



A causal direction test for heterogeneous populations

Vahid Partovi Nia^{a,c,*}, Xinlin Li^a, Masoud Asgharian^d, Shoubo Hu^e, Yanhui Geng^b, Zhitang Chen^b

^a Huawei Noah's Ark Lab, Suit 201, 7101 Park avenue, Montreal, Quebec H3N 1X9, Canada

^b Huawei Noah's Ark Lab, Units 525-530, Core Building 2, Hong Kong Science Park, Shatin, Hong Kong

^c Polytechnique Montreal, Department of Mathematics and Industrial Engineering, 2500 Chemin de Polytechnique, Montreal, QC H3T 1J4, Canada

^d McGill University, Department of Mathematics and Statistics, 805 Rue Sherbrooke West, Montreal, QC H3A 2K6, Canada

^e The Chinese University of Hong Kong, Department of Computer Science and Engineering, Ho Sin-Hang Engineering Building, Shatin N.T., Hong Kong

ARTICLE INFO

MSC:

62H30

62F03

Keywords:

Bayesian hierarchical model

Causal inference

Clustering

Graphical models

Belief network

Probabilistic expert systems

Testing statistical hypotheses

ABSTRACT

A probabilistic expert system emulates the decision-making ability of a human expert through a directional graphical model. The first step in building such systems is to understand data generation mechanism. To this end, one may try to decompose a multivariate distribution into product of several conditionals, and evolving a blackbox machine learning predictive models towards transparent cause-and-effect discovery. Most causal models assume a single homogeneous population, an assumption that may fail to hold in many applications. We show that when the homogeneity assumption is violated, causal models developed based on such assumption can fail to identify the correct causal direction. We propose an adjustment to a commonly used causal direction test statistic by using a k -means type clustering algorithm where both the labels and the number of components are estimated from the collected data to adjust the test statistic. Our simulation result show that the proposed adjustment significantly improves the performance of the causal direction test statistic for heterogeneous data. We study large sample behaviour of our proposed test statistic and demonstrate the application of the proposed method using real data.

1. Introduction

Causal inference is one of the most fundamental concepts in learning. Teaching the machine how to find *cause-and-effect* relationship, often using a mathematical model, is as essential as teaching children how to connect the dots. Causality is perhaps as old as human endeavour for learning and has historically developed along with the development of human knowledge in almost any domain of science. Scientific studies tend to draw cause and effect conclusions after observing associations. However, most observed associations translate to causal conclusions only under certain conditions. Early causal problem formulation based on data observation appears in statistics (Neyman, 1923), economics (Heckman, 1976), medicine (Greenland, Pearl, & Robins, 1999), and computer science (Pearl, 1986) among others. Research on causality typically starts with hypothesizing a cause and setting up an experimental study in which association can be translated into causal relationship (Fisher, 1926). In many applications, however, logistic or other constraints may preclude the possibility of conducting experimen-

tal studies. Methods have therefore been developed for cause-and-effect conclusions to be drawn from data collected in observational studies.

With the advent of technology, modern applications often include a large number of variables, many possibly spurious, measured on each subject under study. Identifying causal direction, especially in such high dimensional settings, albeit challenging, but crucial to gain insights into the data generating mechanism and hence data interpretation, and also a deeper understanding of data structure. There are recent efforts in using causal analysis to interpret large dimensional data in decision making (Wu, 2010), in natural language processing (Dehkharghani, Mercan, Javeed, & Saygin, 2014), in transportation (Kayikci & Stix, 2014), and in genetics (Schadt et al., 2005), among others.

In many applications causal conclusion is drawn for multiple variables, through a cause and effect model also known as belief network. Such models decompose a multivariate distribution into several conditional distributions. Many of such decompositions are theoretically equivalent and detecting which one is the actual data generating mechanism often requires domain knowledge of a human expert, see Fig. 1.

* Corresponding author at: Huawei Noah's Ark Lab, Suit 201, 7101 Park avenue, Montreal, Quebec H3N 1X9, Canada.

E-mail addresses: vahid.partovinia@huawei.com (V. Partovi Nia), xinlinli1@huawei.com (X. Li), masoud.asgharian2@mcgill.ca (M. Asgharian), sbhu@cse.cuhk.edu.hk (S. Hu), geng.yanhui@huawei.com (Y. Geng), chenzhitang2@huawei.com (Z. Chen).




Fig. 1. A human expert (top panel) versus an expert system embedded with automatic reasoning (bottom panel) in which cause and effect model attaches data to the knowledge base.

1.1. Preliminaries

The precursor to establishing the cause and effect relationship between several variables is to identify the causal direction between a pair of variables. It transpires that inferring causal direction between a pair of variables, say X and Y , should be closely tied to the conditional distributions of X given Y and of Y given X . Assume (X, Y) follows a complex joint distribution, say $p(x, y)$. Following the independent causal mechanisms principles postulate (Schölkopf et al., 2012), causal direction inference looks for evidence in the observed data to prefer a certain conditional decomposition, either $p(x, y) = p(y | x)p(x)$ or $p(x, y) = p(x | y)p(y)$. In the former decomposition x causes y or $x \rightarrow y$ and in the latter decomposition y causes x or $y \rightarrow x$, see Fig. 2. In theory, both decompositions are valid and therefore inferring a causal direction without further assumptions is ill-defined and *unidentifiable*. Under more structural assumptions, such as the Additive Noise Model (ANM), the direction of decomposition becomes *identifiable*, and observed data can be used to infer the cause and effect relationship.

1.2. Additive noise model

ANM (Hoyer, Janzing, Mooij, Peters, & Schölkopf, 2009) represents the effect as a function of the cause with an additive independent noise, i.e. $y = f(x) + \epsilon$, in which f is a nonlinear deterministic smooth function and $\epsilon \sim p(\epsilon)$ is an independent noise. There is no backward model of the form $x = g(y) + \epsilon$ that admits an ANM in the anti-causal direction unless the backward noise depend on y . This shows a causal direction can be examined based on testing $\epsilon \perp\!\!\!\perp x$. One may first fit a smooth regression model to predict $\hat{y} = \hat{f}(x)$, predict the noise $\hat{\epsilon} = y - \hat{y}$, and then test if the residuals $\hat{\epsilon}$ is independent of the predictor x , $\hat{\epsilon} \perp\!\!\!\perp x$ (Hoyer et al., 2009). This test of independence can be used as an evidence for inferring the causal direction for a pair of variables empirically.

Most causal inference approaches assume a single causal model for the observed data (Janzing & Schölkopf, 2010; Shimizu, Hoyer, Hyvärinen, & Kerminen, 2006; Zhang & Hyvärinen, 2009) so they are suitable for homogeneous data. In many applications, however, data are collected from several different sources. Due to the unknown data generation process and variability in the data source, sampling scheme, sampling conditions, etc there is no guarantee on the viability of such homogeneity assumptions in practice. Naive use of the existing

clustering algorithms misleads causal direction inference. Moreover, each sub cluster may declare its own causal direction to confuse the ultimate judgment. There has been little research on the effect of heterogeneity in causal direction test. However, there some works that consider population heterogeneity in causal inference, but all of them mostly consider the causal effect rather than causal direction test (Brand & Thomas, 2013; Hong & Raudenbush, 2013; Xie, 2011, 2013).

When observations are generated from a non-homogeneous population that is comprised of several homogeneous sub-populations, the number of homogeneous sub-populations affects the causal direction test statistic performance to a great extent. The distribution of the causal direction test statistic needs to be adjusted when the homogeneity assumption fails to hold.

1.3. Causal direction test

The test statistic for causal direction relies on a measure of dependence between the error and the predictor. Measuring dependence in heterogeneous populations requires careful considerations as the following simple example shows in Fig. 3. Marginal correlation as a measure of dependence which ignores the cluster labels, may contradict the inter-cluster correlation which uses the cluster labels if data are generated from a heterogeneous population. The marginal correlation becomes significantly positive if cluster centres are aligned about the line $y = x$, and is significantly negative if centres are aligned about $y = -x$, see Fig. 3 left panel. Consequently, building a test statistic that ignores cluster labels may lead to misleading judgments and affects the test performance, and type I error probability, see Fig. 3 right panel.

What depicted in Fig. 3 is essentially Simpson's paradox (Simpson, 1951) presented in terms of correlation rather than conditional probability (Rücker & Schumacher, 2008). The intimate tie between Simpson's paradox and causal inference has been well documented by Pearl (2009). This is the main motivation to correct the test of causal direction for heterogeneous data.

Correcting the test statistic requires estimating the cluster labels as well as the number of cluster components. Estimating the number of homogeneous sub-populations is a challenging task. In additive noise models the observed number of cluster components is heavily affected by the range of the observed input x . Fig. 4 provides a visual intuition, where the identified number of cluster components heavily depends on the observed range of x . If x is observed about 0.7, a single component emerges. However, observing x about 0.5 or 1.0 changes the number of observed components to 2 or 3 respectively. To this end, we propose a clustering algorithm that labels and the number of components are estimated from the collected data. The estimated labels are used to adjust the causal direction test statistic.

Recently, Hu, Chen, Partovi Nia, Chan, and Geng (2018) and Liu and Chan (2016) proposed inferring the causal direction on ANMs for discrete and continuous variables respectively. Here we focus on continuous variables. Hu et al. (2018) proposed using the k -means algorithm on causal parameters and used a predetermined number of sub-populations to overcome data heterogeneity. To the best our knowledge, there is nothing in the literature on correcting the causal direction test statistic after clustering. We speculate that this gap is due to a possible misconception, that is clustering has little effect on causal direction test. We show that this is simply wrong, and the cluster labels play a crucial role in causal direction inference.

It is evident that the clustering phase is unjustifiable if it does not help inferring the causal direction. Our work builds on Hu et al. (2018) and extends it in two directions: (i) provides a clustering method with imprecise number of cluster components. (ii) uses the clustering information to adjust the test statistic and re-examine the causal direction using clustering labels.




Fig. 2. Examples of simple cause and effect models for three variables. When y causes x , i.e. $p(x, y, z) = p(x | y)p(y)p(z)$, left panel. When x causes y , i.e. $p(x, y, z) = p(y | x)p(x)p(z)$, middle panel. When y and z cause x , i.e. $p(x, y, z) = p(x | y, z)p(y)p(z)$, right panel.




Fig. 3. Left panel: marginal correlation which ignore cluster labels (grey) may contradict inter-cluster correlation which uses clustering labels (blue and red). Right panel: Empirical type I error probability α for testing inter-cluster correlation with zero while ignoring cluster labels deviates from the theoretical value 0.05 as the number of cluster components increases. Data are simulated from a mixture of $1 \leq k \leq 8$ standard Gaussian while cluster centres are aligned along $y = x$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 4. Illustration of from three clusters, $y = x^3 + \epsilon$ (blue), $y = 0.5x + \epsilon$ (orange), $y = 0.8 - x^3 + \epsilon$ (red). The causal model imposes a different number of observed components depending on observed x , calling for a clustering with flexible component size. Test of causal direction suffers from the same problem in heterogeneous data. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

We use a model in which the causal direction of the mixture of ANMs is identifiable, and adopt Partially Observable Gaussian Processes Model for estimation (Lawrence, 2005) proposed in Hu et al. (2018).

2. Partially observed additive noise model

We assume if $x \rightarrow y$, the distribution of x and the function f mapping x to y are independent (Janzing & Scholkopf, 2010). We note that f is, for instance, the conditional expectation of Y given $X = x$. The independence can perhaps be best understood in a parametric setting where the joint distribution of (X, Y) is known up to finitely

many unknown parameters. The independence then means that the marginal distribution of X and the conditional distribution of Y given X do not have any common unknown parameter. In other words, as far as describing the relationship between X and Y are concerned, the conditional distribution of Y given X and the joint distribution of (X, Y) have the same information. So much so that the marginal distribution of X does not have any pertinent information about the conditional distribution of Y given X .

We interpret the independence between the cause and mechanism only through the cluster-specific model parameter θ_c that captures all properties of the mapping f , while θ_c is independent of the cause x . We assume the model is identifiable, i.e. if $x \rightarrow y$, there is no backward additive noise model $x = g(y; v_c) + \epsilon$ that satisfies $y \perp\!\!\!\perp (\epsilon, v_c)$. This assumption is not restrictive as Hoyer et al. (2009) proved it for a single population additive noise model, and Hu et al. (2018, Theorem 1) extended this result for a several population model. In other words, if x is independent of θ_c in the causal direction, it is very likely that y and v_c are dependent in the anti-causal direction.

We start the estimation process by projecting a set of centred n dimensional data $\mathbf{x} = [x_1, \dots, x_n]^T$ as the observed cause, and $\mathbf{y} = [y_1, \dots, y_n]^T$, as the observed effect onto d hidden dimensions. The projection problem is formalized as the maximization of the Gaussian log-likelihood

$$\mathcal{L}(\mathbf{K}) = -\frac{dn}{2} \log(2\pi) - \frac{d}{2} \log |\mathbf{K}| - \frac{1}{2} \text{tr}(\mathbf{K}^{-1} \mathbf{y} \mathbf{y}^T),$$

where $\mathbf{K}$ is the covariance matrix $\mathbf{K} = \phi \phi^T + \beta^{-1} \mathbf{I}$, β is a positive scale, $\mathbf{I}$ is the identity matrix. The canonical nonlinear feature map $\phi = [\phi(x_1), \dots, \phi(x_n)]^T$ is computed using the kernel trick. In the kernel representation, features have zero mean and the variance covariance of the $\mathbf{y}$ is defined by the cross products of the kernel features which creates $\mathbf{K}$.

The latent variable θ_i is brought in the additive noise model through a concatenated latent predictor $\tilde{\mathbf{x}}_i^T = [x_i, \theta_i]$ and the Hilbert space is re-defined based on the new vector $\tilde{\mathbf{x}}_i$. Therefore, the latent parameters are estimated by maximizing the Gaussian log-likelihood

$$\mathcal{L}(\theta) = -\frac{(d+1)n}{2} \log(2\pi) - \frac{d+1}{2} \log |\tilde{\mathbf{K}}| - \frac{1}{2} \text{tr}(\tilde{\mathbf{K}}^{-1} \mathbf{y} \mathbf{y}^T), \quad (1)$$

where θ is the vector composed of θ_i 's, $\tilde{\mathbf{K}} = \tilde{\Phi}\tilde{\Phi}^\top$, and $\tilde{\Phi}_{n \times (d+1)} = [\phi(\tilde{\mathbf{x}}_1), \dots, \phi(\tilde{\mathbf{x}}_N)]$. This is an augmented Gaussian log likelihood, for $(\mathbf{x}, \theta)$ which is a new feature space of dimension $d + 1$. The variance covariance matrix $(d + 1) \times (d + 1)$ which is constructed through the cross products of the kernel features on the extended vector $(\mathbf{x}_i, \theta_i)$.

The parameter vector θ appears in $\tilde{\mathbf{K}}$ through $\tilde{\mathbf{x}}$. In our developments we focus on univariate θ_i , but the methodology is general and is valid for multivariate projection as well. This approach re-formalizes the additive model $y = f(x, \theta) + \varepsilon$ in terms of the augmented variable $y = f(\tilde{\mathbf{x}}) + \varepsilon$. However, still the log-likelihood of an ill-defined model, because x and θ should be independent.

The Hilbert–Schmidt independence criterion (HSIC) measures the dependence between observations of a pair of random variables by projecting them onto the reproducing kernel Hilbert space. The empirical HSIC is

$$\text{HSIC} = \frac{1}{n^2} \text{tr}(\mathbf{K}\mathbf{H}\mathbf{L}\mathbf{H}), \quad (2)$$

where $\mathbf{K}$ is the kernel element of $x, k(x_i, x_{i'})$, $\mathbf{L}$ is the kernel element of $y, l(y_i, y_{i'})$, $\mathbf{H} = \mathbf{I} - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$ and $\mathbf{1}$ is the unit vector of size n .

The independence between x and θ is encouraged by adding HSIC as a regularizer to the log-likelihood term through the regularization constant $\lambda > 0$

$$\mathcal{J}(\theta) = \mathcal{L}(\theta) - \lambda \log \text{HSIC}(\theta). \quad (3)$$

Causal parameters are estimated by $\hat{\theta} = \arg \max \mathcal{J}(\theta)$ using scaled conjugate gradient maximization (Hu et al., 2018).

We first estimate the model parameters for each subject $\theta_i, i = 1, \dots, n$, we then cluster θ_i while the number of clusters k is imprecise. This maps θ_i to $\theta_c, c = 1, \dots, k$ while k varies $K - \Delta \leq k \leq K + \Delta$ for a given Δ . The cluster component range Δ has little effect on the performance of the clustering algorithm and is determined by the computational budget.

3. Clustering method

The proposed clustering method works via combining additive noise models (Hoyer et al., 2009) with product partition models (Hartigan, 1990). We provide more details about the intuition behind the clustering method below. This intuition provide some insights about how to estimate the number of cluster components.

The clustered version of the additive noise models is composed of several additive noise models of the same causal direction (Hu et al., 2018). Causal models are developed given data labels which are to be estimated, either mutually or after estimation of causal model parameters. Following Hu et al. (2018) we propose the latter approach since it is computationally less demanding.

A cluster additive noise model is a set of causal models of the same causal direction between two continuous random variables x and y with realizations x_i and $y_i, i = 1, \dots, n$

$$y_i = f(x_i | \theta_c) + \varepsilon_i, \quad (4)$$

where x denotes the cause, y denotes the effect, f is a smooth function, nonlinearity parameters θ_c that parametrize the smooth function f , and ε is the statistical noise.

In model (4) we assume

1. $x_i \sim p(x), \theta_c \sim p(\theta)$ are independently and identically drawn from a Gaussian distribution.
2. The statistical error is independent of the covariates and clustering parameters $\varepsilon_i \perp\!\!\!\perp x_{i'}$ and $\varepsilon_i \perp\!\!\!\perp \theta_c$.
3. The clustering parameters θ_c independently and identically drawn from a Gaussian distribution.

The difference between traditional causal models and cluster causal models is the way that causal parameters θ parametrize the smooth function f . The nonlinearity parameter θ_c is drawn randomly from a




Fig. 5. The generative process of the proposed causal model $f(x | \theta)$ and how it interacts with the clustering model $f(x, y | \theta, z)$ of Eq. (6).

probabilistic model independently. In other words, a set of independent generating mechanisms is assumed for each sub-population through θ_c . Our formulation is slightly different from Hu et al. (2018) that assumes the causal parameters θ_c are drawn from a fixed set, but we assume they are generated from an independent Gaussian distribution. Although this modification seems minor, it plays a major role in attaching the causal model to the clustering algorithm in Fig. 5 and allows to aggregate the causal test statistic through the cluster independence assumption. This model is inspired by commonly encountered situations where the data generating process from one independent trial to another may be different due to the influence of uncontrollable factors.

We first re-formalize probabilistic clustering model through Bayesian regression and derive the clustering algorithm using this model. This viewpoint allows us use the marginal posterior as an estimation tool for the cluster component size.

Assume the following Bayesian regression for the latent parameters

$$\begin{aligned} \theta | \mu &\sim \mathcal{N}(\mathbf{Z}\mu, \sigma^2 \mathbf{I}) \\ \mu &\sim \mathcal{N}(\bar{\theta}, \kappa(\tau\sigma)^2 \mathbf{I}), \end{aligned} \quad (5)$$

where σ^2 is the common within-cluster variance, τ^2 is the between-cluster to within-cluster variance ratio, and $\bar{\theta} = [\bar{\theta}_1, \dots, \bar{\theta}_k]$ is the vector of the cluster averages. The over-dispersion parameter $\kappa > 1$ controls the prior information, i.e. a large κ value gives a flat prior with minimal prior information about the parameters. This model is a sort of *empirical Bayes* in which the data statistic is utilized to parametrize the prior.

We adopt the *product partition* model (Hartigan, 1990) for clustering, i.e.

$$f(\mathbf{x}, \mathbf{y} | \theta, \mathbf{z}) = \prod_{c=1}^k \prod_{\{i|z_i=c\}} f(x_i, y_i | \theta_i) \quad (6)$$

The clustering mechanism adds an unobserved label z_i to each observation, i.e. (x_i, y_i, z_i) or equivalently (θ_i, z_i) in which $z_i \in \{1, 2, \dots, k\}$ is the label and k is the uncertain number of sub-populations. The clustering method relies on $\theta_i \in \mathbb{R}$ which is the key to distinguish between different generating mechanisms. Note that for an identifiable mapping f_θ , one can directly cluster generating mechanisms by clustering θ_i 's. Therefore, θ_c in (6) is equivalent to the pair $(\theta_i, z_i = c)$.

A practical causal cluster model should uncover the unknown number of cluster components k as well as the unobserved label z_i . We

therefore focus on devising an algorithm that allows for clustering θ_i with a flexible component size $k = \max(z_i) \in \{K - \Delta, \dots, K + \Delta\}$, given positive integers K and Δ , $K > \Delta + 1$. The algorithm looks like a simple extension of k -means, but the inspiration comes from a probabilistic clustering that satisfies certain conditions to guarantee convergence.

1. Initialization: Set K, Δ , initialize $\mathbf{z}$.
2. Run $2\Delta + 1$ clustering chains in parallel $k \in \{K - \Delta, \dots, K + \Delta\}$
3. For each chain of size k

- (3.1) centre update: $\mu_c = \bar{\theta}_c$
- (3.2) label update: $z_i = \arg \min_c |\theta_i - \mu_c|$.

4. Within-cluster variance computation:

$$\sigma^2 = \frac{1}{N} \sum_{c=1}^k \sum_{\{i|z_i=c\}}^n (\theta_i - \bar{\theta}_c)^2.$$

5. Between-to-within variance ratio computation:

$$\tau^2 = \frac{1}{k\sigma^2} \sum_{c=1}^k (\bar{\theta}_c - \bar{\theta})^2$$

6. Component size estimation: $\hat{k} = \arg \max_k \ell(\mathbf{z} \mid \theta, k)$.

We iterate between 3.1 and 3.2 until convergence and ultimately at step 5 report the labels $\mathbf{z}$ with k that maximize

$$\ell(\mathbf{z} \mid \theta, k) = -\frac{n}{2} \log 2\pi\sigma^2 - \frac{1}{2\sigma^2} \sum_{c=1}^k \sum_{\{i|z_i=c\}} (\theta_i - \bar{\theta}_c)^2 - \frac{1}{2} \sum_{c=1}^k \log(\kappa\tau^2 n_c + 1), \quad (7)$$

see Fig. 6 that visualizes the algorithm in a condensed flowchart. Eq. (7) is derived by marginalizing over the model parameters

$$\ell(\mathbf{z} \mid \theta, k) = \log \int \dots \int f(\mathbf{y} \mid \theta, \mathbf{z}) f(\theta \mid \mathbf{z}) d\theta + \log p(\mathbf{z}),$$

The multi dimensional integral is a Gaussian integral which turns out to be analytically tractable and can be simplified to (7). We assumed a constant prior for labels $p(\mathbf{z})$ but there are some other choices like multinomial Dirichlet (Partovi Nia & Davison, 2015).

The clustering algorithm resembles the k -means to a great extent and only adds a few more steps to estimate the cluster component size using the marginal log likelihood ℓ . The clustering hyperparameter κ is a sort of over-dispersion of cluster centres. Our experiments show $\kappa = e^{kn/4}$ is a good choice. The computational complexity of the clustering algorithm is $\mathcal{O}(N\Delta K)$. The following result (Theorem 1) shows that a stochastic version of the proposed clustering algorithm converges to a stationary distribution independent of K . However, K affects the search space and the convergence rate of the algorithm may differ in practice for different values of K .

Theorem 1. Suppose σ^2 and τ^2 are given.

1. cluster centre update: sample from a Gaussian distribution with mean

$$\mu_c \sim \mathcal{N}\left(\bar{\theta}_c \left(1 + \frac{1}{n_c \kappa \tau^2}\right)^{-1}, \sigma^2 \left(n_c + \frac{1}{\kappa \tau^2}\right)^{-1}\right).$$

2. cluster label update: sample from Multinomial distribution with probability mass

$$z_i \sim \Pr(z_i = c) = \frac{\phi\left(\frac{\theta_i - \mu_c}{\sigma}\right)}{\sum_{c=1}^k \phi\left(\frac{\theta_i - \mu_c}{\sigma}\right)}, c = 1, \dots, k$$

in which $\phi(\cdot)$ is a standard Gaussian density.

3. cluster component update: sample from Multinomial distribution with probabilities proportional to $\exp\{\ell(\mathbf{z} \mid \theta, k)\}$ of (7).




Fig. 6. The condensed flow chart of the proposed clustering algorithm while k -means is used as a subroutine.

Then the proposed clustering algorithm converges to the stationary distribution

$$p(\mathbf{z} \mid \theta) \propto \sum_{k=K-\Delta}^{K+\Delta} \exp\{\ell(\mathbf{z} \mid \theta, k)\}$$

See Appendix for the proof.

This probabilistic clustering method resembles k -means and cluster size estimation resembles BIC scoring (Schwarz et al., 1978). There has been various attempts to use BIC for cluster component selection. Several authors including (Pelleg, Moore, et al., 2000) studied k -means with BIC scoring and report that BIC over-estimates the number of components. Our scoring is developed for the clustering context, in which k -means matches the setting of Bayesian linear regression: labelling update is equivalent to design matrix estimation, and mean update is equivalent to the coefficient estimation. Our formulation shows that the conventional BIC scoring is inappropriate for k -means and requires proper generalization. It also shows the marginal posterior scoring with ℓ becomes BIC scoring of Pelleg et al. (2000) if clusters are balanced $n_c = n_{c'}, c \neq c'$, and $\kappa\tau = 1$.

4. Test statistic adjustment

In additive noise causal models, HSIC is proposed to draw conclusions about the causal direction (Hu et al., 2018). This test statistic is, however, designed for situations in which samples are coming from a homogeneous population, and is highly sensitive to departures from this assumption. Theorem 2 presents the asymptotic distribution of an adjusted empirical HSIC to infer the causal direction in a heterogeneous case for a given set of clustered data.

Theorem 2. Define $\text{HSIC}_c = \frac{1}{n_c} \text{tr}(\mathbf{K}_c \mathbf{H}_c \mathbf{L}_c \mathbf{H}_c)$ to be the cluster-specific empirical HSIC statistic. The aggregated test statistic $t = \sum_{c=1}^k n_c \text{HSIC}_c$ converges in distribution to $\sum_{l=1}^{\infty} \lambda_l z_l^2$, where $z_l, l = 1, 2, \dots$ are independent standard Gaussian random variables and $\lambda_l, l = 1, 2, \dots$ are non-negative constants.

The result of Theorem 2 reduces to the homogeneous case of Gretton et al. (2005) if $k = 1$. The theoretical quantile of the test statistic t




Fig. 7. Top panel: the UN life expectancy dataset is clustered to $k = 5$ using our clustering algorithm. Bottom panel: the density plot of the causal parameters θ_i .

can be calculated using the Gamma basis $\alpha = \frac{\mu_c^2}{\sigma_c^2}$ and $\beta = \frac{\sigma_c^2}{\mu_c}$ with $\mu = \sum_{c=1}^k n_c \mu_c$, $\sigma^2 = \sum_{c=1}^k n_c \sigma_c^2$ (Wood, Booth, & Butler, 1993), in which μ_c and σ_c^2 are the cluster-specific theoretical HSIC mean and variance. Computing the test statistic t is straightforward by clustering observations first, using labels to isolate data in the same cluster, calculate HSIC for each cluster, and then using a weighted average of HSIC to obtain t .

We performed simulation studies based inspired on real data sets. Our results shows that the main role of HSIC correction in Theorem 2 appear in controlling the type I error probability. Occasionally it improves the total error (type I error + type II error probabilities). Computing the corrected HSIC statistic is straightforward given the clustering labels, i.e. HSIC for each cluster is computed separately and then summed up while weighted by the cluster size.

5. Application

The Tübingen cause-effect pairs (Mooij, Peters, Janzing, Zscheischler, & Schölkopf, 2016) is a well-known benchmark in the context of causal direction detection.¹ The database includes 41 data sets arranged in 108 pairs (x, y) with a known causal direction identified for each pair, either $x \rightarrow y$ or $y \rightarrow x$.

5.1. Life expectancy data

First we explore the effect of the number of clusters on the test statistic for the UN life expectancy data by concatenating pairs 56–63 of Tübingen cause-effect pairs. Fig. 7 (left panel) shows the scatter plot of UN data x : life expectancy versus y : latitude; note that the true causal direction is $y \rightarrow x$.

Data are generated from multiple sources, so we expect that data homogeneity assumption fail to hold. The scatter plot in Fig. 7 confirms this visually.

We compute the causal parameters by maximizing the log likelihood (3) with $\lambda = 50$ given the true causal direction $y \rightarrow x$. The test statistic without adjustment is 14.90 and its theoretical 5% quantile is 0.60, so it mistakenly rejects the null. The statistic after adjustment using clustering labels with $k = 2, 3, 4$ still rejects the true direction but with a larger p -value. This is aligned with our observation in the simulated mixture example in Fig. 3 (left panel), i.e. for a large number of components the type I error is more affected, and we expect to see the effect of test statistic correction specially for large k .

In this example $k = 5$ gives the adjusted test statistic $t = 1.73$ with the 5% theoretical quantile 2.21, so infers the causal direction correctly. The direction is inferred correctly also for $k \geq 6$. Our clustering algorithm estimated 7 clusters.

¹ <https://webdav.tuebingen.mpg.de/cause-effect/>.




Fig. 8. Boxplot of type I error probability of Tübingen cause-effect pairs unadjusted (left) versus adjusted (right). The nominal type I error remains on average around the nominal value 0.05 after adjustment.

5.2. Multiple cause-effect pairs

Next we check the performance of our method on all data pairs as well. Following Hu et al. (2018) we exclude pairs 12, 17, 47, 52, 53, 54, 55, 70, 71, 101, and 105. Additionally we excluded pairs 73, 106, and 68 that include outliers. Data that contain outliers violate the Gaussian error assumption and treating outliers is a common practice in smooth modelling. Outliers trouble our clustering method as the clustering algorithm relies on Gaussian assumption. As a consequence, outliers yielded singleton clusters on data sets 73, 106, and 68 that troubled the computation of the cluster specific test statistic numerically. Overall, too much deviation for the Gaussian distributional assumption may lead to inaccurate causal inference direction test and may require special treatment.

We sample 90 data from each pair and repeat this process 50 times independently. Then we estimate the causal parameters by maximizing the log likelihood (3) with $\lambda = 50$ for $x \rightarrow y$ and $y \rightarrow x$ directions. We choose $k = K \pm 2$ with a visually appealing $2 \leq K \leq 6$ for each data set. We used the clustering labels to adjust the HSIC statistic while running our clustering algorithm. Fig. 8 shows the boxplot of type I error. The total error probability (type I error + type II error) remain equal. Theoretically the type I error probability must remain under control at about the significance level. However, the mean of type I error probability for unadjusted statistic is 0.796 while using the adjusted method it is 0.048. The latter preserved the nominal type error probability 0.05. We see a similar behaviour in the simulated toy example of Fig. 3 in which the type I error probability is exceedingly higher than the nominal value for large number of components k .

5.3. Disease symptom data

We analyse the stroke subset of verbal autopsy survey (Murray et al., 2011) benchmark data available in openVA R package. Stroke measures 27 cause and symptoms ranging from vague symptoms such as *ill* to specific symptoms such as *vomit*. Inference about the association




Fig. 9. Causal direction for different clusters, homogeneous case with a single cluster $k = 1$ (left panel), heterogeneous case with two clusters $k = 2$ (middle panel), and with three clusters $k = 3$ (right panel). Edges tend to weaken by increasing k . In $k = 3$, some edges disappear, and a directed edge is reversed.

between variables are made using the spike and slab model of Li, McCormick, and Clark (2019) which only discovers important associations between variables. In contrast, our method can enhance the analysis further by finding the causal direction between the dependent variables with or without the homogeneity assumption, i.e. $k = 1$ or $k > 1$. Scatter plot of data pairs clearly indicates that homogeneity is not a viable assumption. Fig. 9 confirms the inferred direction depends on the number of clusters.

6. Conclusion

We showed that heterogeneity can severely affect causal direction inference. In fact, the distribution of the test statistic deviates from the theoretical distribution obtained under the homogeneity assumption. This deviation considerably affects the type I error probability. The HSIC test statistic used in causal direction identification relies heavily on the homogeneity assumption. It requires proper adjustment using clustering labels when there are reasons to believe that homogeneity assumption is not tenable. To adjust this test statistic, we first developed a clustering method in the context of additive noise models that allows flexible number of clusters. We then used the estimated clustering labels to adjust the test statistic of causal direction for heterogeneous data. Our work can be distinguished from the Hu et al. (2018) in two directions (i) the causal model includes more parameters and incorporates the clustering labels, compare Fig. 5 with Hu et al. (2018, Figure 2), (ii) the causal direction test statistic is adjusted while Hu et al. (2018) only proposes the mixture model without adjusting the test statistic.

It is well-known in hypothesis testing that type I error and type II error work in opposite direction, e.g. decreasing type I error leads to an increase in type II error. For instance, in likelihood ratio tests for composite hypotheses $\alpha + \beta \leq 1$. Total error improves only in a local neighbourhood of the null hypothesis only if the sample size increases. Our method adjusts type I error probability while keeping the total error in the same order. In some of our experiments it even improves the total error so this adjustment builds an empirically more efficient test. The unadjusted HSIC test has a wrong asymptotic distribution when homogeneity assumption fails to hold. This is why, the type I error deviates from its nominal level.

To implement our proposed clustering algorithm we assumed that K and Δ are known. This assumption was only made to facilitate computation. The method, however, works for any value of K and Δ . This restriction can be relaxed by setting $K - \Delta = 1, K + \Delta = N$ if the computational power allows. This modification makes the clustering algorithm $\mathcal{O}(N^3)$.

CRedit authorship contribution statement

Vahid Partovi Nia: Contributed in methodology, theory and algorithm. **Xinlin Li:** Contributed in algorithm and data analysis. **Masoud Asgharian:** Contributed in methodology and theory. **Shoubo Hu:** Contributed in algorithm and data analysis. **Yanhui Geng:** Contributed in writing and proof reading. **Zhitang Chen:** Contributed in data analysis and proof reading.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix

Proof of Theorem 1. The proof is a multi-stage Gibbs sampler adaptation for the clustering case with varying cluster components. First we ensure that sampling from the discrete multivariate posterior $p(\mathbf{z} | \mu, \theta, k)$ and continuous multivariate $p(\mu | \mathbf{z}, \theta, k)$ converges to the joint $p(\mu, \mathbf{z} | \theta, k)$. Note that $p(\mu | \theta, \mathbf{z}, k)$ is multivariate Gaussian and $p(\mathbf{z} | \theta, \mu, k)$ is discrete with support $\{1, \dots, k\}^N$. Define the positive Markov transition kernel

$$k(\mathbf{z} | \mathbf{z}') = \int \dots \int p(\mathbf{z} | \mu, \theta, k) p(\mu | \mathbf{z}', \theta, k) d\mu.$$

This transition kernel is equivalent to taking intermediate samples from $\mu_t \sim p(\mu | \mathbf{z}_{t-1}, \theta, k)$ at iteration t and drawing $\mathbf{z}_t \sim p(\mathbf{z} | \mu_t, \theta, k)$.

It is easy to check that $k(\mathbf{z} | \mathbf{z}')$ is reversible and hence invariant with respect to the marginal $p(\mathbf{z} | \theta, k)$.

Now suppose we sample from the multivariate $p(\mathbf{z} | \mu, \theta, k)$ using univariate multinomial samplers. Let $k_1(\mathbf{z} | \mathbf{z}')$ be the transition kernel of a univariate Gibbs sampler of $p(\mathbf{z} | \mu, \theta, k)$ in increasing order $z_1, \dots, z_N$, i.e.

$$\begin{aligned} k_1(\mathbf{z} | \mathbf{z}') &= p(z'_1 | z_2, \dots, z_N, \mu, \theta, k) \\ &\quad \times p(z'_2 | z'_1, z_3, \dots, z_N, \mu, \theta, k) \dots p(z'_N | z'_1, \dots, z'_{N-1}, \mu, \theta, k) \\ &= \sum_{z_1, \dots, z_N} k_1(\mathbf{z} | \mathbf{z}') p(\mathbf{z} | \mu, \theta, k) \\ &= \sum_{z_1, \dots, z_N} p(z'_1 | z_2, \dots, z_N, \mu, \theta, k) \dots p(z'_N | z'_1, \dots, z'_{N-1}, \mu, \theta, k) \\ &\quad p(z_1 | z_2, \dots, z_N, \mu, \theta, k) p(z_2, \dots, z_N | \mu, \theta, k) \\ &= \sum_{z_2, \dots, z_N} p(z'_2 | z'_1, \dots, z_N, \mu, \theta, k) \\ &\quad \times p(z'_N | z'_1, \dots, z'_{N-1}, \mu, \theta, k) p(z'_1, z_2, \dots, z_N) \end{aligned}$$

in which we integrated over z_1 and decomposed

$$p(z'_1, z_2, \dots, z_N | \mu, \theta, k) = p(z'_1 | z_2, \dots, z_N, \mu, \theta, k) p(z_2, \dots, z_N | \mu, \theta, k).$$

Continue by decomposing

$$\begin{aligned} p(z'_1, z_2, \dots, z_N | \mu, \theta, k) &= p(z_2 | z'_1, z_3, \dots, z_N | \mu, \theta, k) \\ &\quad \times p(z'_1, z_3, \dots, z_N | \mu, \theta, k) \end{aligned}$$

and summing over z_2 . Repeating this re-arranging and summing over $z_3, \dots, z_N$ ends up with $p(z'_1, \dots, z'_N | \mu, \theta, k) = p(\mathbf{z}' | \mu, \theta, k)$.

A similar argument applies to $\mu_t \sim p(\mu | \mathbf{z}, \theta)$ to replace the intermediate sampler with univariate conditional samplers and to show

$$p(\mu' | \mathbf{z}, \theta) = \int \dots \int k_2(\mu | \mu') p(\mu | \mathbf{z}, \theta) d\mu$$

in which

$$k_2(\mu | \mu') = p(\mu'_1 | \mu_2, \dots, \mu_c, \mathbf{z}, \theta) \\ \times p(\mu'_2 | \mu'_1, \mu_3, \dots, \mu_c, \mathbf{z}, \theta) \dots p(\mu'_c | \mu'_1, \dots, \mu'_{c-1}).$$

Simple posterior calculations show these univariate distributions that construct the kernel k_1 are multinomial $z_i \sim \mathcal{M}(1, \pi_i)$ with probabilities

$$\pi_i = \left[\frac{\phi\left(\frac{\theta_i - \mu_1}{\sigma}\right)}{\sum_{c=1}^k \phi\left(\frac{\theta_i - \mu_c}{\sigma}\right)}, \dots, \frac{\phi\left(\frac{\theta_i - \mu_c}{\sigma}\right)}{\sum_{c=1}^k \phi\left(\frac{\theta_i - \mu_c}{\sigma}\right)} \right],$$

and the intermediate univariate samplers are Gaussian with mean $\bar{\theta}_c(1 + \frac{1}{N_c \kappa \tau^2})^{-1}$ and variance $\sigma^2(N_c + \frac{1}{\kappa \tau^2})^{-1}$. The label update step has k_1 and the mean update has k_2 transition kernel. Implementing these two steps sequentially is equivalent to a chain with the composition transition kernel $k_1 \circ k_2$. The last step is to margin over the cluster components k .

Assume a discrete uniform prior on $k \in \{K - \Delta, K + \Delta\}$ which allows to define a posterior proportional to the likelihood. Note that k does not affect the dimension of the marginalized posterior $p(\mathbf{z} | \theta, k)$, otherwise trans-dimensional samplers need to be developed. Marginalizing over k adds another step to the algorithm and implies sampling $k \sim \mathcal{M}(1, \pi)$ in which

$$\pi = \left[\frac{\exp\{\ell(\mathbf{z} | \theta, k = K - \Delta)\}}{\sum_{k=K-\Delta}^{K+\Delta} \exp\{\ell(\mathbf{z} | \theta, k)\}}, \dots, \frac{\exp\{\ell(\mathbf{z} | \theta, k = K + \Delta)\}}{\sum_{k=K-\Delta}^{K+\Delta} \exp\{\ell(\mathbf{z} | \theta, k)\}} \right]. \quad \blacksquare$$

Proof of Theorem 2. The product partition model (6) imposes mutually independent random pairs (x_i, y_i) across clusters. Therefore, given a certain causal direction their projections $\theta_i = g(x_i, y_i)$ are independent across clusters too. A similar argument holds for HSIC as a function of θ_i . This allows us to use the asymptotic results of the homogeneous case (Gretton, Bousquet, Smola and Schölkopf, 2005) in each cluster c and combine them using a product partition independence assumption.

For large n_c inside cluster c

$$n_c \text{HSIC}_c \sim \sum_{l=1}^{\infty} \lambda_l z_l^2 \approx \Gamma(\alpha, \beta)$$

in which λ_l are constants and z_l are independent standard Gaussian random variables with $\mu_c = \mathbb{E}(\text{HSIC}_c)$, $\sigma_c^2 = \mathbb{V}(\text{HSIC}_c)$. Define the aggregated test statistic

$$t = \sum_{c=1}^k n_c \text{HSIC}_c = \sum_{c=1}^k \sum_{l=1}^{\infty} \lambda_{cl} z_{cl}^2$$

which is clearly another countable sum $\sum_{m=1}^{\infty} \gamma_m w_m^2$ after swapping the sum order and re-arranging terms $m = (l-1)k + c$. $\blacksquare$

References

- Brand, J. E., & Thomas, J. S. (2013). Causal effect heterogeneity. In *Handbook of causal analysis for social research* (pp. 189–213). Springer.
- Dehkharghani, R., Mercan, H., Javeed, A., & Saygin, Y. (2014). Sentimental causal rule discovery from Twitter. *Expert Systems with Applications*, 41(10), 4950–4958.
- Fisher, R. A. (1926). The arrangement of field experiments. In *Breakthroughs in statistics* (pp. 82–91). Springer.
- Greenland, S., Pearl, J., & Robins, J. M. (1999). Causal diagrams for epidemiologic research. *Epidemiology*, 37–48.
- Gretton, A., Bousquet, O., Smola, A., & Schölkopf, B. (2005). Measuring statistical dependence with Hilbert-Schmidt norms. In *International conference on algorithmic learning theory* (pp. 63–77). Springer.
- Gretton, A., Smola, A. J., Bousquet, O., Herbrich, R., Belitski, A., Augath, M., et al. (2005). Kernel constrained covariance for dependence measurement. In *AISTATS, Vol. 10* (pp. 112–119).

- Hartigan, J. A. (1990). Partition models. *Communications in Statistics Theory and Methods*, 19(8), 2745–2756.
- Heckman, J. J. (1976). The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. *Annals of Economic and Social Measurement*, 5(4), 475–492. NBER.
- Hong, G., & Raudenbush, S. W. (2013). Heterogeneous agents, social interactions, and causal inference. In *Handbook of causal analysis for social research* (pp. 331–352). Springer.
- Hoyer, P. O., Janzing, D., Mooij, J. M., Peters, J., & Schölkopf, B. (2009). Nonlinear causal discovery with additive noise models. In *Advances in neural information processing systems* (pp. 689–696).
- Hu, S., Chen, Z., Partovi Nia, V., Chan, L., & Geng, Y. (2018). Causal inference and mechanism clustering of a mixture of additive noise models. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in neural information processing systems 31* (pp. 5206–5216).
- Janzing, D., & Schölkopf, B. (2010). Causal inference using the algorithmic Markov condition. *IEEE Transactions on Information Theory*, 56(10), 5168–5194.
- Kayikci, Y., & Stix, V. (2014). Causal mechanism in transport collaboration. *Expert Systems with Applications*, 41(4), 1561–1575.
- Lawrence, N. (2005). Probabilistic non-linear principal component analysis with Gaussian process latent variable models. *Journal of Machine Learning Research*, 6(Nov), 1783–1816.
- Li, Z., McCormick, T., & Clark, S. (2019). Bayesian joint spike-and-slab graphical lasso. In *International conference on machine learning* (pp. 3877–3885). PMLR.
- Liu, F., & Chan, L. (2016). Causal discovery on discrete data with extensions to mixture model. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(2), 21.
- Mooij, J. M., Peters, J., Janzing, D., Zscheischler, J., & Schölkopf, B. (2016). Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research*, 17(1), 1103–1204.
- Murray, C. J., Lopez, A. D., Black, R., Ahuja, R., Ali, S. M., Baqui, A., et al. (2011). Population Health Metrics Research Consortium gold standard verbal autopsy validation study: design, implementation, and development of analysis datasets. *Population Health Metrics*, 9(1), 1–15.
- Neyman, J. S. (1923). On the application of probability theory to agricultural experiments. *Statistical Science*, 5, 465–472.
- Partovi Nia, V., & Davison, A. C. (2015). A simple model-based approach to variable selection in classification and clustering. *The Canadian Journal of Statistics*, 43(2), 157–175.
- Pearl, J. (1986). Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29(3), 241–288.
- Pearl, J. (2009). *Causality*. Cambridge University Press.
- Pelleg, D., Moore, A. W., et al. (2000). X-means: Extending k-means with efficient estimation of the number of clusters. In *International conference on machine learning, Vol. 1* (pp. 727–734).
- Rücker, G., & Schumacher, M. (2008). Simpson's paradox visualized: the example of the rosiglitazone meta-analysis. *BMC Medical Research Methodology*, 8(1), 34.
- Schadt, E. E., Lamb, J., Yang, X., Zhu, J., Edwards, S., GuhaThakurta, D., et al. (2005). An integrative genomics approach to infer causal associations between gene expression and disease. *Nature Genetics*, 37(7), 710–717.
- Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., & Mooij, J. (2012). On causal and anticausal learning. arXiv preprint arXiv:1206.6471.
- Schwarz, G., et al. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461–464.
- Shimizu, S., Hoyer, P. O., Hyvärinen, A., & Kerminen, A. (2006). A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(Oct), 2003–2030.
- Simpson, E. H. (1951). The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 13(2), 238–241.
- Wood, A. T., Booth, J. G., & Butler, R. W. (1993). Saddlepoint approximations to the CDF of some statistics with nonnormal limit distributions. *Journal of the American Statistical Association*, 88(422), 680–686.
- Wu, W. W. (2010). Linking Bayesian networks and PLS path modeling for causal analysis. *Expert Systems with Applications*, 37(1), 134–139.
- Xie, Y. (2011). Causal inference and heterogeneity bias in social science. *Information Knowledge Systems Management*, 10(1–4), 279–289.
- Xie, Y. (2013). Population heterogeneity and causal inference. *Proceedings of the National Academy of Sciences*, 110(16), 6262–6268.
- Zhang, K., & Hyvärinen, A. (2009). On the identifiability of the post-nonlinear causal model. In *Conference on uncertainty in artificial intelligence* (pp. 647–655). AUAI Press.