

Research Article

David Strieder* and Mathias Drton

Confidence in causal inference under structure uncertainty in linear causal models with equal variances

<https://doi.org/10.1515/jci-2023-0030>

received May 16, 2023; accepted September 21, 2023

Abstract: Inferring the effect of interventions within complex systems is a fundamental problem of statistics. A widely studied approach uses structural causal models that postulate noisy functional relations among a set of interacting variables. The underlying causal structure is then naturally represented by a directed graph whose edges indicate direct causal dependencies. In a recent line of work, additional assumptions on the causal models have been shown to render this causal graph identifiable from observational data alone. One example is the assumption of linear causal relations with equal error variances that we will take up in this work. When the graph structure is known, classical methods may be used for calculating estimates and confidence intervals for causal-effects. However, in many applications, expert knowledge that provides an a priori valid causal structure is not available. Lacking alternatives, a commonly used two-step approach first learns a graph and then treats the graph as known in inference. This, however, yields confidence intervals that are overly optimistic and fail to account for the data-driven model choice. We argue that to draw reliable conclusions, it is necessary to incorporate the remaining uncertainty about the underlying causal structure in confidence statements about causal-effects. To address this issue, we present a framework based on test inversion that allows us to give confidence regions for total causal-effects that capture both sources of uncertainty: causal structure and numerical size of non-zero effects.

Keywords: confidence intervals, causal-effects, linear structural equation models, equal error variances, graphical models

MSC 2020: 62D20, 62H22

1 Introduction

Questions about causal relations are at the heart of many research problems. Researchers across various fields, ranging from economics to biology and medicine, are interested in distinguishing causes from effects to predict the outcome of interventions in complex systems. While controlled experiments provide an established method for inferring causal relations, in many applied settings, interventional studies are not feasible for ethical or cost considerations. Thus, inferring causal relationships based solely on observational data is an important task that the field of causal discovery addresses.

In the last few decades, causal discovery has gained popularity, with much fundamental research being done on the topics of identification and estimation [1,2]. Identifiability results clarify when it is theoretically

* **Corresponding author: David Strieder**, TUM School of Computation, Information and Technology, Munich Center for Machine Learning, Technical University of Munich, 80333 Munich, Germany, e-mail: david.strieder@tum.de

Mathias Drton: TUM School of Computation, Information and Technology, Munich Center for Machine Learning, Technical University of Munich, Munich, Germany, e-mail: mathias.drton@tum.de

possible to go beyond mere correlations. More specifically, they clarify under which circumstances, assumptions lead to unique identification of interventional distributions from a given joint probability distribution. Moreover, numerous structure learning algorithms have been proposed that infer the underlying causal structure based on observed data. Knowledge of the underlying causal structure then allows researchers to reason about the effect of interventions in complex systems, and given a causal structure, standard statistical methods may be applied to estimate the involved causal-effects and provide an uncertainty assessment for the numerical size of the effects. However, in the typical applied setting, the exact causal structure is (at least partially) unknown. In order to draw reliable conclusions and make calibrated confidence statements, the remaining uncertainty about the underlying causal structure needs to be incorporated in causal inference.

Without prior expert knowledge about the underlying causal structure, a commonly employed “naïve” approach would involve splitting the task. In the first step, a causal learning algorithm estimates the underlying causal structure. In the second step, within the inferred model, confidence intervals for the causal parameters can be calculated using classical statistical inference methods. However, the resulting intervals are not valid if the causal learning algorithm infers the wrong causal structure. Such a two-step method conditions away the uncertainty in the causal structure arising from the data-driven model choice, and the classical “double-dipping” problem occurs since both causal structure and the effect are estimated from one data set. Thus, the approach is overly optimistic in its conclusion about the strength and existence of causal-effects, and its confidence regions do not achieve the desired coverage probability, especially under high uncertainty with respect to the underlying structure.

Despite the popularity and growing interest in causal discovery, we are unaware of prior attempts to rigorously account for uncertainty in the causal structure when providing confidence statements about causal-effects. Here, we use the term “confidence” in the technical sense for a set with a given desired frequentist coverage probability. A completely different approach for uncertainty quantification that we do not consider here is to form Bayesian credible sets. In fact, a number of authors have used Bayesian approaches for uncertainty quantification in causal discovery, but primarily with an emphasis on the uncertainty in the graphical structure as opposed to causal-effects (see [3–5] for three selected examples).

In this article, we want to shed light on the problem by highlighting the challenges of the task of handling structure uncertainty and by proposing a solution for constructing confidence intervals for total causal-effects that capture both types of uncertainty: the uncertainty regarding the causal structure resulting from a data-driven model selection and the uncertainty about the numerical size of the effect. Our solution is feasible for problems of small to moderate scale – e.g., our simulation experiments will consider problems with up to 12 variables. However, already at this scale, rigorously accounting for structure uncertainty means to account for more than 10^{26} possible structures, reflecting the fast superexponential growth with the number of nodes in the causal graph.

In Section 2 of this article, we review the considered framework of structural causal models [6,7]. More specifically, in order to ensure identifiability from observational data alone, we consider a restricted class of causal models, namely, linear structural equation models (LSEMs) with Gaussian errors and equal variances [8], and review the definition of a total causal-effect, which is the target of interest in this study. We then generalize and improve the ideas introduced for bivariate problems by Strieder et al. [9] and present a general framework for constructing confidence intervals for total causal-effects in multivariate data (Sections 3 and 4). Furthermore, we present computational shortcuts and analyze the performance of the introduced framework in numerical experiments (Section 5). In Section 6, we discuss extensions of the framework to partially quantify the uncertainty over causal structures.

2 Background

This section reviews the LSEMs and the total causal-effect, which is the target of interest in this study.

2.1 LSEMs

Structural causal models constitute a commonly used tool to study causal relationships and represent noisy causal relations among a set of interacting variables. Each variable is modeled as a function of a subset of other variables, its causes, and a stochastic error term. The causal perspective emerges from viewing those relations as making assignments rather than representing mathematical equations. The left-hand side, the effect, is assigned the value specified on the right-hand side, the causes. Externally varying the values on the right-hand side results in a change on the left-hand side, but not vice versa. This reflects the asymmetry inherent in cause–effect relations.

We assume access to observational data in the form of a sample of independent copies of a random vector $X = (X_1, \dots, X_d)$; without loss of generality, we assume the random vector to have zero mean. Without any further assumptions on the structural causal model, e.g., the involved functions, identification of the underlying causal structure is only possible up to a Markov equivalence class. Therefore, to ensure unique identifiability, following a line of research initiated by Peters and Bühlmann [8], we focus on linear relations and normally distributed errors with equal variances by assuming that X solves the equation system:

$$X_j = \sum_{i \neq j} \beta_{j,i} X_i + \varepsilon_j, \quad j = 1, \dots, d, \quad (1)$$

where $B = [\beta_{j,i}]_{j,i=1}^d$ are unknown parameters that represent the direct causal-effects between the variables, and ε_j are independent, normally distributed error terms:

$$\varepsilon_j \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2), \quad j = 1, \dots, d,$$

with common, unknown variance $\sigma^2 > 0$. Each specific LSEM restricts a subset of the direct effects $\beta_{j,i}$ to be zero. Thus, LSEMs are naturally represented by a (minimal) directed graph $G(B)$ with vertex set $V = \{1, \dots, d\}$. In the associated directed graph, the edge set equals the support of B ; thus, a missing edge $i \rightarrow j$ indicates that $\beta_{j,i} = 0$. As in related work, we assume the underlying directed graph to be acyclic (DAG), which entails that B is permutation similar to a triangular matrix and System (1) admits the unique solution $X = (I_d - B)^{-1}\varepsilon$, with I_d denoting the $d \times d$ identity matrix. Incorporating the equal variance assumption, the covariance matrix of X is seen to be

$$\mathbb{E}[XX^T] = \sigma^2(I_d - B)^{-1}(I_d - B)^{-T}.$$

In the sequel, we will use the following graphical concepts. Each DAG admits a (not necessarily unique) topological ordering of its vertices, i.e., a permutation π of $\{1, \dots, d\}$ such that the existence of a directed path from node $\pi(i)$ to node $\pi(j)$ implies $i < j$. We write $i <_G j$ if node i precedes node j in such a causal ordering of the graph. If the DAG contains an edge from node i to node j , then node i is a parent of node j , and we denote the set of all parents of node j with $p(j)$. It has been shown that under the aforementioned modeling assumptions with equal error variances, i.e., homoscedasticity across all interacting variables, the causal ordering is identifiable and corresponds to an ordering of conditional variances [10,11]. Hence, causal-effects can also be uniquely identified.

2.2 Causal-effects

Informally put, we are interested in estimating how much the value of X_j changes if we externally intervene in the system and set the value of X_i to x_i . Such an intervention represents a change in the true data-generating process and, thus, a modification in our probabilistic model. In the structural equation framework, the effect of such an external intervention can be expressed by replacing the i th equation with $X_i = x_i$ and is denoted as $\text{do}(X_i = x_i)$.


In this work, we seek to construct confidence intervals for the total causal-effect $C(i \rightarrow j)$ in LSEMs. This effect is formally defined as the unit change in the expectation of X_j with respect to an intervention in X_i , i.e.,

$$C(i \rightarrow j) := \frac{d}{dx_i} \mathbb{E}[X_j | \text{do}(X_i = x_i)].$$

Remark 2.1. Note that the total causal-effect in an LSEM (B, σ^2) can be intuitively expressed based on the respective underlying DAG $G(B)$. The effect $C(i \rightarrow j)$ equals the sum of all products of edge weights along all directed paths from i to j . Furthermore, the effect is zero if there does not exist a directed path from i to j . This follows from the path properties of the matrix $(I_d - B)^{-1}$ that computes the solution X of the structural equations from the independent errors.

Example 2.2. Consider the following LSEM and the corresponding (minimal) DAG:

$$\begin{aligned} X_1 &= \beta_{1,3}X_3 + \varepsilon_1 \\ X_2 &= \beta_{2,1}X_1 + \beta_{2,4}X_4 + \beta_{2,5}X_5 + \varepsilon_2 \\ X_3 &= \varepsilon_3 \\ X_4 &= \beta_{4,1}X_1 + \varepsilon_4 \\ X_5 &= \beta_{5,3}X_3 + \varepsilon_5 \end{aligned}$$



The total causal-effect is given by the sum of all products of edge weights along all directed paths. For example, we obtain $C(1 \rightarrow 2) = \beta_{2,1} + \beta_{4,1}\beta_{2,4}$, whereas $C(2 \rightarrow 1) = 0$.

Our aim is to construct a confidence interval for the total causal-effect $C(i \rightarrow j)$ when the underlying causal structure is not known and has to be learned. We stress that, under the identifiability assumption of an underlying Gaussian LSEM with equal error variances, this is a well-defined problem. In fact, every admissible distribution can be represented by a single (minimal) DAG. This DAG then uniquely defines the total causal-effect.

A “naïve” two-step approach that first learns a causal structure and then calculates confidence intervals in the inferred model does not incorporate the uncertainty regarding the underlying causal structure. To respect this uncertainty in the calculations of a valid confidence interval, it is necessary to take all possible structures, i.e., graphs, into account. As a result, the target quantity has a composite structure, and thus, obtaining finite sample distributions of estimators is difficult.

In many situations where the sampling distribution of estimators is unobtainable, bootstrapping and other resampling methods can offer an appealing framework to construct confidence intervals. At first sight, resampling may appear to provide a simple solution to the problem under discussion. This solution relies on computing for each resampled data set a causal-effect estimate. To obtain this estimate, one may compose a consistent model selection method that learns the causal graph with a consistent causal-effect estimator in the learned model. Note that this composition is well defined due to the identifiability of LSEMs with equal error variances. However, as a statistic, this composition lacks smoothness, and thus, bootstrapping procedures may fail severely [12,13]. We emphasize that Drton and Williams [13] demonstrate that even in low dimensions, the asymptotic behavior of confidence intervals and bootstrap tests is difficult to predict for complex composite settings as encountered here. The subtleties stem from the fact that the set of covariance matrices associated with at least one possible DAG form a union of smooth manifolds. This union contains singularities at the intersections of the manifolds, which is highlighted in the bivariate case by Strieder et al. [9]. They demonstrated that singular points arise at the intersections of the two models. Consequently, bootstrapping cannot be used to capture model uncertainty correctly. This failure of the bootstrapping methods can also be observed in our simulations in Section 5.3.

In conclusion, there is a need to develop new frameworks for constructing confidence intervals that circumvent the problems that originate from the non-smooth nature of the target of interest.

3 Confidence intervals via test inversion

In the remainder of this article, we develop a framework that circumvents the problems posed by the non-smooth nature of the total causal-effect and provides valid confidence intervals for causal inference under structure uncertainty. In this section, we introduce our main ansatz, the test inversion approach, and derive the concrete hypothesis that subsequently needs to be tested.

3.1 Inversion of tests

Our main idea is to leverage the classical duality between statistical hypothesis tests and confidence regions in the following way. If we perform valid hypothesis tests for all attainable causal-effects, then all values that we cannot reject form a confidence region for the total causal-effect. More specifically, let $\alpha \in (0, 1)$ be a fixed significance level. We perform (asymptotically) valid-level α tests for the hypothesis of a fixed total causal-effect ψ for all attainable values of the causal-effect, i.e., for all $\psi \in \mathbb{R}$, we test

$$H_0 : C(1 \rightarrow 2) = \psi,$$

with a test whose acceptance region we denoted by $A(\psi)$. Then, a (asymptotically) valid $(1 - \alpha)$ -confidence region for the total causal-effect $C(1 \rightarrow 2)$ for the data X is given by:

$$C(X) := \{\psi \in \mathbb{R} : X \in A(\psi)\}.$$

A proof of this classical result can be found, e.g., in [14, Theorem 9.2.2].

With the aforementioned perspective, our problem is shifted to developing suitable hypothesis tests for all attainable causal-effects. Due to the lack of concrete knowledge about the model, all possible causal graphs must be respected. Thus, we have to construct hypothesis tests in a composite setting that remains challenging. In the next section, we explicitly describe the hypotheses that need to be tested to construct confidence regions for the total causal-effect.

3.2 Hypothesis of a fixed causal-effect

Recalling the assumptions of an underlying LSEM with homoscedastic errors across all interacting variables, we consider the following task. We assume that the given data consist of random vectors $X^{(1)}, \dots, X^{(n)}$ drawn independently from a centered multivariate normal distribution $N(0, \Sigma_0)$. The distribution's density $p(x|\Sigma_0)$ factorizes according to an unknown DAG while satisfying the assumption of equal error variances. To account for the uncertainty in the causal structure, we need to test at level α whether any multivariate centered normal distribution that factorizes to any DAG under the equal variance assumption allows for a total causal-effect of size ψ , and repeat that test for all values $\psi \in \mathbb{R}$. In light of the normality assumption, it is natural to parameterize our statistical model in terms of the covariance matrix Σ . The additional assumption of homoscedasticity across all interacting variables restricts the set of possible covariance matrices. Furthermore, as we detail now.

Given a fixed DAG G , under the assumption of equal error variances, all possible normal distributions with a density that factorizes according to G are defined by the family:

$$\mathcal{N}(G) := \{N(0, \Sigma) : \exists B \in \mathbb{R}^G, \sigma^2 > 0, \text{ with } \Sigma = \sigma^2(I_d - B)^{-1}(I_d - B)^{-T}\}, \quad (2)$$

where $\mathbb{R}^G = \{B \in \mathbb{R}^{d \times d} : \beta_{j,i} = 0 \text{ if } i \notin p(j)\}$. Considering and combining all possible DAG models, our overall model corresponds to the following union of sets of covariance matrices:

$$\mathcal{M} := \bigcup_{G \in \mathcal{G}(d)} \{\Sigma \in \text{PD}(d) : N(0, \Sigma) \in \mathcal{N}(G)\}, \quad (3)$$

where $\mathcal{G}(d)$ is the set of all DAGs with d nodes.

Remark 3.1. If G^* is a supergraph of G , then the set $\mathcal{N}(G)$ is a subset of $\mathcal{N}(G^*)$. Hence, in the definition of \mathcal{M} , we only need to consider the union over all complete DAGs with d nodes, i.e., all possible orderings.

The following proposition shows that we can identify the model \mathcal{M} , which is a subset of all positive definite matrices satisfying the assumption of equal error variances, solely based on polynomial constraints in the entries of the covariance matrix. Here, and in the remainder of this article, we write $\Sigma_{i,j|p(i)}$ for the conditional covariance matrix, i.e.,

$$\Sigma_{j,i|p(i)} := \Sigma_{j,i} - \Sigma_{j,p(i)}(\Sigma_{p(i),p(i)})^{-1}\Sigma_{p(i),i}.$$

Proposition 3.2. For a covariance matrix $\Sigma \in \text{PD}(d)$, the following statements are equivalent:

- (i) $\Sigma \in \mathcal{M}$,
- (ii) There exists a complete DAG G and $\sigma^2 > 0$ such that for all $k = 1, \dots, d$

$$\sigma^2 = \Sigma_{k,k|p(k)} = \Sigma_{k,k} - \Sigma_{k,p(k)}(\Sigma_{p(k),p(k)})^{-1}\Sigma_{p(k),k}. \quad (4)$$

Proof. “ \Rightarrow ” Let $\Sigma \in \mathcal{M}$. Then, there exists a DAG G such that $N(0, \Sigma) \in \mathcal{N}(G)$. Considering Remark 3.1, we can assume without loss of generality that G is complete. Straightforward calculations with the multivariate normal density then imply that the variances of each node conditioned on its parents are equal to σ^2 , or see a combinatorial argument as in [15, Equation (7.4)]. Thus, for all $k = 1, \dots, d$,

$$\sigma^2 = \text{Var}(X_k | X_{p(k)}) = \Sigma_{k,k} - \Sigma_{k,p(k)}(\Sigma_{p(k),p(k)})^{-1}\Sigma_{p(k),k}.$$

“ \Leftarrow ” Since G is complete, the multivariate normal distribution $N(0, \Sigma)$ factorizes according to the graph G . This implies the existence of a matrix $B \in \mathbb{R}^G$ and a positive definite diagonal matrix $\Omega \in \mathbb{R}^{d \times d}$ with $\Sigma = (I_d - B)^{-1}\Omega(I_d - B)^{-T}$. Equation (4) implies that $\Omega_{k,k} = \Sigma_{k,k|p(k)}$ is equal for all $k = 1, \dots, d$, and the claim follows. \square

Remark 3.3. Note the identifiability result under homoscedasticity across all interacting variables, e.g., Theorem 1 in [10]. Each $\Sigma \in \mathcal{M}$ corresponds to a unique (B, σ^2) and, thus, a unique minimal DAG $G(B)$. As explained in Remark 3.1, the union in the definition of \mathcal{M} (3) is not necessarily disjoint, and analogously, the (complete) graph satisfying Condition 2 in Proposition 3.2 is not necessarily unique. Nevertheless, since B is unique, all graphs satisfying Condition 2 in Proposition 3.2 share the same edge set with non-zero edge coefficients and correspond to all valid causal orderings of the minimal DAG. Thus, every “additional” parent of node i in a complete DAG G from Condition 2 compared to the unique minimal DAG $G(B)$ is conditionally independent of node i given $p_{G(B)}(i)$. This follows by d-separation since every (undirected) path between node i and an additional parent either traverses the parents $p_{G(B)}(i)$ in the minimal DAG or contains a collider c with $i <_{G(B)} c$, and thus, this collider has no descendent in $p_{G(B)}(i)$.

Our testing problem is the following classical statistical problem. We are concerned with a parametric model \mathcal{M} , parameterized by the covariance matrix, which is an intricate union of subsets of the cone of positive definite matrices given by polynomial constraints. Within that statistical model, we need to test for all attainable values of the total causal-effect ψ . In the following proposition, we show that the parameter of interest is a function of the covariance matrix. Thus, the considered statistical testing problem can be fully parameterized in terms of the covariance matrix.

Proposition 3.4. Let $\Sigma \in \mathcal{M}$. Then, the total causal-effect $C(i \rightarrow j)$ is given by:

$$C(i \rightarrow j) = \Sigma_{j,i|p(i)}(\Sigma_{i,i|p(i)})^{-1},$$

with parent set $p(i)$ based on any corresponding DAG G from Proposition 3.2.

Proof. The case $j <_G i$ is trivially true, since in this case, the total causal-effect $C(i \rightarrow j)$ is zero and, considering G is complete, the parent set $p(i)$ contains the node j , which yields $\Sigma_{j,ilp(i)} = 0$.

Assume $i <_G j$. Since the set $p(i)$ satisfies the back-door criterion with respect to i and j , we have

$$C(i \rightarrow j) = \frac{d}{dx_i} \mathbb{E}[X_j | \text{do}(X_i = x_i)] = \frac{d}{dx_i} \mathbb{E}[X_j | X_i = x_i, X_{p(i)}].$$

It follows that the total causal-effect $C(i \rightarrow j)$ is the regression coefficient of X_i when regressing X_j on $(X_i, X_{p(i)})$, which is given by $\Sigma_{j,ilp(i)}(\Sigma_{i,ilp(i)})^{-1}$.

We emphasize that this expression does not depend on the specific choice of the DAG G . All graphs satisfying Condition 2 in Proposition 3.2 share the same edge set with non-zero edge weights, and thus, the parent set $p(i)$ always satisfies the back-door criterion in the corresponding unique minimal DAG (see Remark 3.3). \square

Using Proposition 3.4, we can define the hypothesis space $\mathcal{M}_\psi \subset \mathcal{M}$, given by all covariance matrices $\Sigma \in \mathcal{M}$ that allow for a total causal-effect of size ψ via

$$\mathcal{M}_\psi = \{\Sigma \in \text{PD}(d) : \exists \text{ complete DAG } G \text{ and } \sigma^2 > 0 \text{ with } \psi\sigma^2 = \Sigma_{j,ilp(i)} \text{ and } \sigma^2 = \Sigma_{k,k|p(k)} \quad \forall k = 1, \dots, d\}.$$

Put together, we have to solve the statistical testing problem

$$H_0^{(\psi)} : \Sigma \in \mathcal{M}_\psi \text{ against } H_1 : \Sigma \in \mathcal{M} \setminus \mathcal{M}_\psi, \quad (5)$$

for all attainable $\psi \in \mathbb{R}$. Then, we invert the tests to construct valid confidence intervals for the total causal-effect. Similar to the union \mathcal{M} , we can write the hypothesis space as a union of single hypotheses over all (complete) DAGs, i.e., $H_0^{(\psi)} = \bigcup_{G \in \mathcal{G}(d)} H_0^{(\psi)}(G)$ with single hypotheses $H_0^{(\psi)}(G) : \Sigma \in \mathcal{M}_\psi(G)$, where


$$\mathcal{M}_\psi(G) = \{\Sigma \in \text{PD}(d) : \exists \sigma^2 > 0 \text{ with } \psi\sigma^2 = \Sigma_{j,ilp(i)} \text{ and } \sigma^2 = \Sigma_{k,k|p(k)} \quad \forall k = 1, \dots, d\}. \quad (6)$$

Remark 3.5. It is clear that for $\psi \neq 0$, we only need to consider the union over all graphs with $i <_G j$. In contrast, for a zero effect, we have to consider all graphs, since the effects of multiple directed paths might “cancel” each other. Furthermore, for graphs with $j <_G i$, the constraint given by the size zero effect is not informative, since $j \in p(i)$ and consequently, $\Sigma_{j,ilp(i)} = 0$. Therefore, the corresponding hypothesis space has higher dimension.

Example 3.6. Consider the following LSEM and the corresponding (minimal) DAG:

There exist two directed paths from node 1 to node 2 in the corresponding (minimal) DAG, and thus, $1 <_G 2$ in all valid causal orderings. However, the effect of the two paths “cancel” each other, such that $C(1 \rightarrow 2) = 0.25 - 0.5^2 = 0$.

$$\begin{aligned} X_1 &= -0.5X_3 + \varepsilon_1 \\ X_2 &= 0.25X_1 + 0.5X_4 + 0.25X_5 + \varepsilon_2 \\ X_3 &= \varepsilon_3 \\ X_4 &= -0.5X_1 + \varepsilon_4 \\ X_5 &= 0.5X_3 + \varepsilon_5 \end{aligned}$$



4 Testing for total causal-effects

For constructing confidence intervals for the total causal-effect, we invert tests of Hypothesis (5). In this section, we present two concrete tests based on likelihood ratio theory: (i) a classical constrained likelihood ratio test [16] and (ii) an approach based on the recently proposed split likelihood ratio test [17].

4.1 Constrained likelihood ratio test

The first option we consider is to form a classical likelihood ratio statistic for the testing Problem (5). Writing ℓ_n for the Gaussian log-likelihood function based on the sample $X^{(1)}, \dots, X^{(n)}$, the likelihood ratio statistic for (5) is

$$\check{\lambda}_n^{(\psi)} := 2 \left(\sup_{\Sigma \in \mathcal{M}} \ell_n(\Sigma) - \sup_{\Sigma \in \mathcal{M}_\psi} \ell_n(\Sigma) \right). \quad (7)$$

In regular settings, the statistic may be compared to a chi-square quantile for an asymptotic test. Here, however, asymptotic distribution theory is difficult because both the null hypothesis and the alternative are intricate unions and do not satisfy the usual regularity assumptions. In the simplified bivariate setting, Strieder et al. [9] avoid the difficulties posed by the existence of singularities by testing modified problems in two different approaches. In the first approach, the hypothesis and alternative are relaxed to test an ordering of (conditional) variances. However, it is not feasible to optimize this modified testing problem in problems beyond the bivariate case. Furthermore, the mixture weights of the arising limit distribution, which is a mixture of chi-square distributions, are difficult to be determined in higher dimensions. The second approach relaxes the alternative to an unconstrained Gaussian model. In this modification, the statistic (7) is changed to a larger statistic for which the optimization of \mathcal{M} is replaced by an optimization over the positive definite cone $\text{PD}(d)$. This approach, however, may cause problems in the test inversion approach, as the union of tested null hypotheses no longer coincides with the alternative. In particular, small confidence regions may arise due to model misspecification.

In this article, we will derive our confidence region for the total causal-effect from tests that reject for too large values of the original statistic $\check{\lambda}_n^{(\psi)}$ from (7). However, we will consider relaxing the alternative \mathcal{M} and use the theory of intersection union tests to obtain a simple yet effective upper bound on the distribution of $\check{\lambda}_n^{(\psi)}$.

To obtain this upper bound, first note that the hypothesis we need to test is an intricate union of single hypotheses given by $\bigcup_{G \in \mathcal{G}(d)} \mathcal{M}_\psi(G)$. While we will later show that every single hypothesis defines a smooth submanifold, their union is not necessarily a smooth submanifold again. Nevertheless, the original test statistic (7) can then be written as $\check{\lambda}_n^{(\psi)} = \min_{G \in \mathcal{G}(d)} \check{\lambda}_n^{(\psi)}(G)$, with

$$\check{\lambda}_n^{(\psi)}(G) := 2 \left(\sup_{\Sigma \in \mathcal{M}} \ell_n(\Sigma) - \sup_{\Sigma \in \mathcal{M}_\psi(G)} \ell_n(\Sigma) \right). \quad (8)$$

Thus, the original test statistic can be upper-bounded by every single test statistic $\check{\lambda}_n^{(\psi)}(G)$, and a valid-level α test of the union of single hypotheses can be constructed by testing every single hypothesis. We then reject the union null hypothesis if we reject all single hypotheses. A formal proof of this classical theory of intersection union tests can be found, e.g., in [14, Theorem 8.3.23].

In the second step, we relax the alternative to an unconstrained Gaussian model to obtain the following larger test statistic:

$$\lambda_n^{(\psi)}(G) := 2 \left(\sup_{\Sigma \in \text{PD}(d)} \ell_n(\Sigma) - \sup_{\Sigma \in \mathcal{M}_\psi(G)} \ell_n(\Sigma) \right).$$

Relaxing the alternative to an optimization over the entire cone of positive definite matrices $\text{PD}(d)$ potentially increases the achieved likelihood and consequently upper-bounds each single test statistic $\check{\lambda}_n^{(\psi)}(G)$.

Combining both approximations yields an upper bound of the original test statistic $\check{\lambda}_n^{(\psi)}$ and, thus, a simple way to construct asymptotically conservative hypothesis tests. In the following, we state our main result by turning them into conservative confidence intervals for the total causal-effect $C(i \rightarrow j)$ via test inversion.

Theorem 4.1. *Let $\alpha \in (0, 1)$. Then, an asymptotic $(1 - \alpha)$ -confidence set for the causal-effect $C(i \rightarrow j)$ is given by:*

$$C = \{\psi \in \mathbb{R} : \check{\lambda}_n^{(\psi)}(i <_G j) \leq \chi_{d,1-\alpha}^2\} \cup \{0 : \check{\lambda}_n^{(0)}(j <_G i) \leq \chi_{d,1-\alpha}^2\},$$

where $\check{\lambda}_n^{(\psi)}(i <_G j) = \min_{G \in \mathcal{G}(d) : i <_G j} \check{\lambda}_n^{(\psi)}(G)$ and $\check{\lambda}_n^{(0)}(j <_G i) = \min_{G \in \mathcal{G}(d) : j <_G i} \check{\lambda}_n^{(0)}(G)$.

Remark 4.2. Note that we only need to consider complete DAGs, which drastically reduces the number of single hypothesis tests that need to be evaluated and thus the computational burden. Furthermore, we need to respect graphs with $j <_G i$ only for zero-sized effects. Therefore, our proposed confidence regions may contain a non-zero interval and an isolated zero. In contrast to the bivariate case, however, zero effects can also correspond to orderings with $i <_G j$ due to cancellation (see also Remark 3.5).

Proof. Let $\psi \in \mathbb{R}$ and $\Sigma \in \mathcal{M}_\psi$. Since $\mathcal{M}_\psi = \bigcup_{G \in \mathcal{G}(d)} \mathcal{M}_\psi(G)$, there exists a complete graph G , such that $\Sigma \in \mathcal{M}_\psi(G)$. If $i <_G j$, it follows from the introduced upper bound and Lemma 4.3

$$P_\Sigma(\check{\lambda}_n^{(\psi)} > \chi_{d,1-\alpha}^2) \leq P_\Sigma(\check{\lambda}_n^{(\psi)}(G) > \chi_{d,1-\alpha}^2) \leq P_\Sigma(\lambda_n^{(\psi)}(G) > \chi_{d,1-\alpha}^2) \rightarrow \alpha.$$

We obtain an analogous result for $j <_G i$ with Lemma 4.4. Using the test inversion approach yields the claim. \square

In the following, we derive the asymptotic distribution of the upper bound $\lambda_n^{(\psi)}(G)$ of our test statistic under every single hypothesis $\mathcal{M}_\psi(G)$. As previously indicated, the cases $i <_G j$ versus $j <_G i$ define different dimensional hypothesis spaces; thus, we consider both instances separately.

First, we consider the case $i <_G j$ and assume without loss of generality that $(1, 2, \dots, d)$ is the causal ordering. In this case, the single hypothesis can be written as $\mathcal{M}_\psi(G) = \{\Sigma \in \text{PD}(d) : f_\psi(\Sigma|G) = 0\}$, with

$$f_\psi(\Sigma|G) = \begin{pmatrix} \Sigma_{j,1|p(i)} - \psi \Sigma_{1,1} \\ \Sigma_{2,2|p(2)} - \Sigma_{1,1} \\ \vdots \\ \Sigma_{d,d|p(d)} - \Sigma_{1,1} \end{pmatrix}.$$

The relaxed model space of all positive definite matrices $\text{PD}(d)$ can be identified with an open subspace in $\mathbb{R}^{\frac{1}{2}(d^2+d)}$. The hypothesis space $\mathcal{M}_\psi(G)$ then defines a $\frac{1}{2}(d^2 - d)$ -dimensional submanifold of $\mathbb{R}^{\frac{1}{2}(d^2+d)}$, since the Jacobian of f_ψ has full rank d . To see that, note its (non-zero) triangular structure for derivatives for $(\Sigma_{k,k})_{k=2,\dots,d}$ in all rows except the first row. The first row, however, has a non-zero derivative for $\Sigma_{j,i}$, and zero derivatives for $(\Sigma_{k,k})_{k=j,\dots,d}$.

Using the fact that this single hypothesis defines a smooth submanifold, we determine the limit distribution of the upper bound $\lambda_n^{(\psi)}(G)$.

Lemma 4.3. *Let $G \in \mathcal{G}(d)$ be a DAG with $i <_G j$ and $\psi \in \mathbb{R}$. Under the hypothesis $H_0^{(\psi)}(G)$, the likelihood ratio test statistic $\lambda_n^{(\psi)}(G)$ satisfies*

$$\lambda_n^{(\psi)}(G) \xrightarrow{\mathcal{D}} \chi_d^2.$$

Proof. Since the hypothesis space $\mathcal{M}_\psi(G)$ defines a $\frac{1}{2}(d^2 - d)$ -dimensional submanifold of $\mathbb{R}^{\frac{1}{2}(d^2+d)}$, the tangent cone at every $\Sigma \in \mathcal{M}_\psi(G)$ is a $\frac{1}{2}(d^2 - d)$ -dimensional linear subspace of $\mathbb{R}^{\frac{1}{2}(d^2+d)}$ and the claim follows. For details on this classical result, we refer to the study of Drton [18]. \square

In the case $j <_G i$, the causal effect from i to j can only be zero. Therefore, we solely need to test the single hypothesis $H_0^{(0)}(G)$. We assume again without loss of generality that $(1, \dots, d)$ is the underlying causal order. The polynomial constraint in Hypothesis (6) corresponding to the fixed-size effect is not informative, as the node j is included in $p(i)$, and thus, $\Sigma_{j, i|p(i)} = 0$. Therefore, the hypothesis $H_0^{(0)}(G)$ is solely defined by $d - 1$ polynomial constraints corresponding to the assumption of equal error variances, i.e., in this case, we have $\mathcal{M}_0(G) = \{\Sigma \in \text{PD}(d) : f_0(\Sigma|G) = 0\}$, where

$$f_0(\Sigma|G) := \begin{pmatrix} \Sigma_{2,2|p(2)} - \Sigma_{1,1} \\ \vdots \\ \Sigma_{d,d|p(d)} - \Sigma_{1,1} \end{pmatrix}.$$

Thus, similar to the previous case, we obtain a chi-square distribution as the limit distribution, however with fewer degrees of freedom.

Lemma 4.4. *Let $G \in \mathcal{G}(d)$ be a DAG with $j <_G i$. Under Hypothesis $H_0^{(0)}(G)$, the likelihood ratio test statistic $\lambda_n^{(0)}(G)$ satisfies*

$$\lambda_n^{(0)}(G) \xrightarrow{\mathcal{D}} \chi_{d-1}^2.$$

Proof. Analogously to Lemma 4.3. Note that the Jacobian has full rank $d - 1$ due to its triangular structure. \square

With Theorem 4.1, we defined an asymptotically valid confidence interval for the total causal-effect that is able to capture both types of uncertainty: the uncertainty about the causal structure and the uncertainty about the numerical size of the effect. Besides the theoretical guarantees for the asymptotic coverage of the proposed confidence set, we demonstrate in the simulation section, that even in finite samples, the confidence set achieves the desired coverage probability. We emphasize again that our new approach uses the original test statistic (7) with an upper bound on its distribution. Thus, we avoid the issues that may arise under model misspecification in the (bivariate) method by Strieder et al. [9] due to contrasting a hypothesis restricted to equal error variances with a general alternative.

In the following, we propose an alternative framework for constructing hypothesis tests for (5) based on the theory of universal inference [17]. While this framework leads to more conservative intervals, it comes with a finite sample guarantee.

4.2 Split likelihood ratio test

Likelihood ratio tests provide a powerful tool for constructing hypothesis tests that are calibrated via asymptotic distribution theory. However, in this context, the needed distribution-theoretic insights are difficult to obtain and we adopted stochastic upper bounds. An interesting alternative is to conduct the needed tests in the framework of universal inference that was introduced by Wasserman et al. [17]. Their method uses a modification of the likelihood ratio statistic, called split likelihood ratio, based on a data-splitting approach. Due to the independence of the split data, one can apply a universal critical value, instead of relying on asymptotic distribution theory. Type I error control is then guaranteed by an application of Markov's inequality. In the following, we use their methodology for constructing conservative confidence regions for total causal-effects with finite sample guarantee.

To construct a split likelihood ratio test for the test Problem (5), we proceed as follows. First, we split the data into two subsets $D_0 = \{X^{(1)}, \dots, X^{(k)}\}$ and $D_1 = \{X^{(k+1)}, \dots, X^{(n)}\}$ and define the log-likelihood functions:

$$\ell_j(\Sigma) := \sum_{X^{(i)} \in D_j} \log p(X^{(i)}|\Sigma), \quad j = 0, 1,$$

based on D_0 and D_1 , respectively. Let $\tilde{\Sigma}_1 := \operatorname{argmax}_{\Sigma \in \mathcal{M}} \ell_1(\Sigma)$ be the maximum-likelihood estimator of Σ restricted to LSEMs with equal error variances based on D_1 . Then, the split likelihood ratio test statistic is defined as:

$$\tilde{\lambda}_n^{(\psi)} = 2(\ell_0(\tilde{\Sigma}_1) - \sup_{\Sigma \in \mathcal{M}_\psi} \ell_0(\Sigma)).$$

The insight of Wasserman, Ramdas, and Balakrishnan [17] is that for any fixed $\Sigma^* \in \text{PD}(d)$, it holds that

$$\mathbb{E}_\Sigma \left[\prod_{X^{(i)} \in D_0} \frac{p(X^{(i)}|\Sigma^*)}{p(X^{(i)}|\Sigma)} \right] \leq \int \prod_{X^{(i)} \in D_0} p(X^{(i)}|\Sigma^*) = 1,$$

and, thus, Markov's inequality implies that for any $\alpha \in (0, 1)$, under $H_0^{(\psi)} : \Sigma \in \mathcal{M}_\psi$,

$$\begin{aligned} P_\Sigma(\tilde{\lambda}_n^{(\psi)} > -2\log \alpha) &\leq \alpha \mathbb{E}_\Sigma \left[\prod_{X^{(i)} \in D_0} \frac{p(X^{(i)}|\tilde{\Sigma}_1)}{\sup_{\Sigma \in \mathcal{M}_\psi} p(X^{(i)}|\Sigma)} \right] \\ &\leq \alpha \mathbb{E}_\Sigma \left[\mathbb{E}_\Sigma \left[\prod_{X^{(i)} \in D_0} \frac{p(X^{(i)}|\tilde{\Sigma}_1)}{\sup_{\Sigma \in \mathcal{M}_\psi} p(X^{(i)}|\Sigma)} \middle| D_1 \right] \right] \leq \alpha. \end{aligned} \quad (9)$$

Thus, the decision rule

$$\text{Reject } H_0^{(\psi)} \quad \text{if } \tilde{\lambda}_n^{(\psi)} > -2\log \alpha$$

constitutes a valid-level α test. This test holds level in finite samples and without any regularity conditions. Thus, we can define the following confidence region for the total causal-effect based on the split likelihood ratio test.

Theorem 4.5. *Let $\alpha \in (0, 1)$. Then, a $(1 - \alpha)$ -confidence set for the causal-effect $C(i \rightarrow j)$ is given by:*

$$C = \{\psi \in \mathbb{R} : \tilde{\lambda}_n^{(\psi)}(i <_G j) \leq -2\log(\alpha)\} \cup \{0 : \tilde{\lambda}_n^{(0)}(j <_G i) \leq -2\log(\alpha)\},$$

where $\tilde{\lambda}_n^{(\psi)}(i <_G j) = \min_{G \in \mathcal{G}(d) : i <_G j} \tilde{\lambda}_n^{(\psi)}(G)$ and $\tilde{\lambda}_n^{(0)}(j <_G i) = \min_{G \in \mathcal{G}(d) : j <_G i} \tilde{\lambda}_n^{(0)}(G)$ and

$$\tilde{\lambda}_n^{(\psi)}(G) = 2(\ell_0(\tilde{\Sigma}_1) - \sup_{\Sigma \in \mathcal{M}_\psi(G)} \ell_0(\Sigma)).$$

Remark 4.6. Similar to the asymptotic confidence set given by Theorem 4.1, we need to consider all graphs when testing for zero effects. Thus, the confidence region may consist of a non-zero interval and a disconnected zero that reflects the remaining uncertainty about the existence of an effect.

Proof. This result follows directly from the test inversion approach and (9). Note that for non-zero effects, we only need to consider graphs with $i <_G j$. \square

5 Computational details and numerical experiments

In this section, we present computational details as well as shortcuts in our proposed algorithm that significantly improve the computation time. We also present a simulation study in which we analyze and compare the coverage probabilities and performance of the different proposed testing procedures.

5.1 Algorithm and computational shortcuts

We focus on the algorithm for calculating the likelihood ratio-based confidence regions introduced in Theorem 4.1, denoted with LRT. By adjusting the critical values in the algorithm and incorporating the data-splitting procedure, the split likelihood ratio-based confidence regions can be computed analogously. As the number of possible underlying causal structures, i.e., the number of DAGs, grows superexponentially with the number of nodes, we need to carefully use combinatorial shortcuts in order to be able to compute our proposed confidence intervals for the total causal-effect.

Our algorithm, presented as Algorithm 1, consists of two subroutines. The first routine `possible.order` (Algorithm 2) immediately rejects all implausible causal orderings and, thus, reduces the set of possible causal orderings in which we subsequently test for all attainable effects with the second routine `testeffect` (Algorithm 3). Carefully reducing the space of plausible causal orderings drastically decreases the computation time and allows us to compute our proposed confidence intervals, which take uncertainty over all possible DAGs into account, in a reasonable time for a moderate (but far from trivial) number of involved variables. Precise computation times can be found in Section 5.3, which presents the results of a simulation study.

Algorithm 1 Returns LRT confidence region for $C(i \rightarrow j)$

Input: data, level α , stepsize s
`(orderings, likelihood, effect) \leftarrow possible.order(data, α)` \triangleright all plausible orderings
`L1 \leftarrow likelihood[1]` \triangleright maximum likelihood alternative
`startvalues \leftarrow unique(effect)`
while startvalues is not empty **do** \triangleright start with effects from unrestricted MLEs
 `(leftbound, rightbound) \leftarrow min(startvalues)`
 `(left, right) \leftarrow TRUE`
 while left **do**
 `leftbound \leftarrow leftbound $- s$`
 `left \leftarrow test.effect(data, leftbound, orderings, L1, α)` \triangleright test for size leftbound effect
 end while
 while right **do**
 `rightbound \leftarrow rightbound $+ s$`
 `right \leftarrow test.effect(data, rightbound, orderings, L1, α)`
 end while
 `append (leftbound $+ \frac{s}{2}$, rightbound $- \frac{s}{2}$) to intervals` \triangleright possibly multiple intervals
 remove all values smaller than rightbound from startvalues
end while
`zero \leftarrow test.effect(data, 0, orderings, L1, α)` \triangleright test for size zero effect
return (zero, intervals)

In order to reduce the space of plausible graphical structures as much as possible, we note again that we do not need to consider all possible DAGs but merely complete ones, represented by all possible causal orderings. This immensely decreases the worst-case number of single hypotheses we need to perform to reject an effect of size ψ to d factorial. Our subroutine `possible.order` (Algorithm 2) further reduces the number of plausible causal orderings as follows.

First, we quickly approximate the maximum likelihood under the alternative with the maximum likelihood $\widehat{L1}$ obtained under the ordering obtained from the method in the study by Chen et al. [10] that sorts conditional variances. Subsequently, we calculate the maximum likelihood of all possible orderings without restrictions on the causal-effect and immediately reject implausible orderings using the (largest) critical value $\chi^2_{d,1-\alpha}$. Note that we can search through the space of orderings recursively starting from the first node and immediately reject (partial) orderings if the corresponding (partial) maximum likelihood already exceeds the threshold. Furthermore, for the subsequent testing procedure, the specific ordering of the parent set $p(i)$ is not relevant. Thus, for all orderings that only differ in the specific ordering of $p(i)$, we subsequently just test with the parent order that achieves the maximum likelihood. Additionally, we only need to test for a zero effect in the maximum-likelihood ordering out of all orderings with $j <_{\mathcal{G}} i$. Finally, we include an attainable effect in our confidence set if it cannot be rejected for at least one of the possible causal orderings using the subroutine `testeffect`, Algorithm 3. Therefore, we do not need to continue testing for a specific value of the effect if we find an ordering such that we cannot reject the value. Thus, first sorting the causal orderings by their maximum likelihood further improves the computation time.

Algorithm 2 possible.order() returns plausible causal orderings

Input: data, level α
crit.val \leftarrow qchisq($1 - \alpha, d$) $\triangleright 1 - \alpha$ quantile of χ_d^2 -distr.
 $\hat{\Sigma} \leftarrow$ cov(data)
for $1 : d$ **do**
 append which.min($\hat{\Sigma}_{\cdot, \cdot | \text{var.order}}$) to var.order \triangleright ordering of cond. variances
end for
 $\widehat{L1} \leftarrow$ get.likelihood(var.order) \triangleright (unrestricted) maximum likelihood
for $1 : d$ **do**
 for order in orderings **do**
 for v in setdiff($1 : d, \text{order}$) **do**
 part.order \leftarrow c(order, v)
 likelihood \leftarrow get.likelihood(part.order) \triangleright (partial) maximum likelihood
 if $2(\widehat{L1} - \text{likelihood}) \leq \text{crit.val}$ **then**
 append part.order to orderings
 if v equals i **then**
 effect $\leftarrow \hat{\Sigma}_{i,j|\text{order}} / \hat{\Sigma}_{i,i|\text{order}}$ \triangleright causal-effect of maximum likelihood
 end if
 end if
 end for
 remove order from orderings \triangleright order only relevant after i
 choose order with max likelihood out of all orders in orderings that do not contain i and are defined on the same set of nodes
end for \triangleright keep only max likelihood for zero effect
choose order with max likelihood out of all orders in orderings with $j <_{\mathcal{G}} i$
sort orderings for decreasing likelihood
return (orderings, likelihood, effect)

Algorithm 3. test.effect() tests causal-effect $C(i \rightarrow j) = \psi$

Input: data, causal-effect ψ , orderings, likelihood alternative L1, level α
for order in orderings **do**
 if $j <_{\mathcal{G}} i$ **then**
 crit.val \leftarrow qchisq($1 - \alpha, d - 1$)
 if ψ equals zero **then**
 likelihood \leftarrow get.likelihood(order) \triangleright (unrestricted) maximum likelihood
 else
 likelihood $\leftarrow -\infty$ \triangleright reject
 end if
 else
 crit.val \leftarrow qchisq($1 - \alpha, d$)
 likelihood \leftarrow get.likelihood(order, ψ) \triangleright (restricted) maximum likelihood
 end if
 if $2(L1 - \text{likelihood}) \leq \text{crit.val}$ **then**
 return TRUE
 end if
end for
return FALSE

We start testing within the plausible causal orderings with the minimal causal-effect that corresponds to the maximum-likelihood estimates. Then, we step-wise decrease the value until we reject in all orderings. The last value we did not reject is a lower bound for our confidence region. We repeat the procedure with increasing steps until we reject in all orderings and the last value we did not reject is an upper bound for this part of our confidence region. If there remain causal-effects from the maximum-likelihood estimates that exceed this upper bound, we repeat this procedure. Thus, using this strategy, we might obtain a confidence region that consists of multiple intervals. Finally, we test whether we have to include zero in our confidence interval, i.e., whether there remains uncertainty about the existence of an effect.

5.2 Maximum-likelihood estimation

Implementing the proposed algorithm involves the routine `get.likelihood`, which maximizes the centered Gaussian (log-)likelihood function $\ell_n(\Sigma)$ for a given causal ordering. The function is given by:

$$\frac{2}{n}\ell_n(\Sigma) = -\log \det(2\pi\Sigma) - \text{tr}(\Sigma^{-1}\hat{\Sigma}),$$

where $\hat{\Sigma} = \frac{1}{n}\sum_{l=1}^n X^{(l)}(X^{(l)})^T$ is the empirical covariance. Due to the underlying assumption of homoscedasticity across all interacting variables, the set of possible distributions can be uniquely identified via (2). Therefore, equivalently, we maximize

$$\frac{2}{n}\ell_n(B, \sigma^2) = -d \log(2\pi\sigma^2) - \frac{1}{\sigma^2} \text{tr}((I_d - B)^T(I_d - B)\hat{\Sigma}),$$

using the edge matrix $B \in \mathbb{R}^G$ and the common variance $\sigma^2 > 0$. The zero entries of the matrix B are implied by the given causal ordering. Furthermore, due to the acyclic structure, B is permutation similar to a lower triangular matrix. Straightforward calculations yield that the maximum of $\ell_n(B, \sigma^2)$ over the equal variance $\sigma^2 > 0$ is given by:

$$\sup_{\sigma^2 > 0} \ell_n(B, \sigma^2) = -\frac{nd}{2} \log \left(\frac{2\pi}{d} \text{tr}((I_d - B)^T(I_d - B)\hat{\Sigma}) \right) - \frac{nd}{2}. \quad (10)$$

Maximizing equation (10) over $B \in \mathbb{R}^G$ is then equivalent to minimizing

$$\text{tr}((I_d - B)^T(I_d - B)\hat{\Sigma}) = \frac{1}{n} \sum_{k=1}^d \sum_{p \in p(k)}^n (X_k^{(l)} - \sum_{p \in p(k)} \beta_{k,p} X_p^{(l)})^2. \quad (11)$$

Thus, in the unrestricted case without constraints for a fixed-size causal-effect, the problem is equivalent to solving d separate linear least-squares problems with minimizer:

$$\hat{\beta}_{k,p(k)}^T = (\hat{\Sigma}_{p(k),p(k)})^{-1} \hat{\Sigma}_{p(k),k}, \quad \text{for } k = 1, \dots, d, \quad (12)$$

and corresponding minimum

$$\frac{1}{n} \sum_{k=1}^d \sum_{l=1}^n \left(X_k^{(l)} - \sum_{p \in p(k)} \hat{\beta}_{k,p} X_p^{(l)} \right)^2 = \sum_{k=1}^d \hat{\Sigma}_{k,k|p(k)}.$$

Furthermore, in this case, we can calculate the maximum likelihood of any partial ordering that starts from the source node, since it collapses into subproblems. Thus, we can reject all orderings that start with a given partial ordering if the corresponding partial maximum likelihood already exceeds the threshold.

In addition, we need to calculate the maximum-likelihood estimates under each single hypothesis $H_0^{(\psi)}(G)$. For $j <_G i$, we only test Hypothesis $H_0^{(0)}(G)$. However, in that case, the constraint of a size zero causal-effect $C(i \rightarrow j) = 0$ is not restrictive, and thus the minimizer is similarly given by (12). Therefore, we assume in the following $i <_G j$. Analogously to the previous calculations, maximizing the Gaussian likelihood for a given

causal ordering and a fixed total causal-effect can be solved by minimizing the linear least-squares Problem (11). However, the constraint of a fixed total causal-effect $C(i \rightarrow j) = \psi$ links a subset of the regression parameters, and we cannot solve all d linear least-squares problems separately.

Recalling the path properties of the total causal-effect (see Remark 2.1), the assumption $C(i \rightarrow j) = (I_d - B)_{j,i}^{-1} = \psi$ restricts all $\beta_{k,p(k)}$ for $k \in p(j) \cup \{j\} \setminus p(i) \cup \{i\}$. This follows by observing that

$$(I_d - B)_{j,i}^{-1} = \sum_{k=1}^{d-1} (B^k)_{j,i}$$

and further inductively

$$(B^k)_{j,i} = \sum_{p \in p(j) \setminus p(i)} \beta_{j,p} (B^{k-1})_{p,i}.$$

For all other $k \notin p(j) \cup \{j\} \setminus p(i) \cup \{i\}$, the corresponding part of the optimization Problem (11) can be solved separately with the (unrestricted) minimizer given by (12). However, it remains to minimize the following non-trivial joint least-squares problem:

$$\sum_{k \in p(j) \cup \{j\} \setminus p(i) \cup \{i\}} \sum_{l=1}^n (X_k^{(l)} - \sum_{p \in p(k)} \beta_{k,p} X_p^{(l)})^2, \quad \text{where } \beta_{j,i} = \psi - \sum_{m=2}^{d-1} (B^m)_{j,i}.$$

Note that we can reduce the constraint to the submatrix B_s , which is formed by selecting all rows and columns corresponding to the set $p(j) \cup \{j\} \setminus p(i)$, since $(B^m)_{j,i} = (B_s^m)_{j,i}$. Then, the constraint yields $\beta_{j,i}$ as a function of the remaining $\beta_{k,p(k)}$ for $k \in p(j) \cup \{j\} \setminus p(i) \cup \{i\}$, and we solve the remaining least-squares problem numerically with quasi-Newton methods.

5.3 Simulations

In this section, we present the results of a simulation study that compares the coverage frequencies and the widths of the proposed confidence regions as well as their computing times. Our simulation experiments are designed as follows. We generate synthetic data sets based on LSEMs with Gaussian errors and equal variances associated with randomly selected directed acyclic graphs in a three-step procedure. First, we randomly select a permutation of d nodes, representing the causal ordering of (X_1, \dots, X_d) . Second, we generate edge weights for all edges in the complete graph that corresponds to the chosen causal ordering according to a normal distribution $N(\beta, 0.1)$. In that way, β reflects the average strength of the direct effects in the graph, and thus, smaller values indicate higher uncertainty about the underlying causal structure. In the next step, we prune the complete graph, i.e., we include an edge with a probability of 0.9 for dense and 0.5 for sparse graphs. Finally, we generate n samples according to the LSEM, represented by the chosen graph, with standard normal distributed errors. For a range of edge weights β , sample sizes n and dimensions d , we generated multiple independent data sets and determined the confidence sets for the total causal-effect of X_1 on X_2 with confidence level $\alpha = 0.05$ using the different proposed methods (LRT given by Theorem 4.1 and SLRT given by Theorem 4.5). We repeated this procedure twice, for the case of a true non-zero effect and the case of no effect. We chose the tuning parameter of the SLRT interval, the splitting ratio, optimal in the sense of Strieder and Drton [19].

We report the observed empirical coverage probabilities for six-dimensional graphs in Table 1. All proposed methods achieve the desired coverage frequency of 0.95 and seem to be conservative. This is expected for the split likelihood ratio method since it is based solely on Markov's inequality. However, it is also not surprising that the LRT method is conservative, considering the use of intersection union testing and the fact that we set critical values based on the relaxed alternative. For comparison, we also included classical bootstrapping results in our simulation study. The bootstrapping method uses GDS, an established causal discovery algorithm for LSEMs with equal variance, proposed by Peters and Bühlmann [8]. For each resampled data set, the GDS method uses a greedy search algorithm to maximize the likelihood and estimate the total causal-effect. The results validate the theoretical findings in the studies by Strieder et al. [9] and Drton and Williams [13].

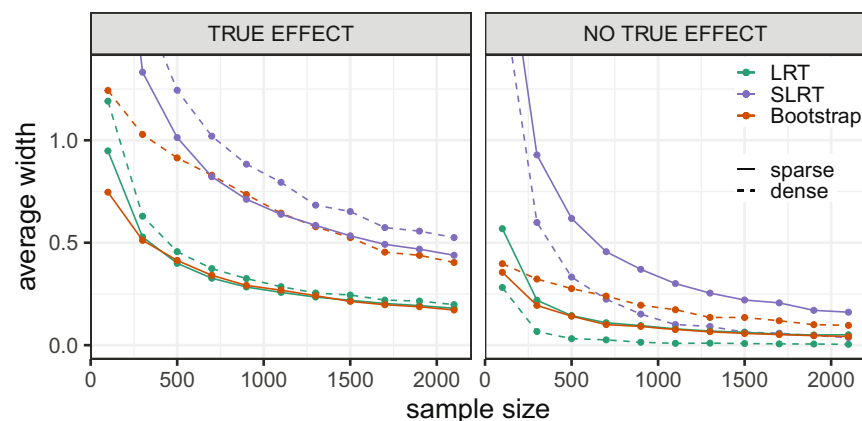
Table 1: Empirical coverage of 95% confidence intervals for the total causal-effect of X_1 on X_2 in randomly generated 6-dim. DAGs (1,000 replications)

| Method | $n \backslash \beta$ | True effect | | | | | | No true effect | | | | | |
|-----------|----------------------|-------------|-------------|------|-------------|-------------|------|----------------|------|------|-------|------|------|
| | | Sparse | | | Dense | | | Sparse | | | Dense | | |
| | | 0.05 | 0.1 | 0.5 | 0.05 | 0.1 | 0.5 | 0.05 | 0.1 | 0.5 | 0.05 | 0.1 | 0.5 |
| LRT | 100 | 1.00 | 1.00 | 0.99 | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | 500 | 0.99 | 0.99 | 0.99 | 0.99 | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | 1,000 | 0.98 | 0.98 | 1.00 | 0.99 | 1.00 | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| SLRT | 100 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | 500 | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | 1,000 | 0.99 | 0.99 | 1.00 | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| Bootstrap | 100 | 0.70 | 0.75 | 0.97 | 0.75 | 0.80 | 0.97 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | 500 | 0.72 | 0.77 | 0.97 | 0.78 | 0.85 | 0.98 | 1.00 | 1.00 | 1.00 | 1.00 | 0.99 | 1.00 |
| | 1,000 | 0.76 | 0.80 | 0.95 | 0.79 | 0.88 | 0.98 | 0.99 | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 |

Bootstrapping methods do not correctly account for uncertainty in the causal structure and do not achieve the required coverage frequency, as highlighted in bold.

In addition to the coverage frequencies, we investigated how conclusive the confidence regions are, i.e., how informative are intervals for practitioners in deciding about the existence and size of the causal-effects, while still providing a valid uncertainty assessment. To this, we show in Figure 1 the average width of the non-zero part of the confidence sets. We report the results for dimension $d = 6$, where the data were generated according to random DAGs with expected edge weight $\beta = 0.5$, against the sample size. In contrast to the bootstrapping method, our proposed LRT method constitutes a valid confidence region. Nevertheless, even in the considered setting with edge weight $\beta = 0.5$, where all variants achieved the desired coverage, our method significantly outperforms bootstrapping in the dense setting, while performing similarly in the sparse setting. In contrast, the split likelihood ratio tests yield wider confidence intervals. We also note that the confidence intervals are generally wider in the dense setting, i.e., the remaining uncertainty about the numerical size of the effects is higher for dense DAGs. Our intuition here is that in dense DAGs, more directed paths contribute to the total causal-effect; thus, more edge weights are involved with remaining uncertainty. Figure 1 shows that our proposed methods successfully help to locate the numerical size of existing causal-effects while correctly quantifying remaining uncertainty.

In addition to locating the numerical size of existing effects, an informative confidence interval should help in deciding whether there exists remaining uncertainty about the existence of an effect. Therefore, to

**Figure 1:** Mean width of 95% confidence intervals for the total causal-effect of X_1 on X_2 in randomly generated 6-dim. DAGs (1,000 replications).

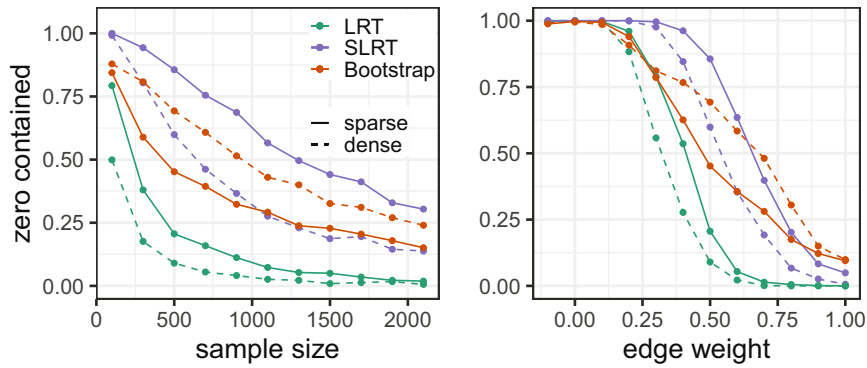


Figure 2: Percentages of times zero contained in 95% confidence intervals in randomly generated 6-dim. DAGs with non-zero effect (1,000 replications): (Left) against sample size and (Right) against expected edge weights.

investigate whether the proposed methods pick up on the causal direction of the effect between two variables, we plotted the proportions of times zero is included in the confidence sets when there is a true non-zero effect. Figure 2 shows the results for dimensions $d = 6$ with expected edge weight $\beta = 0.5$ against the sample size and for the sample size $n = 500$ against the expected edge weight. The LRT method eliminates the remaining uncertainty about the direction of the causal-effect the quickest with increasing sample size, significantly outperforming the bootstrapping method and the split likelihood ratio tests. As intuitively expected, for higher average edge weights, there is less uncertainty about the underlying causal structure, and thus, the methods more often exclude the possibility of no causal-effect.

Furthermore, one observes that in contrast to the bootstrapping method, the confidence regions of our proposed methods contain the zero less often in the dense setting compared to those in the sparse setting. This stems from the fact that there exists less uncertainty about the causal ordering for dense DAGs, and thus, the remaining uncertainty about the direction of the involved effects is lower than in the sparse setting. While there might be less uncertainty about the numerical size of the effect and thus, narrower confidence regions in the sparse setting, the uncertainty about the causal structure is higher.

We obtain similar results for larger causal structures with 12 involved variables. Figure 3 shows the percentages of times our proposed LRT-confidence regions contain zero as well as the mean width of the non-zero part. The data were generated with an expected edge weight $\beta = 0.5$ and increasing sample sizes. In the case of no true effect, the confidence intervals always correctly contain zero. The average width of the non-zero part is close to zero, even for low sample sizes. If there is a non-zero total causal-effect, the remaining uncertainty about the direction of the effect is already eliminated in samples of size around 2,000. We highlight again that the remaining uncertainty about the causal structure (reflected by the percentage of times zero is




Figure 3: Mean width and percentages of times zero contained in 95% confidence intervals for the total causal-effect of X_1 on X_2 in randomly generated 12-dim. DAGs (100 replications).

contained in the confidence interval) is higher in the sparse setting. In contrast, the remaining uncertainty about the numerical size of the effect, reflected by the average width of the non-zero part of the confidence intervals, is lower.

In summary, the proposed framework for constructing confidence intervals for total causal-effects successfully detects the direction and the numerical size of causal-effects and correctly quantifies the remaining uncertainty in causal structure and effect size.

Figure 4 compares the average computation times over 100 randomly drawn DAGs. As the number of possible causal structures grows superexponentially with the dimension, the computation times for our proposed confidence intervals naturally increase. Since the intervals take uncertainty over all possible structures into account, the computation times similarly increase superexponentially. However, Figure 4 shows that it is feasible to compute this complete uncertainty over all causal structures for moderate dimensions with lower computation times than the bootstrapping method. We plotted the computation times for an expected edge weight $\beta = 0.5$ and sample size $n = 1,000$ against the dimension (up to $d = 12$). Contrasting the bootstrapping procedure based on the greedy search algorithm GDS, the computation times for our proposed confidence intervals are lower in the dense compared to those in the sparse setting. This is not unexpected, as dense DAGs generally coincide with less uncertainty in the causal structure and thus lead to fewer plausible causal orderings passed on to the testing procedure. Similarly, if there is no true effect, computation times are lower since we need to test within fewer plausible causal orderings. In general, the computation times are mainly determined by reducing the set of plausible orderings with Algorithm 2 and the uncertainty in the causal structure within the data.

Furthermore, Figure 4 also includes the average computation times for increasing sample sizes with expected edge weight $\beta = 0.5$ and dimension $d = 6$. In contrast to the bootstrapping procedure, increasing the sample size reduces the computation times for our proposed confidence intervals. More available data reduce the uncertainty about the underlying structure, which reduces the computation time.

Our framework is based on the assumption of an underlying Gaussian LSEM with equal error variances. We emphasize that if the true data-generating mechanism follows a Gaussian LSEM with differing error variances, the causal structure and, consequently, the true causal-effects are generally not identifiable. Therefore, our task of constructing a confidence interval for this causal-effect is not well defined. However, in this setting, our method targets the causal-effect in an approximation of the data-generating mechanism under the equal error variances assumption. With practical applications in mind, we investigated this behavior and, thus, the sensitivity of our method toward deviations from equal error variances. We generated data according to the same procedure as described earlier (for $d = 6$, $\beta = 0.1$, and $n = 500$), but with error variances sampled uniformly from $[1 - 0.5v, 1 + 0.5v]$. In that way, v quantifies the degree of deviation from homoscedasticity across the interacting variables, ranging from 0 to 1.8. The empirical coverage probabilities in Figure 5 show that our framework is rather robust to small deviations from equal error variances. For small




Figure 4: Mean computation times of 95% confidence intervals for causal-effect in randomly generated DAGs in seconds (100 replications): (Left) against dimension and (Right) against sample size.

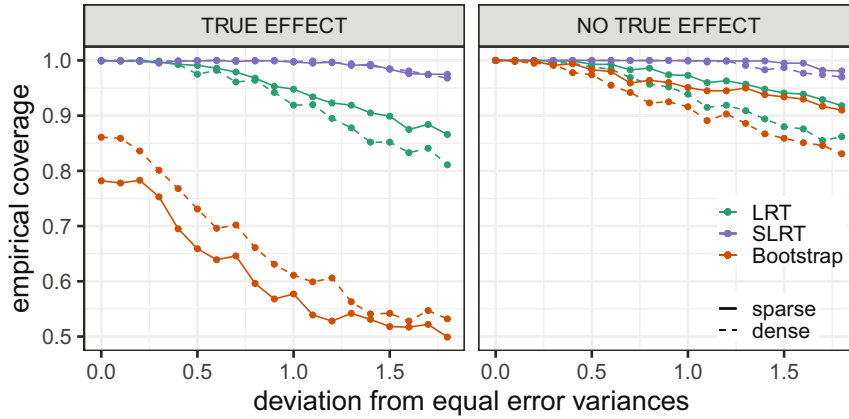


Figure 5: Empirical coverage of 95% confidence intervals for the total causal-effect of X_1 on X_2 in randomly generated 6-dim. DAGs under departure from equal error variances (1,000 replications).

deviations, the targeted causal-effect in the approximated equal error variance model is close to the true causal-effect, and thus, our confidence intervals, nevertheless, cover the (unidentifiable) true effect.

In conclusion, the likelihood ratio method performs best in our simulation studies and yields the most informative confidence intervals. Nonetheless, while the split likelihood ratio tests are more conservative, we emphasize that it is the only method with theoretical finite sample guarantees.

5.4 Real data example

To illustrate the importance of correctly accounting for uncertainty in the causal structure, we consider the frequently studied Sachs protein data [20]. This collection of 14 data sets comprises expression levels for 11 proteins and phospholipids in human T-cells obtained under different experimental conditions. In applications involving variables from similar domains and measurements obtained by similar techniques, the assumption of equal error variances offers a simple way to conduct exploratory analyzes of cause-effect relations, also in light of the robustness of our method toward small deviations from equal variance that was seen in our simulations. In our data analysis, we compared our LRT method against naive bootstrapping with the GDS method [8] as well as bootstrapping with the LiNGAM method [21]. The LiNGAM method explicitly assumes non-Gaussian errors to ensure identifiability. We considered the first of the 14 data sets with a sample size of 853 and focused on a subset of 10 proteins (excluding PKA) that have measurements on a similar scale. The protein PKA is a highly variable source node in the ground truth causal DAG reported in the study by Sachs et al. [20]. To account for its influence, we first performed linear regressions of each protein on PKA and continued our analysis using the residuals.

Figure 6 shows the calculated 95% confidence intervals for six exemplary causal-effects. The regression coefficient of X_i when regressing X_j on $(X_i, X_{p(i)})$, where the conventionally accepted ground truth structure defines the parent set, is seen as the true target effect $C(i \rightarrow j)$ and marked in red. In total, our proposed confidence intervals cover 84 of the 90 pairwise causal-effects among the 10 involved variables. In contrast, naive bootstrapping with GDS and with LiNGAM cover only 66 and 72, respectively. Our confidence intervals indicate that, even within our strict modeling assumptions, there is substantial structure uncertainty. Out of the 90 calculated intervals, 67 contain zero as well as non-zero effects, thus reflecting uncertainty about the direction and existence of effects. In comparison, the naive methods with GDS and with LiNGAM contain both directions in only 35 and 55 cases, respectively. Thus, while the bootstrapping methods yield narrower intervals on average, they do not correctly account for structure uncertainty and tend to miss more minor target effects.




Figure 6: 95% confidence intervals for the total causal-effect between proteins from the real world expression data.

6 Discussion

In this article, we raise the problem of rigorously quantifying uncertainty in causal inference when the causal structure is unknown. We demonstrate the intricacies of this task and propose a precise solution for constructing confidence regions for total causal-effects, which capture both types of uncertainty: numerical size of effects and causal structure. Using a test inversion approach with intersection union tests and exploiting combinatorial shortcuts, our proposed framework mathematically rigorously quantifies the remaining uncertainty in the concrete setting of LSEMs with homoscedastic Gaussian errors across all interacting variables.

Our proposal of leveraging test inversions of joint tests for causal structure and size of effect to obtain confidence regions for causal-effect estimates under structure uncertainty has a general nature and with suitable statistical tests could be extended to other modeling assumptions (see Strieder et al.'s study [9] for results from a heuristic for the example of bivariate LSEMs with non-Gaussian errors). Moreover, our framework could be adjusted to target different (causal) quantities by reformulating the corresponding constraint.

Assuming LSEMs with equal error variances ensures identifiability. Thus, a confidence region corresponds to a set of covariance matrices, which all uniquely entail a corresponding total causal-effect. Without equal error variances (or rather identifiability in general), this unique correspondence does not hold. While our computational framework of focusing on causal orders and complete graphs might be adapted to consider a set of plausible covariance matrices, every covariance matrix then entails a set of possible causal-effects that always contains zero. Thus, the corresponding confidence regions are not conclusive about the direction of effects. In contrast, we demonstrated that under the assumption of equal error variances, our proposed confidence regions not only correctly quantify the remaining uncertainty in causal structure and numerical effect size, but also provide conclusive information that can help practitioners to decide about existence and numerical size of effects.

The computational shortcuts we adopt allow us to quantify structure uncertainty over all possible DAGs among a moderate number of nodes. For higher-dimensional systems, it is not feasible to consider all possible DAGs given that there already exist more DAGs on 22 nodes (10^{87}) than estimated atoms in the universe. However, our method can still be used by practitioners as a valuable tool for analyzing higher-dimensional causal systems by partially quantifying the uncertainty in the causal structure.

In the *status quo* scenario, practitioners would either confer with experts that provide the underlying causal structure or use structure learning algorithms to estimate the DAG from available data. Starting from this inferred model, classical statistical methods can then be used to calculate confidence intervals for the involved causal parameters. These confidence intervals quantify uncertainty in the numerical size of the causal-effect conditional on the inferred model. However, they do not incorporate uncertainty about the underlying causal structure and the (data-driven) model choice. While in higher-dimensional systems, it is not computationally tractable to quantify structure uncertainty by considering all DAGs, our framework can still be used as a compromise in between. Rather than solely relying on the full structure provided by experts, our method can be used to consider a set of plausible causal structures. This way practitioners can incorporate varying degrees of belief of experts in different parts of the structure by fixing parts with high confidence but still considering the remaining uncertainty over the rest. Along similar lines, instead of focusing on a single

output of causal learning algorithms, practitioners can use our method to a set of (most likely) causal structures and incorporate some degree of structure uncertainty in causal-effect estimates.

In summary, we consider our study of total causal-effect in linear causal models with equal error variances as the beginning of a new line of research on confidence sets under structure uncertainty. We anticipate fruitful generalizations in many directions including not only other model settings such as the linear non-Gaussian case, or additive noise models, but also different causal parameters other than total effects. From this perspective, while this article develops our ideas in a very concrete setting, it sets the scene for various follow-up works in providing rigorously justified confidence regions in causal-effect estimation when the causal structure itself is learned from data.

Funding information: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 83818). Furthermore, this work has been funded by the German Federal Ministry of Education and Research and the Bavarian State Ministry for Science and the Arts. The authors of this work take full responsibility for its content.

Conflict of interest: The authors state no conflict of interest.

References

- [1] Pearl J. Causality. Models, reasoning, and inference. 2nd ed. Cambridge: Cambridge University Press; 2009.
- [2] Spirtes P, Glymour C, Scheines R. Causation, prediction, and search. Adaptive computation and machine learning. Cambridge, MA: MIT Press; 2000.
- [3] Hoyer PO, Hyttinen A. Bayesian discovery of linear acyclic causal models. In: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence. UAI'09. Arlington, Virginia, USA: AUAI Press; 2009. p. 240–8.
- [4] Claassen T, Heskes T. A Bayesian approach to constraint based causal inference. In: Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence. UAI'12. Arlington, Virginia, USA: AUAI Press; 2012. p. 207–16.
- [5] Cao X, Khare K, Ghosh M. Posterior graph selection and estimation consistency for high-dimensional Bayesian DAG models. *Ann Statist.* 2019;47(1):319–48.
- [6] Peters J, Janzing D, Schölkopf B. Elements of causal inference. Adaptive Computation and Machine Learning. Foundations and learning algorithms. Cambridge, MA: MIT Press; 2017.
- [7] Maathuis M, Drton M, Lauritzen S, Wainwright M, editors. Handbook of graphical models. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. Boca Raton, FL: CRC Press; 2019.
- [8] Peters J, Bühlmann P. Identifiability of Gaussian structural equation models with equal error variances. *Biometrika.* 2014;101(1):219–28.
- [9] Strieder D, Freidling T, Haffner S, Drton M. Confidence in causal discovery with linear causal models. In: Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence. UAI'21. Vol. 161. PMLR; 2021. p. 1217–26.
- [10] Chen W, Drton M, Wang YS. On causal discovery with an equal-variance assumption. *Biometrika.* 2019;106(4):973–80.
- [11] Ghoshal A, Honorio J. Learning linear structural equation models in polynomial time and sample complexity. In: Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics. Vol. 84. PMLR; 2018. p. 1466–75.
- [12] Andrews DWK, Guggenberger P. Asymptotic size and a problem with subsampling and with the m out of n bootstrap. *Econ Theory.* 2010;26(2):426–68.
- [13] Drton M, Williams B. Quantifying the failure of bootstrap likelihood ratio tests. *Biometrika.* 2011;98(4):919–34.
- [14] Casella G, Berger RL. Statistical inference. The Wadsworth & Brooks/Cole Statistics/Probability Series. Pacific Grove, CA: Wadsworth & Brooks/Cole Advanced Books & Software; 1990.
- [15] Drton M. Algebraic problems in structural equation modeling. In: The 50th anniversary of Gröbner bases. Vol. 77 of Adv. Stud. Pure Math. Math. Soc. Japan, Tokyo; 2018. p. 35–86.
- [16] Silvapulle MJ, Sen PK. Constrained statistical inference. Wiley Series in Probability and Statistics. Hoboken, NJ: Wiley-Interscience [John Wiley & Sons]; 2005.
- [17] Wasserman L, Ramdas A, Balakrishnan S. Universal inference. *Proc Natl Acad Sci USA.* 2020;117(29):16880–90.
- [18] Drton M. Likelihood ratio tests and singularities. *Ann Statist.* 2009;37(2):979–1012.
- [19] Strieder D, Drton M. On the choice of the splitting ratio for the split likelihood ratio test. *Electron J Stat.* 2022;16(2):6631–50.
- [20] Sachs K, Perez O, Pe'er D, Lauffenburger DA, Nolan GP. Causal protein-signaling networks derived from multiparameter single-cell data. *Science.* 2005;308(5721):523–9.
- [21] Shimizu S, Hoyer PO, Hyvärinen A, Kerminen A. A linear non-Gaussian acyclic model for causal discovery. *J Mach Learn Res.* 2006;7:2003–30.