Biometrika (2014), **101**, 1, *pp.* 219–228
Printed in Great Britain
doi: 10.1093/biomet/ast043
Advance Access publication 8 November 2013

Identifiability of Gaussian structural equation models with equal error variances

By J. Peters and P. Bühlmann

Seminar für Statistik, ETH Zürich, Rämistrasse 101, 8092 Zürich, Switzerland
peters@stat.math.ethz.ch buhlmann@stat.math.ethz.ch

Summary

We consider structural equation models in which variables can be written as a function of their parents and noise terms, which are assumed to be jointly independent. Corresponding to each structural equation model is a directed acyclic graph describing the relationships between the variables. In Gaussian structural equation models with linear functions, the graph can be identified from the joint distribution only up to Markov equivalence classes, assuming faithfulness. In this work, we prove full identifiability in the case where all noise variables have the same variance: the directed acyclic graph can be recovered from the joint Gaussian distribution. Our result has direct implications for causal inference: if the data follow a Gaussian structural equation model with equal error variances, then, assuming that all variables are observed, the causal structure can be inferred from observational data only. We propose a statistical method and an algorithm based on our theoretical findings.

Some key words: Causal inference; Graphical model; Structural equation model.

1. Introduction

1·1. *Graphical and structural equation models*

For random variables $X_1, \ldots, X_p$, we define a graphical model to be a pair $\{\mathcal{G}, \mathcal{L}(\mathcal{X})\}$, where $\mathcal{L}(\mathcal{X}) = \mathcal{L}(X_1, \ldots, X_p)$ is a joint probability distribution that is Markov with respect to a directed acyclic graph $\mathcal{G}$ (Lauritzen, 1996, Ch. 3.2). Structural equation models, also referred to as functional models, are related to graphical models. They are specified by a collection $\mathcal{S} = \{S_1, \ldots, S_p\}$ of p equations

$$S_j: \quad X_j = f_j(X_{\mathrm{PA}_j}, N_j) \quad (j = 1, \ldots, p) \tag{1}$$

and a joint distribution $\mathcal{L}(\mathcal{N}) = \mathcal{L}(N_1, \ldots, N_p)$ of the noise variables. Here, $\mathrm{PA}_j \subset \{1, \ldots, p\} \setminus \{j\}$ denotes the parents of j. We require the noise terms to be jointly independent, so $\mathcal{L}(\mathcal{N})$ is a product distribution. The graph $\mathcal{G}$ of a structural equation model is obtained by drawing directed edges from each variable X_k, $k \in \mathrm{PA}_j$, occurring on the right-hand side of (1) to X_j. The graph $\mathcal{G}$ is required to be acyclic. Furthermore, given a structural equation model, the joint distribution $\mathcal{L}(\mathcal{X})$ is fully determined and $\mathcal{L}(\mathcal{X})$ is Markov with respect to the graph $\mathcal{G}$ (Pearl, 2009, Theorem 1.4.1).

1·2. *Identifiability from the distribution*

We address the following problem. Given the joint distribution $\mathcal{L}(\mathcal{X}) = \mathcal{L}(X_1, \ldots, X_p)$ from a graphical model or from a structural equation model with directed acyclic graph $\mathcal{G}_0$, can we

© 2013 Biometrika Trust

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

recover the graph $\mathcal{G}_0$? By first considering graphical models, one can easily see that the answer is negative: the joint distribution $\mathcal{L}(\mathcal{X})$ is Markov with respect to different directed acyclic graphs, such as all fully connected directed acyclic graphs. Thus, there are many possible graphical models $\{\mathcal{G}, \mathcal{L}(\mathcal{X})\}$ for the same distribution $\mathcal{L}(\mathcal{X})$. Similarly, there are structural equation models with different structures that could have generated the distribution $\mathcal{L}(\mathcal{X})$. By making additional assumptions, one can obtain restricted graphical models and restricted structural equation models for which the graph is identifiable from the joint distribution. It is precisely here that the difference between graphical and functional models becomes apparent.

Given a graphical model, the distribution $\mathcal{L}(\mathcal{X})$ is faithful with respect to the directed acyclic graph $\mathcal{G}_0$ if each conditional independence found in $\mathcal{L}(\mathcal{X})$ is implied by the Markov condition. If faithfulness holds, one can obtain the Markov equivalence graph of the true directed acyclic graph $\mathcal{G}_0$ (Spirtes et al., 2000). But the Markov equivalence class may still be large (cf. Andersson et al., 1997) and the directed acyclic graph $\mathcal{G}_0$ is not identifiable. Furthermore, faithfulness in its full generality cannot be tested from data (Zhang & Spirtes, 2008). Since both the Markov condition and faithfulness only restrict the conditional independences in the joint distribution, it is not surprising that two graphs entailing the same conditional independences cannot be distinguished from one another.

Structural equation models enable us to exploit a different type of restriction. First, a general Gaussian structural equation model is equivalent to a Gaussian graphical model $\{\mathcal{G}_0, \mathcal{L}(\mathcal{X})\}$, so the structure $\mathcal{G}_0$ is not identifiable from $\mathcal{L}(\mathcal{X})$. Recently, however, it has been shown that this case is exceptional: (i) if we consider linear functions and non-Gaussian noise, we can identify the underlying directed acyclic graph $\mathcal{G}_0$ (Shimizu et al., 2006); (ii) if one restricts the functions to be additive in the noise component and excludes the linear Gaussian case as well as a few other pathological function-noise combinations, one can show that $\mathcal{G}_0$ is identifiable from $\mathcal{L}(\mathcal{X})$ (Hoyer et al., 2009; Peters et al., 2011). In this article, we prove that there is a third way to deviate from the general linear Gaussian case: (iii) Gaussian structural equation models where all functions are linear but the normally distributed noise variables which have equal variance σ^2 are again identifiable. The identifiability results (i) and (ii) require a condition called causal minimality. In its original form, Zhang & Spirtes (2008) define causal minimality as follows: for the true causal graph $\mathcal{G}_0$, $\mathcal{L}(\mathcal{X})$ is not Markov with respect to any proper subgraph of $\mathcal{G}_0$. Causal minimality is therefore a weak form of faithfulness. Remark 2 shows that for proving (iii) we assume causal minimality.

It may come as a surprise that for a class of Gaussian structural equation models the underlying directed acyclic graph is identifiable. The assumption of equal error variances seems natural for applications with variables from a similar domain and is commonly used in time series models.

1·3. *Causal interpretation*

Our result has implications for causal inference. If $\mathcal{G}_0$ is interpreted as the causal graph of the data-generating process for $X_1, \ldots, X_p$, the problem considered here is to infer the causal structure from the joint distribution. This is particularly interesting when the causal graph is of interest but interventional experiments are too expensive, unethical, or even impossible to perform. In the causal setting, our result reads as follows. If the observational data are generated by a Gaussian structural equation model that represents the causal relationships and has equal error variances, then the causal graph is identifiable from the joint distribution. Despite the potentially important application to causal inference, we present the main statement and its proof without using causal terminology; in particular, equations (1) and (2) can be interpreted as holding in distribution.

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

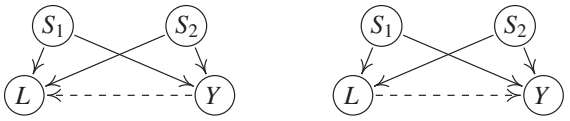

Fig. 1. The situation dealt with in the second part of case (ii) in the proof of Theorem 1, with $\mathcal{S} = \{S_1, S_2\}$ and $\mathcal{D} = \emptyset$. It contains the proof's main argument.

2. Identifiability for Gaussian models with equal error variances

We first introduce some notation. The index set $\mathcal{J} = \{1, \ldots, p\}$ corresponds to a set of vertices in a graph. Associated with $j \in \mathcal{J}$ are random variables X_j from $\mathcal{X} = (X_1, \ldots, X_p)$. Given a directed acyclic graph $\mathcal{G}$, we denote the parents of a node j by $\mathrm{PA}_j^{\mathcal{G}}$, the children by $\mathrm{CH}_j^{\mathcal{G}}$, the descendants by $\mathrm{DE}_j^{\mathcal{G}}$ and the nondescendants by $\mathrm{ND}_j^{\mathcal{G}}$.

We consider a structural equation model with directed acyclic graph $\mathcal{G}_0$ of the form

$$X_j = \sum_{k \in \mathrm{PA}_j^{\mathcal{G}_0}} \beta_{jk} X_k + N_j \quad (j = 1, \ldots, p), \tag{2}$$

where all the N_j are independent and identically distributed according to $N(0, \sigma^2)$ with $\sigma^2 > 0$. Additionally, for each $j \in \{1, \ldots, p\}$, we require that $\beta_{jk} \neq 0$ for all $k \in \mathrm{PA}_j^{\mathcal{G}_0}$.

Theorem 1. *Let $\mathcal{L}(\mathcal{X})$ be generated from model (2). Then $\mathcal{G}_0$ is identifiable from $\mathcal{L}(\mathcal{X})$ and the coefficients β_{jk} can be reconstructed for all j and $k \in \mathrm{PA}_j^{\mathcal{G}_0}$.*

Remark 1. The idea of the proof is to assume that there are two structural equation models with distinct graphs $\mathcal{G}$ and $\mathcal{G}'$ that lead to the same joint distribution. We exploit the Markov condition and causal minimality, see Remark 2, to find variables L and Y that have the same set of parents $\mathcal{S} = \{S_1, S_2\}$ in both graphs but have reversed edges between each other in $\mathcal{G}$ and $\mathcal{G}'$, as shown in Fig. 1. Defining $L^* = L\,|_{\mathcal{S}=s}$ for some value $s \in \mathbb{R}^2$, we can use the equal error variances to show that L^* has different variances in both graphs. This leads to a contradiction.

Remark 2. Theorem 1 assumes that the coefficients $\beta_{jk} \neq 0$ do not vanish for any $k \in \mathrm{PA}_j^{\mathcal{G}_0}$. Lemma A4 below and Proposition 2 in Peters et al. (2011) show that this condition implies causal minimality. From our point of view, causal minimality is a natural condition that is in accordance with the intuitive understanding of a causal influence between variables.

Remark 3. Theorem 1 can be generalized to the case where the error covariance matrix has the form $\mathrm{cov}(N_1, \ldots, N_p) = \sigma^2\,\mathrm{diag}(\alpha_1, \ldots, \alpha_p)$ with prespecified $\alpha_1, \ldots, \alpha_p$ and unknown σ^2.

3. Penalized maximum likelihood estimator

Consider data which are independent and identically distributed realizations of $X^{(1)}, \ldots, X^{(n)}$ from model (2) with true coefficients β_{jk}^0. The representation in vector form is $X = BX + N$, where B is the $p \times p$ matrix with entries $B_{jk} = \beta_{jk}$. To make the article easier to read, we write

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

B or β whenever we think of a matrix or vector of parameters, respectively. As estimator for the coefficients $B^0 = (\beta^0_{jk})_{j,k}$ and the error variance σ^2, we consider

$$\{\hat{\beta}(\lambda), \hat{\sigma}^2(\lambda)\} = \underset{\beta \in \mathcal{B}, \sigma^2 \in \mathbb{R}^+}{\arg\min} \; -\ell(\beta, \sigma^2; X^{(1)}, \ldots, X^{(n)}) + \lambda \|\beta\|_0, \tag{3}$$

where

$$-\ell(\beta, \sigma^2; X^{(1)}, \ldots, X^{(n)}) = \frac{np}{2} \log(2\pi\sigma^2) + \frac{n}{2\sigma^2} \mathrm{tr}\{(I - B)^{\mathrm{T}}(I - B)\hat{\Sigma}\},$$

with sample covariance matrix $\hat{\Sigma}$, is the negative loglikelihood assuming equal error variances σ^2, and $\|\beta\|_0 |\{(j,k) : \beta_{jk} \neq 0\}|$. Furthermore, $\mathcal{B} = \{B \in \mathbb{R}^{p \times p} : \mathrm{Adj}(B) \text{ has only zero eigenvalues}\}$ contains only those coefficient matrices whose corresponding graphs do not have cycles (Cvetković et al., 1995, p. 81). Here, $\mathrm{Adj}(B)_{kj} = 1_{\beta_{jk} \neq 0}$ is the adjacency matrix. Minimizing over all $\beta \in \mathcal{B}$ includes optimizing over all directed acyclic graphs; see § 4. The induced directed acyclic graph from $\hat{\beta}(\lambda)$ is denoted by $\hat{\mathcal{G}}$. For $\lambda = \log(n)/2$, the objective function in equation (3) is the BIC score.

The convergence rate and consistency of the penalized maximum likelihood estimator for the true coefficients β^0_{jk} and the true structure $\mathcal{G}_0$ follow from an analysis in van de Geer & Bühlmann (2013, Theorem 5.1) under regularity conditions. More precisely, for $\lambda_n = \log(n)/2$ we have, as $n \to \infty$,

$$\sum_{j,k=1}^{p} \{\hat{\beta}_{jk}(\lambda_n) - \beta^0_{jk}\}^2 = O_{\mathrm{P}}\{\log(n) n^{-1}\}, \quad \mathrm{pr}(\hat{\mathcal{G}}_n = \mathcal{G}_0) \to 1.$$

The results in van de Geer & Bühlmann (2013, § 5) also cover the high-dimensional sparse setting where $p = p_n = O\{n/\log(n)\}$.

One could use a combination of the PC-algorithm and minimization of the penalized likelihood in equation (3): the former, which is computationally very efficient, could be employed for estimating the Markov equivalence class, while the latter could be used for orienting the remaining undirected edges. A related approach has been suggested by Tillman et al. (2010). For consistency in the first step, one necessarily requires a version of the strong faithfulness assumption, which can be very restrictive (Uhler et al., 2013). Penalized maximum likelihood estimation does not need such an assumption (van de Geer & Bühlmann, 2013) but pays a price in terms of computational complexity.

4. GREEDY SEARCH ALGORITHM

Because the optimization in (3) is over the space of all directed acyclic graphs, the estimator is hard to compute. Already for $p = 20$ there are $2{\cdot}3 \times 10^{72}$ directed acyclic graphs (OEIS Foundation Inc., 2011), which makes an exhaustive search infeasible. Instead we propose a greedy procedure, which we call greedy directed acyclic graph search with equal error variances. At each iteration t, we are given a directed acyclic graph $\mathcal{G}_t$ and move to the neighbouring directed acyclic graph with the largest drop in BIC score. If all neighbours have a higher BIC score in (3) than $\mathcal{G}_t$, the algorithm terminates. Here, we say that two directed acyclic graphs are neighbours if they can be transformed into each other by one edge addition, removal or reversal. Chickering (2002) proposed a similar search strategy but with the search done in the space of Markov equivalence classes rather than over directed acyclic graphs.

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

In order to shorten the runtime, we randomly search through neighbouring directed acyclic graphs until we find a directed acyclic graph with a better score than $\mathcal{G}_t$, and then we use this directed acyclic graph for $\mathcal{G}_{t+1}$. We consider at least k neighbours; if there are several directed acyclic graphs among the first k that have better scores than $\mathcal{G}_t$, we take the best one. The whole procedure improves further if we increase the probability of changing edges that point into nodes whose residuals have a high variance. This modification and the score function are the only parts of the algorithm that make use of the equal error variances. Additionally, we reinitiate the method five times starting from a random sparse graph with $k = p$, $k = 2p$, $k = 3p$, $k = 5p$ and $k = 300$. This choice is ad hoc but works well in practice, as it decreases the risk of getting stuck in a local optimum. R code for this method is available as Supplementary Material.

5. Experiments

5·1. *Existing methods*

We compare our method against the PC-algorithm (Spirtes et al., 2000) and the greedy equivalence search (Chickering, 2002). The latter approximates the BIC-regularized maximum likelihood estimator for nonrestricted Gaussian structural equation models. Both methods can only recover the Markov equivalence class, see § 1·2, and therefore leave some arrows undirected. The Markov equivalence class can be represented by a completed partially directed acyclic graph. In the experiments, we report the structural Hamming distance between the true and estimated partially directed acyclic graphs; this assigns a distance of 2 for each pair of reversed edges, for example, $\rightarrow$ in the true and $\leftarrow$ in the estimated graph; all other edge mistakes count as 1.

5·2. *Random graphs*

For varying n and p, we compare the three methods. For a given value of p, we randomly choose an ordering of the variables with respect to the uniform distribution and include each of the $p(p-1)/2$ possible edges with a probability of p_{edge}. All noise variances are set to 1, since scaling all noise variables by a common factor yields exactly the same estimates $\hat{\beta}$ and $\hat{\mathcal{G}}$. The coefficients β^0_{jk} are uniformly chosen from $[-1, -0{\cdot}1] \cup [0{\cdot}1, 1]$. We consider a sparse setting with $p_{\text{edge}} = 3/(2p-2)$, which results in an expected number of $3p/4$ edges, and a dense setting with $p_{\text{edge}} = 0{\cdot}3$. Table 1 shows, for the sparse setting, the average structural Hamming distance to the true directed acyclic graph and to the true completed partially directed acyclic graph over 100 simulations. Except for $p = 40$ and $n = 100$, the graphs estimated by the proposed method are closer to the true directed acyclic graph than are the graphs resulting from state-of-the-art methods, which can only recover the true Markov equivalence class; greedy directed acyclic graph search also performs better when comparing the distance to the true completed partially directed acyclic graph. Table 2 shows the analogous results for the dense setting, in which the improvement obtained by greedy directed acyclic graph search with equal error variances is even greater.

As a proof of concept, we also simulate data with $n = 500$ from a non-faithful distribution: $X_1 = N_1$, $X_2 = -X_1 + N_2$ and $X_3 = X_1 + X_2 + N_3$. As expected from the theory, the PC-algorithm and greedy equivalence search fail here; in all 100 experiments, they output $X_1 \rightarrow X_2 \leftarrow X_3$, which is not the correct Markov equivalence class. In contrast, greedy directed acyclic graph search always identified the correct directed acyclic graph.

5·3. *Deviation from equal error variances*

When the data are generated by a Gaussian structural equation model with different error variances, the method is not guaranteed to find the correct directed acyclic graph or the correct

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

Table 1. *Average structural Hamming distance between the estimated and true directed acyclic graphs and between the estimated and true Markov equivalence classes, for sparse graphs with p nodes and sample size n*

		$n = 100$			$n = 500$			$n = 1000$		
p		GDS_{EEV}	PC	GES	GDS_{EEV}	PC	GES	GDS_{EEV}	PC	GES
5	DAG	1·5	3·9	3·6	0·5	2·9	2·8	0·4	3·0	2·5
	CPDAG	1·5	2·9	2·3	0·5	1·4	1·2	0·3	1·0	0·7
20	DAG	12·2	14·1	18·0	4·5	11·1	10·3	2·7	10·1	8·7
	CPDAG	13·9	10·9	17·0	5·2	7·7	7·6	3·0	6·9	5·6
40	DAG	44·7	29·6	53·0	15·7	22·6	26·1	10·7	20·1	21·9
	CPDAG	50·0	24·4	53·1	18·9	15·9	23·4	13·4	13·3	17·5

DAG, directed acyclic graph; CPDAG, completed partially directed acyclic graph; GDS_{EEV}, greedy directed acyclic graph search with equal error variances; PC, PC-algorithm; GES, greedy equivalence search.

Table 2. *Average structural Hamming distance between the estimated and true directed acyclic graphs and between the estimated and true Markov equivalence classes, for dense graphs with p nodes and sample size n*

		$n = 100$			$n = 500$			$n = 1000$		
p		GDS_{EEV}	PC	GES	GDS_{EEV}	PC	GES	GDS_{EEV}	PC	GES
5	DAG	1·2	2·9	3·0	0·6	2·4	2·2	0·3	2·1	2·1
	CPDAG	1·3	2·1	1·9	0·5	1·2	0·7	0·2	0·8	0·5
20	DAG	30·0	56·6	63·9	12·5	55·7	66·3	8·2	57·6	69·1
	CPDAG	31·0	56·1	63·2	13·1	55·5	66·2	8·8	57·5	68·5
40	DAG	216·1	242·8	323·1	185·2	247·2	430·4	172·0	248·9	470·6
	CPDAG	217·1	242·4	323·0	185·7	247·0	430·1	172·2	248·5	470·4

DAG, directed acyclic graph; CPDAG, completed partially directed acyclic graph; GDS_{EEV}, greedy directed acyclic graph search with equal error variances; PC, PC-algorithm; GES, greedy equivalence search.

Markov equivalence class. When the true data-generating process follows such a Gaussian structural equation model with different variances, we can always represent it as a model with equal error variances if we apply a fine-tuned rescaling of the variables, $X_i \mapsto a_i X_i$, with a_i equal to the inverse of the standard deviation of the error in the ith structural equation. Of course, such a rescaling is only possible when we know the error variances, hence the term fine-tuned. In the hypothetical case where the data would be scaled with such a deceptive fine-tuned standardization, the graph identified by our method would belong to the correct Markov equivalence class. We emphasize, however, that this is for an artificial scenario, which differs from the situation of having raw data from a Gaussian structural equation model with different error variances.

An important question is how sensitive our method is to deviations from the assumption of equal error variances. We investigate this empirically. For $p = 10$ and $n = 500$, we sample the noise variances uniformly from $[1 - a, 1 + a]$ and vary a between 0 and 0·9. Theorem 1 establishes identifiability of the graph only for $a = 0$. As before, the coefficients β^0_{jk} are uniformly chosen from $[-1, -0{\cdot}1] \cup [0{\cdot}1, 1]$. The parameter p_{edge} is chosen to be $2/(p - 1)$, on average resulting in p edges; this is between the sparse and dense settings. Figure 2 shows that the performance of greedy directed acyclic graph search is relatively robust as the parameter a changes. Even for large values of a, the method does not perform worse than the PC-algorithm. The best-score method reports the result of greedy directed acyclic graph search or greedy equivalence search depending on which method obtained the better score. Greedy directed acyclic graph search was chosen in 100%, 100%, 88%, 36%, 7%, 1%, 2%, 0%, 0% and 0% of the cases, for $a = 0, 0{\cdot}1, \ldots, 0{\cdot}9$, respectively.

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

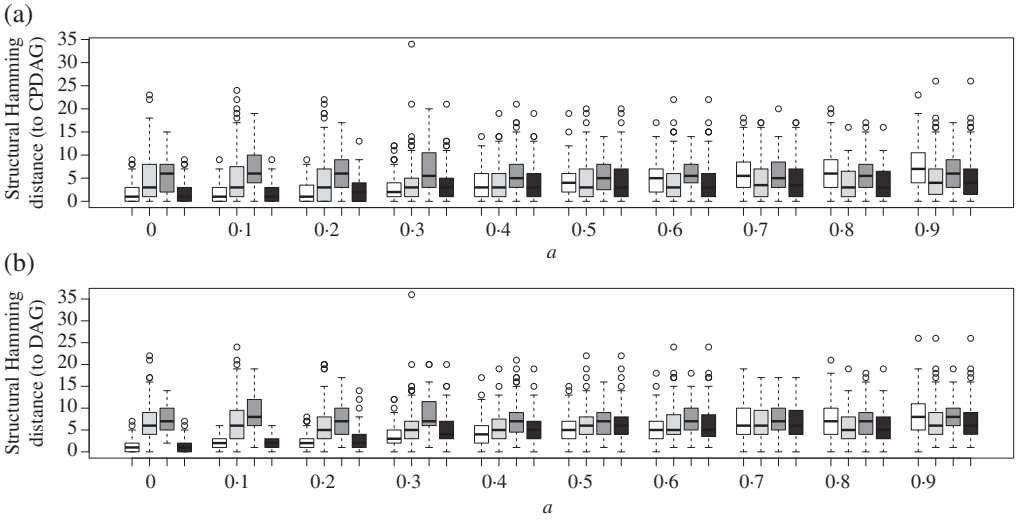

Fig. 2. Box plots for the structural Hamming distance of greedy directed acyclic graph search (white), greedy equivalence search (light grey), the PC-algorithm (grey) and a best-score method (dark grey) to: (a) the true directed acyclic graph, DAG; and (b) the true partially directed acyclic graph, CPDAG. The horizontal axis shows values of a parameter a which measures the perturbation from the case of equal error variances; $a = 0$ corresponds to equal error variances.

Table 3. BIC *scores of greedy equivalence search and greedy directed acyclic graph search with equal error variances obtained for different types of microarray data; smaller is better*

	Prostate	Lymphoma	Riboflavin	Leukaemia	Brain	Cancer	Colon
GES	4095	4560	2711	5456	1411	5891	3224
GDS_{EEV}	6057	5404	3236	5481	1343	6288	3201

GES, greedy equivalence search; GDS_{EEV}, greedy directed acyclic graph search with equal error variances.

5·4. *Real data*

We now apply the greedy equivalence search and greedy directed acyclic graph search to seven datasets containing microarray data, described by Dettling & Bühlmann (2003) and Bühlmann et al. (2014), and compare their BIC scores. When greedy equivalence search gives the better score, this indicates that the assumption of equal error variances is not justified. In Fig. 2 we saw that even then, it might sometimes be useful to look at the greedy directed acyclic graph search solution. If, on the other hand, greedy directed acyclic graph search yields a better score than greedy equivalence search, we prefer the solution obtained by greedy directed acyclic graph search, which furthermore is a graph rather than a Markov equivalence class. To avoid a high-dimensional setting with $p > n$, we always chose the $0{\cdot}8n$ genes with the highest variance. Table 3 shows that in two out of the seven datasets, greedy directed acyclic graph search produced a better score than greedy equivalence search.

For the Colon example, greedy directed acyclic graph search proposed a directed acyclic graph with 192 edges, while greedy equivalence search proposed a graph with 217 edges. There are 91 edges in both solutions, 61 with the same orientation. The graphs therefore differ on roughly half of the edges.

ACKNOWLEDGEMENT

We thank R. Tanase for fruitful discussions. The research leading to these results received funding from the European Union's Seventh Framework Programme.

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

SUPPLEMENTARY MATERIAL

Supplementary material available at *Biometrika* online includes R code for the greedy directed acyclic graph search with equal error variances.

APPENDIX

Some lemmata

In the following we consider different subsets of the set of variables $\mathcal{X}$; to simplify notation, we do not distinguish between indices and variables, since the context should make the meaning clear. This way, we can also speak of the parents $\mathrm{PA}_B^{\mathcal{G}}$ of a variable $B \in \mathcal{X}$. We also consider sets of variables $\mathcal{S} \subset \mathcal{X}$ to be a single multivariate variable.

The following four statements are all plausible and their proofs mostly involve technicalities. The reader may skip to the next section and use the lemmata whenever needed.

LEMMA A1. *Let* $(A_1, \ldots, A_m) \sim N\{(\mu_1, \ldots, \mu_m)^{\mathrm{T}}, \Sigma\}$ *with strictly positive definite* Σ*, and define* $A_1^* = A_1 \,|_{(A_2,\ldots,A_m)=(a_2,\ldots,a_m)}$*, in distribution. Then* $\mathrm{var}(A_1^*) \leqslant \mathrm{var}(A_1)$ *for all* $(a_2, \ldots, a_m) \in \mathbb{R}^{m-1}$.

We use the notation of conditional variables rather than conditional distributions to improve readability.

Proof. Let us decompose Σ into

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \Sigma_{12}^{\mathrm{T}} \\ \Sigma_{12} & \Sigma_{22} \end{pmatrix}$$

with an $(m-1) \times 1$ vector Σ_{12}. Since Σ_{22}^{-1} is positive definite, $\mathrm{var}(A_1^*) = \sigma_1^2 - \Sigma_{12}^{\mathrm{T}} \Sigma_{22}^{-1} \Sigma_{12} \leqslant \sigma_1^2$. □

LEMMA A2 (Peters et al., 2011). *Let* Y, N, Q *and* R *be random variables taking values in* $\mathcal{Y}$, $\mathcal{N}$, $\mathcal{Q}$ *and* $\mathcal{R}$*, respectively, and suppose that their joint distribution is absolutely continuous with respect to some product measure; we denote the density by* $p_{Y,Q,R,N}(y, q, r, n)$*. The variables* Q *and* R *can be multivariate. Let* $f : \mathcal{Y} \times \mathcal{Q} \times \mathcal{N} \to \mathbb{R}$ *be a measurable function. If* $N \perp\!\!\!\perp (Y, Q, R)$*, then for all* $q \in \mathcal{Q}$ *and* $r \in \mathcal{R}$ *with* $p_{Q,R}(q, r) > 0$*,* $f(Y, Q, N) \,|_{Q=q,R=r} = f(Y \,|_{Q=q,R=r}, q, N)$ *in distribution.*

LEMMA A3 (Peters et al., 2011). *Let* $\mathcal{L}(\mathcal{X})$ *be generated by a structural equation model as in* (2) *with corresponding directed acyclic graph* $\mathcal{G}$ *and consider a variable* $X \in \mathcal{X}$*. If* $\mathcal{S} \subseteq \mathrm{ND}_X^{\mathcal{G}}$*, then* $N_X \perp\!\!\!\perp \mathcal{S}$.

LEMMA A4. *Let* $\mathcal{L}(\mathcal{X})$ *be generated from a structural equation model as in* (2) *with directed acyclic graph* $\mathcal{G}$*. Consider a variable* $B \in \mathcal{X}$ *and one of its parents* $A \in \mathrm{PA}_B^{\mathcal{G}}$*. For all sets* $\mathcal{S}$ *with* $\mathrm{PA}_B^{\mathcal{G}} \setminus \{A\} \subseteq \mathcal{S} \subseteq \mathrm{ND}_B^{\mathcal{G}} \setminus \{A\}$ *we have* $B \not\perp\!\!\!\perp A \mid \mathcal{S}$.

Proof. Define $Q = \mathrm{PA}_B^{\mathcal{G}} \setminus \{A\}$ such that we have $\mathcal{S} = (Q, R)$ for some R. From Lemma A2 we obtain that for some linear function f

$$B \,|_{Q=q,R=r} = f(q) + \beta A \,|_{Q=q,R=r} + N_B$$

in distribution, with $N_B \perp\!\!\!\perp A \,|_{Q=q,R=r}$. But since $\beta \neq 0$, $A \,|_{Q=q,R=r} \not\perp\!\!\!\perp B \,|_{Q=q,R=r}$. □

Proof of Theorem 1

If we assumed faithfulness, we could recover the correct Markov equivalence class, which itself implies the existence of an L and Y as in Remark 1 (Chickering, 1995, Theorem 2). But since we are not assuming faithfulness, proving existence of a situation similar to that in Fig. 1 requires more work. This part of the proof, due to not assuming faithfulness, is taken from Peters et al. (2011) and remains almost the same. The difference from the argument in Peters et al. (2011) is that we can prove causal minimality and need not assume it. Also new are Lemmata A1 and A4, as well as the proof's main argument given in the second part of case (ii).

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

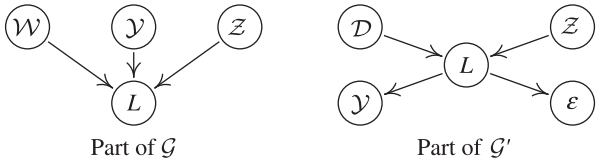

Fig. A1. Nodes adjacent to L in $\mathcal{G}$ and $\mathcal{G}'$.

Proof. We assume that there are two structural equation models as in (2) which both induce $\mathcal{L}(\mathcal{X})$, one with graph $\mathcal{G}$ and the other with graph $\mathcal{G}'$. We will show that $\mathcal{G} = \mathcal{G}'$. Since directed acyclic graphs do not contain any cycles, we can always find nodes that have no descendants. To see this, start a directed path at some node; after at most $\#\mathcal{X} - 1$ steps, we reach a node without a child. Eliminating such a node from the graph leads to a directed acyclic graph again; we can discard further nodes without children in the new graph. We repeat this process for all nodes that have no children in both $\mathcal{G}$ and $\mathcal{G}'$ and which have the same parents in both graphs. If we end up with no nodes left, the two graphs are identical and the result is proved. Otherwise, we end up with a smaller set of variables which we again call $\mathcal{X}$, two smaller graphs which we again call $\mathcal{G}$ and $\mathcal{G}'$, and a node L that has no children in $\mathcal{G}$ and is such that either $\mathrm{PA}_L^{\mathcal{G}} \neq \mathrm{PA}_L^{\mathcal{G}'}$ or $\mathrm{CH}_L^{\mathcal{G}'} \neq \emptyset$. We will show that this leads to a contradiction. Importantly, because of the Markov property of the distribution with respect to $\mathcal{G}$, all other nodes are independent of L given $\mathrm{PA}_L^{\mathcal{G}}$:

$$L \perp\!\!\!\perp \mathcal{X} \setminus (\mathrm{PA}_L^{\mathcal{G}} \cup \{L\}) \mid \mathrm{PA}_L^{\mathcal{G}}. \tag{A1}$$

To make the arguments easier to understand, we introduce the following notation; see also Fig. A1. We partition $\mathcal{G}$-parents of L into $\mathcal{Y}$, $\mathcal{Z}$ and $\mathcal{W}$. Here, the $\mathcal{Z}$ are also $\mathcal{G}'$-parents of L, the $\mathcal{Y}$ are $\mathcal{G}'$-children of L, and the $\mathcal{W}$ are not adjacent to L in $\mathcal{G}'$. Let $\mathcal{D}$ be the $\mathcal{G}'$-parents of L that are not adjacent to L in $\mathcal{G}$ and let $\mathcal{E}$ be the $\mathcal{G}'$-children of L that are not adjacent to L in $\mathcal{G}$. Thus we have $\mathrm{PA}_L^{\mathcal{G}} = \mathcal{Y} \cup \mathcal{Z} \cup \mathcal{W}$, $\mathrm{CH}_L^{\mathcal{G}} = \emptyset$, $\mathrm{PA}_L^{\mathcal{G}'} = \mathcal{Z} \cup \mathcal{D}$ and $\mathrm{CH}_L^{\mathcal{G}'} = \mathcal{Y} \cup \mathcal{E}$. Consider $\mathcal{T} = \mathcal{W} \cup \mathcal{Y}$. We distinguish two cases.

Case (i): $\mathcal{T} = \emptyset$. Then there must be a node $D \in \mathcal{D}$ or a node $E \in \mathcal{E}$; otherwise L would have been discarded. If there is a $D \in \mathcal{D}$, then (A1) implies $L \perp\!\!\!\perp D \mid \mathcal{S}$ for $\mathcal{S} = \mathcal{Z} \cup \mathcal{D} \setminus \{D\}$, which contradicts Lemma A4 applied to $\mathcal{G}'$. If $\mathcal{D} = \emptyset$ and there is an $E \in \mathcal{E}$, then $E \perp\!\!\!\perp L \mid \mathcal{S}$ holds for $\mathcal{S} = \mathcal{Z} \cup \mathrm{PA}_E^{\mathcal{G}'} \setminus \{L\}$, which also contradicts Lemma A4. To avoid cycles it is necessary that $\mathcal{Z} \subseteq \mathrm{ND}_E^{\mathcal{G}'}$.

Case (ii): $\mathcal{T} \neq \emptyset$. Then $\mathcal{T}$ contains a $\mathcal{G}'$-youngest node with the property that there is no directed $\mathcal{G}'$-path from this node to any other node in $\mathcal{T}$. This node may not be unique. Suppose that $W \in \mathcal{W}$ is such a youngest node. Consider the directed acyclic graph $\tilde{\mathcal{G}}'$ that equals $\mathcal{G}'$ with additional edges $Y \to W$ and $W' \to W$ for all $Y \in \mathcal{Y}$ and $W' \in \mathcal{W} \setminus \{W\}$. In $\tilde{\mathcal{G}}'$, L and W are not adjacent. Thus we can find a set $\tilde{\mathcal{S}}$ that d-separates L and W in $\tilde{\mathcal{G}}'$; indeed, one can take $\tilde{\mathcal{S}} = \mathrm{PA}_L^{\tilde{\mathcal{G}}'}$ if $W \notin \mathrm{DE}_L^{\tilde{\mathcal{G}}'}$ and $\tilde{\mathcal{S}} = \mathrm{PA}_W^{\tilde{\mathcal{G}}'}$ if $L \notin \mathrm{DE}_W^{\tilde{\mathcal{G}}'}$. Then $\mathcal{S} = \tilde{\mathcal{S}} \cup \{\mathcal{Y}, \mathcal{Z}, \mathcal{W} \setminus \{W\}\}$ d-separates L and W in $\tilde{\mathcal{G}}'$.

We now prove this claim. All $Y \in \mathcal{Y}$ are already in $\tilde{\mathcal{S}}$ in order to block $L \to Y \to W$. Suppose there is a $\tilde{\mathcal{G}}'$-path that is blocked by $\tilde{\mathcal{S}}$ and unblocked if we add Z and W' nodes to $\tilde{\mathcal{S}}$. How can we unblock a path by including more nodes? The path $L \cdots V_1 \cdots U_1 \cdots W$, see Fig. A2, must contain a collider V_1 that is an ancestor of a Z with $V_1, \ldots, V_m$, $Z \notin \tilde{\mathcal{S}}$ and corresponding nodes U_i for a W' node. Choose V_1 and U_1 on the given path so close to each other that there is no such collider in between. If there is no V_1, choose U_1 close to L; if there is no U_1, choose V_1 close to W. Now the path $L \leftarrow Z \cdots V_1 \cdots U_1 \cdots W' \to W$ is unblocked given $\tilde{\mathcal{S}}$, which contradicts the fact that $\tilde{\mathcal{S}}$ d-separates L and W. This ends the proof of the claim.

The set $\mathcal{S}$ d-separates L and W also in $\mathcal{G}'$, because $\mathcal{G}'$ contains fewer paths. We have $L \perp\!\!\!\perp W \mid \mathcal{S}$, which contradicts Lemma A4 applied to $\mathcal{G}$. In summary, $W \in \mathcal{W}$ cannot be the $\mathcal{G}'$-youngest node.

Therefore, the $\mathcal{G}'$-youngest node in $\mathcal{T}$ must be some $Y \in \mathcal{Y}$. We have that

$$\sigma_{\mathcal{G}}^2 = \sigma_{\mathcal{G}'}^2 = \min_{X \in \mathcal{X}} \mathrm{var}(X) = \sigma^2. \tag{A2}$$

Define $\mathcal{S} = \mathrm{PA}_L^{\mathcal{G}} \setminus \{Y\} \cup \mathcal{D}$. Clearly, $\mathcal{S} \subseteq \mathrm{ND}_L^{\mathcal{G}}$ since L does not have any descendants in $\mathcal{G}$. Define $Q = \mathrm{PA}_L^{\mathcal{G}} \setminus \{Y\}$ and take any $s = (q, d)$. Define $L^* = L \mid_{\mathcal{S}=s}$ in distribution, and also $Y^* = Y \mid_{\mathcal{S}=s}$ in distribution. Then, from $\mathcal{G}$ and by using Lemma A2, we find that for some linear function f, $L^* = f(q) + \beta Y^* + N_L$

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014

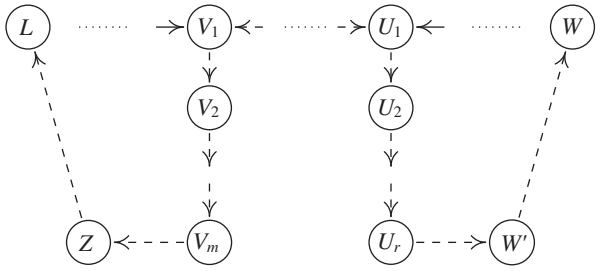

Fig. A2. Assume that the path $L \cdots V_1 \cdots U_1 \cdots W$ is blocked by $\tilde{S}$ but unblocked if we include Z and W'; then the dashed path is unblocked given $\tilde{S}$.

in distribution, with $N_L \perp\!\!\!\perp Y|_{S=s}$. The independence holds because $S \subseteq \mathrm{ND}_L^{\mathcal{G}}$. We then have

$$\mathrm{var}(L^*) = \beta^2 \mathrm{var}(Y^*) + \sigma^2 > \sigma^2. \tag{A3}$$

Since $\mathrm{PA}_L^{\mathcal{G}'} \subseteq S$, we obtain from $\mathcal{G}'$ and Lemma A1 that

$$\mathrm{var}(L^*) \leqslant \sigma^2 \tag{A4}$$

since $\det\{\mathrm{cov}(\mathcal{X})\} \neq 0$. Equations (A3) and (A4) contradict each other.

To prove Remark 3, replace $\mathrm{var}(X)$ by $\mathrm{var}(X)/\alpha_X$ in (A2) and σ^2 by $\sigma^2\alpha_L$ in (A3) and (A4). □

REFERENCES

ANDERSSON, S. A., MADIGAN, D. & PERLMAN, M. D. (1997). A characterization of Markov equivalence classes for acyclic digraphs. *Ann. Statist.* **25**, 505–41.

BÜHLMANN, P., KALISCH, M. & MEIER, L. (2014). High-dimensional statistics with a view towards applications in biology. *Ann. Rev. Statist. Appl.* **1**, 255–78.

CHICKERING, D. M. (1995). A transformational characterization of equivalent Bayesian network structures. In *Proc. 11th Conf. Uncertainty Artificial Intel. (UAI 1995)*. San Francisco, California: Morgan Kaufmann, pp. 87–98.

CHICKERING, D. M. (2002). Optimal structure identification with greedy search. *J. Mach. Learn. Res.* **3**, 507–54.

CVETKOVIĆ, D., DOOB, M. & SACHS, H. (1995). *Spectra of Graphs: Theory and Application*. Heidelberg: Barth, 3rd ed.

DETTLING, M. & BÜHLMANN, P. (2003). Boosting for tumor classification with gene expression data. *Bioinformatics* **19**, 1061–9.

HOYER, P., JANZING, D., MOOIJ, J., PETERS, J. & SCHÖLKOPF, B. (2009). Nonlinear causal discovery with additive noise models. In *Adv. Neural Info. Proces. Syst. 21 (NIPS 2008)*. Red Hook, New York: Curran Associates, pp. 689–96.

LAURITZEN, S. (1996). *Graphical Models*. New York: Oxford University Press.

OEIS FOUNDATION INC. (2011). The on-line encyclopedia of integer sequences. http://oeis.org/A003024.

PEARL, J. (2009). *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press, 2nd ed.

PETERS, J., MOOIJ, J. M., JANZING, D. & SCHÖLKOPF, B. (2011). Identifiability of causal graphs using functional models. In *27th Conf. Uncertainty Artificial Intel. (UAI 2011)*. Corvallis, Oregon: AUAI Press, pp. 589–98.

SHIMIZU, S., HOYER, P. O., HYVÄRINEN, A. & KERMINEN, A. (2006). A linear non-Gaussian acyclic model for causal discovery. *J. Mach. Learn. Res.* **7**, 2003–30.

SPIRTES, P., GLYMOUR, C. & SCHEINES, R. (2000). *Causation, Prediction, and Search*. Cambridge, Massachusetts: MIT Press, 2nd ed.

TILLMAN, R., GRETTON, A. & SPIRTES, P. (2010). Nonlinear directed acyclic structure learning with weakly additive noise models. In *Adv. Neural Info. Proces. Syst. 22 (NIPS 2009)*. Red Hook, New York: Curran Associates, pp. 1847–55.

UHLER, C., RASKUTTI, G., BÜHLMANN, P. & YU, B. (2013). Geometry of the faithfulness assumption in causal inference. *Ann. Statist.* **41**, 436–63.

VAN DE GEER, S. & BÜHLMANN, P. (2013). ℓ_0-penalized maximum likelihood for sparse directed acyclic graphs. *Ann. Statist.* **41**, 536–67.

ZHANG, J. & SPIRTES, P. (2008). Detection of unfaithfulness and robust causal inference. *Minds Mach.* **18**, 239–71.

[*Received May* 2012. *Revised July* 2013]

Downloaded from http://biomet.oxfordjournals.org/ at ETH ZÃ¼rich on March 3, 2014