

Dual Likelihood for Causal Inference under Structure Uncertainty

David Strieder DAVID.STRIEDER@TUM.DE and **Mathias Drton** MATHIAS.DRTON@TUM.DE
Munich Center for Machine Learning and Department of Mathematics, TUM School of Computation, Information and Technology, Technical University of Munich, Germany

Editors: Francesco Locatello and Vanessa Didelez

Abstract

Knowledge of the underlying causal relations is essential for inferring the effect of interventions in complex systems. In a widely studied approach, structural causal models postulate noisy functional relations among interacting variables, where the underlying causal structure is then naturally represented by a directed graph whose edges indicate direct causal dependencies. In the typical application, this underlying causal structure must be learned from data, and thus, the remaining structure uncertainty needs to be incorporated into causal inference in order to draw reliable conclusions. In recent work, test inversions provide an ansatz to account for this data-driven model choice and, therefore, combine structure learning with causal inference. In this article, we propose the use of dual likelihood to greatly simplify the treatment of the involved testing problem. Indeed, dual likelihood leads to a closed-form solution for constructing confidence regions for total causal effects that rigorously capture both sources of uncertainty: causal structure and numerical size of nonzero effects. The proposed confidence regions can be computed with a bottom-up procedure starting from sink nodes. To render the causal structure identifiable, we develop our ideas in the context of linear causal relations with equal error variances.

Keywords: linear structural causal models, graphical models, dual likelihood, causal effects, uncertainty quantification

1. Introduction

Reasoning about causal relations and inferring the effect of interventions in complex systems is a central task of scientific research. The field of causal discovery addresses this challenge with much fundamental research being done over the last decades (Pearl, 2009; Spirtes et al., 2000). In particular, in many applied settings, access to interventional data is limited, and thus, performing classical controlled experiments to investigate causal relations is not feasible. To tackle this problem, various identifiability results clarify under which conditions interventional distributions can theoretically be identified from observational data alone. Furthermore, many structure learning algorithms have been proposed that estimate causal structures using only observed data. On the other hand, when the underlying causal structure is known, causal inference results enable researchers to estimate causal quantities in complex systems and assess remaining uncertainty. However, when the underlying causal structure is learned from data, the remaining structure uncertainty needs to be incorporated into causal inference in order to draw reliable conclusions.

Recently, Strieder et al. (2021) introduced a first ansatz for constructing confidence regions for causal effects that account for both types of uncertainty, uncertainty about the causal structure and uncertainty about the numerical size of the effect. They consider bivariate linear structural causal models (SCMs) and develop a solution based on test inversion. Their strategy was improved and generalized to higher-dimensional linear SCMs in Strieder and Drton (2023). In this work, we

take up their idea of using a test inversion ansatz to construct confidence regions for total causal effects (see Section 2). However, in contrast to the earlier work, we suggest here to employ dual likelihood for the arising testing problem. An important consequence of the use of dual likelihood is that we are able to provide a closed-form solution for our proposed confidence regions (Section 3) and can, thus, avoid the numerical optimization algorithms deployed by [Strieder and Drton \(2023\)](#). Furthermore, while the earlier methods rely on a top-down procedure starting from the source node, by employing dual likelihood, we obtain a bottom-up procedure starting from the sink node. Thus, our proposed dual likelihood method provides an interesting alternative perspective on strategies to save on computation. Finally, in Section 4.2, we present computational details and analyze the performance in a simulation study. Further simulation results can be found in Appendix A. In the concluding Section 5, we discuss extensions and possible future work.

2. Background

A common tool for studying causal relationships among interacting variables are structural causal models ([Peters et al., 2017](#); [Maathuis et al., 2019](#)), where each variable is represented as a function of other variables (its causes) and a random error term. The causal perspective views these relationships as assignments rather than mathematical equations, and thus, changes in the causes result in changes in the effects but not vice versa, reflecting the inherent asymmetry in cause-effect relationships. In this section, we review a restricted class of causal models, namely linear structural causal models with Gaussian errors and equal error variances ([Peters and Bühlmann, 2014](#)), as well as the definition of the total causal effect, which is the causal target of interest in this study. Further, we review the test inversion ansatz introduced by [Strieder et al. \(2021\)](#) to construct confidence regions for causal effects that account for structure uncertainty.

2.1. Causal Effects in Linear Structural Causal Models

We assume access to observational data in the form of n independent copies of a random vector $X = (X_1, \dots, X_d)$ with zero mean. To ensure unique identifiability, we follow a line of research introduced by [Peters and Bühlmann \(2014\)](#) that focuses on linear relations and normal distributed errors with equal variances, represented by the equation system

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

where $B := [\beta_{j,i}]_{j,i=1}^d$ represents the direct causal effects between variables, and ε_j are independent, normally distributed error terms with common variance σ^2 . Such a linear structural causal model is naturally represented by a directed graph, where a missing edge $i \rightarrow j$ indicates no direct effects, that is, $b_{j,i} = 0$. Further, we assume the underlying graph to be acyclic, which entails the unique solution $X = (I_d - B)^{-1} \varepsilon$ of the equation system (1), where I_d denotes the $d \times d$ identity matrix. The corresponding covariance matrix is given by $\Sigma = \sigma^2 (I_d - B)^{-1} (I_d - B)^{-T}$.

In the remainder of the article, we use the following notation and graphical concepts. We write $i <_G j$ if node i precedes node j in a causal ordering of the corresponding directed acyclic graph (DAG). If the DAG contains an edge from node i to node j , then node i is called a parent of node j , and we denote the set of all parents of node j with $p(j)$. Further, if the DAG contains a directed path from node i to node j , then node j is called a descendant of node i , and we denote the set of

all descendants of node i with $d(i)$. Finally, we write $\Sigma_{j,i|p(i)}$ for the conditional covariance matrix, that is,

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

Our aim is to estimate the total causal effect $\mathcal{C}(i \rightarrow j)$ in linear SCMs and rigorously quantify the remaining uncertainty. Formally, the total causal effect is defined as the unit change in the expectation of X_j with respect to an intervention in X_i

$$\mathcal{C}(i \rightarrow j) := \frac{d}{dx_i} \mathbb{E}[X_j | \text{do}(X_i = x_i)] = \Sigma_{j,i|p(i)}(\Sigma_{i,i|p(i)})^{-1}.$$

While this parameter of interest is given as a simple function of the covariance matrix, difficulties arise in practice when the underlying causal structure is unknown. In such situations, the conditioning set is unknown and has to be inferred from data, which introduces structure uncertainty. We emphasize that due to the identifiability of Gaussian SCMs with equal error variances (e.g., [Chen et al., 2019](#); [Ghoshal and Honorio, 2018](#)), estimating the total causal effect is nevertheless a well-defined problem.

A naive two-step approach that treats the two tasks of causal discovery and causal inference separately by learning the causal structure first and then calculating confidence intervals in the inferred model does not account for uncertainty in the structure. Furthermore, classical bootstrapping methods also fail to correctly account for model uncertainty due to singularities at the intersection of models. To address this issue, [Strieder et al. \(2021\)](#) introduced a framework for constructing confidence intervals that rigorously account for structure uncertainty, which was then generalized in [Strieder and Drton \(2023\)](#). However, thoroughly accounting for structure uncertainty proves a challenging task, given the enormous number of possible structures, even in moderate dimensions.

2.2. Confidence Intervals via Test Inversion

The main idea is to leverage the classical duality between statistical hypothesis tests and confidence regions ([Casella and Berger, 1990](#)). To this end, we consider all possible causal structures in the following way. If we perform valid hypothesis tests for all attainable causal effects, then all values that cannot be rejected form a confidence region for the total causal effect. This shifts the task to constructing suitable hypothesis tests for all attainable causal effects.

Under the assumption of an underlying linear structural causal model with equal error variances, every possible distribution $N(0, \Sigma)$ is given by a set of covariance matrices $\mathcal{M} := \bigcup_{G \in \mathcal{G}(d)} \mathcal{M}(G)$, where $\mathcal{G}(d)$ is the set of all possible complete DAGs on d nodes and

$$\mathcal{M}(G) = \left\{ \Sigma \in \text{PD}(d) : \exists \sigma^2 > 0 \text{ with } \sigma^2 = \Sigma_{k,k|p(k)} \quad \forall k = 1, \dots, d \right\}.$$

Here, we write $\text{PD}(d)$ for the cone of positive definite $d \times d$ matrices. The hypothesis of a fixed causal effect $\mathcal{C}(i \rightarrow j)$ of size ψ further restricts the set of possible distributions, where the specific structure of the constraint depends on the DAG G . Thus, with

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

and $\mathcal{M}_\psi := \bigcup_{G \in \mathcal{G}(d)} \mathcal{M}_\psi(G)$, the task is to invert the statistical testing problem

$$H_0^{(\psi)} : \Sigma \in \mathcal{M}_\psi \quad \text{against} \quad H_1 : \Sigma \in \mathcal{M} \setminus \mathcal{M}_\psi \quad (3)$$

for all $\psi \in \mathbb{R}$.

3. Testing for Total Causal Effects with the Dual Likelihood

Earlier methods (Strieder and Drton, 2023) propose to use constrained likelihood ratio tests for the testing problem (3). However, the direct maximization of the Gaussian likelihood with constraints on the total causal effect has no closed-form solution, which leads to algorithms that rely on numerical optimization routines and grid searches. Instead, our proposal is to employ dual likelihood theory for Gaussian models (e.g., Brown, 1986; Kauermann, 1996) to solve the testing problem (3) and, thus, obtain a simple closed-form solution.

To obtain our main result, we first introduce the dual likelihood for Gaussian graphical models. Maximizing the Gaussian likelihood corresponds to the minimization problem

$$\arg \min_{\Sigma \in \mathcal{M}} \text{KL}(P_{\hat{\Sigma}}, P_{\Sigma}) = \arg \min_{\Sigma \in \mathcal{M}} (\text{tr}(\Sigma^{-1} \hat{\Sigma}) + \log \det(\Sigma)),$$

where KL is the Kullback–Leibler divergence with sample covariance $\hat{\Sigma}$. As the Kullback–Leibler divergence is not symmetrical in its arguments, the minimization problem

$$\begin{aligned} \arg \min_{\Sigma \in \mathcal{M}} \text{KL}(P_{\Sigma}, P_{\hat{\Sigma}}) &= \arg \min_{\Sigma \in \mathcal{M}} (\text{tr}(\hat{\Sigma}^{-1} \Sigma) - \log \det(\Sigma)) \\ &= \arg \min_{\Sigma \in \mathcal{M}} (\text{tr}(\Sigma \hat{\Sigma}^{-1}) + \log \det(\Sigma^{-1})) \end{aligned}$$

yields a different estimator, the dual maximum likelihood estimator. Along similar lines, the dual (log-)likelihood of Σ for the observed parameter $\hat{\Sigma}^{-1}$ is defined as

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

We note that this dual likelihood is reciprocal to the Gaussian likelihood in the sense that for observed sample covariance $\hat{\Sigma}^{-1}$, we have

$$\ell_n^{dual}(\Sigma | \hat{\Sigma}^{-1}) = \ell_n(\Sigma^{-1} | \hat{\Sigma}^{-1}), \quad (4)$$

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

3.1. Dual Likelihood Ratio Test

In order to construct confidence intervals for the total causal effects, our idea is to invert dual likelihood ratio tests of the testing problem (3) with the dual likelihood ratio test statistic

$$\begin{aligned} \text{dual-}\check{\lambda}_n^{(\psi)} &:= 2 \left(\sup_{\Sigma \in \mathcal{M}} \ell_n^{dual}(\Sigma | \hat{\Sigma}^{-1}) - \sup_{\Sigma \in \mathcal{M}_{\psi}} \ell_n^{dual}(\Sigma | \hat{\Sigma}^{-1}) \right) \\ &= 2 \left(\sup_{\Sigma \in \mathcal{M}} \ell_n(\Sigma^{-1} | \hat{\Sigma}^{-1}) - \sup_{\Sigma \in \mathcal{M}_{\psi}} \ell_n(\Sigma^{-1} | \hat{\Sigma}^{-1}) \right). \end{aligned} \quad (5)$$

First, we observe that by employing equation (4), this dual likelihood ratio test statistic equals the classical likelihood ratio test statistic for a modified problem. More precisely, we redefine the model space $\mathcal{M}^{-1} := \bigcup_{G \in \mathcal{G}(d)} \mathcal{M}^{-1}(G)$ with

$$\mathcal{M}^{-1}(G) := \left\{ \Omega \in \text{PD}(d) : \exists \tilde{\sigma}^2 > 0 \text{ with } \tilde{\sigma}^2 = \Omega_{k,k|d(k)} \quad \forall k = 1, \dots, d \right\},$$

where $\Omega = \Sigma^{-1}$. Similarly, we redefine the hypothesis space $\mathcal{M}_\psi^{-1} = \bigcup_{G \in \mathcal{G}(d)} \mathcal{M}_\psi^{-1}(G)$ with

$$\mathcal{M}_\psi^{-1}(G) := \left\{ \Omega \in \text{PD}(d) : \exists \tilde{\sigma}^2 > 0 \text{ with } \tilde{\sigma}^2 = \Omega_{k,k|d(k)} \quad \forall k = 1, \dots, d \right. \\ \left. \text{and } \psi = \tau_j(\Omega_{i,d(i)}(\Omega_{d(i),d(i)})^{-1}) \right\},$$

where τ_j projects the $|d(i)|$ -dimensional vector onto the component corresponding to j if $j \in d(i)$ and zero otherwise. Then, we obtain the following Lemma.

Lemma 1 *Let $\Sigma \in \text{PD}(d)$. Then the following equivalences hold.*

1. $\Sigma \in \mathcal{M}$ if and only if $\Sigma^{-1} \in \mathcal{M}^{-1}$.
2. $\Sigma \in \mathcal{M}_\psi$ if and only if $\Sigma^{-1} \in \mathcal{M}_\psi^{-1}$.

Proof This result follows from straightforward calculations using the following observation. If $\Sigma = \sigma^2(I_d - B)^{-1}(I_d - B)^{-T}$ for some DAG G with edge weights B and equal error variance $\sigma^2 > 0$, then $\Sigma^{-1} = \sigma^{-2}(I_d - T)^{-1}(I_d - T)^{-T}$, where T represents the negative total causal effect between variables. Thus, Σ^{-1} corresponds to the covariance matrix of a DAG G' , where the parent sets $p(i)$ in G' are the descendants $d(i)$ in G for all $i = 1, \dots, d$, the edge weights are given by T and equal error variance σ^{-2} . Constraints on the total causal effects in Σ and G thus correspond to constraints on direct effects in Σ^{-1} and G' . \blacksquare

Using Lemma 1, we immediately obtain that the dual likelihood ratio test statistic (5) corresponds to

$$\text{dual-}\check{\lambda}_n^{(\psi)} = 2 \left(\sup_{\Omega \in \mathcal{M}^{-1}} \ell_n(\Omega | \hat{\Sigma}^{-1}) - \sup_{\Omega \in \mathcal{M}_\psi^{-1}} \ell_n(\Omega | \hat{\Sigma}^{-1}) \right),$$

which is the classical likelihood ratio test statistic for the modified testing problem

$$\tilde{H}_0^{(\psi)} : \Omega \in \mathcal{M}_\psi^{-1} \quad \text{against} \quad \tilde{H}_1 : \Omega \in \mathcal{M}^{-1} \setminus \mathcal{M}_\psi^{-1}, \quad (6)$$

with observed sample covariance $\hat{\Sigma}^{-1}$. Thus, we can employ similar strategies to [Strieder and Drton \(2023\)](#) in order to obtain confidence regions for total causal effects.

3.2. Confidence Intervals with Dual Likelihood

In this subsection, we state our main result, a confidence region for total causal effects $\mathcal{C}(i \rightarrow j)$ that captures structure uncertainty as well as uncertainty about the numerical size of the effect. Our first step to tackle the testing problem (6) is to relax the alternative \mathcal{M}^{-1} and employ the theory of intersection union tests, see e.g. [Casella and Berger \(1990\)](#), to obtain a simple upper bound on the distribution of the dual likelihood ratio test statistic via

$$\text{dual-}\check{\lambda}_n^{(\psi)} \leq \text{dual-}\lambda_n^{(\psi)}(G) := 2 \left(\sup_{\Omega \in \text{PD}(d)} \ell_n(\Omega | \hat{\Sigma}^{-1}) - \sup_{\Omega \in \mathcal{M}_\psi^{-1}(G)} \ell_n(\Omega | \hat{\Sigma}^{-1}) \right).$$

Further, we define $\text{dual-}\check{\lambda}_n^{(\psi)}(i <_G j) := \min_{G \in \mathcal{G}(d) : i <_G j} \text{dual-}\check{\lambda}_n^{(\psi)}(G)$, where

$$\text{dual-}\check{\lambda}_n^{(\psi)}(G) := 2 \left(\sup_{\Omega \in \mathcal{M}^{-1}} \ell_n(\Omega | \hat{\Sigma}^{-1}) - \sup_{\Omega \in \mathcal{M}_\psi^{-1}(G)} \ell_n(\Omega | \hat{\Sigma}^{-1}) \right).$$

Then we obtain our following main result via test inversions of asymptotically conservative hypothesis test based on the asymptotic distribution of the upper bound dual- $\lambda_n^{(\psi)}(G)$ under every single hypothesis $\mathcal{M}_\psi^{-1}(G)$, see Lemma 3, as well as the closed-form solution of the dual likelihood estimation under constraints on the total causal effect, see Section 3.3.

Theorem 2 *Let $\alpha \in (0, 1)$. Then an asymptotic $(1 - \alpha)$ -confidence set for the total causal effect $\mathcal{C}(i \rightarrow j)$ is given by*

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

This confidence set can be constructed explicitly as follows. Define

$$K := \min_{G \in \mathcal{G}(d)} \sum_{k=1}^d (\widehat{\Sigma}^{-1})_{k,k|d(k)}, \quad Z := \min_{G \in \mathcal{G}(d): j <_G i} \sum_{k=1}^d (\widehat{\Sigma}^{-1})_{k,k|d(k)},$$

and for all $G \in \mathcal{G}(d)$ with $i <_G j$,

$$D_G := (\widehat{\Sigma}^{-1})_{i,j|d(i)\setminus\{j\}}^2 - (\widehat{\Sigma}^{-1})_{j,j|d(i)\setminus\{j\}} \left(\sum_{k \neq i}^d (\widehat{\Sigma}^{-1})_{k,k|d(k)} + (\widehat{\Sigma}^{-1})_{i,i|d(i)\setminus\{j\}} - K \exp\left(\frac{\chi_{d,1-\alpha}^2}{dn}\right) \right).$$

Moreover, for $D_G \geq 0$, define

$$L_G := \frac{-(\widehat{\Sigma}^{-1})_{i,j|d(i)\setminus\{j\}} - \sqrt{D_G}}{(\widehat{\Sigma}^{-1})_{j,j|d(i)\setminus\{j\}}}, \quad U_G := \frac{-(\widehat{\Sigma}^{-1})_{i,j|d(i)\setminus\{j\}} + \sqrt{D_G}}{(\widehat{\Sigma}^{-1})_{j,j|d(i)\setminus\{j\}}}.$$

Then the closed-form solution for the confidence set C is given by

$$C = \bigcup_{G \in \mathcal{G}(d): D_G \geq 0} [L_G, U_G] \cup \{0 : Z \leq K \exp\left(\frac{\chi_{d-1,1-\alpha}^2}{dn}\right)\}.$$

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 with the introduced upper bound and Lemma 3

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

Furthermore, plugging in the dual likelihood estimates (7) and (8) derived in Section 3.3, we view

$$\text{dual-}\check{\lambda}_n^{(\psi)}(G) - \chi_{d,1-\alpha}^2 = 2 \left(\sup_{\Omega \in \mathcal{M}^{-1}} \ell_n(\Omega | \widehat{\Sigma}^{-1}) - \sup_{\Omega \in \mathcal{M}_\psi^{-1}(G)} \ell_n(\Omega | \widehat{\Sigma}^{-1}) \right) - \chi_{d,1-\alpha}^2$$

as a strictly convex quadratic polynomial in ψ . With $K := \min_{G \in \mathcal{G}(d)} \sum_{k=1}^d (\widehat{\Sigma}^{-1})_{k,k|d(k)}$ and

$$D_G := (\widehat{\Sigma}^{-1})_{i,j|d(i)\setminus\{j\}}^2 - (\widehat{\Sigma}^{-1})_{j,j|d(i)\setminus\{j\}} \left(\sum_{k \neq i}^d (\widehat{\Sigma}^{-1})_{k,k|d(k)} + (\widehat{\Sigma}^{-1})_{i,i|d(i)\setminus\{j\}} - K \exp\left(\frac{\chi_{d,1-\alpha}^2}{dn}\right) \right),$$

this quadratic polynomial has real roots if $D_G \geq 0$. Thus, the inequality $\text{dual-}\check{\lambda}_n^{(\psi)}(G) \leq \chi_{d,1-\alpha}^2$ holds if and only if

$$\psi \in \left[\frac{-(\widehat{\Sigma}^{-1})_{i,j|d(i)\setminus\{j\}} - \sqrt{D_G}}{(\widehat{\Sigma}^{-1})_{j,j|d(i)\setminus\{j\}}}, \frac{-(\widehat{\Sigma}^{-1})_{i,j|d(i)\setminus\{j\}} + \sqrt{D_G}}{(\widehat{\Sigma}^{-1})_{j,j|d(i)\setminus\{j\}}} \right].$$

We obtain an analogous result for $j <_G i$, and the test inversion approach yields the claim. \blacksquare

Note that for non-zero effects, we do not need to test among graphs with $j <_G i$. Thus, our confidence regions may consist of a non-zero interval and an isolated zero, reflecting the remaining uncertainty about the direction of the causal relation. In contrast, for zero-sized effects, we need to consider all causal orderings since multiple paths might cancel.

Our main result employs the following asymptotic distribution of the upper bound $\text{dual-}\lambda_n^{(\psi)}(G)$ under every single hypothesis $\mathcal{M}_\psi^{-1}(G)$.

Lemma 3 *Let $G \in \mathcal{G}(d)$ be a DAG and let $\psi \in \mathbb{R}$. Then, the relaxed dual likelihood ratio test statistic $\text{dual-}\lambda_n^\psi(G)$ satisfies one of the following properties.*

a) *If $i <_G j$, then under the hypothesis $H_0^{(\psi)}(G)$*

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

b) *If $j <_G i$, then under the hypothesis $H_0^{(0)}(G)$*

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

Proof We assume $i <_G j$. Without loss of generality, let $(1, 2, \dots, d)$ be the causal ordering of G . Then $\mathcal{M}_\psi^{-1}(G) = \{\Omega \in \text{PD}(d) : f_\psi(\Omega|G) = 0\}$, with

$$f_\psi(\Omega|G) := \begin{pmatrix} \tau_j(\Omega_{i,\{i+1,\dots,d\}}(\Omega_{\{i+1,\dots,d\},\{i+1,\dots,d\}}^{-1}) - \psi) \\ \Omega_{1,1|\{2,\dots,d\}} - \Omega_{d,d} \\ \vdots \\ \Omega_{d-1,d-1|d} - \Omega_{d,d} \end{pmatrix}.$$

The Jacobian of f_ψ has full rank d , which follows due to its (nonzero) triangular structure in derivatives for $\Omega_{k,k}$ for all $k = 1, \dots, d-1$, while the first row has nonzero derivative for $\Omega_{i,i}$ but zero derivatives for $\Omega_{k,k}$ for all $k = 1, \dots, i$. Thus, $\mathcal{M}_\psi^{-1}(G)$ defines a $\frac{1}{2}(d^2 - d)$ -dimensional submanifold of $\mathbb{R}^{\frac{1}{2}(d^2+d)}$. The case $j <_G i$ follows similarly, and thus, the claim follows. For details, we refer to [Drton \(2009\)](#). \blacksquare

Note that the constraint of a zero-sized effect does not restrict the submanifold $\mathcal{M}_0^{-1}(G)$ for $j <_G i$. All edge weights in the corresponding linear SCMs can vary freely, and thus, the two different degrees of freedom for the chi-square limits arise.

3.3. Dual Likelihood Estimation

Maximizing the classical Gaussian likelihood under constraints on total causal effects leads to complex optimization problems with polynomial constraints that need to be solved with numerical optimization routines. In contrast to that, we can explicitly maximize the dual likelihood under constraints on total causal effects. Thus, as an important consequence of the use of dual likelihood, we are able to provide a closed-form solution for our proposed confidence regions in Theorem 2.

In order to derive this closed-form solution, we need to calculate the dual likelihood ratio test statistic (5). This involves maximizing the dual likelihood in two cases, the unrestricted case $\Sigma \in \mathcal{M}$, and under the constraint of a fixed-size total causal effect $\Sigma \in \mathcal{M}_\psi$. In the following, we solve the dual likelihood equations for a fixed graph G . In order to consider full uncertainty over all possible causal structures, we subsequently need to cleverly optimize over the space of DAGs, see Section 4.1. In the case $\Sigma \in \mathcal{M}(G)$, we obtain

$$\begin{aligned} \frac{2}{n} \sup_{\Sigma \in \mathcal{M}(G)} \ell_n^{dual}(\Sigma | \hat{\Sigma}^{-1}) &= \sup_{\Omega \in \mathcal{M}^{-1}(G)} -\log \det(2\pi\Omega) - \text{tr}(\Omega^{-1} \hat{\Sigma}^{-1}) \\ &= \sup_{T \in \mathbb{R}^{-G}, \sigma^2 > 0} -d \log(2\pi\sigma^{-2}) - \sigma^2 \text{tr}((I - T)^T (I - T) \hat{\Sigma}^{-1}), \end{aligned}$$

where $\mathbb{R}^{-G} := \{T \in \mathbb{R}^{d \times d} : t_{k,l} = 0 \text{ if } l \notin d(k)\}$. This immediately follows if we recapitulate

$$\Sigma = \sigma^2 (I_d - B)^{-1} (I_d - B)^{-T} = \sigma^2 (I - T)^T (I - T),$$

where $T \in \mathbb{R}^{-G}$ represents the negative total causal effect between variables. Maximizing over the equal error variance σ^2 then leads to a linear least squares problem with the optimal solution

$$\min_{T \in \mathbb{R}^{-G}} \text{tr}((I - T)^T (I - T) \hat{\Sigma}^{-1}) = \sum_{k=1}^d (\hat{\Sigma}^{-1})_{k,k|d(k)}. \quad (7)$$

In the second case of an additional constraint of a fixed size total causal effect $\mathcal{C}(i \rightarrow j) = \psi$, we can similarly first optimize over the equal error variance σ^2 to derive a linear least squares problem. That is, for $\Sigma \in \mathcal{M}_\psi(G)$, maximizing the dual likelihood then corresponds to the linear least squares problem

$$\min_{T \in \mathbb{R}^{-G, \psi}} \text{tr}((I - T)^T (I - T) \hat{\Sigma}^{-1}).$$

Here, the search space is additionally restricted by the fixed size total causal effect $\mathcal{C}(i \rightarrow j) = \psi$ and given by $\mathbb{R}^{-G, \psi} := \{T \in \mathbb{R}^{d \times d} : t_{k,l} = 0 \text{ if } l \notin d(k), t_{i,j} = -\psi\}$. Solving this linear least squares problem with the constraint of one fixed parameter leads to the solution

$$\sum_{k \neq i}^d (\hat{\Sigma}^{-1})_{k,k|d(k)} + (\hat{\Sigma}^{-1})_{i,i|d(i) \setminus \{j\}} + \psi^2 (\hat{\Sigma}^{-1})_{j,j|d(i) \setminus \{j\}} + 2\psi (\hat{\Sigma}^{-1})_{i,j|d(i) \setminus \{j\}}. \quad (8)$$

Note that maximizing dual likelihood under constraints on total causal effects corresponds with the reciprocal property (4) to maximizing classical Gaussian likelihood under modified constraints on direct causal effects. Thus, while directly maximizing the classical Gaussian likelihood under constraints on the total causal effect is complicated due to the arising polynomial constraint on the parameters, we are able to obtain a closed-form solution for the maximum dual likelihood estimate since the constraint only pertains to one parameter.

4. Computation of Confidence Regions for Total Causal Effects

In the following section, we present important computational shortcuts as well as the results of a simulation study that analyzes the performance and computation times of our proposed confidence regions.

4.1. Computational Shortcuts

Our proposal of employing the dual likelihood to construct confidence regions for total causal effects resolves the need for numerical optimization routines of earlier methods and leads to a closed-form solution that significantly reduces computation times (see Figure 3). However, the bottleneck of the task itself is the superexponential growth of the number of possible causal structures with the number of nodes. Without any knowledge about the underlying causal relation, we have to consider all possible causal structures in order to make calibrated confidence statements that include the structure uncertainty. This, already at the scale of systems with 12 involved variables, entails accounting for more than 10^{26} possible graphs. Using combinatorial shortcuts that quickly reduce the search space of plausible causal structures is thus extremely important to compute confidence regions that solve this task.

First, we highlight that with our methodology, we only need to consider complete DAGs, thus reducing the search space to d factorial structures for d involved variables. Namely, we have to search through all permutations of d nodes, which then correspond to all possible topological orderings of the involved variables. Furthermore, we employ a branch and bound type search algorithm to quickly reduce the space of plausible causal orderings before performing the testing procedure for all possible causal effects $\psi \in \mathbb{R}$. The main idea of this procedure in order to immediately reject implausible causal orderings is the following. We quickly approximate the maximum dual likelihood under the alternative $\Omega \in \mathcal{M}^{-1}$ with the dual likelihood estimate (7) under the causal ordering obtained via recursively minimizing $(\hat{\Sigma}^{-1})_{k,k|d(k)}$ starting from the sink node. Then, recursively starting from the sink node, we search through the space of causal orderings in reverse and reject a partial ordering if the unrestricted partial dual likelihood estimate (7) already exceeds the threshold of the alternative. This is possible because the dual likelihood estimate (7) collapses into subproblems which only depend on descendants. Moreover, looking at (8), the specific ordering among the descendants of i is not relevant for the subsequent testing procedure for the precise total causal effects $\psi \in \mathbb{R}$. Therefore, in our branch and bound type search algorithm, it suffices to proceed with the partial ordering that achieves the highest dual likelihood among all plausible orderings with the same set of descendants $d(i)$. More specifically, at each step of our algorithm, we reject all partial orderings that do not achieve the highest dual likelihood among all plausible partial orderings of the same set of nodes and do not include i .

Combined, all these computational shortcuts vastly improve computation times and lead to a bottom-up procedure that allows us to compute the proposed confidence region in a reasonable time for moderate dimensions, see Section 4.2. We emphasize that considering the fast superexponential growth of the number of possible DAGs with the number of nodes, rigorously accounting for structure uncertainty over all DAGs is an intrinsically difficult task already in moderate dimensions. Further, we highlight the interesting distinction to the LRT-method of [Strieder and Drton \(2023\)](#), that by employing dual likelihood, our method leads to a bottom-up procedure, starting from the sink node, while their method is a top-down procedure, starting from the source node.

4.2. Simulation Study

In the following simulation study, we compare the width of the proposed confidence region as well as the computation times to the LRT-method by [Strieder and Drton \(2023\)](#). For more simulation results, we refer the reader to the [Appendix A](#). Our experiments are designed as follows. We generate 1000 synthetic data sets based on linear SCMs corresponding to randomly selected DAGs. The edge weights are sampled from a normal distribution $N(\beta, 0.1)$, and the error terms are sampled from a standard normal distribution. Then we compute the confidence regions for the total causal effect $\mathcal{C}(1 \rightarrow 2)$ with confidence level $\alpha = 0.05$. We repeat this procedure for a range of different average direct effect strengths β , sample sizes n , dimensions d , sparse or dense graphs, and for true non-zero total effects as well as for no total effect.

First, we note that the theoretical coverage guarantees of our proposed confidence regions are based on (conservative) asymptotic critical values. Nevertheless, the confidence regions achieve the desired coverage in all our simulation settings, even in low sample sizes. We explore this in [Table 1](#) in the [Appendix A](#), which shows the empirical coverage frequencies against the sample size. Furthermore, while the employed dual likelihood estimators are asymptotically efficient in the sense of having the same asymptotic variance as ordinary maximum likelihood estimates, differences in finite samples are possible. Thus, we compare the performance of our confidence regions to the existing maximum likelihood approach. In [Figure 1](#), we compare the average width of the non-zero part of the confidence regions, reflecting the remaining uncertainty about the numerical size of the effect. We show the results for $d = 10$ and $\beta = 0.5$ against the sample size. The main observation is that the difference in performance between both methods seems negligible. Similar observations can be made for other performance measures, which we defer to the [Appendix A](#). For example, we investigate how often the zero is included in the confidence regions when there is a true nonzero effect and observe no significant differences. Thus, both methods seem to be equally conclusive about the existence as well as the numerical size of the total causal effects. Further, the two methods also perform similarly under model misspecification of equal error variances and linearity.

However, the big advantage of the proposed dual likelihood method is the significantly reduced computation time. Employing the dual likelihood avoids time-consuming numerical optimization routines, and we immensely benefit from the closed-form solution calculated in [Theorem 2](#). [Figure 2](#) shows the computation times for $\beta = 0.5$ and sample size 1000 against the dimension. Our proposed confidence regions are faster to compute by a factor of up to 10^2 , without any compromises in terms of information value. This is even more apparent in [Figure 3](#), which compares the computation times for $\beta = 0.1$ and sample size 1000 against the dimension. In this setting, the generated data encompasses more structure uncertainty, given the lower average edge strength. Thus, the time difference between repeatedly applying numerical optimization algorithms versus directly computing the closed-form solution is more fatal, such that we are not able to compute the LRT-confidence regions in a reasonable time in dimensions beyond $d = 8$. In contrast, the computation times of our proposed method seem to vary less with the amount of structure uncertainty inherent in the data set.

Furthermore, we highlight the observation that dense systems generally lead to wider confidence regions compared to sparse systems. Intuitively, for dense graphs more paths contribute to the total causal effects on average, and thus, more uncertainty about the numerical size of the effect remains in the data. In contrast, data generated by sparse DAGs generally exhibits more uncertainty about the underlying causal structure, which explains the observed increased computation times.

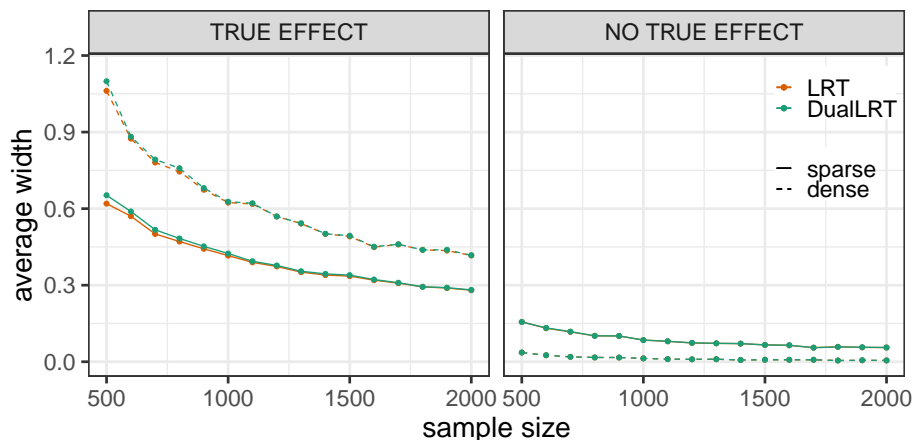


Figure 1: Mean width of 95%-confidence regions for the total causal effect in randomly generated 10-dim. DAGs (1000 replications).

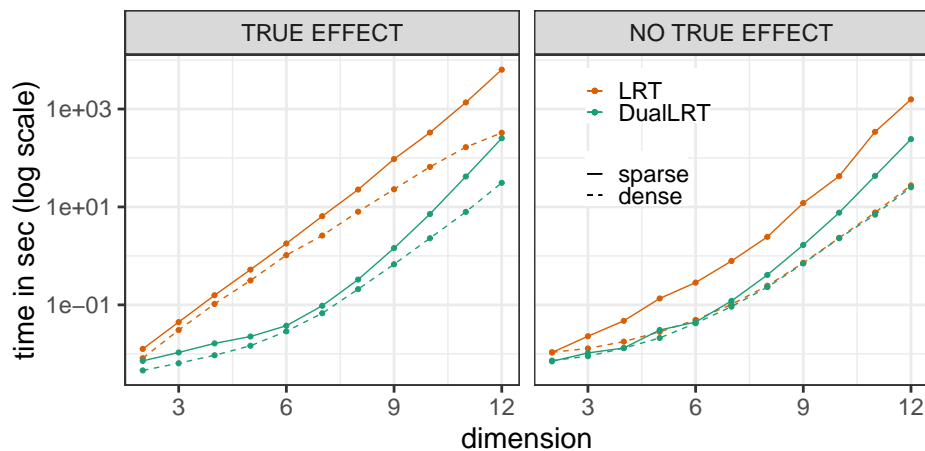


Figure 2: Mean computation times of 95%-confidence regions for the total causal effect in randomly generated DAGs with average edge strength $\beta = 0.5$ (1000 replications).

In summary, our simulation study validates that our proposed confidence regions based on the dual likelihood successfully pick up on the direction and the numerical size of the total causal effect while correctly quantifying the remaining uncertainty in structure as well as effect size. Furthermore, the computation times of our proposed confidence regions, computed with the proposed bottom-up procedure, are significantly lower than the competition. Here, we immensely benefit from the available closed-form solution due to employing dual likelihood. Especially under high structure uncertainty, avoiding repeated numerical optimization algorithms is crucial in order to compute the confidence regions in a reasonable time.

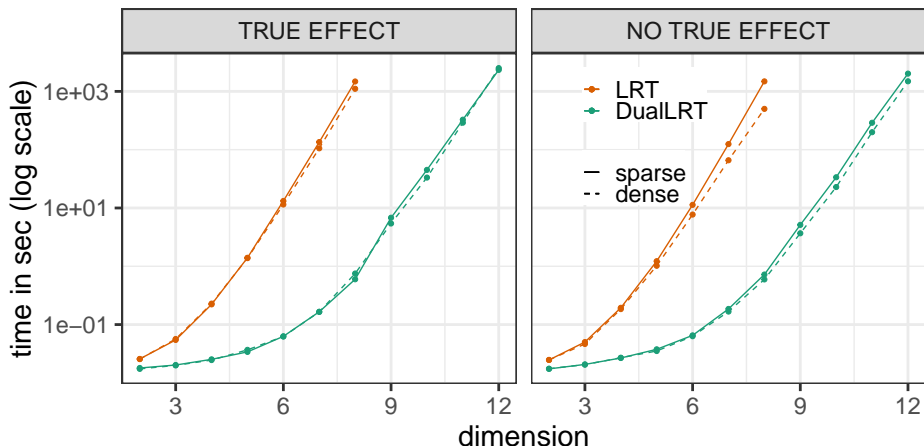


Figure 3: Mean computation times of 95%-confidence regions for the total causal effect in randomly generated DAGs with average edge strength $\beta = 0.1$ (1000 replications).

5. Conclusion

In this article, we address the important problem of rigorously accounting for uncertainty in causal inference when the causal structure itself has to be learned from data. We provide a closed-form solution for constructing confidence regions for the total causal effect that captures structure uncertainty as well as uncertainty about the numerical size of the effect. The confidence regions are developed for the concrete context of linear SCMs with equal error variances, where the underlying assumption of homoscedasticity across all interacting variables renders the causal structure identifiable. Thus, causal inference of the total causal effect without knowledge of the underlying structure is a well-defined task. An interesting conclusion of our work is that even in this specialized setting, we already observe high levels of uncertainty about causal directions and size of effects, which highlights the importance of methods that properly account for structure uncertainty in causal inference.

One thing to note is that our general approach of leveraging test inversions of joint tests for causal structure and effect size is generalizable to other settings. In particular, if one deals with parametric models for which the causal DAG is identifiable, then likelihood ratio tests may be deployed with the same stochastic approximation strategies as in this paper, and we anticipate that the matrix inversion interplay between direct and total effects that is behind our use of dual likelihood can be similarly exploited.

While unique structure identifiability is required in order to consistently resolve uncertainty about nonzero causal effects (e.g., when facing full Markov equivalence in linear Gaussian models without variance assumption, the confidence region would always include zero), our framework could in principle be adjusted to target all effects that correspond to graphs within a Markov equivalence class. However, in cases where the causal structure is merely identifiable up to some faithfulness-type assumption, more research needs to be done to fully understand which causal effects are entailed by a given (possibly unfaithful) distribution.

The observed reciprocal correspondence between dual likelihood with total effects and classical likelihood with direct effects could be exploited for other causal inference tasks. The idea is to leverage tools developed for direct effects and classical likelihood to causal inference results for total effects via the dual likelihood. In our simulation study, the differences in performance and explanatory power that stem from employing dual likelihood ratio tests instead of classical likelihood ratio tests are negligible and heavily outweighed by the significantly reduced computation times due to the available closed-form solution.

Finally, given the fast superexponential growth of the number of possible causal structures with the dimension of the system, it is essential to quickly reject implausible orderings and reduce the search space in order to compute confidence regions that account for uncertainty over all structures. The introduced computational shortcuts lead to a bottom-up procedure that, starting from sink nodes, recursively reduces the search space by immediately rejecting implausible partial orderings. In contrast, using the classical likelihood leads to a top-down procedure that starts with source nodes. Cleverly alternating between both procedures in a first step to reduce the search space might even further decrease the computation times of our proposed confidence regions.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 83818). Further, 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.

References

- Lawrence D. Brown. *Fundamentals of statistical exponential families with applications in statistical decision theory*, volume 9 of *Institute of Mathematical Statistics Lecture Notes—Monograph Series*. Institute of Mathematical Statistics, Hayward, CA, 1986.
- George Casella and Roger L. Berger. *Statistical inference*. The Wadsworth & Brooks/Cole Statistics/Probability Series. Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, CA, 1990.
- Wenyu Chen, Mathias Drton, and Y. Samuel Wang. On causal discovery with an equal-variance assumption. *Biometrika*, 106(4):973–980, 2019.
- Mathias Drton. Likelihood ratio tests and singularities. *Ann. Statist.*, 37(2):979–1012, 2009.
- Asish Ghoshal and Jean Honorio. Learning linear structural equation models in polynomial time and sample complexity. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84, pages 1466–1475. PMLR, 09–11 Apr 2018.
- Göran Kauermann. On a dualization of graphical Gaussian models. *Scand. J. Statist.*, 23(1):105–116, 1996.

Marloes Maathuis, Mathias Drton, Steffen Lauritzen, and Martin Wainwright, editors. *Handbook of graphical models*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL, 2019.

Judea Pearl. *Causality*. Cambridge University Press, Cambridge, second edition, 2009. Models, reasoning, and inference.

Jonas Peters and Peter Bühlmann. Identifiability of Gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, 2017.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, 2000.

David Strieder and Mathias Drton. Confidence in causal inference under structure uncertainty in linear causal models with equal variances. *J. Causal Inference*, 11(1), 2023.

David Strieder, Tobias Freidling, Stefan Haffner, and Mathias Drton. Confidence in causal discovery with linear causal models. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence. UAI'21*, volume 161, pages 1217–1226. PMLR, 27–30 Jul 2021.

Appendix A. Additional Simulation Results

In this section, we present additional results of our simulation study. The data generation process and general simulation setup are outlined in Section 4.2. Table 1 shows the observed coverage probabilities for 10-dimensional graphs. Both methods seem to be conservative and achieve the desired coverage frequency. Considering that we employ critical values based on an upper bound as well as the theory of intersection union tests for the arising testing problem, it is not surprising that the resulting confidence regions are conservative. Further, note that under high structure uncertainty, it is not possible to compute the competitor LRT in a reasonable time, which we denoted with NA.

As an additional performance measure, in Figure 4, we compare the proportions of times zero is included in the confidence regions when there is a true nonzero effect. The inclusion of zero in the confidence regions reflects the remaining uncertainty about the existence of the effect. We show the results for $d = 10$ and $\beta = 0.5$ against the sample size n as well as for $d = 10$ and sample size $n = 1000$ against the average direct effect strength β . In terms of performance, we observe no significant differences. However, we note again that for high structure uncertainty, which is the case for low average direct effect strength β , it is not possible to compute the competitor LRT in a reasonable time.

Furthermore, with practical applications in mind, Figure 5 investigates the robustness of the methods towards small deviations from equal error variances. Here, to generate data, we sample the error variances uniformly from $[1 - 0.5v, 1 + 0.5v]$, where v indicates the degree of deviation from homoscedasticity across all variables. We show the empirical coverage probabilities for $d = 10, \beta = 0.5$, and sample size $n = 1000$ against the degree of deviation v . Once again, we only observe negligible differences, and both methods indicate some robustness to small deviations from equal error variances.

Finally, to further explore aspects of model misspecification, we also investigate the robustness of the methods towards deviations from linear relations. We generate data as outlined before, however, with an increasing emphasis on an additional quadratic dependence structure via the relation $X_j = \sum_{i \in p(j)} \beta_i((1 + 3v)X_i + 0.02vX_i^2) + \varepsilon_j$, for all $j = 1, \dots, d$. Thus, v indicates the degree of deviation from linearity. The scaling ensures similar magnitudes of true causal effects across all investigated simulation settings. Note that for nonlinear relations, generally, the true total causal effect is not given by a single parameter but rather by a functional relation. However, when applying linear models it is natural to consider as a parameter of interest the slope parameter determining the best linear approximation within the true causal structure. In Figure 6, we consider this ‘pseudo-linear’ total effect and show the empirical coverage probabilities for $d = 10, \beta = 0.5$, and sample size $n = 1000$ against the degree of deviation v . We observe that even under nonlinearity, the methods pick up on the direction of the causal relation, and thus, for no causal effect, zero is always covered by the confidence intervals. Note that in the considered nonlinear additive equal variance setting, the causal ordering is still implied by ordering conditional variances. Further, the performance differences between both methods are negligible, and for nonzero effects, both methods indicate some robustness to small deviations from linearity.

Table 1: Empirical coverage of 95%-confidence regions for the total causal effect in randomly generated 10-dim. DAGs (1000 replications).

method	$n \setminus \beta$	TRUE EFFECT				NO TRUE EFFECT			
		SPARSE		DENSE		SPARSE		DENSE	
		0.1	0.5	0.1	0.5	0.1	0.5	0.1	0.5
DualLRT	500	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	1000	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	2000	1.00	0.99	0.99	1.00	1.00	1.00	1.00	1.00
LRT	500	NA	0.99	NA	1.00	NA	1.00	NA	1.00
	1000	NA	1.00	NA	1.00	NA	1.00	NA	1.00
	2000	NA	1.00	NA	1.00	NA	1.00	NA	1.00

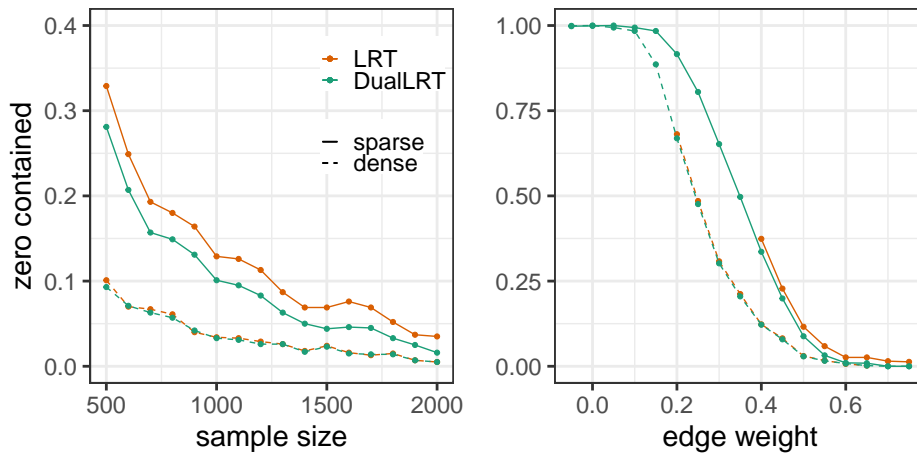


Figure 4: Proportion of times zero contained in 95%-confidence regions for the total causal effect in randomly generated 10-dim. DAGs with true non-zero effect (1000 replications).

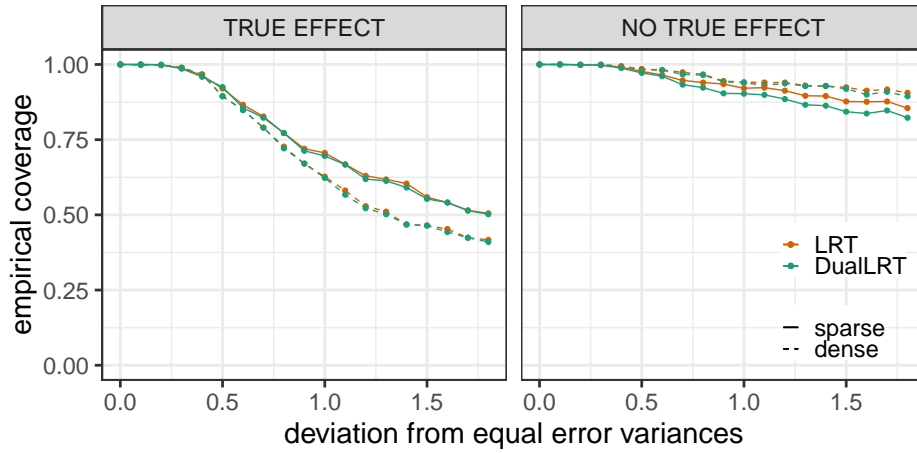


Figure 5: Empirical coverage of 95%-confidence regions for the total causal effect in randomly generated 10-dim. DAGs under departure from equal error variances (1000 replications).

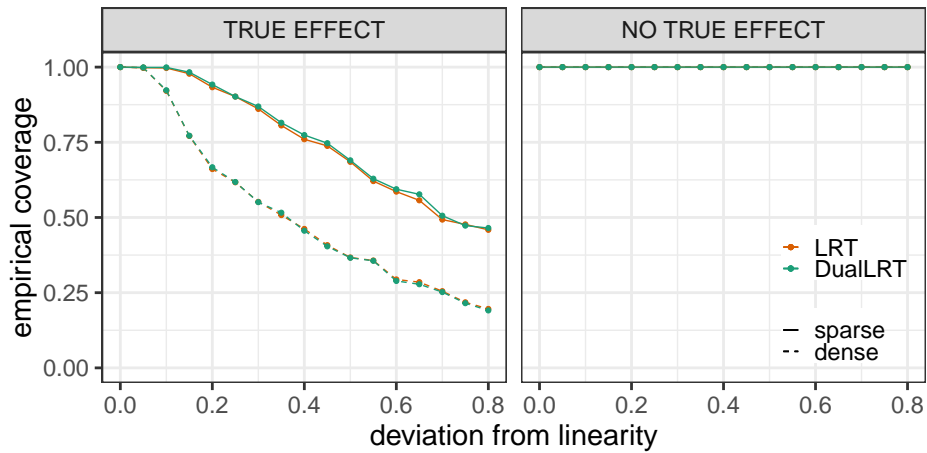


Figure 6: Empirical coverage of 95%-confidence regions for the total causal effect in randomly generated 10-dim. DAGs under departure from linearity (1000 replications).