

Research Article

Christoph Schultheiss* and Peter Bühlmann

On the pitfalls of Gaussian likelihood scoring for causal discovery

<https://doi.org/10.1515/jci-2022-0068>

received October 20, 2022; accepted March 29, 2023

Abstract: We consider likelihood score-based methods for causal discovery in structural causal models. In particular, we focus on Gaussian scoring and analyze the effect of model misspecification in terms of non-Gaussian error distribution. We present a surprising negative result for Gaussian likelihood scoring in combination with nonparametric regression methods.

Keywords: graphical models, model misspecification, non-parametric regression, structural causal models

MSC 2020: 62D20

1 Introduction

We consider the problem of finding the causal structure of a set of random variables X_1, \dots, X_p . We assume that the data can be represented by a structural equation model whose structure is given by a directed acyclic graph (DAG), say, G^0 , with nodes $1, \dots, p$ and denote by $\text{PA}(j)$ the parents of j . Without further assumptions on the structural causal model, one can only find G^0 up to its Markov equivalence class. This can be achieved, for example, with the PC algorithm [1]. Its generality comes at the price of requiring conditional independence tests, which are a statistically hard problem.

Here, we focus on the so-called additive noise model (ANM)

$$X_j \leftarrow f_j(X_{\text{PA}(j)}) + \mathcal{E}_j \quad \forall j \in 1, \dots, p, \quad (1)$$

where the \mathcal{E}_j are mutually independent, centered random variables. This is a popular modeling assumption that allows for better identifiability guarantees [2,3].

For an arbitrary DAG G , let $\text{PA}^G(j)$ be the nodes from which a directed edge toward j starts. Define

$$\mathcal{E}_j^G = X_j - \mathbb{E}[X_j | X_{\text{PA}^G(j)}] \quad \forall j \in 1, \dots, p.$$

Obviously, $\mathcal{E}_j^{G^0} = \mathcal{E}_j$. Under not overly restrictive assumptions, it holds that $\mathcal{E}_1^G, \dots, \mathcal{E}_p^G$ are mutually independent only if $G \supseteq G^0$ [3]. Thus, an obvious approach to find G^0 is to loop over all possible graphs and test for independence of the residuals. Of course, this becomes infeasible when the dimensionality p grows.

A more reasonable algorithm based on greedy search is presented in the study by Peters et al. [3]. They also introduce regression with subsequent independence test (RESIT) that iteratively detects sink nodes. Finding the true DAG is guaranteed, assuming perfect regressors and independence tests. It involves $O(p^2)$ nonparametric independence tests, which might be computationally involved or lack power when the sample size is small.

* **Corresponding author: Christoph Schultheiss**, Seminar for Statistics, Department of Mathematics, ETH Zürich, Switzerland, e-mail: christoph.schultheiss@stat.math.ethz.ch

Peter Bühlmann: Seminar for Statistics, Department of Mathematics, ETH Zürich, Switzerland, e-mail: buhlmann@stat.math.ethz.ch

Instead of performing independence tests, one can compare the likelihood score of different graphs

$$\mathcal{L}(G) = \mathbb{E} \left[\log \left(\prod_{j=1}^p p_j^G(\mathcal{E}_j^G) \right) \right] = \sum_{j=1}^p \mathbb{E} [\log(p_j^G(\mathcal{E}_j^G))],$$

where p_j^G denotes the density of \mathcal{E}_j^G . Only for independent \mathcal{E}_j^G , their multivariate density factorizes as suggested. Therefore, the true DAG maximizes this quantity due to the properties of the Kullback-Leibler divergence. In practice, this comes with the additional difficulty of estimating the densities $p_j^G(\cdot)$, and one needs to add some penalization to prefer simpler graphs and avoid selecting complete graphs [4].

If one additionally assumes that the \mathcal{E}_j are marginally normally distributed $\mathcal{N}(0, \sigma_j^2)$, one can instead consider the Gaussian likelihood

$$\mathcal{L}^N(G) := \sum_{j=1}^p \left(-\log(\sigma_j^G) - \frac{1}{2} - \frac{1}{2} \log(2\pi) \right) = -\sum_{j=1}^p \log(\sigma_j^G) + C,$$

where $(\sigma_j^G)^2 = \mathbb{E}[(\mathcal{E}_j^G)^2]$ [5]. It holds, $\mathcal{L}(G) \geq \mathcal{L}^N(G)$ and $\mathcal{L}^N(G^0) = \mathcal{L}(G^0)$. Thus, the problem simplifies to finding the graph that leads to the lowest sum of log-variances.

Under such a normality assumption for \mathcal{E}_j and the $f_j(\cdot)$ in (1) being additive in their arguments, Bühlmann et al. [5] presents a causal discovery method that is consistent for high-dimensional data. To justify this approach for a broader class of error distributions, define

$$\Delta := \min_{G \not\supseteq G^0} \sum_{j=1}^p \log(\sigma_j^G) - \log(\sigma_j^{G^0}). \quad (2)$$

Then, one has to assume that

$$\Delta > 0. \quad (A1)$$

That is, the lowest possible expected negative Gaussian log-likelihood with any graph G that is not a superset of the true G^0 must be strictly larger than the expected negative Gaussian log-likelihood with the true graph G^0 . An argument for that assumption is that a true causal model should be easier to fit in some sense and thus also obtain a lower error variance. Using simple “non-pathological” examples, we demonstrate how this can easily be a fallacy when the true error distribution is misspecified. Thus, we advocate the need to be very careful when using Gaussian scoring with flexible nonparametric regression functions in causal discovery. The main part of this article considers theoretical population properties. We discuss some data applications in Section 5.1.

2 Data-generating linear model

We begin with data-generating linear models in which

$$f_j(X_{\text{PA}(j)}) = \sum_{k \in \text{PA}(j)} \beta_{jk} X_k.$$

For these, we find the explicit Theorem 1. The intuition for this result carries over to a range of nonlinear ANMs (1), especially when the causal effects are close to linear. We present according examples in Section 3.

If all \mathcal{E}_j in a data-generating linear model are Gaussian, i.e., X_1, \dots, X_p are jointly Gaussian, it is known that any causal order could induce the given multivariate normal distribution and obtain an optimal Gaussian score. Assuming that the distribution is faithful with respect to the true DAG G^0 , the set of the most sparse DAGs obtaining the optimal Gaussian score corresponds to the Markov equivalence class of G^0 [6]. Thus, one can obtain this Markov equivalence class by preferring more sparse DAGs in case of equal scores. In general, the single true DAG cannot be determined even if the full multivariate distribution is known. On the contrary, the linear non-Gaussian acyclic model introduced by Shimizu et al. [7] is known to

be identifiable. In such linear non-Gaussian models, algorithms designed for linear Gaussian models, e.g., the PC algorithm using partial correlation to assess conditional independence, do not use all the available information, but typically still provide the same guarantees since they depend on the covariance structure only. The covariance matrix of data generated by a linear causal model does not change when replacing a Gaussian \mathcal{E}_j by an additive error of the same variance but otherwise arbitrary distribution. Thus, assuming the faithfulness condition for the data-generating distribution, Gaussian scoring for linear causal models in the infinite sample limit leads to the true underlying Markov equivalence class even under misspecification of the error distribution.

If the data-generating model is not known to be linear such that nonparametric regression methods, or the conditional mean as their population version, are applied, this generalization does not hold true anymore, as laid out in the following theorem. Let π be a permutation on $\{1, \dots, p\}$ and G^π be the full DAG according to π , i.e., $\pi(k) \in \text{PA}^{G^\pi}(\pi(j))$ if and only if $k < j$.

Theorem 1. *Let X_1, \dots, X_p come from a linear model:*

$$X_j \leftarrow \sum_{k \in \text{PA}(j)} \beta_{jk} X_k + \mathcal{E}_j, \quad \text{with } \mathbb{E}[\mathcal{E}_j] = 0, \quad \mathbb{E}[\mathcal{E}_j^2] < \infty \quad \forall j \in 1, \dots, p, \quad (3)$$

with mutually independent \mathcal{E}_j . Then, for every possible permutation π ,

$$\sum_{j=1}^p \log(\sigma_j^{G^\pi}) - \log(\sigma_j^{G^0}) \leq 0.$$

That is, for every causal order, the corresponding full graph scores as least as well as the true causal graph. Furthermore,

$$\begin{aligned} & \sum_{j=1}^p \log(\sigma_j^{G^\pi}) - \log(\sigma_j^{G^0}) < 0 \quad \text{if} \\ & \exists j \in \{1, \dots, p\} : \quad \mathbb{E}[X_j | X_{\text{PA}^{G^\pi}(j)}] \neq X_{\text{PA}^{G^\pi}(j)}^\top \beta^{\pi,j}, \quad \text{where} \\ & \beta^{\pi,j} = \mathbb{E}[X_{\text{PA}^{G^\pi}(j)} X_{\text{PA}^{G^\pi}(j)}^\top]^{-1} \mathbb{E}[X_j X_{\text{PA}^{G^\pi}(j)}^\top] \quad \text{is the least squares parameter.} \end{aligned}$$

That is, for every causal order, the corresponding full graph scores strictly better than the true causal graph if at least one conditional expectation is nonlinear in the parental variables.

Apart from some pathological cases, the last condition holds for permutations that are not conformable with the true DAG unless all the \mathcal{E}_j are Gaussian. Thus, the gap condition (A1) does not hold true for this whole set of distributions, and, without Gaussianity, a wrong model would not only score equivalently but also be preferred. This is in stark contrast to results around Gaussian scoring when fitting linear models.

2.1 Illustrative examples

For illustrative purposes, we restrict ourselves to the two-variable case with $X_2 \leftarrow \beta X_1 + \mathcal{E}_2$. We have the true DAG $G^0 = \{X_1 \rightarrow X_2\}$ and define $G^\perp = \{X_1 \leftarrow X_2\}$. If $\beta X_1 \stackrel{\mathcal{D}}{=} \mathcal{E}_2$, it holds

$$X_2 = \mathbb{E}[X_2 | X_1] = \mathbb{E}[\beta X_1 + \mathcal{E}_2 | X_1] = \beta X_1, \quad \text{implying that} \quad \mathbb{E}[X_1 | X_2] = \frac{1}{\beta} X_2$$

such that $\Delta = 0$; see the definition in (2). In the two variable linear model, $\exp(\Delta)^2$ equals the ratio of the attainable mean squared error for the backward direction between the best-fitting unrestricted model and the best-fitting linear model.

First, we consider the analytically tractable case where $X_1 \stackrel{\mathcal{D}}{=} \mathcal{E}_2 \sim \text{Unif}[-1, 1]$. Then, for $\beta = \pm 1$, $\mathbb{E}[X_1 | X_2] = \pm X_2 / 2$ and $\Delta = 0$. For every other nonzero and bounded β , $\Delta < 0$ so a causal discovery method based on Gaussian scoring would – assuming a correct regression function estimator – wrongly claim $X_2 \rightarrow X_1$ in the large sample limit. We make things more explicit.

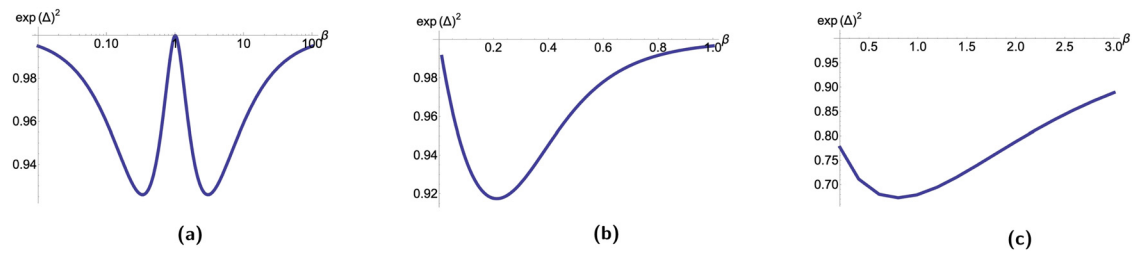


Figure 1: Two-variable linear model: effect of changing β . (a) $X_1 \stackrel{\mathcal{D}}{=} \mathcal{E}_2 \sim \text{Unif}[-1, 1]$. (b) $X_1 \sim \mathcal{N}(0, 1)$, $\mathcal{E}_2 \sim \text{Unif}[-1, 1]$. (c) $X_1 \sim \mathcal{N}(0, 1)$, $\mathcal{E}_2 + 1 \sim \chi_1^2$.

Proposition 1. Let $X_2 = \beta X_1 + \mathcal{E}_2$ with $X_1 \stackrel{\mathcal{D}}{=} \mathcal{E}_2 \sim \text{Unif}[-1, 1]$. Define $\gamma = \max\{|\beta|, |1/\beta|\}$. Then, the ratio of the variances is given as follows:

$$\exp(\Delta)^2 = \left(\prod_{j=1}^2 (\sigma_j^{G^\perp})^2 \right) / \left(\prod_{j=1}^2 (\sigma_j^{G^0})^2 \right) = \frac{(\gamma^2 + 1)(2\gamma - 1)}{2\gamma^3} := r(\gamma).$$

As argued above, $r(1) = 1$. For $|\beta| \rightarrow 0$, the ratio between the variance products approaches 1 as X_1 and X_2 become independent, and hence, both models perform equally well. For $|\beta| \rightarrow \infty$, the ratio approaches 1 as X_2 becomes a deterministic linear map of X_1 , and hence, the linear model is invertible. The ratio $r(\gamma)$ is minimized for $\gamma = 3$, with $r(3) \approx 0.93$ strictly decreasing for $\gamma \in [1, 3)$ and strictly increasing for $\gamma \in (3, \infty)$. This is shown in Figure 1(a).

Next, we consider a similar example but with $X_1 \sim \mathcal{N}(0, 1)$. Analytic expressions for $\mathbb{E}[X_1|X_2]$ and $\text{Var}(X_1|X_2)$ can be found in terms of the Gaussian cumulative distribution function. For $(\sigma_1^{G^\perp})^2 = \mathbb{E}[\text{Var}(X_1|X_2)]$, we invoke numerical integration. The obtained $\exp(\Delta)^2$ is shown in Figure 1(b). As elaborated before, it approaches 1 for $\beta \rightarrow 0$ and $\beta \rightarrow \infty$. In between, it is strictly less than 1 since $\mathbb{E}[X_1|X_2]$ is not linear in X_2 . The minimum of $\exp(\Delta)^2 \approx 0.92$ is obtained for $\beta \approx 0.21$.

Finally, we use an asymmetric error distribution instead, namely, $\mathcal{E}_2 \sim \chi_1^2 - 1$. Figure 1(c) shows that this leads to more extreme values of $\Delta < 0$. Notably, $\mathbb{E}[X_1|X_2]$ is not monotone in this case, so the best linear fit is not a good approximation.

3 Beyond a data-generating linear model

If all \mathcal{E}_j are Gaussian, we know that $\Delta = 0$ for linear conditional expectations in (1), but $\Delta > 0$ otherwise. Thus, involving nonlinearities enables the identifiability of the model.

For non-Gaussian \mathcal{E}_j in a linear model, $\Delta < 0$ holds true apart from some special cases, see Theorem 1. However, the intuition is that nonlinearities could be beneficial for identifiability. As the lower bound for Δ is negative and not 0, presumably a higher degree of nonlinearity might be necessary to achieve $\Delta > 0$. We analyze this with the following simple model:

$$X_2 \leftarrow \beta \text{Sign}(X_1) \frac{|X_1|^\nu}{\sqrt{\mathbb{E}[|X_1|^{2\nu}]}} + \mathcal{E}_2, \quad \nu > 0. \quad (4)$$

The normalization ensures that the variance of X_2 depends only on β . For $\nu = 1$, we obtain a linear model.

Consider the case $X_1 \sim \mathcal{N}(0, 1)$ and $\mathcal{E}_2 \sim \text{Unif}[-1, 1]$. Analytic expressions for $\mathbb{E}[X_1|X_2]$ and $\text{Var}(X_1|X_2)$ can be found in terms of the Gaussian cumulative distribution function and the Γ -function. For $(\sigma_1^{G^\perp})^2 = \mathbb{E}[\text{Var}(X_1|X_2)]$, we invoke numerical integration. The obtained $\exp(\Delta)^2$ is shown in Figure 2(a) for different values of β . It confirms the intuition, that sufficiently strong nonlinearity leads to $\Delta > 0$ even for non-Gaussian errors. For $\beta = 1$ and $\beta = 2$, the model becomes identifiable for most $\nu \neq 1$. For $\beta = 0.5$, stronger nonlinearities are necessary.

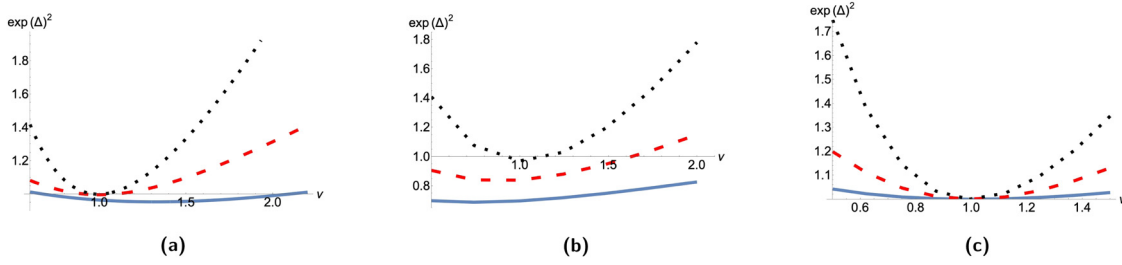


Figure 2: Two-variable nonlinear model (4): effect of changing ν for $\beta = 0.5$ (solid blue curve), $\beta = 1$ (dashed red curve), and $\beta = 2$ (dotted black curve). (a) $X_1 \sim \mathcal{N}(0, 1)$, $\mathcal{E}_2 \sim \text{Unif}[-1, 1]$. (b) $X_1 \sim \mathcal{N}(0, 1)$, $\sqrt{6}\mathcal{E}_2 + 1 \sim \chi_1^2$. (c) $X_1 \stackrel{\mathcal{D}}{=} \sqrt{3}\mathcal{E}_2 \sim \mathcal{N}(0, 1)$.

Furthermore, the minimum Δ is not obtained for $\nu = 1$ but at $\nu \approx 1.32$. Thus, not every nonlinear model is better identifiable than the corresponding linear model.

As in the linear case, we consider the effect of an asymmetric error distribution, namely, a scaled and centered chi-squared distribution. We show this in Figure 2(b). The factor $\sqrt{6}$ leads to the different lines corresponding to the same signal-to-noise ratios. As expected, higher degrees of nonlinearity are necessary to obtain a positive Δ .

We show the behavior for a correctly specified model in Figure 2(c). For $\nu = 1$, $\Delta = 0$, as implied by the unidentifiability result; otherwise, $\Delta > 0$. For $\nu \neq 1$, higher signal-to-noise ratios lead to a more distinct $\Delta > 0$.

3.1 Monotonicity of $f_2(\cdot)$

The nonlinearities discussed here are designed to be slight deviations from the linear model and, thus, chosen to be strictly monotone. Notably, for non-monotone functions, the intuition that the anti-causal model is harder to fit is more applicable. In particular, if X_1 is centered and symmetric and $f_2(\cdot)$ is an even function, it holds $\mathbb{E}[X_1|X_2] \equiv 0$. Then,

$$(\sigma_1^G)^2 = \text{Var}(X_1) \quad \text{and} \quad \exp(\Delta)^2 = \frac{\text{Var}(X_1)\text{Var}(X_2)}{\text{Var}(X_1)\text{Var}(\mathcal{E}_2)} > 1. \quad (5)$$

Thus, the gap condition (A1) is satisfied regardless of the distribution of \mathcal{E}_2 as long as X_1 and X_2 have a finite variance.

4 Heteroskedastic noise model

A simple extension of model (1) that has recently gained some attention, is the heteroskedastic noise model, also referred to as location-scale noise model

$$X_j \leftarrow f_j(X_{\text{PA}(j)}) + g_j(X_{\text{PA}(j)})\mathcal{E}_j \quad \text{with} \quad \mathbb{E}[\mathcal{E}_j] = 0, \quad \mathbb{E}[\mathcal{E}_j^2] = 1 \quad \forall j \in 1, \dots, p,$$

with some nonnegative functions $g_j(\cdot)$ [8–10]. It comes with similar identifiability guarantees as the ANM when testing for mutual independence between the \mathcal{E}_j . Let accordingly

$$\begin{aligned} f_j^G(X_{\text{PA}^G(j)}) &= \mathbb{E}[X_j|X_{\text{PA}^G(j)}], \\ g_j^G(X_{\text{PA}^G(j)})^2 &= \mathbb{E}[(X_j - f_j^G(X_{\text{PA}^G(j)}))^2|X_{\text{PA}^G(j)}], \quad \text{and} \\ \mathcal{E}_j^G &= \frac{X_j - f_j^G(X_{\text{PA}^G(j)})}{g_j^G(X_{\text{PA}^G(j)})} \end{aligned}$$

be the conditional means, conditional variances, and residuals according to any potentially wrong DAG G .

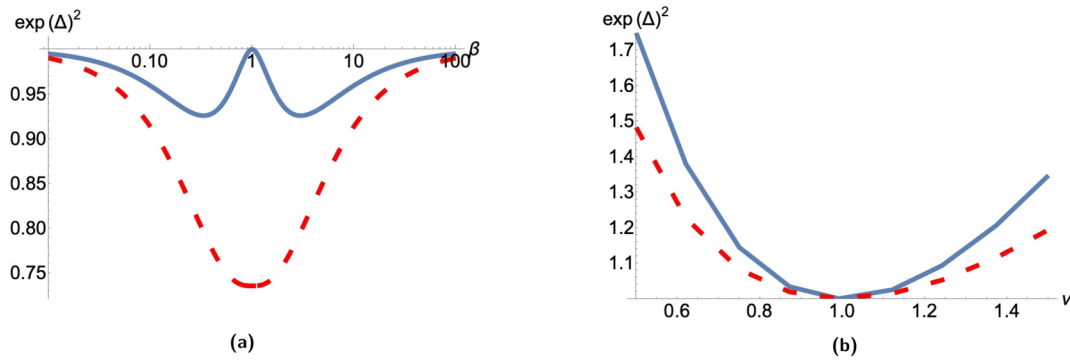


Figure 3: Two variable additive model: fitting homoskedastic models (solid blue curve) versus heteroskedastic models (dashed red curve). (a) Model (3) with $X_1 \stackrel{\mathcal{D}}{=} \varepsilon_2 \sim \text{Unif}[-1, 1]$. (b) Model (4) with $X_1 \stackrel{\mathcal{D}}{=} \sqrt{3} \varepsilon_2 \sim \mathcal{N}(0, 1)$ and $\beta = 2$.

Then, one obtains

$$\begin{aligned} \mathcal{L}(G) &= \sum_{j=1}^p \mathbb{E} \left[\log \left(\frac{p_j^G(\varepsilon_j^G)}{g_j^G(X_{\text{PA}^G(j)})} \right) \right] \\ \mathcal{L}^N(G) &:= - \sum_{j=1}^p \mathbb{E} [\log(g_j^G(X_{\text{PA}^G(j)}))] + C = - \sum_{j=1}^p \frac{1}{2} \mathbb{E} [\log(g_j^G(X_{\text{PA}^G(j)})^2)] + C \\ &\geq - \sum_{j=1}^p \frac{1}{2} \log(\mathbb{E}[g_j^G(X_{\text{PA}^G(j)})^2]) + C = - \sum_{j=1}^p \log(\sigma_j^G) + C. \end{aligned}$$

Thus, when fitting heteroskedastic models, the score can only be increased compared to the homoskedastic fit. This can further increase the difficulty of finding the correct direction under non-Gaussian noise. Even if the true forward model is homoskedastic, i.e., $g_j(\cdot) \equiv \sigma_j$, the backward model is typically heteroskedastic and profits from this new score. For example, in the setup of Figure 1(a), Δ would be negative even for $\beta = 1$. If one allows to fit heteroskedastic models, a result analogous to Theorem 1 exists. A negative gap is obtained unless all conditional expectations are linear and all conditional variances are constant for the wrong causal order.

In Figure 3, we review the examples from Figures 1(a) and 2(c) and see how allowing for a heteroskedastic fit makes the problem harder. For the sake of comparison, we look at $\exp(\Delta)^2$ although it does not have the same simple interpretation in the heteroskedastic case.

In terms of the location-scale noise model, the data-generating model, as shown in Figure 3(a) is unidentifiable as $X_1|X_2$ is uniformly distributed, i.e., its distribution is independent of X_2 apart from location and scale. This does not contradict the identifiability results as they are derived for random variables with full support in \mathbb{R} .

5 Discussion

5.1 Data applications

For an extensive comparison between methods relying on Gaussian scoring and nonparametric independence tests in ANMs or heteroskedastic noise models, we refer to the study by Immer et al. [10], where several fitting methods are considered, combined with both approaches, and evaluated on a variety of benchmark cause and effect pairs. Those pairs include both real and artificial data. For some of the considered data sources, using independence tests clearly improved the success rate for inferring the causal direction as compared to using the Gaussian score.

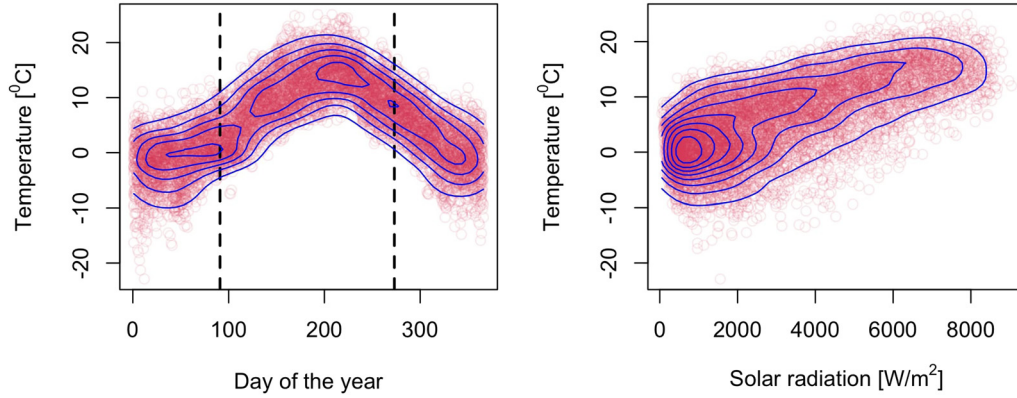


Figure 4: Scatter plot and contour lines of the density estimate for two selected pairs from the Tübingen data.

Let us consider two specific examples of the Tübingen data by Mooij et al. [11]. Details on the data can be found in Section D.11 of their article. Both examples have the temperature as the effect variable, while the cause is the day of the year or the intensity of the solar radiation, respectively. The corresponding scatter plots and the contour lines of the density estimates are shown in Figure 4. It is evident that in neither case the cause variable is normally distributed, the days are perfectly uniformly distributed while solar radiation is right-skewed. Therefore, the assumptions for Gaussian scoring to infer the true causal direction are not fulfilled. For the first dataset, we restricted the numerical analysis to the time frame 1st April to 30th September (31st March to 29th September in leap years) to circumvent the issue that the data are circular. This is indicated by the black dotted lines.

To evaluate the Gaussian scores, we estimate the conditional expectation for either direction with a smoothing spline. In the first case, the causal effect is non-monotone, and the conditional mean in the anti-causal direction is not very informative to predict the day of the year. Therefore, we obtain the correct causal direction with Gaussian scoring even though the assumptions are not fulfilled. We obtain the data estimate $\exp(\Delta)^2 = 1.48$.

The effect of solar radiation on the temperature appears to be monotone, which makes the conditional expectation in the anti-causal direction more informative. In addition, it seems that the conditional expectation in the causal direction is not so far from being linear. This indeed makes the Gaussian scoring algorithm prefer the wrong direction. The estimate is $\exp(\Delta)^2 = 0.91$. Similarly, it fails for the first dataset when considering the first 183 days of the year instead ($\exp(\Delta)^2 = 0.99$).

With RESIT relying on independence testing, we see for both datasets that the hypothesis of residuals being independent of the predictor is rejected in either direction. This indicates that the ANM in (1) is not rich enough to explain the data. However, applying Algorithm 1 from the study by Peters et al. [3], which minimizes the estimated dependence between predictor and residuals, finds the true causal direction for both data pairs.

5.2 Conclusion

We discuss causal discovery in structural causal models using Gaussian likelihood scoring and analyze the effect of model misspecification.

In the case where the data-generating distribution comes from a linear structural equation model and linear regression functions are used for estimation, the following holds. When the true error distribution is Gaussian, one can only identify the Markov equivalence class of the underlying data-generating DAG. The same holds true when the error distribution is non-Gaussian, but one wrongly relies on a Gaussian error distribution for estimation.

Thus, popular algorithms like the greedy equivalence search [12] for Gaussian models or the PC algorithm [1] assessing partial correlation are potentially conservative and only infer the Markov equivalence class when the error distributions are non-Gaussian, as they do not exploit the maximal amount of information. But they are safe to use within the domain of data-generating linear structural equation models. We prove here that this fact does not necessarily hold true when invoking nonparametric regression estimation. Especially, if the true causal model is linear or just “slightly nonlinear,” one would systematically obtain the wrong causal direction under error misspecification. As optimizing Gaussian scores is the same as optimizing ℓ_2 -loss, regressors that are more flexible than necessary for the causal model support anti-causal decisions. The intuition carries over when allowing for the flexibility of heteroskedastic error terms. If the true causal model has homoskedastic additive errors, fitting heteroskedastic models will increase the range of setups where misspecified Gaussian scoring chooses anti-causal relationships.

To overcome these issues, one could rely on general nonparametric independence tests, either between the different residuals or between residuals and predictors. Of course, this comes at a higher computational cost and potentially lower sample efficiency in cases where Gaussian scoring works, including in the presence of non-monotonic causal effects.

Funding information: The project leading to this application has received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation program (grant agreement no. 786461).

Author contributions: All authors have accepted responsibility for the entire content of this manuscript and approved its submission. The authors applied the SDC approach for the sequence of authors.

Conflict of interest: The authors state that there is no conflict of interest.

Data availability statement: The datasets analyzed during this study are available in the database with cause-effect pairs, <https://webdav.tuebingen.mpg.de/cause-effect/>. We consider pairs 42 and 77.

References

- [1] Spirtes P, Glymour CN, Scheines R, Heckerman D. Causation, prediction, and search. Cambridge, MA, USA: MIT Press; 2000.
- [2] Hoyer P, Janzing D, Mooij JM, Peters J, Schölkopf B. Nonlinear causal discovery with additive noise models. *Adv Neur In*. 2008;21:689–96.
- [3] Peters J, Mooij J, Janzing D, Schölkopf B. Causal discovery with continuous additive noise models. *J Mach Learn Res*. 2014;15:2009–53.
- [4] Nowzohour C, Bühlmann P. Score-based causal learning in additive noise models. *Statistics*. 2016;50(3):471–85.
- [5] Bühlmann P, Peters J, Ernest J. CAM: Causal additive models, high-dimensional order search and penalized regression. *Ann Statist*. 2014;42(6):2526–56.
- [6] Zhang J, Spirtes P. Detection of unfaithfulness and robust causal inference. *Minds Mach*. 2008;18(2):239–71.
- [7] Shimizu S, Hoyer PO, Hyvärinen A, Kerminen A, Jordan M. A linear non-Gaussian acyclic model for causal discovery. *J Mach Learn Res*. 2006;7(10):2003–30.
- [8] Strobl EV, Lasko TA. Identifying patient-specific root causes with the heteroscedastic noise model. 2022. arXiv: <http://arXiv.org/abs/arXiv:220513085>.
- [9] Xu S, Mian OA, Marx A, Vreeken J. Inferring cause and effect in the presence of heteroscedastic noise. In: *International Conference on Machine Learning*. PMLR; 2022. p. 24615–630.
- [10] Immer A, Schultheiss C, Vogt JE, Schölkopf B, Bühlmann P, Marx A. On the identifiability and estimation of causal location-scale noise models. 2022. arXiv: <http://arXiv.org/abs/arXiv:221009054>.
- [11] Mooij JM, Peters J, Janzing D, Zscheischler J, Schölkopf B. Distinguishing cause from effect using observational data: methods and benchmarks. *J Mach Learn Res*. 2016;17(1):1103–204.
- [12] Chickering DM. Optimal structure identification with greedy search. *J Mach Learn Res*. 2002;3:507–54.

Appendix

A Proofs

A.1 Proof of Theorem 1

It is well known that for jointly Gaussian variables X_1, \dots, X_p , every possible causal order can induce the multivariate distribution with a suitable linear model. As the sum of log-variances determines the Kullback-Leibler divergence in the Gaussian case, this sum must be equal for all these linear models that induce the same multivariate distribution.

For every possible multivariate distribution with existing and bounded moment matrix $\Sigma^X = \mathbb{E}[\mathbf{X}\mathbf{X}^\top]$, which the assumed model has, the linear least squares parameter and corresponding residual variances using arbitrary sets of regressor covariates are completely determined by the moment matrix. Therefore, the residual variances correspond to those of multivariate Gaussian data. Accordingly, one can obtain the same sum of log-variances as for the true model using the best linear predictors for every possible permutation. This proves the non-strict inequality in the theorem.

If for some variable j the conditional expectation given its parents (in the DAG G^π) is not a linear function, the linear model cannot be optimal in terms of residual variance. Therefore, $\sigma_j^{G^\pi}$ in an unrestricted model is lower than the residual standard error of the best fitting linear model, and the score is further improved. Hence, the inequality is strict as soon as there exists at least one such variable.

A.2 Proof of Proposition 1

The variances of X_1 and X_2 as well as \mathcal{E}_2 follow directly from the properties of the uniform distribution.

$$\begin{aligned} (\sigma_1^{G^0})^2 &= \text{Var}(X_1) = \frac{1}{3}, \quad (\sigma_2^{G^0})^2 = \text{Var}(\mathcal{E}_2) = \frac{1}{3}, \quad \text{and} \\ (\sigma_2^{G^\perp})^2 &= \text{Var}(X_2) = \beta^2 \text{Var}(X_1) + \text{Var}(\mathcal{E}_2) = \frac{1}{3}(\beta^2 + 1). \end{aligned}$$

The last term requires some more work. Due to the symmetry, we can assume without loss of generality that $\beta \geq 0$. For the densities, we obtain

$$\begin{aligned} f_{X_1}(x_1) &= \frac{1}{2} \mathbb{1}_{\{|x_1| \leq 1\}}, \quad f_{X_2|X_1}(x_2, x_1) = \frac{1}{2} \mathbb{1}_{\{|x_2 - \beta x_1| \leq 1\}}, \quad \text{and} \\ f_{X_2}(x_2) &= \int \frac{1}{4} \mathbb{1}_{\{|x_1| \leq 1\}} \mathbb{1}_{\{|x_2 - \beta x_1| \leq 1\}} dx_1 = \frac{1}{4} \left(\min \left\{ 1, \frac{x_2 + 1}{\beta} \right\} - \max \left\{ -1, \frac{x_2 - 1}{\beta} \right\} \right) := \frac{1}{4}(a - b). \end{aligned}$$

For notational simplicity, we define random variables A and B with realizations a and b . We obtain the moments

$$\begin{aligned} \mathbb{E}[X_1|X_2] &= \int x_1 f_{X_1|X_2}(x_1, X_2) dx_1 = \frac{\int \frac{x_1}{4} \mathbb{1}_{\{|x_1| \leq 1\}} \mathbb{1}_{\{|X_2 - \beta x_1| \leq 1\}} dx_1}{\int \frac{1}{4} \mathbb{1}_{\{|x_1| \leq 1\}} \mathbb{1}_{\{|X_2 - \beta x_1| \leq 1\}} dx_1} = \frac{A^2 - B^2}{2(A - B)} = \frac{1}{2}(A + B), \\ \mathbb{E}[X_1^2|X_2] &= \int x_1^2 f_{X_1|X_2}(x_1, X_2) dx_1 = \frac{\int \frac{x_1^2}{4} \mathbb{1}_{\{|x_1| \leq 1\}} \mathbb{1}_{\{|X_2 - \beta x_1| \leq 1\}} dx_1}{\int \frac{1}{4} \mathbb{1}_{\{|x_1| \leq 1\}} \mathbb{1}_{\{|X_2 - \beta x_1| \leq 1\}} dx_1} = \frac{A^3 - B^3}{3(A - B)} \\ &= \frac{1}{3}(A^2 + AB + B^2), \\ \text{Var}(X_1|X_2) &= \mathbb{E}[X_1^2|X_2] - \mathbb{E}[X_1|X_2]^2 = \frac{1}{12}(A - B)^2. \end{aligned}$$

Finally, we are interested in

$$\begin{aligned}(\sigma_1^{G^\perp})^2 &= \mathbb{E}[\text{Var}(X_1|X_2)] = \mathbb{E}\left[\frac{1}{12}(A-B)^2\right] = \int_{-1-\beta}^{1+\beta} \frac{1}{12}(a-b)^2 f_{X_2}(x_2) dx_2 \\&= \int_{-1-\beta}^{1+\beta} \frac{1}{48}(a-b)^3 dx_2 = 2 \int_0^{1+\beta} \frac{1}{48}(a-b)^3 dx_2.\end{aligned}$$

Assume first $\beta \geq 1$. Then, $b = (x_2 - 1)/\beta$ in the area of integration. For $x_2 \geq \beta - 1$, it holds $a = 1$.

$$\begin{aligned}(\sigma_1^{G^\perp})^2 &= \int_0^{1+\beta} \frac{1}{24}(a-b)^3 dx_2 = \int_0^{\beta-1} \frac{1}{24}\left(\frac{2}{\beta}\right)^3 dx_2 + \int_{\beta-1}^{1+\beta} \frac{1}{24}\left(\frac{\beta+1-x_2}{\beta}\right)^3 dx_2 \\&= \frac{1}{24\beta^3} \left(8(\beta-1) + \int_0^2 u^3 du \right) = \frac{1}{24\beta^3} (8\beta - 4) = \frac{1}{6\beta^3} (2\beta - 1),\end{aligned}$$

where we applied the change of variable $u = \beta + 1 - x_2$ to simplify the integration. Inserting residual variances with $\gamma = \beta$, the proposition's statement follows.

Alternatively, if $\beta < 1$, $a = 1$ in the interval of integration. For $x_2 < 1 - \beta$, it holds $b = -1$.

$$\begin{aligned}(\sigma_1^{G^\perp})^2 &= \int_0^{1+\beta} \frac{1}{24}(a-b)^3 dx_2 = \int_0^{1-\beta} \frac{1}{24}(2)^3 dx_2 + \int_{1-\beta}^{1+\beta} \frac{1}{24}\left(\frac{\beta+1-x_2}{\beta}\right)^3 dx_2 \\&= \frac{1}{24} \left(8(1-\beta) + \frac{1}{\beta^3} \int_0^{2\beta} u^3 du \right) = \frac{1}{24} (8 - 4\beta) = \frac{1}{6} (2 - \beta),\end{aligned}$$

where we applied the change of variable $u = \beta + 1 - x_2$ to simplify the integration. Inserting all the residual variances with $\gamma = 1/\beta$, the proposition's statement follows.

B Derivations for the figures

Assume model (4), which has model (3) as a special case for $\nu = 1$. With $X_1 \sim \mathcal{N}(0, 1)$, the normalization is

$$V(\nu) := \mathbb{E}[|X_1|^{2\nu}] = \frac{2^\nu \Gamma\left(\nu + \frac{1}{2}\right)}{\sqrt{\pi}}.$$

As before,

$$\begin{aligned}(\sigma_1^{G^0})^2 &= \text{Var}(X_1), \quad (\sigma_2^{G^0})^2 = \text{Var}(\mathcal{E}_2) \quad \text{and} \\(\sigma_2^{G^\perp})^2 &= \text{Var}(X_2) = \beta^2 + \text{Var}(\mathcal{E}_2).\end{aligned}$$

B.1 Gaussian and uniform

For $\mathcal{E}_2 \sim \text{Unif}[-1, 1]$, X_1 is constrained to lie within

$$\text{Sign}((X_2 - 1)\sqrt{V(\nu)}/\beta)(X_2 - 1)\sqrt{V(\nu)}/\beta^{1/\nu} \quad \text{and} \quad \text{Sign}((X_2 + 1)\sqrt{V(\nu)}/\beta)(X_2 + 1)\sqrt{V(\nu)}/\beta^{1/\nu},$$

i.e., given X_2 , it is a truncated standard Gaussian. The conditional mean and variance follow from standard theory about truncated Gaussian random variables. For simplicity, call the upper bound A and the lower bound B . The density of X_2 is then given as follows:

$$f_2(x_2) = \int_b^a \frac{1}{2} \phi(x_1) dx_1 = \frac{1}{2} (\Phi(a) - \Phi(b)).$$

Finally,

$$(\sigma_1^{G^\perp})^2 = \mathbb{E}[\text{Var}(X_1|X_2)]$$

is obtained by numerically integrating over (a sufficient part of) the real line.

B.2 Two Gaussian random variables, or Gaussian and χ_1^2

Except for $V(v)$, all quantities are obtained by brute force numerical integration.

B.3 Two uniform random variables with heteroskedastic fitting

We can mainly follow the derivation in Appendix A.2. Instead of $(\sigma_1^{G^\perp})^2$, we need $\exp(\mathbb{E}[\log(\text{Var}(X_1|X_2))])$, which is obtained by numerical integration.