

Telling Cause from Effect using MDL-based Local and Global Regression

Alexander Marx and Jilles Vreeken
Max Planck Institute for Informatics and Saarland University
Saarland Informatics Campus, Saarbrücken, Germany
{amarx, jilles}@mpi-inf.mpg.de

Abstract—We consider the fundamental problem of inferring the causal direction between two univariate numeric random variables X and Y from observational data. The two-variable case is especially difficult to solve since it is not possible to use standard conditional independence tests between the variables.

To tackle this problem, we follow an information theoretic approach based on Kolmogorov complexity and use the Minimum Description Length (MDL) principle to provide a practical solution. In particular, we propose a compression scheme to encode local and global functional relations using MDL-based regression. We infer X causes Y in case it is shorter to describe Y as a function of X than the inverse direction. In addition, we introduce SLOPE, an efficient linear-time algorithm that through thorough empirical evaluation on both synthetic and real world data we show outperforms the state of the art by a wide margin.

Index Terms—Kolmogorov Complexity, MDL, Causal Inference, Regression, Hypercompression

I. INTRODUCTION

Telling cause from effect from observational data is one of the fundamental problems in science [26], [18]. We consider the problem of inferring the most likely direction between two univariate numeric random variables X and Y . That is, we are interested in identifying whether X causes Y , whether Y causes X , or whether they are merely correlated.

Traditional methods, that rely on conditional independence tests, cannot decide between the Markov equivalent classes of $X \rightarrow Y$ and $Y \rightarrow X$ [18]. Recently, it has been postulated however that if $X \rightarrow Y$, there exists an independence between the marginal distribution of the cause, $P(X)$, and the conditional distribution of the effect given the cause, $P(Y | X)$ [25], [9]. The state of the art exploits this asymmetry in various ways, and overall obtain up to 70% accuracy on a well-known benchmark of cause-effect pairs [24], [8], [20], [10], [17]. In this paper we break this barrier, and give an elegant score that is computable in linear-time and obtains over 82% accuracy on the same benchmark.

To illustrate its strength, we show the results of our method, called SLOPE, and four state-of-the-art methods over this benchmark in Fig. 1. In particular, we show the accuracy over the top- k pairs ranked according to the score at hand, the so-called decision rate. The plot shows that SLOPE leads by large margin over its competitors; it is 100% accurate over the 32 pairs it is most certain about, and is 90% accurate over its top-74 out of 98 pairs. Unlike its competitors, its accuracies are strongly significant with regard to the 95% confidence




Figure 1: [Higher is better] Weighted accuracy of our method, SLOPE, versus the state of the art in causal inference for univariate numeric pairs as identified in a recent survey [17], i.e., CURE [24], RESIT [20], IGCI [10] and ANM-pHSIC [8] on the Tuebingen benchmark data set (98 pairs). The gray area indicates the 95% confidence interval of a fair coin toss.

interval of a fair coin flip. Moreover, our score comes with a natural significance test that allows us to weed out insignificant results; when we consider the 83 significant pairs only we find that SLOPE is even 85% accurate.

We base our method on the algorithmic Markov condition, a recent postulate by Janzing and Schölkopf [9], that states that if X causes Y , the factorization of the joint distribution $P(X, Y)$ in the causal direction has a simpler description—in terms of Kolmogorov complexity—than that in the anti-causal direction. That is, if $X \rightarrow Y$, $K(P(X)) + K(P(Y | X)) \leq K(P(Y)) + K(P(X | Y))$. As any physical process can be modelled by a Turing machine, this ideal score can detect any causal dependence that can be explained by a physical process. However, Kolmogorov complexity is not computable, so we need a practical instantiation of this ideal. In this paper we do so using the Minimum Description Length (MDL) principle, which provides a statistically well-founded approximation of Kolmogorov complexity.

Simply put, we propose to fit a regression model from X to Y , and vice versa, measuring both the complexity of the function, as well as the error it makes in bits, and infer that causal direction by which we can describe the data most succinctly. There is a little bit more to it, of course. We carefully construct an MDL score such that we can meaningfully compare between different types of functional dependencies, including linear, quadratic, cubic, reciprocal, and exponential functions. This way, for example, we will find

that we can more succinctly describe the data in Fig. 2a by a cubic function than with a linear function, as while it takes fewer bits to describe the latter function, it will take many more bits to describe the large error it makes.

We do not only consider models that try to explain all the data with a single, global, deterministic regression function, but also allow for non-deterministic models. That is, we consider compound regression functions that in addition to a global deterministic function additionally include regression functions for local parts of the data corresponding to specific, duplicated X values. For example, consider the data in Fig. 2b, where the Y values belonging to a single X value clearly show more structure than the general linear trend. In contrast, if we rotate the plot by 90 degrees, we do not observe the same regularities for the X values mapped to a single Y value. In many cases, e.g., $Y = 1$ there is only one mapping X value. We can exploit this asymmetry by considering local regression functions per value of X , each individually fitted but as we assume all noise to be of the same type, all should be of the same function class. In this particular example, we therewith correctly infer that X causes Y . The MDL principle prevents us from overfitting, as such local functions are only included if they aid global compression. Last, but not least, we give a linear-time algorithm, SLOPE, to compute this score.

As we model Y as a function of X and noise, our approach is somewhat reminiscent to causal inference based on Additive Noise Models (ANMs) [25], where one assumes that Y is generated as a function of X plus additive noise, $Y = f(X) + N$ with $X \perp N$. In the ANM approach, we infer $X \rightarrow Y$ if we can find a function from X to Y that admits an ANM, but cannot do so in the opposite direction. In practice, ANM methods often measure the independence between the presumed cause and the noise in terms of p-values, and infer the direction of the lowest p-value. As we will see, this leads to unreliable confidence scores—not the least because p-values are often strongly influenced by sample size [1], but also as that a lower p-value does not necessarily mean that H_1 is more true, just that H_0 is very probably not true [1]. We will show that our score, on the other hand, is robust against sample size, and correlates strongly with accuracy. Moreover, it admits an elegant and effective analytical statistical test on the *difference* in score between the two causal directions based on the no-hypercompressibility inequality [3], [7].

Our key contributions can be summarised as follows, we

- (a) propose an MDL score for causal inference on pairs of univariate numeric random variables,
- (b) formulate an analytic significance test based on compression,
- (c) show to model unobserved mechanisms via compound deterministic and non-deterministic functions,
- (d) introduce the linear-time SLOPE algorithm,
- (e) give extensive empirical evaluation on synthetic and real-world data, including a case study
- (f) make all code, data generators, and data available.

The remainder of this paper is organised as usual. We first give a brief primer to Kolmogorov complexity and the




Figure 2: Example deterministic and non-deterministic data. In both cases the ground truth is X causes Y . The left-hand data is generated using a quadratic function with Gaussian noise, whereas the right-hand data is generated using a non-deterministic function.

Minimum Description Length principle in Sec. II. In Sec. III we introduce our score based on the algorithmic independence of conditional, as well as a practical instantiation based on the MDL principle. We introduce the linear-time SLOPE algorithm to efficiently compute the conditional score in Sec. IV. Sec. V discusses related work. We empirically evaluate SLOPE in Sec. VI and discuss the results in Sec. VII. We round up with conclusions in Sec. VIII.

II. PRELIMINARIES

In causal inference, the goal is to determine for two random variables X and Y that are statistically dependent whether it is more likely that X causes Y , denoted by $X \rightarrow Y$, or whether it is more likely that Y causes X , $Y \rightarrow X$. In this paper we consider the case where X and Y are univariate and numeric. We work under the common assumption of causal sufficiency [4], [17], [20], [28]. That is, we assume there is no hidden confounder variable Z that causes both X and Y .

We base our causal inference score on the notion of Kolmogorov complexity, which we will approximate via MDL. Below we give brief primers to these two main concepts.

A. Kolmogorov Complexity

The Kolmogorov complexity of a finite binary string x is the length of the shortest binary program p^* for a universal Turing machine \mathcal{U} that outputs x and then halts [11], [13]. Formally,

$$K(x) = \min\{|p| \mid p \in \{0, 1\}^*, \mathcal{U}(p) = x\}.$$

Simply put, p^* is the most succinct *algorithmic* description of x , and therewith Kolmogorov complexity of x is the length of its ultimate lossless compression. Conditional Kolmogorov complexity, $K(x \mid y) \leq K(x)$, is then the length of the shortest binary program p^* that generates x , and halts, given y as input.

The Kolmogorov complexity of a probability distribution P , $K(P)$, is the length of the shortest program that outputs $P(x)$ to precision q on input $\langle x, q \rangle$ [13]. More formally, we have

$$K(P) = \min \{ |p| : p \in \{0, 1\}^*, |\mathcal{U}(\langle x, \langle q, p \rangle \rangle) - P(x)| \leq 1/q \} .$$

The conditional, $K(P \mid Q)$, is defined similarly except that the universal Turing machine \mathcal{U} now gets the additional information Q . The algorithmic mutual information between two distributions P and Q is $I(P : Q) = K(P) - K(P \mid Q^*)$, where Q^* is the shortest binary program for Q . For more details on Kolmogorov complexity see [13].

B. Minimum Description Length Principle

Kolmogorov complexity is, however, not computable [13], but we can approximate it in a well-founded and computable way through the Minimum Description Length (MDL) principle [22], [7]. Conceptually, instead of all programs, Ideal MDL considers only those for which we know that they output x and halt, i.e., lossless compressors. Formally, given a model class \mathcal{M} , MDL identifies the best model $M \in \mathcal{M}$ for data D as the one minimizing

$$L(D, M) = L(M) + L(D \mid M) ,$$

where $L(M)$ is the length in bits of the description of M , and $L(D \mid M)$ is the length in bits of the description of data D given M . This is known as two-part, or *crude* MDL. There also exists one-part, or *refined* MDL. Although refined MDL has theoretically appealing properties, it is only efficiently computable for a small number of model classes.

To use MDL in practice we need to define a model class, and how to encode a model, resp. the data given a model, into bits. Note that in MDL we are only concerned with optimal code *lengths*, not actual codes—our goal is to measure the *complexity* of a dataset under a model class, after all [7].

III. INFORMATION THEORETIC CAUSAL INFERENCE

In this section, we first introduce how to infer causal directions using Kolmogorov complexity. Thereupon, we show how to obtain a computable score based on the MDL principle.

A. Causal Inference by Kolmogorov Complexity

A central postulate in causal inference concerns the algorithmic independence of conditionals. For multiple random variables, this postulate is defined as follows [9].

Algorithmic Independence of Conditionals: *A causal hypothesis is only acceptable if the shortest description of the joint density P is given by the concatenation of the shortest description of the Markov kernels. Formally, we write*

$$K(P(X_1, \dots, X_n)) \stackrel{\pm}{=} \sum_j K(P(X_j \mid PA_j)) , \quad (1)$$

which holds up to an additive constant independent of the input, and where PA_j corresponds to the parents of X_j in a causal directed acyclic graph (DAG).

As we consider two variables, X and Y , either X is the parent of Y or the other way round. That is, either

$$K(P(X, Y)) \stackrel{\pm}{=} K(P(X)) + K(P(Y \mid X)) , \text{ or} \\ K(P(X, Y)) \stackrel{\pm}{=} K(P(Y)) + K(P(X \mid Y)) .$$

In other words, two valid ways to describe the joint distribution of X and Y include to first describe the marginal distribution $P(X)$ and then the conditional distribution $P(Y \mid X)$, or first to describe $P(Y)$ and then $P(X \mid Y)$.

Thereupon, Janzing and Schölkopf formulated the postulate for algorithmic independence of Markov kernels [9].

Algorithmic Independence of Markov Kernels: *If $X \rightarrow Y$, the marginal distribution of the cause $P(X)$ is algorithmically independent of the conditional distribution of the effect given the cause $P(Y \mid X)$, i.e., the algorithmic mutual information between the two will be zero,*

$$I(P(X) : P(Y \mid X)) \stackrel{\pm}{=} 0 , \quad (2)$$

while this is not the case in the other direction.

Simply put, for the true causal direction, the marginal distribution of the cause is algorithmically independent of the conditional distribution of the effect given the cause. Building upon Eqs. (1) and (2), Mooij et al. [16] derived an inference rule stating that if X causes Y ,

$$K(P(X)) + K(P(Y \mid X)) \leq K(P(Y)) + K(P(X \mid Y)) \quad (3)$$

holds up to an additive constant. This means that if $X \rightarrow Y$, the description of the joint distribution $K(P(X, Y))$ of first describing the marginal distribution of the cause $K(P(X))$ and then describing the conditional distribution of the effect given the cause $K(P(Y \mid X))$, will be shorter than the other way around.

Although Eq. (3) already allows for inferring the causal direction for a given pair, we obtain a more robust score, allowing for fair comparison of results independent of data sizes, when we normalise the result. In particular, Budhathoki and Vreeken [4] recently proposed to normalise the scores with the sum of the description lengths for the marginal distributions. We therefore define our causal indicator as

$$\Delta_{X \rightarrow Y} = \frac{K(P(X)) + K(P(Y \mid X))}{K(P(X)) + K(P(Y))} ,$$

and $\Delta_{Y \rightarrow X}$ in the same manner. Consequently, we infer $X \rightarrow Y$, if $\Delta_{X \rightarrow Y} < \Delta_{Y \rightarrow X}$, and $Y \rightarrow X$, if $\Delta_{X \rightarrow Y} > \Delta_{Y \rightarrow X}$ and do not decide if $\Delta_{X \rightarrow Y} = \Delta_{Y \rightarrow X}$.

The confidence of our score is $\mathbb{C} = |\Delta_{X \rightarrow Y} - \Delta_{Y \rightarrow X}|$. The higher, the more certain we are that the inferred causal direction is correct. To avoid confusion, we want to emphasise that \mathbb{C} has nothing to do with a confidence interval, but can be used to rank results of several tests. Below, after introducing our practical score, we will show how we can in addition define a practical statistical test.

B. Causal Inference by MDL

As Kolmogorov complexity is not computable, we will instantiate $\Delta_{X \rightarrow Y}$ and $\Delta_{Y \rightarrow X}$ using the Minimum Description Length principle [13], [7]. In practice this means we will estimate $\Delta_{X \rightarrow Y}$ as

$$\hat{\Delta}_{X \rightarrow Y} = \frac{L(X) + L(Y | X)}{L(X) + L(Y)}$$

where $L(X)$ is the length in bits of the description of the marginal distribution of X , $L(Y)$ that of the marginal distribution of Y , and $L(Y | X)$ that of the conditional distribution of Y given X . We define $\hat{\Delta}_{Y \rightarrow X}$ analogue to $\hat{\Delta}_{X \rightarrow Y}$, and we infer $X \rightarrow Y$, if $\hat{\Delta}_{X \rightarrow Y} < \hat{\Delta}_{Y \rightarrow X}$, $Y \rightarrow X$, if $\hat{\Delta}_{X \rightarrow Y} > \hat{\Delta}_{Y \rightarrow X}$ and do not decide if $\hat{\Delta}_{X \rightarrow Y} = \hat{\Delta}_{Y \rightarrow X}$ or below a user-defined threshold. Like above, confidence \mathbb{C} is simply the absolute difference between $\hat{\Delta}_{X \rightarrow Y}$ and $\hat{\Delta}_{Y \rightarrow X}$.

Considering the difference between the encoded lengths is related to, but not the same as considering the ratio of the posteriors; we also include the complexity of the model, which helps against overfitting. Intuitively, if the functions we find for the two directions both explain the data equally well, we prefer that direction that explains it using the simplest function.

This leaves us to explain how we encode the data, and, most importantly, how we encode $L(Y | X)$.

Intuition of the Conditional Encoding: The general idea is simple: we use regression to model the data of Y given X . That is, we model Y as a function f of X and independent noise N , i.e. $Y = f(X) + N$. We do so by fitting a regression function f over X and treating the error it makes as Gaussian distributed noise. Naturally, the better $f(X)$ fits Y , the fewer bits we will have to spend on encoding errors. The more parameters $f(X)$ has, however, the more bits we will have to spend on encoding these. This way, MDL naturally balances the complexity of the model to that of the data [7]. For example, while a linear function is more simple to describe than a cubic one, the latter will fit the data plotted in Fig. 2a so much better that MDL decides it is the better choice.

A key idea in our approach is to consider not only single global deterministic regression functions f_g , which works well for deterministic data, but to also consider non-deterministic, or compound functions as models. That is, we consider models that besides the global regression function f_g may additionally consist of *local* regression functions f_l that model Y for those values x of X that non-deterministically map to multiple values of Y . That is, per such value of X , we take the associated values of Y , sort these ascending, and uniformly re-distribute them on X over a fixed interval. We now see how well, just for these re-distributed points, we can fit a local regression model f_l . This way, we will for example be able to much more succinctly describe the data in Fig. 2b than with a single global deterministic regression function, as we can now exploit the structure that the values of Y have given a value of X , namely, being approximately equally spaced. To avoid overfitting we use MDL, and only allow a local function for a value of X into our model if it provides a gain in overall compression. Since we assume that for the true causal model

the data in the local components follows the same pattern, we only allow models in which all local functions are of the same type, e.g., all are linear, all are quadratic, etc.

In the following paragraphs, we formalise these ideas and define our cost functions. All logarithms are to base 2, and we use the common convention that $0 \log 0 = 0$.

Complexity of the Marginals: We start by defining the cost for the marginal distributions, $L(X)$ and $L(Y)$, which mostly serve to normalise our causal indicators $\hat{\Delta}_{X \rightarrow Y}$ and $\hat{\Delta}_{Y \rightarrow X}$. As we beforehand do not know how X or Y are distributed, and do not want to incur any undue bias, we encode both using a uniform prior with regard to the data resolution τ of X and Y . That is, we have $L(X) = -n \log \tau$, where τ is the resolution of the data of X . Note that resolution τ can be different between X and Y —we specify how we choose τ in the next section. We define $L(Y)$ analogue.

Complexity of the Conditional Model: Formally, we write F for the set of regression functions, or model, we use to encode the data of Y given X . A model F consists of at least one global regression function $f_g \in \mathcal{F}$, and up to $dom(X)$ local regression functions $f_l \in \mathcal{F}$, associated with individual values of X . We write F_l for the set of local regression functions $f_l \in F_l$, and require that all $f_l \in F_l$ are of the same type. The description length, or encoded size, of F is

$$L(F) = L_{\mathbb{N}}(|F|) + \log \binom{|X| - 1}{|F_l| - 1} + 2 \log(|\mathcal{F}|) + L(f_g) + \sum_{f_l \in F_l} L(f_l),$$

where we first describe the number of local functions using $L_{\mathbb{N}}$, the MDL optimal encoding for integers $z \geq 1$ [23], then map each f_l to its associated value of X , after which we use $\log |\mathcal{F}|$ bits to identify the type of the global regression function f_g , and whenever F_l is non-empty also $\log |\mathcal{F}|$ bits to identify the type of the local regression functions f_l , finally, we encode the functions themselves. Knowing the type of a function, we only need to encode its parameters, and hence

$$L(f) = \sum_{\phi \in \Phi_f} L_{\mathbb{N}}(s) + L_{\mathbb{N}}(\lceil \phi \cdot 10^s \rceil) + 1,$$

where we encode each parameter ϕ up to a certain precision p . We shift ϕ by the smallest integer number s such that $\phi \cdot 10^s \geq 10^p$, i.e. $p = 3$ means that we consider three digits. What remains is that we need to encode the shift, the shifted digit and the sign.

Complexity of the Conditional Data: Reconstructing the data of Y given $f(X)$ corresponds to encoding the errors, or deviations, the model makes. Since we fit our regression functions by minimizing the sum of squared errors, which corresponds to maximizing the likelihood under a Gaussian, it is a natural choice to encode the errors using a Gaussian distribution with zero-mean.

Since we have no assumption on the standard deviation, we use the empirical estimate $\hat{\sigma}$ to define the standard deviation

of the Gaussian. By doing so, the encoded size of the error of $F(X)$ with respect to the data of Y corresponds to

$$L(Y | F, X) = \sum_{f \in F} \left(\frac{n_f}{2} \left(\frac{1}{\ln 2} + \log 2\pi\hat{\sigma}^2 \right) - n_f \log \tau \right),$$

where n_f is the number of data points for which we use a specific function $f \in F$. Intuitively, this score is higher the less structure of the data is described by the model and increases proportionally to the sum of squared errors.

Complexity of the Conditional: Having defined the data and model costs above, we can now proceed and define the total encoded size of the conditional distribution of Y given X as

$$L(Y | X) = L(F) + L(Y | F, X). \quad (4)$$

By MDL we are after that model F that minimises Eq. (4). After discussing a significance test for our score, we will present the SLOPE algorithm to efficiently compute the conditional score in the next section.

Significance by Hypercompression: Ranking based on confidence works well in practice. Ideally, we would additionally like to know the significance of an inference. It turns out we can define an appropriate hypothesis test using the no-hypercompressibility inequality [3], [7]. In a nutshell, under the hypothesis that the data was sampled from the null-model, the probability that any other model can compress it k bits better is $P_0(L_0(x) - L(x) \geq k) \leq 2^{-k}$.

This means that if we assume the null complexity, L_0 , to be the least-well compressed causal direction, we can evaluate the probability of gaining k bits by instead using the most-well compressed direction. Formally, if we write $L(X \rightarrow Y)$ for $L(X) + L(Y | X)$, and vice-versa for $L(Y \rightarrow X)$, L_0 would be $\max\{L(X \rightarrow Y), L(Y \rightarrow X)\}$. The probability that the data can be compressed $k = |L(X \rightarrow Y) - L(Y \rightarrow X)|$ bits better in the opposite direction then simply is 2^{-k} .

In fact, we can construct a stronger test by assuming that the data is not causated, but merely correlated. That is, we assume *both* directions are wrong; the one compresses too well, the other compresses too poor. Following, if we assume these two to be equal in terms of exceptionality, the null complexity is the mean between the complexities of the two causal directions, i.e., $L_0 = \min\{L(X \rightarrow Y), L(Y \rightarrow X)\} + |L(X \rightarrow Y) - L(Y \rightarrow X)|/2$. The probability of the best-compressing direction is then 2^{-k} with $k = |L(X \rightarrow Y) - L(Y \rightarrow X)|/2$. We can now set a significance threshold α as usual, such as $\alpha = 0.001$, and use this to prune out those cases where the difference in compression between the two causal directions is insignificant. We will evaluate this procedure, in addition to our confidence score, in the experiments.

IV. THE SLOPE ALGORITHM

With the framework defined in the previous section, we can determine the most likely causal direction and the corresponding confidence value. In this section we present the SLOPE algorithm to efficiently compute the causal indicators. To keep the computational complexity of the algorithm linear,

Algorithm 1: CONDITIONALCOSTS(Y, X)

input : random variables Y and X
output: score $L(Y | X)$

- 1 $F =$ empty model;
- 2 $f_g =$ FITDETERMINISTIC($Y \sim X, \mathcal{F}$);
- 3 $F = F \cup f_g$;
- 4 $s = s_g = L(F) + L(Y | F, X)$;
- 5 $X_u = \{x \in X \mid \text{count}(x) \geq 2\}$;
- 6 **foreach** $\mathcal{F}_c \in \mathcal{F}$ **do**
- 7 $s_c = s_g, F_c = F$;
- 8 **foreach** $x_i \in X_u$ **do**
- 9 $Y_i = \{y \in Y \mid y \text{ maps to } x_i\}$;
- 10 $X_i = \text{norm}(1 : |Y_i|, \min = -t, \max = t)$;
- 11 $f_l =$ FITDETERMINISTIC($Y_i \sim X_i, \mathcal{F}_c$);
- 12 $\hat{s} = L(F_c \cup f_l) + L(Y | F_c \cup f_l, X)$;
- 13 **if** $\hat{s} < s_c$ **then** $s_c = \hat{s}, F_c = F_c \cup f_l$;
- 14 **if** $s_c < s$ **then** $s = s_c$;
- 15 **return** s ;

we restrict ourselves to linear, quadratic, cubic, exponential and reciprocal functions—although at the cost of extra computation this class may be expanded arbitrarily. We start by introducing the subroutine of SLOPE that computes the conditional complexity of Y given X .

A. Calculating the Conditional Scores

Algorithm 1 describes the subroutine to calculate the conditional costs $L(Y | X)$ or $L(X | Y)$. We start with fitting a global function f_g for each function class $c \in \mathcal{F}$ and choose the one f_g with the minimum sum of data and model costs (line 2). Next, we add f_g to the model F and store the total costs (3–4). For purely deterministic functions, we are done.

If X includes duplicate values, however, we need to check whether fitting a non-deterministic model leads to a gain in compression. To this end we have to check for each value x_i of X that occurs at least twice, whether we can express the ascendingly ordered corresponding Y values, Y_i , as a function f_l of uniformly distributed data X_i between $[-t, t]$, where t is a user-determined scale parameter (lines 9–12). If the model costs of the new local function f_l are higher than the gain on the data side, we do not add f_l to our model (13). As it is fair to assume that for truly non-deterministic data the generating model for each local component is the same, we hence restrict all local functions to be of the same model class $c \in \mathcal{F}$. As final result, we return the costs according to the model with the smallest total encoded size. In case of deterministic data, this will be the model containing only f_g .

B. Causal Direction and Confidence

In the previous paragraph, we described Algorithm 1, which is the main algorithmic part of SLOPE. Before applying it, we first normalise X and Y to be from the same domain and then determine the data resolutions τ_X and τ_Y for X and Y . To obtain the data resolution, we calculate the smallest

difference between two instances of the corresponding random variable. Next, we apply Algorithm 1 for both directions to obtain $L(Y | X)$ and $L(X | Y)$. Subsequently, we estimate the marginals $L(X)$ and $L(Y)$ based on their data resolutions. This we do by modelling both as a uniform prior with $L(X) = -n \log \tau_X$ and $L(Y) = -n \log \tau_Y$. In the last step, we compute $\hat{\Delta}_{X \rightarrow Y}$ and $\hat{\Delta}_{Y \rightarrow X}$ and report the causal direction as well as the corresponding confidence value \mathbb{C} .

C. Computational Complexity

To assess the computational complexity, we have to consider the score calculation and the fitting of the functional relations. The model costs are computed in linear time according to the number of parameters, whereas the data costs need linear time with regard to the number of data points n . Since we here restrict ourselves to relatively simple functions, we can fit these in time linear to the number of data points. To determine the non-deterministic costs, in the worst case we perform $n/2$ times $|\mathcal{F}|$ fits over two data points, which is still linear. In total the runtime complexity of SLOPE hence is $O(n|\mathcal{F}|)$. In practice, this means that SLOPE is fast, and only takes seconds up to a few minutes depending on the size of the data.

V. RELATED WORK

Causal inference from observational data is an important open problem that has received a lot of attention in recent years [4], [17], [24], [18]. Traditional constraint based approaches, such as conditional independence test, require at least three random variables and can not decide between Markov equivalent causal DAGs [18], [28]. In this work, we focus specifically on those cases where we have to decide between the Markov equivalent DAGs $X \rightarrow Y$ and $Y \rightarrow X$.

A well studied framework to infer the causal direction for the two-variable case relies on the additive noise assumption [25]. Simply put, it makes the strong assumption that Y is generated as a function of X plus additive noise, $Y = f(X) + N$, with $X \perp\!\!\!\perp N$. It can then be shown that while such a function is admissible in the causal direction, this is not possible in the anti-causal direction. There exist many approaches based on this framework that try to exploit linear [25] or non-linear functions [8] and can be applied to real valued [25], [8], [31], [20] as well as discrete data [19]. Recently, Mooij et al. [17] reviewed several ANM-based approaches from which ANM-pHSIC, a method employing the Hilbert-Schmidt Independence Criterion (HSIC) to test for independence, performed best. For ANMs the confidence value is often expressed as the negative logarithm of the p-value from the used independence test [17]. P-values are, however, quite sensitive to the data size [1], which leads to a less reliable confidence value. As we will show in the experiments, our score is robust and nearly unaffected by the data size.

Another class of methods rely on the postulate that if $X \rightarrow Y$ the marginal distribution of the cause $P(X)$ and the conditional distribution of the effect given the cause $P(Y | X)$ are independent of each other. The same does not hold for the opposite direction [9]. The authors of IGCI define

this independence via orthogonality in the information space. Practically, they define their score using the entropies of X and Y [10]. Liu and Chan implemented this framework by calculating the distance correlation for discrete data between $P(X)$ and $P(Y | X)$ [14]. A third approach based on this postulate is CURE [24]. Here, the main idea is to estimate the conditional using unsupervised inverse Gaussian process regression on the corresponding marginal and compare the result to the supervised estimation. If the supervised and unsupervised estimation for $P(X | Y)$ deviate less than those for $P(Y | X)$, an independence of $P(X | Y)$ and $P(X)$ is assumed and causal direction $X \rightarrow Y$ is inferred. Although well formulated in theory, the proposed framework is only solvable for data of up to 200 data points and otherwise relies strongly on finding a good sample of the data.

Recently, Janzing and Schölkopf postulated that if $X \rightarrow Y$, the complexity of the description of the joint distribution in terms of Kolmogorov complexity, $K(P(X, Y))$, will be shorter when first describing the distribution of the cause $K(P(X))$ and then describing the distribution of the effect given the cause $K(P(Y | X))$ than vice versa [9], [12]. To the best of our knowledge, Mooij et al. [16] were the first to propose a practical instantiation of this framework based on the Minimum Message Length principle (MML) [30] using Bayesian priors. Vreeken [29] proposed to approximate the Kolmogorov complexity for numeric data using the cumulative residual entropy, and gave an instantiation for multivariate continuous-valued data. Perhaps most related to SLOPE is ORIGO [4], which uses MDL to infer causal direction on binary data, whereas we focus on univariate numeric data.

VI. EXPERIMENTS

In this section we empirically evaluate SLOPE. In particular, we consider synthetic data, a benchmark data set, and a real-world case study. We implemented SLOPE in *R* and make both the code, the data generators, and real world data publicly available for research purposes.¹ We compare SLOPE to the state of the art for univariate causal inference. These include CURE [24], IGCI [10] and RESIT [20]. From the class of ANM-based methods we compare to ANM-pHSIC [8], [17], which a recent survey identified as the most reliable ANM inference method [17]. We use the implementations by the authors, sticking to the recommended parameter settings.

To run SLOPE, we have to define the parameter t , which is used to normalise the data X_i within a local component, on which the data Y_i is fitted. Generally, the exact value of t is not important for the algorithm, since it only defines the domain of the data points X_i , which can be compensated by the parameters of the fitted function. In our experiments, we use $t = 5$ and set the precision p for the parameters to three.

A. Synthetic Data

We first consider data with known ground truth. To generate such data, we follow the standard scheme of Hoyer et al. [8].

¹<http://eda.mmci.uni-saarland.de/slope/>




Figure 3: [Higher is better] Accuracies of SLOPE, RESIT and IGCI on synthetic data. The first letter of the labels corresponds to the distribution of X (u : uniform, g : sub-Gaussian, b : binomial and p : Poisson), the second letter to that of the noise (u : uniform, g : Gaussian and n : non-additive).

That is, we first generate X randomly according to a given distribution, and then generate Y as $Y = f(X) + N$, where f is a function that can be linear, cubic or reciprocal, and N is the noise term, which can either be additive or non-additive.

Accuracy: First, we evaluate the performance of SLOPE under different distributions. Following the scheme above, we generate X randomly from either

- 1) a uniform distribution with $\min = -t$ and $\max = t$, where $t \sim \text{unif}(1, 10)$,
- 2) a sub-Gaussian distribution by sampling data with $\mathcal{N}(0, s)$, where $s \sim \text{unif}(1, 10)$ and taking each value to the power of 0.7 maintaining its sign,²
- 3) a binomial distribution with $p \sim \text{unif}(0.1, 0.9)$ and the number of trials $t \sim \lceil \text{unif}(1, 10) \rceil$, or
- 4) a Poisson distribution with $\lambda \sim \text{unif}(1, 10)$.

Note that the binomial and Poisson distribution generate discrete data points, which with high probability results in non-deterministic pairs. To generate Y we first apply either a linear, cubic or reciprocal function on X , with fixed parameters, and add either additive noise using a uniform or Gaussian distribution with $t, s \sim \text{unif}(1, \max(x)/2)$ or non-additive noise with $\mathcal{N}(0, 1)|\sin(2\pi\nu X)| + \mathcal{N}(0, 1)|\sin(2\pi(10\nu)X)|/4$ according to [24], where we choose $\nu \sim \text{unif}(0.25, 1.1)$. For every combination we generate 100 data sets of 1000 samples each.

Next, we apply SLOPE, RESIT, and IGCI and record how many pairs they correctly infer. As they take up to hours to process a single pair, we do not consider CURE and ANM here. We give the averaged results over all three function types in Fig. 3. In general, we find that SLOPE and IGCI perform on par and reach 100% for most setups, whereas SLOPE performs better on the sub-Gaussian data. If we consider the single results for linear, cubic and reciprocal, we find that on the linear data with sub-Gaussian distributed X , SLOPE performs on average 7% better than IGCI.

Confidence: Second, we investigate the dependency of the RESIT, IGCI, and SLOPE scores on the size of the data. In an

²We consider sub-Gaussian distributions since linear functions with both X and N Gaussian distributed are not identifiable by ANMs [8].

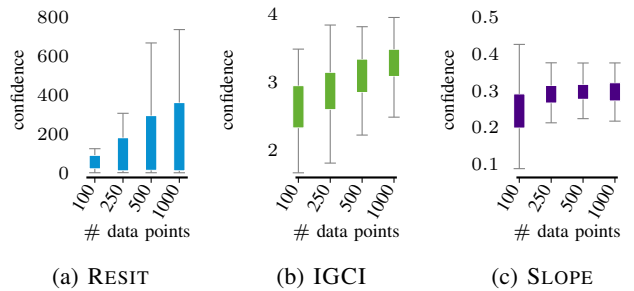


Figure 4: [The more stable the better] Confidence values on a cubic function for different sample sizes. Unlike RESIT and IGCI, the SLOPE scores can be meaningfully compared between different sample sizes.

ideal world, a confidence score is not affected by the size of the data, as this allows easy comparison and ranking of scores.

To analyse this, we generate 100 datasets of 100, 250, 500 and 1000 samples each, where X is Gaussian distributed and Y is a cubic function of X with uniform noise. Subsequently, we apply RESIT, IGCI and SLOPE and record their confidence values. We show the results per sample size in Fig. 4. As each method uses a different score, the scale of the Y-axis is not important. What is important to note, is the trend of the scores over different sample sizes. We see SLOPE has very consistent confidence values that are independent of the number of samples, in strong contrast to RESIT and IGCI. For RESIT the standard deviation of the confidence values grows with the sample size, while for IGCI, we observe that the average confidence increases with the number of samples. In other words, while it is easy to compare and rank SLOPE scores, this is not the case for the two others—which, as we will see below results in comparatively bad decision rates.

Non-Determinacy: Local regression on non-deterministic data adds to the modelling power of SLOPE, yet, it may also lead to overfitting. Here we evaluate whether MDL protects us from picking up spurious structure.

To control non-determinacy, we sample X uniformly from k equidistant values over $[0, 1]$, i.e., $X \in [\frac{0}{k}, \frac{1}{k}, \dots, \frac{k}{k}]$. To obtain Y , we apply a linear function and additive Gaussian noise as above. Per data set we sample 1000 data points.

In Fig. 5 we plot the non-determinism of the model, i.e. the average number of used bins divided by the average number of bins SLOPE could have used, against the number of distinct X values. As a reference, we also include the average number of values of Y per value of X . We see that for at least 75 unique values, SLOPE does not infer non-deterministic models. Only at 40 distinct values, i.e., an average of 25 duplicates per X , SLOPE consistently starts to fit non-deterministic models. This shows that if anything, rather than being prone to overfit, SLOPE is conservative in using local models.




Figure 5: [SLOPE does not overfit] Percentage of non-deterministic models SLOPE chooses, resp. the expected number of Y values per X value, for the number of unique values of X .

B. Real World Data

Next we evaluate SLOPE on real world benchmark data. In particular, we consider the Tuebingen cause-effect data set.³ At the time of writing the data set included 98 univariate numeric cause effect pairs. We compare SLOPE to IGCI, RESIT, ANM, and CURE, using the suggested parameter settings for this benchmark. In addition, we include SLOPED in our comparison, which is an ablated version of SLOPE that only fits a single global deterministic function.

Decision Rate and Accuracy: We first consider the accuracy and decision rates over the benchmark data. To determine the decision rate for an approach, we order the results according to its confidence values. We then determine, from $k = 1$ to 98, the accuracy over the top- k ranked pairs. To determine the accuracy at k we weigh each pair according to its specification in the database. In case an algorithm does not decide, we weigh this result as one half of the corresponding weight.

We plot the results in Fig. 6, where in addition we show the 95% confidence interval for the binomial distribution with success probability of 0.5 in gray. We observe that SLOPE strongly outperforms its competitors in both decision rate and overall accuracy; it identifies the correct result for top-ranked 32 data sets, over the top-74 pairs (which correspond to 73.8% of the weights) it has an accuracy of 90%, while when overall it obtains an accuracy of 82.4%.

In Fig. 7 we show the corresponding confidence values of SLOPE for the benchmark pairs. The plot emphasises not only the predictive power of SLOPE, but also the strong correlation between confidence value and accuracy. In comparison to the other approaches the decision rate (Fig. 6) of SLOPE is stable and only decreases slightly at the very end. Moreover, our competitors obtain much worse overall accuracies, between 56% (CURE) and 71% (RESIT), which for the most part are insignificant.

We identify three reasons why SLOPE performs better than other approaches based on functional learning. First, we consider also piecewise regression, which leads to an improvement of 10% for SLOPE compared to SLOPED. Second, we consider the complexity of the functions, and incorporate the model cost to prevent overfitting; as they are so powerful, methods




Figure 6: [Higher is better] Decision rates of SLOPE, CURE, RESIT, IGCI and ANM on the Tuebingen benchmark data set (98 pairs). SLOPED is an ablated version of SLOPE, which fits the data with a single deterministic function (more detailed version of Fig. 1).

based on Gaussian process learning, for example, are likely to overfit. Last, the calculation of our confidence value is less dependent on the data size.

If we not only consider the confidence values, but also our proposed statistical test, we can improve our results even further. After adjusting the p-values using the Benjamini-Hochberg correction [2] to control the false discovery rate (FDR), 83 out of the 98 decisions are significant w.r.t. $\alpha = 0.001$. As shown in Fig. 7 the pairs rated as insignificant correspond to small confidence values. In addition, from the 15 insignificant pairs, 9 were inferred incorrect from SLOPE and 6 correct. Over the significant pairs the weighted accuracy increases to 84.9%.

To provide further evidence that the confidence values and the p-values are indeed related, we plot the adjusted p-values and confidence values in Fig. 8a. We observe that high confidence values correspond to highly significant p-values. We also computed the decision rate for SLOPE when ranking by p-values, and find it is only slightly worse than that ranked by confidence. We posit that confidence works better as it is more independent of the data size. To test this, we calculate the correlation between data size and corresponding measures using the maximal information coefficient (MIC) [21]. We find a medium correlation between confidence and p-values (0.64), and between p-values and data size (0.55), and only a weak correlation between confidence and data size (0.31).

To show that the strength of our method comes from its ability to fit non-deterministic models, we also consider SLOPED, an ablated version of SLOPE that only considers deterministic models, i.e., only considers models that consist of a single global regression function. As it uses the same score, we see in Fig. 6 that its decision rate is still good. Although the overall accuracy drops by 10% to 72.1%, it is still as good as the best competing method. This also shows that the full approach is much better able to deal with a broader spectrum of cause effect pairs.

Apart from the accuracies, we also tracked which functional dependencies SLOPE found on the benchmark data. We found that most of the time (54.6%), it fits linear functions. For 23.7% of the data it fits exponential models, and for 15.5%

³<https://webdav.tuebingen.mpg.de/cause-effect/>




Figure 7: Confidence values of SLOPE for the Tuebingen benchmark pairs, in descending order, corresponding to Fig. 6. Correct inferences marked in green, errors in red, and inferences insignificant at $\alpha = 0.001$ marked with a grey arrow.




Figure 8: (left) Confidence and significance of SLOPE on the Tuebingen benchmark pairs. Only samples with low confidence are also insignificant. (right) Runtime in seconds over all 98 pairs, in log-scale. SLOPE is more accurate than all, and faster than all except for IGCI.

cubic models. Quadratic and reciprocal models are rarely fitted (6.2%). This shows that working with a richer class of functions allows us to pick up more structure. In addition, we observe that although we allow to fit complex models, in many cases a simple model is preferred since it has sufficient explanatory power at lower model costs.

C. Case Study: Octet Binary Semi Conductors

To evaluate real-world performance we conduct a case study on octet binary semi-conductors [5], [27]. In essence, these are 82 different materials for which we have data on the radii of electron orbitals, and important other energy quantities that together determine their chemical properties. The target of interest is the energy difference δ_E between rocksalt and zincblende crystal structures. It is still an open challenge to find that combination of physical properties that can fully explain δ_E [5], [6]. For causal inference, it suffices to know that the energy difference is influenced by the descriptor variables and not the other way round.

As there is no known combination of physical properties that fully describes δ_E , we mine the top 10 features that each explain as much of δ_E as possible [15], and therewith obtain 10 cause effect pairs where we set δ_E as X and one of the mined features as Y . After consulting with domain experts, we assume $Y \rightarrow X$ as ground truth for all pairs.

We find that SLOPE infers the correct direction for 9 out of 10 pairs. The only error is also the only insignificant score ($p = 0.199$) at $\alpha = 0.001$. In comparison, we find that CURE infers all pairs correctly, whereas IGCI makes the same decisions as SLOPE. RESIT and ANM, on the other hand, only get 4 resp. 5 pairs correct.

D. Runtime

Last, we evaluate the computational efficiency of SLOPE. To this end we report, per method, the wall-clock time needed to decide on all 98 pairs of the benchmark data set. We ran these experiments on Linux servers with two six-core Intel Xenon E5-2643v2 processors and 64GB RAM. The implementations of SLOPE, IGCI and RESIT are single-threaded, whereas ANM and CURE are implemented in Matlab and use the default number of threads. We give the results in Fig. 8. We see that IGCI is fastest, followed by SLOPE, which takes 1 475 seconds to processes all pairs. The other competitors are all at least one order of magnitude slower. Taking 13 days, CURE has the longest runtime. The large gain in runtime of SLOPE compared to RESIT, ANM and CURE rises from the fact that those methods employ Gaussian process regression to fit the functions.

VII. DISCUSSION

The experiments clearly show that SLOPE works very well. It performs well in a wide range of settings, both on synthetic and real world data. In particular on the latter it outperforms the state of the art, obtaining highly stable decision rates and an overall accuracy of more than 10% better than the state of the art. Our case study showed it makes sensible decisions. Most importantly, SLOPE is simple and elegant. Its models are easy to interpret, it comes with a stable confidence score, a natural statistical test, and is computationally efficient.

The core idea of SLOPE is to decide for the causal direction by the simplest, best fitting regression function. To deal with non-deterministic data, we allow our model to additionally use local regression functions for non-deterministic values of X , which the experiments show leads to a large increase in performance. Importantly, we employ local regression within an MDL framework; without this, fitting local regressors would not make sense, as it would lead to strong overfitting.

A key advantage of our MDL-based instantiation of the algorithmic Markov condition, compared to HSIC-based independence tests and IGCI, is that our score is not dependent on the size of the data. This makes it possible to meaningfully compare results among different tests; this is clearly reflected in the stable decision rates. Another advantage is that it allows us to define a natural statistical test based on compression rates, which allows us to avoid insignificant inferences.

Although the performance of SLOPE is impressive, there is always room for improvement. With regard to confidence and significance, we are highly interested in investigating whether we can define a test that directly infers the significance of a confidence score (without resorting to permutation testing).

In addition, it is possible to improve the search for local components by considering alternate re-distributions of X' , apart from uniformly ascending values. This is not trivial, as there exist $n!$ possible orders, and it is not immediately clear how to efficiently optimise regression fit over this space. More obviously, there is room to expand the set of function classes that we use at the moment—kernel based, or Gaussian-process based regression are powerful methods that, at the expense of computation, will likely improve performance further.

For future work, we additionally aim to consider the detection of confounding variables—an open problem that we believe our information theoretic framework naturally lends itself to—as well as to extend SLOPE to multivariate and possibly mixed-type data. We are perhaps most enthusiastic about leveraging the high accuracy of SLOPE towards inferring causal networks from biological processes without the need of conditional independence tests.

VIII. CONCLUSION

We studied the problem of inferring the causal direction between two univariate numeric random variables X and Y . To model the causal dependencies we proposed an MDL-based framework employing local and global regression. Further, we proposed SLOPE, an efficient linear-time algorithm, to instantiate this framework. In addition, we introduced 10 new cause effect pairs from a material science data set.

Empirical evaluations on synthetic and real world data show that SLOPE reliably infers the correct causal direction with an high accuracy. On benchmark data, at over 82% accuracy SLOPE outperforms the state of the art by more than 10%, provides a more robust decision rate, while additionally also being computationally more efficient. In future research, we plan to refine our statistical test, consider detecting confounding, causal inference on multivariate setting, and use SLOPE to infer causal networks directly from data.

ACKNOWLEDGMENT

The authors wish to thank Panagiotis Mandros and Mario Boley for help with the octet binary causal pairs. Alexander Marx is supported by the International Max Planck Research School for Computer Science (IMPRS-CS). Both authors are supported by the Cluster of Excellence Multimodal Computing and Interaction within the Excellence Initiative of the German Federal Government.

REFERENCES

[1] D. Anderson, K. Burnham, and W. Thompson, “Null Hypothesis Testing: Problems, Prevalence, and an Alternative,” *J. Wildl. Manage.*, vol. 64, no. 4, pp. 912–923, 2000.

[2] Y. Benjamini and Y. Hochberg, “Controlling the false discovery rate: a practical and powerful approach to multiple testing,” *J. R. Statist. Soc. B*, vol. 57(1), pp. 289–300, 1995.

[3] P. Bloem and S. de Rooij, “Large-scale network motif learning with compression,” *CoRR*, vol. abs/1701.02026, 2017.

[4] K. Budhathoki and J. Vreeken, “Causal inference by compression,” in *ICDM*. IEEE, 2016, pp. 41–50.

[5] L. M. Ghiringhelli, J. Vybiral, S. V. Levchenko, C. Draxl, and M. Scheffler, “Big data of materials science: Critical role of the descriptor,” *PRL*, vol. 114, no. 10, pp. 1–5, 2015.

[6] B. R. Goldsmith, M. Boley, J. Vreeken, M. Scheffler, and L. M. Ghiringhelli, “Uncovering structure-property relationships of materials by subgroup discovery,” *New J. Phys.*, vol. 19, no. 1, p. 013031, 2017.

[7] P. Grünwald, *The Minimum Description Length Principle*. MIT Press, 2007.

[8] P. Hoyer, D. Janzing, J. Mooij, J. Peters, and B. Schölkopf, “Nonlinear causal discovery with additive noise models,” in *NIPS*, 2009, pp. 689–696.

[9] D. Janzing and B. Schölkopf, “Causal inference using the algorithmic markov condition,” *IEEE TIT*, vol. 56, no. 10, pp. 5168–5194, 2010.

[10] D. Janzing, J. Mooij, K. Zhang, J. Lemeire, J. Zscheischler, P. Daniušis, B. Stuedel, and B. Schölkopf, “Information-geometric approach to inferring causal directions,” *AIJ*, vol. 182-183, pp. 1–31, 2012.

[11] A. Kolmogorov, “Three approaches to the quantitative definition of information,” *Problemy Peredachi Informatsii*, vol. 1, no. 1, pp. 3–11, 1965.

[12] J. Lemeire and E. Dirx, “Causal models as minimal descriptions of multivariate systems,” 2006.

[13] M. Li and P. Vitányi, *An Introduction to Kolmogorov Complexity and its Applications*. Springer, 1993.

[14] F. Liu and L. Chan, “Causal inference on discrete data via estimating distance correlations,” *Neur. Comp.*, vol. 28, no. 5, pp. 801–814, 2016.

[15] P. Mandros, M. Boley, and J. Vreeken, “Discovering Reliable Approximate Functional Dependencies,” in *KDD*. ACM, 2017, pp. 355–364.

[16] J. Mooij, O. Stegle, D. Janzing, K. Zhang, and B. Schölkopf, “Probabilistic latent variable models for distinguishing between cause and effect,” *NIPS*, no. 26, pp. 1–9, 2010.

[17] J. M. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf, “Distinguishing cause from effect using observational data: Methods and benchmarks,” *JMLR*, vol. 17, no. 32, pp. 1–102, 2016.

[18] J. Pearl, *Causality: Models, Reasoning and Inference*, 2nd ed. New York, NY, USA: Cambridge University Press, 2009.

[19] J. Peters, D. Janzing, and B. Schölkopf, “Identifying cause and effect on discrete data using additive noise models,” in *AISTATS*. JMLR, 2010, pp. 597–604.

[20] J. Peters, J. Mooij, D. Janzing, and B. Schölkopf, “Causal discovery with continuous additive noise models,” *JMLR*, vol. 15, pp. 2009–2053, 2014.

[21] D. N. Reshef, Y. A. Reshef, H. K. Finucane, S. R. Grossman, G. McVean, P. J. Turnbaugh, E. S. Lander, M. Mitzenmacher, and P. C. Sabeti, “Detecting novel associations in large data sets,” *Science*, vol. 334, no. 6062, pp. 1518–1524, 2011.

[22] J. Rissanen, “Modeling by shortest data description,” *Automatica*, vol. 14, no. 1, pp. 465–471, 1978.

[23] —, “A universal prior for integers and estimation by minimum description length,” *Annals Stat.*, vol. 11, no. 2, pp. 416–431, 1983.

[24] E. Sgouritsa, D. Janzing, P. Hennig, and B. Schoelkopf, “Inference of Cause and Effect with Unsupervised Inverse Regression,” *AISTATS*, vol. 38, pp. 847–855, 2015.

[25] S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen, “A linear non-gaussian acyclic model for causal discovery,” *JMLR*, vol. 7, pp. 2003–2030, 2006.

[26] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search*. MIT press, 2000.

[27] J. A. Van Vechten, “Quantum dielectric theory of electronegativity in covalent systems. i. electronic dielectric constant,” *Physical Review*, vol. 182, no. 3, p. 891, 1969.

[28] T. Verma and J. Pearl, “Equivalence and synthesis of causal models,” in *UAI*, 1991, pp. 255–270.

[29] J. Vreeken, “Causal inference by direction of information,” in *SDM*. SIAM, 2015, pp. 909–917.

[30] C. S. Wallace and D. M. Boulton, “An information measure for classification,” *Comput. J.*, vol. 11, no. 1, pp. 185–194, 1968.

[31] K. Zhang and A. Hyvärinen, “On the identifiability of the post-nonlinear causal model,” in *UAI*. AUAIU Press, 2009, pp. 647–655.