need to be applied for each variable separately and the resulting p -values would need to be aggregated with a Bonferroni correction.

Algorithm 5 Invariant residual distribution test for nonlinear ICP.**Input:** i.i.d. sample of (Y, X_S, E) , α , subroutine for test in step 4.

- 1: Pool the data from all environments and fit a model to predict Y with X_S .
- 2: Initialize $pv \leftarrow 1$, $t \leftarrow 0$.
- 3: **for each** $e \in \mathcal{E}$ **do**
- 4: Use a two-sample test to assess whether the residuals of samples from environment e have the same distribution as the residuals of samples from environments in the index set $\mathcal{E}'$ where $\mathcal{E}' = \mathcal{E} \setminus \{e\}$, yielding the p -value pv_e .
- 5: $t \leftarrow t + 1$
- 6: $pv \leftarrow \min(pv, pv_e)$.
- 7: **if** $|\mathcal{E}| = 2$ **then**
- 8: **break**
- 9: **end if**
- 10: **end for**
- 11: Apply a Bonferroni correction for the number of performed tests t : $pv \leftarrow t \cdot pv$.

Output: Decision about $H_{0,S}$

Step 1

For instance, one could use a random forest or a GAM as the estimation procedure. The latter implicitly assumes that the components in f in Eq. (1) are additive.

Step 4

For instance, a nonparametric test such as Kolmogorov-Smirnov can be used here. Alternatively, we can limit the test to assess equality of first and second moments by first using a Wilcoxon test for the expectation with an one-vs-all scheme as described in the algorithm. Subsequently, Levene's test for homogeneity of variance across groups can be used to test for equality of the second moments of the residual distributions. In this case, the final p -value would be twice the minimum of (a) the Bonferroni-corrected p -value from the one-vs-all Wilcoxon test and (b) the p -value from Levene's test.

Parameter settings used in simulations

In step 1, we use a GAM or a random forest. In step 4, we use both approaches described above, using (a) `ks.test()` from the `stats` package in R [52] and (b) `wilcox.test()` and `levene.test()` (the latter being contained in the `lawstat` package in R [53, 54]). These combinations yield approaches (E)(i)–(iv) in Section 4.

II.6 Invariant conditional quantile prediction

Setting and assumptions

For continuous and/or higher-dimensional environmental variables E the test described in Steps 4–11 which assesses whether Exceedance $\perp\!\!\!\perp E$ would need to be modified according to the structure of E .

Algorithm 6 Invariant conditional quantile prediction for nonlinear ICP.

Input: i.i.d. sample of (Y, X_S, E) , α , set of quantiles $\mathcal{B}$, subroutine for test in step 7.

```

1: Initialize  $pv \leftarrow 1, t \leftarrow 0$ .
2: for each  $\beta \in \mathcal{B}$  do
3:   Predict  $1 - \beta$  quantile  $Q_{1-\beta}(x)$  of  $Y|X_S = x$ .
4:   for each  $e \in \mathcal{E}$  do
5:     Define one-vs-all environment  $I = \mathbb{1}_{\{E=e\}}$ 
6:     Define exceedance  $E_{1-\beta} = \mathbb{1}_{\{Y > \hat{Q}_{1-\beta}(x)\}}$ 
7:     Test whether  $E_{1-\beta}$  is independent of  $I$ :  $pv_{e,\beta} \leftarrow \text{StatTest}(E_{1-\beta}, I, \alpha)$ 
8:      $t \leftarrow t + 1, pv \leftarrow \min(pv, pv_{e,\beta})$ 
9:     if  $|\mathcal{E}| = 2$  then
10:       break
11:   end if
12: end for
13: end for
14: Apply a Bonferroni correction for the number of performed tests  $t$ :  $pv \leftarrow t \cdot pv$ 
```

Output: Decision about $H_{0,S}$

Step 3

For instance, a Quantile Regression Forest [55] can be used here.

Step 7

For instance, Fisher's exact test can be used here by computing the 2×2 contingency table of the exceedance of the residuals for the quantile $1 - \beta$ for $I = 0$ and for $I = 1$.

Parameter settings used in simulations

In step 3, we use a quantile regression forest for $\mathcal{B} = \{0.1, 0.5, 0.9\}$. In step 7, we use `fisher.test()` from the `stats` package in R.

II.7 Overview of conditional independence tests in CondIndTests package

The described conditional independence tests are available in the R package `CondIndTests`. A wrapper function `CondIndTest()` is provided which takes the respective test as the argument `method`. The package supports the estimation procedures, subroutines and statistical tests shown in Table 2. The column E indicates whether the environmental variables can be discrete ('D'), continuous ('C'), or both; the column d shows the supported dimensionality of E .

As described at the beginning of Appendix II, a Bonferroni correction is applied when calling the function `CondIndTest()` with a conditional independence test that does not support a multidimensional environment variable. Similarly, a Bonferroni correction is applied when the first input argument Y to the respective test is multidimensional and if the specified test does not support this internally.

Table 2: Overview of implemented test combinations in `CondIndTests` package.

CIT	R function name/METHOD	TEST	E	d
(A)	KCI() KCI (without GP support)	–	D/C	≥ 1
(B)	ResidualPredictionTest() Residual prediction test with – Nyström random features (RBF and polynomial kernel) – Fourier random features – fixed basis expansion	–	D/C	≥ 1
(C)	InvariantEnvironmentPrediction() Random forest	χ^2 test (<code>prop.test()</code>) Wilcoxon test (<code>wilcox.test()</code>)	D/C	1
(D)	InvariantTargetPrediction() Random forest GAM	F-Test Wilcoxon test (<code>wilcox.test()</code>)	D/C	≥ 1
(E)	InvariantResidualDistributionTest() Random forest GAM	Kolmogorov-Smirnov (<code>ks.test()</code>) Levene's test + Wilcoxon test (<code>levene.test()</code> , <code>wilcox.test()</code>)	D	1
(F)	InvariantConditionalQuantilePrediction() Quantile regression forest	Fisher's exact test (<code>fisher.test()</code>)	D	1

Appendix III. Experimental settings for numerical studies

For each simulation, we compare the performance of all methods and conditional independence tests while choosing the following parameters randomly (but keeping them constant for one simulation): In total, there are 27478 simulations from 1240 distinct settings that are evaluated for each of the 22 considered methods.

- (i) **Sample size.** Sample size 'n' is chosen randomly in the set $\{100, 200, 500, 2000, 5000\}$.
- (ii) **Target variable.** We sample one of the variables in the graph in Figure 6 at random as a target variable (variable 'target' is chosen uniformly from $\{1, 2, \dots, 6\}$ in other words).

- (iii) **Tail behavior of the noise.** The noise η_k for $k = 1, \dots, 6$ is sampled from a t-distribution and the degrees of freedom are chosen at random from $\text{df} \in \{2, 3, 5, 10, 20, 50, 100\}$, where the latter is very close to a Gaussian distribution.
- (iv) **Multiplicative or additive effects.** For each simulation setting, we determine whether $g_k(\cdot)$ has additive or multiplicative components. We sample additive components of the form $g_k(Z_{\text{pa}_k}) = \sum_{j \in \text{pa}_k} f(\epsilon_{j,k} \cdot Z_j)$ (*multiplic* = FALSE) and multiplicative components of the form $g_k(Z_{\text{pa}_k}) = \prod_{j \in \text{pa}_k} f(\epsilon_{j,k} \cdot Z_j)$ (*multiplic* = TRUE) with equal probability, where the signs $\epsilon_{j,k} \in \{-1, 1\}$ are as shown in Figure 6 along the relevant arrows.
- (v) **Shift- or do-Interventions.** The variable ‘shift’ is set with equal probability to either TRUE (shift-interventions) or FALSE (do-interventions). For do-interventions we replace the structural equation of the intervened variable $k \in \{1, \dots, q\}$ by

$$Z_k \leftarrow e_k,$$

where e_k is the randomly chosen intervention value which is sampled i.i.d for each observation as described under (vi). For shift-interventions, the value e_k is added as

$$Z_k \leftarrow g_k(Z_{\text{pa}_k}) + \eta_k + e_k.$$

See for example Section 5 of [1] for a more detailed discussion of shift interventions.

- (vi) **Strength of interventions** The intervention values e_k are chosen independently for all variables from a t-distribution with ‘df’ degrees of freedom, shifted by a constant ‘meanshift’ (chosen uniformly at random in $\{0, 0.1, 0.2, 0.5, 1, 2, 5, 10\}$, and scaled by a factor ‘strength’, chosen uniformly at random in $\{0, 0.1, 0.2, 0.5, 1, 2, 5, 10\}$).
- (vii) **Non-linearities.** For the functions $f = f_{\text{id}}$ we consider the following four nonlinear functions, where the index ‘id’ is sampled uniformly from $\{1, 2, 3, 4\}$ and the same nonlinearity is used throughout the graph:

$$\begin{aligned} f_1(x) &= x, \\ f_2(x) &= \max\{0, x\}, \\ f_3(x) &= \text{sign}(x) \cdot \sqrt{|x|}, \\ f_4(x) &= \sin(2\pi x), \end{aligned}$$

- (viii) **Location of interventions.** Each sample is independently assigned into one environment in $\mathcal{E} = \{1, 2, 3\}$, where $\{1\}$ corresponds to observational data, that is all samples in environment $\{1\}$ are sampled as observational data, where samples in environments $\{2, 3\}$ are intervention data. The intervention targets and strengths for samples in environment $\{2\}$ are drawn as per the description below and kept constant for all samples in environment $\{2\}$ and then analogously for environment $\{3\}$, where intervention targets are drawn independently and identically to environment $\{2\}$. The variable ‘interv’ is set uniformly at random to one of the values {‘all’, ‘rand’, ‘close’}. If it is equal to ‘all’, then interventions in environments $\{2, 3\}$ occur at all variables except for the target variable. If it is equal to ‘rand’, then interventions occur at one ancestor, chosen uniformly at random, and one descendant of the target variable, chosen again uniformly at random (the empty set is chosen in case there are no ancestors or descendants). Finally, if it is equal to ‘close’, interventions occur at one parent, chosen uniformly at random, and one child of the target variable, chosen again uniformly at random (and again no interventions occur if these sets are empty).

Appendix IV. Additional experimental results

Figures 10–13 show the equivalent to Figure 8, separately for target variables 2, 3, 4 and 6. For the sample size n , we observe that increasing it from 2000 to 5000 decreases power in case of target variable 4. This behavior

Figure 10: Average Jaccard similarity over the conditional independence tests (A)–(F) (y-axis) against average FWER (x-axis) when estimating the parents of variable 2. The figure is otherwise generated analogously to Figure 8.

Figure 11: Average Jaccard similarity over the conditional independence tests (A)–(F) (y-axis) against average FWER (x-axis) when estimating the parents of variable 3. The figure is otherwise generated analogously to Figure 8.

Figure 12: Average Jaccard similarity over the conditional independence tests (A)–(F) (y-axis) against average FWER (x-axis) when estimating the parents of variable 4. The figure is otherwise generated analogously to Figure 8.

Figure 13: Average Jaccard similarity over the conditional independence tests (A)–(F) (y-axis) against average FWER (x-axis) when estimating the parents of variable 6. The figure is otherwise generated analogously to Figure 8.

can be explained by the fact that when testing S^* in Eq. (3), the null is rejected too often as the bias in the estimation performed as part of the conditional independence test yields deviations from the null that become significant with increasing sample size. For the nonlinearity, we find that the function $f_4(x) = \sin(2\pi x)$ is the most challenging one among the nonlinearities considered. It is associated with very low Jaccard similarity values for the target variables that do have parents. For the intervention type, it may seem surprising that ‘all’ does not yield the largest power. A possible explanation is that intervening on all variables except for the target yields more similar intervention settings—the intervention targets do not differ between environments 2 and 3, even though the strength of the interventions is different. So more heterogeneity between the intervention environments, i. e. also having different intervention targets, seems to improve performance in terms of Jaccard similarity. Lastly, we see that power is often higher for additive parental contributions than for multiplicative ones.

Appendix V. Example

Here we illustrate the methods presented in this manuscript by considering a causal DAG $X_1 \rightarrow X_2 \rightarrow X_3$. Figure 14 visualizes the generated data. There are six environments with shift interventions. The latter act

Figure 14: Visualization of the sample considered in the example in Section V.

on X_1 in two environments (green, yellow) and on X_3 in four environments (green, cyan, blue, magenta). The red environment consists of observational data. We run the proposed approaches (A)–(F) to retrieve the parents of X_2 , i. e. $S^* = \{X_1\}$. Below we give an overview of which sets were accepted by the respective methods with $\alpha = 0.05$. We see that approaches (A), (B)(i)+(ii), (E)(i)–(iii) and (F) retrieve S^* correctly, while the other approaches return the empty set.

CIT	$S_0 = \{\}$	$S_1 = \{X_1\}$	$S_2 = \{X_3\}$	$S_3 = \{X_1, X_3\}$	$\hat{S}$
(A)		✓		✓	$\{X_1\}$
(B)(i)		✓		✓	$\{X_1\}$
(B)(ii)		✓		✓	$\{X_1\}$
(B)(iii)					$\{\}$
(B)(iv)					$\{\}$
(C)		✓	✓	✓	$\{\}$
(D)(i)	✓	✓		✓	$\{\}$
(D)(ii)	✓	✓	✓	✓	$\{\}$
(D)(iii)					$\{\}$
(D)(iv)	✓			✓	$\{\}$
(E)(i)		✓		✓	$\{X_1\}$
(E)(ii)		✓		✓	$\{X_1\}$
(E)(iii)		✓			$\{X_1\}$
(E)(iv)					$\{\}$
(F)		✓		✓	$\{X_1\}$

References

- Peters J, Bühlmann P, Meinshausen N. Causal inference using invariant prediction: identification and confidence intervals. *J R Stat Soc, Ser B (with discussion)*. 2016;78(5):947–1012.
- Pearl J. *Causality: Models, Reasoning, and Inference*. 2nd ed. New York, USA: Cambridge University Press; 2009.
- Spirtes P, Glymour C, Scheines R. *Causation, Prediction, and Search*. 2nd ed. Cambridge, USA: MIT Press; 2000.
- Peters J, Janzing D, Schölkopf B. *Elements of Causal Inference: Foundations and Learning Algorithms*. Cambridge, MA, USA: MIT Press; 2017.
- Chickering DM. Optimal structure identification with greedy search. *J Mach Learn Res*. 2002;3:507–54.
- Peters J, Mooij JM, Janzing D, Schölkopf B. Causal discovery with continuous additive noise models. *J Mach Learn Res*. 2014;15:2009–53.
- Heckerman D. A Bayesian approach to causal discovery. Technical report, Microsoft Research (MSR-TR-97-05). 1997.
- Hauser A, Bühlmann P. Jointly interventional and observational data: estimation of interventional Markov equivalence classes of directed acyclic graphs. *J R Stat Soc B*. 2015;77:291–318.
- Pearl J. A constraint propagation approach to probabilistic reasoning. In: *Proceedings of the 4th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 1985. p. 31–42.
- Pearl J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA: Morgan Kaufmann Publishers Inc.; 1988.
- Heinze-Deml C, Maathuis MH, Meinshausen N. Causal structure learning. *Annu Rev Stat App*. 2018;5(1):371–91.
- Maathuis M, Kalisch M, Bühlmann P. Estimating high-dimensional intervention effects from observational data. *Ann Stat*. 2009;37:3133–64.
- Colombo D, Maathuis M, Kalisch M, Richardson T. Learning high-dimensional directed acyclic graphs with latent and selection variables. *Ann Stat*. 2012;40:294–321.
- Claassen T, Mooij JM, Heskes T. Learning sparse causal models is not NP-hard. In: *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 2013.
- Shimizu S, Hoyer PO, Hyvärinen A, Kerminen AJ. A linear non-Gaussian acyclic model for causal discovery. *J Mach Learn Res*. 2006;7:2003–30.
- Peters J, Bühlmann P. Identifiability of Gaussian structural equation models with equal error variances. *Biometrika*. 2014;101:219–28.
- Hoyer PO, Janzing D, Mooij JM, Peters J, Schölkopf B. Nonlinear causal discovery with additive noise models. *Adv Neural Inf Process Syst*. 2009;21:689–96.

18. Haavelmo T. The probability approach in econometrics. *Econometrica*. 1944;12:S1–115. (supplement).
19. Aldrich J. *Autonomy* Oxford Economic Papers. 1989;41:15–34.
20. Hoover KD. The logic of causal inference. *Econ Philos*. 1990;6:207–34.
21. Schölkopf B, Janzing D, Peters J, Sgouritsa E, Zhang K, Mooij J. On causal and anticausal learning. In: *Proceedings of the 29th International Conference on Machine Learning (ICML)*. 2012. p. 1255–62.
22. Bergsma W, Dassios A. A consistent test of independence based on a sign covariance related to Kendall's tau. *Bernoulli*. 2014;20:1006–28.
23. Hoeffding W. A non-parametric test of independence. *Ann Math Stat*. 1948;19:546–57. 12.
24. Blum JR, Kiefer J, Rosenblatt M. Distribution free tests of independence based on the sample distribution function. *Ann Math Stat*. 1961;32:485–98.
25. Rényi A. On measures of dependence. *Acta Math Acad Sci Hung*. 1959;10:441–51. ISSN 1588-2632.
26. Székely GJ, Rizzo ML, Bakirov NK. Measuring and testing dependence by correlation of distances. *Ann Stat*. 2007;35:2769–94.
27. Zhang K, Peters J, Janzing D, 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)*. 2011. p. 804–13.
28. Goeman JJ, Solari A. Multiple testing for exploratory research. *Statistical Science*. 2011. 584–597.
29. R Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing; 2017. <https://www.R-project.org>.
30. Hirschman C. Why fertility changes. *Annu Rev Sociol*. 1994;20(1):203–33.
31. Huinink J, Kohli M, Ehrhardt J. Explaining fertility: The potential for integrative approaches. *Demogr Res Monogr*. 2015;33: 93.
32. Raftery A, Lewis S, Aghajanian A. Demand or ideation? Evidence from the Iranian marital fertility decline. *Demography*. 1995;32(2):159–82.
33. Angrist JD, Imbens GW, Rubin DB. Identification of causal effects using instrumental variables. *J Am Stat Assoc*. 1996;91:444–55.
34. Imbens GW. Instrumental variables: An econometrician's perspective. *Stat Sci*. 2014;29(3):323–58.
35. Nations U. *World population prospects: The 2012 revision*. Population Division, Department of Economic and Social Affairs. New York: United Nations; 2013. <https://esa.un.org/unpd/wpp/Download/Standard/ASCII/>.
36. Lauritzen SL. *Graphical Models*. New York, USA: Oxford University Press; 1996.
37. Ernest J, Rothenhäusler D, Bühlmann P. Causal inference in partially linear structural equation models. *Ann Stat*. 2018;46:2904–38.
38. Shah RD, Peters J. The hardness of conditional independence testing and the generalised covariance measure. *ArXiv e-prints*. 2018. 1804.07203.
39. Künsch HR. The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*. 1989. 1217–1241.
40. Breiman L. Random forests. *Mach Learn*. 2001;45:5–32.
41. Fukumizu K, Gretton A, Sun X, Schölkopf B. Kernel measures of conditional dependence. In: *Advances in Neural Information Processing Systems 20*. 2008. p. 489–96.
42. Shah RD, Bühlmann P. Goodness-of-fit tests for high dimensional linear models. *J R Stat Soc, Ser B, Stat Methodol*. 2018;80(1):113–35. <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssb.12234>.
43. Williams CKI, Seeger M. Using the nyström method to speed up kernel machines. In: *Advances in Neural Information Processing Systems 13*. MIT Press; 2001. p. 682–8.
44. Rahimi A, Recht B. Random features for large-scale kernel machines. In: *Advances in Neural Information Processing Systems 20*. Curran Associates, Inc.; 2008. p. 1177–84.
45. Zeileis A, Hothorn T, Hornik K. Model-based recursive partitioning. *J Comput Graph Stat*. 2008;17(2):492–514.
46. Hothorn T, Zeileis A. partykit: A modular toolkit for recursive partytioning in R. *J Mach Learn Res*. 2015;16:3905–9.
47. Bühlmann P, Peters J, Ernest J. CAM: Causal additive models, high-dimensional order search and penalized regression. *Ann Stat*. 2014;42:2526–56.
48. Gretton A, Bousquet O, Smola A, Schölkopf B. Measuring statistical dependence with Hilbert-Schmidt norms. In: *Algorithmic Learning Theory*. Springer; 2005. p. 63–78.
49. Gretton A, Fukumizu K, Teo CH, Song L, Schölkopf B, Smola AJ. A kernel statistical test of independence. *Proc Neural Inf Process Syst*. 2007;20:1–8.
50. Pfister N, Bühlmann P, Schölkopf B, Peters J. Kernel-based tests for joint independence. *J R Stat Soc, Ser B*. 2017;80:5–31.
51. Wilson EB. Probable Inference, the Law of Succession, and Statistical Inference. *J Am Stat Assoc*. 1927;22:209–12.
52. Conover WJ. *Practical nonparametric statistics*. New York: John Wiley & Sons; 1971.
53. Levene H. Robust tests for equality of variances. In: Olkin I, editor. *Contributions to Probability and Statistics*. Palo Alto, CA: Stanford University Press; 1960. p. 278–92.
54. Gastwirth JL, Gel YR, Wallace Hui WL, Lyubchich V, Miao W, Noguchi K. lawstat: Tools for Biostatistics, Public Policy, and Law. 2015. <https://CRAN.R-project.org/package=lawstat>. R package version 3.0.
55. Meinshausen N. Quantile regression forests. *J Mach Learn Res*. 2006;7:983–99.