

TSLINGAM

Citation for published version (APA):

Leyder, S., Raymaekers, J., & Verdonck, T. (2023). TSLiNGAM: DirectLiNGAM under heavy tails. Cornell University - arXiv. arXiv.org No. 2308.05422

Document status and date: Published: 10/08/2023

Document Version: Publisher's PDF, also known as Version of record

Please check the document version of this publication:

 A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.

• The final author version and the galley proof are versions of the publication after peer review.

 The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
 You may not further distribute the material or use it for any profit-making activity or commercial gain
 You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

TSLiNGAM: DirectLiNGAM under heavy tails

Sarah Leyder*

Department of Mathematics, University of Antwerp, Jakob Raymaekers Department of Quantitative Economics, Maastricht University, Tim Verdonck Department of Mathematics, University of Antwerp - imec

Department of Mathematics, KU Leuven

August 11, 2023

Abstract

One of the established approaches to causal discovery consists of combining directed acyclic graphs (DAGs) with structural causal models (SCMs) to describe the functional dependencies of effects on their causes. Possible identifiability of SCMs given data depends on assumptions made on the noise variables and the functional classes in the SCM. For instance, in the LiNGAM model, the functional class is restricted to linear functions and the disturbances have to be non-Gaussian.

In this work, we propose TSLiNGAM, a new method for identifying the DAG of a causal model based on observational data. TSLiNGAM builds on DirectLiNGAM, a popular algorithm which uses simple OLS regression for identifying causal directions between variables. TSLiNGAM leverages the non-Gaussianity assumption of the error terms in the LiNGAM model to obtain more efficient and robust estimation of the causal structure. TSLiNGAM is justified theoretically and is studied empirically in an extensive simulation study. It performs significantly better on heavy-tailed and skewed data and demonstrates a high small-sample efficiency. In addition, TSLiNGAM also shows better robustness properties as it is more resilient to contamination.

Keywords: Causal discovery, Efficiency, LiNGAM, Structural causal models

^{*}First author is supported by Fonds Wetenschappelijk onderzoek - Vlaanderen (FWO) as a PhD fellow Fundamental Research (PhD fellowship 11K5523N).

1 Introduction

Over the last decades, statistics and machine learning have proven to be very strong tools for finding and modeling associations in data sets. However, in recent years, it has become clear that when analyzing and using data, causal relations are often more valuable than associations. As a result, there has been a growing interest in causal inference in statistics and machine learning, and it has become a crucial tool in many empirical sciences including medicine, social sciences, neuroinformatics and biology.

One of the established approaches to causal inference builds on the directed acyclic graph (DAG) framework, studied in great depth by Pearl [2009]. DAGs represent the variables of interest as nodes in a graph where directed edges between nodes correspond with causal relations. As DAGs are acyclic, i.e. there are no cycles in the network, they represent a one-directional causal order of the variables, such that no variable with a later causal order can influence an earlier variable. DAGS are typically complemented by structural causal models (SCMs), which are used to describe the functional dependence of an effect on its causes in a DAG. This combined theoretical framework then bestows us with the necessary tools to compute observational, interventional and counterfactual distributions for answering all possible causal queries.

SCMs describe the causal relationships among random variables X_1, \ldots, X_p as a set of equations

$$X_i = f_i(\mathbf{P}_i, e_i), \qquad i = 1, \dots, p$$

where the variable set $P_i \subset \{X_1, \ldots, X_p\}$ denotes the parent variables of X_i , i.e. the variables that determine X_i , and e_i represents the random noise variable disturbing X_i . A central question in causal inference is whether the SCM can be identified from data. In other words, given enough data, can we find the DAG (and the functional dependencies) underlying the data generating process? It is known that this is impossible in full generality. However, under appropriate conditions on the behavior of the noise variables and the functional class of the f_i , it is indeed possible to recover the SCM from observational data alone [Peters et al., 2017]. This process is called causal discovery. It provides a very valuable addition to (randomized) controlled experiments, which are often difficult or impossible due

to prohibitive costs or ethical objections.

One prime example of conditions under which identifiability is possible, is the LiNGAM model [Shimizu et al., 2006]. In LiNGAM, the functional class is restricted to linear functions and the disturbances have to be non-Gaussian and mutually independent. Under these assumptions, identifiability is provable and given continuous data, the complete causal structure can be recovered. The above described model is known as the linear, non-Gaussian, acyclic model (LiNGAM) and the original discovery algorithm is based on independent component analysis (ICA), justified by the assumption of non-Gaussianity [Shimizu et al., 2006]. Several extension to the LiNGAM model have been made. LvLiNGAM [Hoyer et al., 2008] includes the presence of hidden variables or latent confounders by using overcomplete ICA, other methods resilient against latent confounders are ParceLiNGAM [Tashiro et al., 2014] and MLCLiNGAM [Chen et al., 2022]. Dai et al. [2022] consider causal discovery of the LiNGAM model in the presence of measurement errors. Hyvärinen et al. [2010] discuss the integration of LiNGAM in autoregressive models for time series.

In this work, we propose TSLiNGAM, a new algorithm for estimating the causal structure in a LiNGAM model. TSLiNGAM builds on DirectLiNGAM [Shimizu et al., 2011], which is a popular method to obtain the causal LiNGAM model based on simple OLS regressions, but relies on regression estimators which are more efficient under heavy tails and skewness. These alternative regression estimators are more natural given the nongaussianity assumption in the LiNGAM model, and their appropriateness is further motivated theoretically and empirically.

The remainder of the article is organized as follows. Section 2 briefly reviews DirectLiNGAM and introduces TSLiNGAM. It also discusses the theoretical and computational properties of TSLiNGAM. Section 3 then demonstrates the advantage of TSLiNGAM over DirectLiNGAM in an extensive simulation study. Lastly, in Section 4, we compare TSLiNGAM to DirectLiNGAM on four real data examples. Finally, Section 5 concludes.

2 Method

We start by reviewing the LiNGAM model and the DirectLiNGAM algorithm, before introducing TSLiNGAM.

2.1 Preliminaries

The LiNGAM structural causal model postulates that the functional dependencies are linear and the external influences are independent and non-Gaussian. More precisely, it relies on the following three assumptions:

- The generating process can be described by a directed acyclic graph such that the variables {X₁,...,X_p} can be arranged in a causal order. The causal order of the variable X_i is denoted by k(i).
- 2. Each variable is a linear combination of other variables with a lower causal order, plus an external influence:

$$X_i = \sum_{k(j) < k(i)} b_{ij} X_j + e_i$$

The coefficients b_{ij} , called the connection strengths, can be arranged into a matrix B, which can be permuted to strict lower triangularity since the generating process concerns a DAG. The noise terms e_i can be placed into a vector e. Hence we obtain the matrix notation:

$$X = BX + e \tag{1}$$

We call X_i an exogenous variable if X_i is equal to e_i , so no variable X_j has a directed path to X_i . In the DAG framework, there is always at least one exogenous variable. Non-exogenous variables are called endogenous variables.

3. The external influences e_i are continuous random variables following a non-Gaussian distribution with zero mean and non-zero variance and all the e_i for $i \in \{1, \ldots, p\}$ are independent of each other.

Given a data set, the underlying LiNGAM structure can be recovered by rewriting Equation (1) as:

$$X = Ae \tag{2}$$

where $A = (I - B)^{-1}$. Since the disturbance vector e contains mutually independent, non-Gaussian variables, Equation (2) corresponds to the well-known linear independent component analysis model (ICA). The matrix A is called the mixing matrix and efficient ICA-algorithms exist to estimate it for a given data set. Subsequently scaling and permutation steps can be performed to produce a strictly lower triangular matrix B, from which the corresponding causal order can then easily be derived. More information on the LiNGAM discovery algorithm can be found in the original LiNGAM paper [Shimizu et al., 2006].

The LiNGAM algorithm using ICA-estimation does, however, have some drawbacks. First, the optimization used for ICA can get trapped in a local minimum and hence we have no guaranteed computational stability for the method. Second, for the gradient-based algorithm, appropriate parameters must be selected which is not easily done.

In 2011, a direct method was proposed to estimate causal ordering in the linear non-Gaussian context, namely DirectLiNGAM [Shimizu et al., 2011]. In contrast to ICA-LiNGAM, this new method has guaranteed convergence and requires no parameter specification. DirectLINGAM uses two main ingredients. The first is OLS regression to remove the effect of an exogenous variable from the other variables. The second is an independence measure to identify the next exogenous variable. Denote with $r_i^{(j)} := X_i - \frac{\operatorname{cov}(X_i, X_j)}{\operatorname{var}(X_j)}X_j$ the ordinary least squares residual when X_i is regressed on X_j . Further denote the kernel-based estimator of mutual information [Bach and Jordan, 2002] with $\widehat{MI}_{\text{kernel}}$. For each variable X_j we sum the mutual information of it with each of its ordinary least squares residuals $r_i^{(j)}$ to obtain the kernel-based independence measure (KBI):

$$T_{\text{kernel}}(X_j, U) = \sum_{i \in U, \ i \neq j} \widehat{MI}_{\text{kernel}}(X_j, r_i^{(j)})$$
(3)

Here U is the set of indices of the remaining variables. The variable with the lowest T_{kernel} is then the most independent and will be used as the next exogenous variable.

In summary, the DirectLiNGAM algorithm proceeds as follows:

Require: $n \times p$ data set X	
$U \leftarrow \{1, \dots, p\}$	\triangleright Initialize the set of variable subscripts
$K \leftarrow \emptyset$	\triangleright Initialize an empty ordered list of variable subscripts
while K contains less than $p-1$ indices do	
for $j \in U \setminus K$ do	\triangleright Cycle through the variables in $U \setminus K$
for $i \in U \setminus (K \cup j)$ do	$\triangleright Cycle through the variables in U \setminus (K \cup j)$
$R_{\cdot,i}^{(j)} \leftarrow r_i^{(j)}$	\triangleright Store the OLS residuals of variable X_i on X_j
end for	
$T_j \leftarrow T_{\text{kernel}}(X_j, U \setminus K)$	
end for	
$m = \operatorname{argmin}_{j \in U \setminus K} T_j$	\triangleright Find the next exogenous variable
$K \leftarrow \{K,m\}$	\triangleright Append <i>m</i> to <i>K</i>
$X \leftarrow r^{(m)}, \mathbf{X} \leftarrow R^{(m)}$	\triangleright Consider the residuals as new input
end while	
$K \leftarrow \{K, (U \setminus K)\}$	\triangleright Append the final variable to obtain the complete causal ordering
$B \leftarrow \text{OLS}(\mathbf{X}, K)$	\triangleright perform OLS on ${\bf X}$ following the order in K

Algorithm 1 DirectLiNGAM algorithm [Shimizu et al., 2011]

The algorithm above can be extended to make use of prior knowledge on the structure if this is available. For more details, we refer to the DirectLiNGAM paper [Shimizu et al., 2011]. As is clear from the pseudo-code description in Algorithm 1, DirectLiNGAM relies crucially on least squares regression. In addition, the proofs for the identification of the LiNGAM structure by DirectLiNGAM also rely on the use of the least squares estimator [Shimizu et al., 2011].

The DirectLiNGAM algorithm as introduced so far only identifies the causal ordering and returns a fully connected DAG, which is the focus of this paper. In order to drop redundant edges, it can be followed by a sparse regression estimator, for which the adaptive lasso [Zou, 2006] was used by Shimizu et al. [2011].

2.2 TSLiNGAM

The reliance of DirectLiNGAM on OLS regression is counterintuitive. OLS is known to perform extremely well under independent Gaussian errors, but loses its superiority when the errors are skewed, heavy tailed or heteroscedastic, especially when data samples are small [Wilcox, 1998]. Given that the LiNGAM model assumes non-Gaussianity of the error terms, OLS is potentially a weak point of the algorithm.

In order to study this hypothesis, we propose the use of a different slope estimator to identify exogenous variables, namely the Theil-Sen regression estimator. This is motivated by its favorable properties on heavy-tailed and skewed distributions. Theil-Sen regression was first introduced by Theil [1950] and later extended by Sen [1968]. It is defined as follows:

Definition 1 (Theil-Sen slope). For the linear regression of a random variable Y on X, $Y = \beta X + e$, the Theil-Sen slope estimator is defined as

$$\hat{\beta} = med_{i,j} \frac{y_j - y_i}{x_j - x_i}, \qquad for \ x_j \neq x_i \tag{4}$$

for data pairs $\{(x_i, y_i) : i = 1, ..., n\}$.

The Theil-Sen slope estimator is unbiased, regression equivariant, robust with a breakdown value of 0.293 and a bounded influence function [Sen, 1968, Peng et al., 2008]. Compared to OLS, it has a high small-sample efficiency and it is super-efficient when combined with discontinuous or discrete errors [Wilcox, 1998, Peng et al., 2008]. Also, when the errors are (close to) normal, Theil-Sen only loses little efficiency compared to OLS.

To apply the Theil-Sen slope in the DirectLiNGAM algorithm, we need to justify its use theoretically by generalizing the lemmas in Shimizu et al. [2011]. For this, we need the functional form of the Theil-Sen given by

$$T(X,Y) = \operatorname{med}_{X,X',Y,Y'}\left(\frac{Y-Y'}{X-X'}\right) = F^{-1}(0.5),$$
(5)

where F denotes the distribution of $\frac{Y-Y'}{X-X'}$ with $Y \stackrel{d}{=} Y'$ and $X \stackrel{d}{=} X'$. Using this form, we first proof the Fisher Consistency of the Theil-Sen estimator, a property we will need later for generalizing the lemmas. It is defined as follows

Definition 2 (Fisher consistency). The functional T estimating a parameter Θ is Fisher consistent for a distribution F if, when calculating the functional on the whole population, it equals the estimated parameter:

for the distribution
$$F: T(F) = \Theta$$

It turns out that the Theil-Sen estimator is indeed Fisher consistent as shown in the result below

Theorem 1 (Fisher consistency of Theil-Sen slope). For a simple linear regression model $Y = \beta X + \varepsilon$ such that X and ε are independent, continuous random variables, the Theil-Sen slope is a Fisher consistent slope estimator.

Before proving the validity of the Theil-Sen slope in the LiNGAM model we need the additional concept of correlation-faithfulness [Shimizu et al., 2009]:

Definition 3 (Correlation-faithfulness). The distribution of (X_1, \ldots, X_p) is said to be correlation-faithful to the underlying graph if and only if the (conditional) correlations of the X_i 's are implicated by the graph structure.

Now suppose the data are realizations of a *p*-variate random vector $(X_1, \ldots, X_p) \sim F_p$. Lemma 1 then states that the Theil-Sen slope can successfully identify exogenous variables and generalizes Lemma 1 of Shimizu et al. [2011].

Lemma 1 (Generalization of Lemma 1 of Shimizu et al. [2011]). Suppose that the random variables X_1, \ldots, X_p strictly follow the LiNGAM assumptions and that their distribution is correlation-faithful. We consider slope estimators as functionals T acting on bivariate distributions $(X, Y) \sim F_2$. Assume that following properties hold for these slope functionals $T(F_2) = T(X, Y)$ when Y is regressed on X:

- 1. T is regression equivariant: $\forall \gamma \in \mathbb{R} : T(X, Y + X\gamma) = T(X, Y) + \gamma$
- 2. If X and Y are independent: $T(F_2) = T(X, Y) = 0$ (6)
- 3. $T(X,Y) \neq 0 \implies T(Y,X) \neq 0$

Define the residual when X_i is regressed on X_j using slope functional T as the following random variable: $r_i^{(j)} \coloneqq X_i - T(X_j, X_i)X_j$ $(i \neq j)$. Then the variable X_j is exogenous if and only if X_j is independent of $r_i^{(j)}$ for all $i \neq j$. In particular, this holds for the Theil-Sen slope.

Next, Lemma 2 shows that the LiNGAM model holds on the residuals after an exogenous variable is regressed out using the Theil-Sen slope.

Lemma 2 (Generalization of Lemma 2 of Shimizu et al. [2011]). Suppose that the random variables $X = (X_1, \ldots, X_p)^T$ strictly follow the LiNGAM assumptions and that their distribution is correlation-faithful. Assume that the variable X_j is exogenous and denote by $r^{(j)}$ the (p-1)-dimensional vector holding all the residuals when the X_i , $i \neq j$, are regressed on X_j using the a slope estimator satisfying the properties in (6). Then a LiNGAM holds for the residual random variables $r^{(j)}$: $r^{(j)} = B^{(j)}r^{(j)} + e^{(j)}$ or $r^{(j)} = A^{(j)}e^{(j)}$. Moreover, the causal order is preserved: $k_{r^{(j)}}(l) < k_{r^{(j)}}(m) \iff k(l) < k(m)$.

We conclude that using the Theil-Sen slope is effective at identifying exogenous variables and regressing out their effect, and hence our new method correctly identifies the underlying causal model under the LiNGAM assumptions. We will refer with TSLiNGAM (Theil-Sen LiNGAM) to the resulting algorithm which uses Theil-Sen regression for identifying exogenous variables.

2.3 Robustness

Theil-Sen regression is not only more efficient than OLS at heavy-tailed and skewed distributions, it is also more robust in the sense that it is more resilient against contamination in the data. As discussed, the Theil-Sen slope has a bounded influence function [Hampel et al., 1986], implying that the effect that a single outlying observation can have on the measure is limited. Furthermore, the considered slope has a breakdown value of 0.293, meaning that the slope is robust up to 29.3% contamination in the data [Rousseeuw and Leroy, 2005]. In contrast, the OLS has an unbounded influence function and a breakdown value of 0%. To further explore the effect of the robustness-efficiency trade-off, we additionally consider the repeated median, defined by Siegel [1982], for the identification of the exogenous variables. The repeated median is defined as:

Definition 4 (Repeated median slope). For the linear regression of a random variable Y on X, $Y = \beta X + e$, the repeated median slope estimator is defined as

$$\hat{\beta} = med_i med_{j \neq i} \frac{y_i - y_j}{x_i - x_j}, \qquad for \ x_j \neq x_i \tag{7}$$

for data pairs $\{(x_i, y_i) : i = 1, ..., n\}$.

The functional form of the repeated median is given by:

$$T(X,Y) = \underset{X,Y}{\operatorname{med}} \left(\underset{X',Y'}{\operatorname{med}} \left(\frac{Y-Y'}{X-X'} \right) \right)$$
with $Y \stackrel{\mathrm{d}}{=} Y', \ X \stackrel{\mathrm{d}}{=} X'$

$$(8)$$

Just as the Theil-Sen slope, the repeated median is unbiased and regression equivariant. However, the repeated median is more robust, with a bounded influence function and a breakdown value of 0.5. In order to use the repeated median for the identification of exogenous variables, we need to verify its plug-in theoretically. As the Fisher consistency of the repeated median has only been proven for symmetric errors, see Siegel [1982], we proof this property for general error distributions.

Theorem 2 (Fisher consistency of the repeated median). For a simple linear regression model $Y = \beta X + \varepsilon$ such that X and ε are independent, continuous random variables, the repeated median slope estimator is Fisher consistent.

As the repeated median is regression equivariant, Fisher consistent and makes use of medians, Lemma 1 and 2 also hold for this slope estimator. Hence it can thus also be used to identify exogenous variables in a LiNGAM structure.

We now illustrate the effect that a single outlier can have on LiNGAM methods in the following robustness experiment. The objective is to estimate the causal order of a simple two node DAG in the LiNGAM family:

$$\begin{cases} X_1 = e_1 \\ X_2 = X_1 + e_2 \end{cases}$$

Here e_1 and e_2 are distributed according to a Student-*t* distribution with 5 degrees of freedom. We generate n = 500 observations of this bivariate causal model and replace one observation by an outlier of value $(\pm 2^i, \pm 2^j)$ for $i, j \in \{0, 1, 2, ..., 10\}$. The causal direction is then estimated from the contaminated data using the original ICA-LiNGAM algorithm from the *pcalg* package [Kalisch, 2022], the DirectLiNGAM algorithm and our adapted versions using Theil-Sen and the repeated median. We iterate this process 100 times to get representative results. The outcome is shown in Figure 1. For ICA-LiNGAM and DirectLiNGAM, we observe that a single outlying observation has the ability to distort the discovery of the causal order, even in the simplest of causal models. Both plots show different regions for the outlier such that the obtained causal order is the inverse of the true causal direction. In contrast, when Theil-Sen (TSLiNGAM) or the repeated median are used to identify the exogenous variables, we notice that the causal order is always correctly estimated, regardless of the values of the added outlier.



Figure 1: Number of correctly estimated causal orders out of 100 runs per discovery method given the (x, y)-coordinate of added outlier.

According to Figure 1, using the repeated median to identify the exogenous variables seems promising. As discussed, it is a robust slope estimator, but at a cost of lower efficiency. Notwithstanding, the repeated median is still more efficient than OLS for various error distributions, e.g. skewed, heavy-tailed or discrete distributions. Therefore from now on, we will also consider the repeated median as an alternative slope estimator and we will compare its use to TSLiNGAM in an extensive simulation study in Section 3.

2.4 Computational considerations

Finally, we discuss the computational cost of DirectLiNGAM and the proposed TSLiNGAM. As Theil-Sen regression and the repeated median can be computed in $\mathcal{O}(n \log(n))$ time [Cole et al., 1989, Katz and Sharir, 1993, Stein and Werman, 1992, Matoušek et al., 1993], they do not add much to the computational cost when used instead of simple OLS (which requires $\mathcal{O}(n)$ time). Therefore, TSLiNGAM has a computational cost that is similar to that of DirectLiNGAM. The computational costs of both algorithms is in fact dominated by the independence measure for each remaining variable per iteration in the algorithm. In the original paper [Shimizu et al., 2011], the independence measure used in Equation (3) is the kernel-based estimator of mutual information defined by Bach and Jordan [2002]. Although this measure performs well, its accumulated computational cost can become rather large, as lots of Gram matrices with Gaussian kernels have to be computed.

Therefore, as an alternative independence measure, we consider the use of the distance correlation (dcorr) between random variables [Székely et al., 2007]. This measure, unlike Pearson's correlation, is zero if and only if the variables are independent. One of the advantages is that it can be computed in $O(n\log(n))$ time [Chaudhuri and Hu, 2019]. In contrast, the kernel-based independence measure computed on two variables has a computational complexity of $O(nM^2 + M^3)$, where $M (\ll n)$ is the maximal rank found by the low-rank approximations of the Gram matrices used in the algorithm for the independence measure. Therefore, only if $M = O(\sqrt{\log(n)})$, we obtain the same computational complexity as dcorr. Empirically, M often seems to grow substantially faster than this rate. Hence, when the data sets are too large to use the kernel-based independence measure, it can be beneficial to use the distance correlation instead to speed up DirectLiNGAM and TSLiNGAM.

3 Simulation

In this section we compare the proposed method with direct competitors.

3.1 Setup and methods

We compare TSLiNGAM with the original DirectLiNGAM algorithm and with extremal ancestral search (EASE) [Gnecco et al., 2021]. The latter method is designed for heavytailed data, and is thus a highly relevant competitor to TSLiNGAM. In addition, we also compare with three variations of the proposed algorithm. The first uses TS regression with dcorr as independence measure. The other two use RM regression paired with KBI and dcorr respectively.

The data is generated by the following procedure, inspired by the simulation setup of Gnecco et al. [2021]:

- 1. We simulate data of dimension p = 2 with sample sizes $n = \{5, 10, 25, 50, 100\}$, of dimension p = 5 with sample sizes $\{n = 30, 50, 100, 200\}$ and of dimension p = 10 with sample sizes $\{n = 50, 100, 200, 300\}$.
- 2. We generate a linear structural causal model X = BX + e with $X = (X_1, \ldots, X_p)^T$, $e = (e_1, \ldots, e_p)^T$ and $B \in \mathbb{R}^{p \times p}$ as follows:
 - (a) First, we generate a random causal order between the variables X_1, \ldots, X_p as a permutation π of $\{1, \ldots, p\}$.
 - (b) Per variable X_i with $\pi(i) > 1$, the number of parents is distributed as $Bin(\pi(i) 1, q)$ with $q = \{1, 0.6, 0.5\}$ respectively for dimension $p = \{2, 5, 10\}$.
 - (c) Next, we select those parents randomly from the variables with a lower causal order, such that cycles are ruled out.
 - (d) Then, we sample the connection strengths B_{ij} per variable per parent uniformly from $[-0.9, -0.1] \cup [0.1, 0.9]$. This yields the matrix B.
 - (e) Finally, we sample the noise variables randomly from following distributions: Student-t with 1, 2 or 5 degrees of freedom, a centered lognormal distribution, a centered Pareto distribution and a centered exponential distribution. Combining B and e, we obtain X.

For each setting, we generate 1000 data sets. To compare performance among the different methods, we count the number of times the algorithm returns the correct causal order. All implementations are done in R. For the Theil-Sen slope and the repeated median we use the corresponding functions from the *robslopes* package [Raymaekers, 2022, 2023]. For the distance correlation we use the *dccpp* package Berrisch [2022] and for EASE we use the implementation in the *causalXtreme* package [Gnecco, 2021].

3.2 Results

We discuss the results for p = 10 variables here. The results for p = 2 and p = 5 are qualitatively similar and can be found in Section B of the Appendix. The results for p = 10 are presented in Table 1.

We observe that in almost every scenario the proposed TSLiNGAM achieves the best result. TSLiNGAM strongly outperforms DirectLiNGAM when the data is heavy-tailed. For skewed distributions, TSLiNGAM also performs better than DirectLiNGAM, although the difference is somewhat smaller. As the distribution moves closer to normality, such as for the t_5 distribution, DirectLiNGAM becomes the preferred method. However, note that the difference in performance is almost negligible and perhaps more importantly, the absolute performance is very poor. This is explained by the fact that the identifiability of the LiNGAM structure is lost when there are Gaussian errors, and as we move closer to that scenario, it becomes increasingly difficult to identify the underlying structure. EASE does not perform well here. This is probably explained by the fact that EASE only looks at the tails and therefore needs bigger sample sizes in order to perform well. Finally, we consider the variations of the TSLiNGAM algorithm. The repeated median performs well on heavytailed distributions, but does not offer an improvement over TSLiNGAM. Using dcorr as independence measure becomes a viable strategy when the sample size is reasonably large.

3.2.1 Computation time

In addition to the results on the recovery of the underlying LiNGAM structure, we study the computation times of the methods. We discuss the computation times for the simulation study with p = 10 and for the t_5 distribution here. The other distributions had similar computational costs. The computation times for p = 2 and p = 5 are qualitatively similar and can be found in Section B of the Appendix.

Table 2 presents the computation times for p = 10 and the t_5 distribution. The first thing to note is that TSLiNGAM and its variants have essentially the same computational cost as DirectLiNGAM. This is explained by the fact that the computation time of both algorithms is dominated by the kernel-based independence measure. As a result, when

sample size	50	100	200	300	50	100	200	300
		t	1					
TSLiNGAM	477	806	942	984	539	892	983	1000
DirectLiNGAM	286	432	555	640	368	552	733	840
EASE	21	124	307	450	3	10	17	27
Repeated Median & KBI	379	733	915	969	365	759	961	981
Theil-Sen & dcorr	232	530	747	841	361	740	942	982
Repeated Median & dcorr	152	425	657	791	178	548	877	947
		t	2					
TSLiNGAM	167	387	601	711	457	798	959	992
DirectLiNGAM	149	333	539	642	351	649	859	924
EASE	3	13	52	81	0	3	8	12
Repeated Median & KBI	128	348	592	663	309	675	919	981
Theil-Sen & dcorr	38	189	519	733	273	703	932	980
Repeated Median & dcorr	21	133	427	651	156	514	848	937
		t	5			expo	nential	
TSLiNGAM	17	35	88	151	236	601	858	942
DirectLiNGAM	16	37	90	158	245	575	812	903
EASE	0	1	5	1	0	0	1	1
Repeated Median & KBI	11	34	72	132	155	482	818	896
Theil-Sen & dcorr	2	1	7	11	140	542	867	941
Repeated Median & dcorr	0	2	4	16	66	360	746	867

Table 1: Number of correct causal orders out of 1000 trials for p = 10.

using dcorr as a measure of independence, we see a substantial speedup of about one order of magnitude. This suggests that dcorr is useful when the sample size gets larger, which is precisely the scenario in which its performance is also competitive with TSLiNGAM. Finally, we note that EASE is by far the fastest method here. However, as pointed out before, it is not competitive in these relatively small-sample scenarios.

sample size	50	100	200	300
TSLiNGAM	15.68	19.32	24.68	30.19
DirectLiNGAM	14.64	17.98	22.90	27.28
EASE	0.03	0.03	0.04	0.04
Repeated Median & KBI	15.98	19.83	26.05	32.83
Theil-Sen & dcorr	0.96	1.54	2.62	3.62
Repeated Median & dcorr	1.45	2.64	4.43	6.61

Table 2: Computational time in minutes for 1000 runs for p = 10 for a Student-*t* distribution with 5 degrees of freedom.

4 Real data applications

In this section we illustrate the application of TSLiNGAM on four data sets from medical and social sciences.

4.1 Physician data

As a first real data example, we consider data originating from the US National Medical Expenditure Survey conducted in 1987 and 1988. This data contains health-related information on 4406 individuals and can be found at http://qed.econ.queensu.ca/jae/1997-v12.3/deb-trivedi/ or in the R-package *AER* [Kleiber and Zeileis, 2008] as the data set NMES1988. We work with the following variables: age, school (years of education), income (family income), chronic (number of chronic conditions), visits (number of physician office visits) and hospital (number of hospital stays).

We compare TSLiNGAM with the standard DirectLingam to find the causal structure. To prune redundant edges in the resulting adjacency matrices B, we perform Adaptive Lasso, as is done in Shimizu et al. [2011]. This results in the directed acyclic graphs shown in Figure 2.

The causal order found by DirectLiNGAM is (hospital, chronic, visits, age, income, school). This order is not very logical. We would for example expect that the number of

chronic conditions and a person's age are causes of the number of hospital stays. Additionally, years of schooling should have an impact on a persons income. In contrast, the causal order found by TSLiNGAM is (age, school, income, chronic, visits, hospital). This order is very logical and corresponds with our intuition. Furthermore, the causal graph found by TSLiNGAM consists of edges that match our understanding of the variables.



Figure 2: Directed acyclic graphs found by DirectLiNGAM (left) and TSLiNGAM (right).

The better result obtained by TSLiNGAM can be explained by studying the underlying variables of the data set. We know that TSLiNGAM tends to outperform DirectLiNGAM on heavy-tailed and skewed data, and it turns out that the variables visits, hospital and income are leptokurtic and have very fat tails, see for instance the boxplots in Figure 3.

4.2 GAGurine data

The second data set considered in this work is the GAG data from the package *MASS* in R [Venables and Ripley, 2002]. This data contains the concentration of the chemical GAG in the urine of 314 children between 0 and 17 years old, where it is known that age is a dominant cause of the concentration of GAG.



Figure 3: Boxplots of the heavy-tailed variables in the physician data set.



Figure 4: Ground truth: age is a cause of GAG concentration in urine.

On this data set, both DirectLiNGAM and TSLiNGAM succeed in discovering the right causal order. However, as we would like to demonstrate the small-sample efficiency of TSLiNGAM, we sample 1000 data subsets of sizes $\{5, 10, 15, \ldots, 50\}$ from the original data. The number of times TSLiNGAM and DirectLiNGAM find the right causal order on these subsamples are shown in Figure 5. Overall we observe that TSLiNGAM has a 10% higher small-sample efficiency compared to DirectLiNGAM. For example, for sample size n = 45, TSLiNGAM finds the right causal order 762 times, while DirectLiNGAM only succeeds 657 times. This increase in efficiency can be explained by observing that the distribution of the variable GAG is right-skewed and tailed, a scenario where the Theil-Sen slope is better suited than OLS.

4.3 FMRI data

As a third data set, we study the functional magnetic resonance imaging (FMRI) data simulated in Smith et al. [2011]. This data was previously studied within the LiNGAM



Figure 5: Number of correctly found causal orders by TSLiNGAM and DirectLiNGAM for 1000 runs on subsamples of the GAGurine data with specified sample sizes.

framework [Smith et al., 2011, Hyvärinen and Smith, 2013, Tashiro et al., 2014], however we now use the FMRI data in a robustness context. We take the first simulation data set from the paper which contains 10.000 continuous observations of 5 variables and has a causal structure as demonstrated in Figure 6.



Figure 6: The causal order of the FMRI data is (1,2,3,4,5).

If we run DirectLiNGAM or TSLiNGAM on the original data set, both methods succeed in discovering the correct causal order. However, to demonstrate the robustness of

TSLiNGAM, we artificially contaminate 20 observations by changing the value of the first variable to a number generated from a Student-t distribution with 3 degrees of freedom, centered around 25. We then run both algorithms 100 times again and observe that DirectLiNGAM never finds the right causal order. TSLiNGAM, in contrast, recovers the ground truth 72 times and is therefore much less influenced by the 20 outlying observations. This shows that TSLiNGAM using the Theil-Sen slope instead of OLS is more resilient towards contamination.

4.4 Sociological data

As a final real data example, we try to replicate the example performed in the original DirectLiNGAM paper on sociology data. The data are publicly available at the General Social Survey (GSS: https://gssdataexplorer.norc.org/gss_data). We study the variables father's occupation level, son's income, father's education, son's occupation level, son's education and number of siblings. We take the same subset of the data as studied by Shimizu et al. [2011]: non-farm background, ages 35 to 44, white, male, in the labor force at the time of the survey and years 1972 to 2006, see Table 3 for details. After omitting observations which contain missing values, this results in a data set with 2117 observations.

	studied variables	GSS codebook name	selection variables	GSS codebook name
X_1	Father's occupation level	PAPRES10	non-farm background	RES16
X_2	Son's income	REALRINC	age	AGE
X_3	Father's education	PAEDUC	white	RACE
X_4	Son's occupation level	PRESTG10	sex	SEX
X_5	Son's education	EDUC	in the labor force	WRKSTAT
X_6	Number of siblings	SIBS	year	YEAR

Table 3: Studied variables taken from the GSS repository and which variables we selected our sample on.

Domain knowledge on the causal relations between these variables suggests the causal structure shown in Figure 7 [Shimizu et al., 2011].



Figure 7: Ground truth based on domain knowledge: a directed arrow indicates a possible causal relation, a bidirected purple arrow signifies an unknown causal relation (there might be a relation, a latent confounder or nothing).

On this data we then run the DirectLiNGAM algorithm and our TSLiNGAM. To prune redundant edges in the resulting adjacency matrices B, we again perform Adaptive Lasso. This yields the directed acyclic graphs shown in Figures 8 and 9. The two discovered DAGs are fairly good. DirectLiNGAM finds 5 correct edges, 2 wrongly directed edges and 2 redundant edges. TSLiNGAM finds 4 correct edges, 1 wrongly directed edge and no redundant edges. Overall, both methods perform equally well on the sociological data.

Additionally we remark that the Theil-Sen estimator combined with the distance correlation gave a better outcome, see the resulting DAG in Figure 10 of the Appendix. This combination of slope estimator and independence measure yields 6 correct edges, 1 wrongly directed edge and 2 redundant edges, the best result yet.

5 Conclusion

In this work, we proposed TSLiNGAM, which builds on the popular DirectLiNGAM algorithm for causal discovery in LiNGAM structures. We proved that TSLiNGAM recovers the underlying causal LiNGAM structure by building on a new Fisher consistency result of the



Figure 8: Causal graph found by DirectLiNGAM with correct arrows in black, wrongly directed arrows in red, unverifiable arrows in purple and redundant arrows in blue.



Figure 9: Causal graph found by TSLiNGAM with correct arrows in black, wrongly directed arrows in red and unverifiable arrows in purple (there are no redundant arrows).

Theil-Sen slope. By leveraging the attractive properties of the Theil-Sen slope estimator, we obtain improved recovery under heavy tailed and skewed data models, without sacrificing performance in models with symmetric distributions which are close to normal. This improved performance was illustrated in an extensive simulation study. Furthermore, we suggested considering a different independence measure in the algorithm. More precisely, using the distance correlation instead of the original kernel-based independence measure reduces the overall computational cost of the method without sacrificing performance, provided the data set is large enough.

We additionally illustrated the TSLiNGAM on four real data sets. These applications confirm better performance under skewness and heavy tails, an improved small-sample efficiency and increased robustness to outliers compared to the original DirectLiNGAM algorithm.

In summary, we conclude that the newly developed method can be considered a premium alternative to DirectLiNGAM. TSLiNGAM performs significantly better on heavy-tailed data and discovers the right causal order on smaller sample sizes, without increasing the computational cost. In addition, TSLiNGAM also showed better robustness properties as it is more resilient to contamination.

Disclosure statement

The authors report there are no competing interests to declare.

SUPPLEMENTARY MATERIAL

appendix: Document containing proofs, additional simulation results and figures. (.pdf file)

References

- F. R. Bach and M. I. Jordan. Kernel independent component analysis. Journal of Machine Learning Research, 3:1–48, 07 2002.
- J. Berrisch. dccpp: Fast Computation of Distance Correlations, 2022. URL https://CRAN. R-project.org/package=dccpp. R package version 0.0.2.
- A. Chaudhuri and W. Hu. A fast algorithm for computing distance correlation. *Computational Statistics & Data Analysis*, 135:15–24, 2019.

- W. Chen, R. Cai, K. Zhang, and Z. Hao. Causal discovery in linear non-gaussian acyclic model with multiple latent confounders. *IEEE Transactions on Neural Networks and Learning Systems*, 33(7):2816–2827, 2022.
- R. Cole, J. S. Salowe, W. L. Steiger, and E. Szemerédi. An optimal-time algorithm for slope selection. SIAM Journal on Computing, 18(4):792–810, 1989. URL https://doi. org/10.1137/0218055.
- H. Dai, P. Spirtes, and K. Zhang. Independence testing-based approach to causal discovery under measurement error and linear non-gaussian models. https://arxiv.org/abs/ 2210.11021, 2022.
- N. Gnecco. *causalXtreme: Causal discovery in heavy-tailed models*, 2021. URL https://github.com/nicolagnecco/causalXtreme. R package version 0.0.0.9000.
- N. Gnecco, N. Meinshausen, J. Peters, and S. Engelke. Causal discovery in heavy-tailed models. *The Annals of Statistics*, 49, 06 2021.
- F. R. Hampel, E. M. Ronchetti, P. Rousseeuw, and W. A. Stahel. *Robust statistics: the approach based on influence functions*. Wiley-Interscience; New York, 1986.
- P. O. Hoyer, S. Shimizu, A. J. Kerminen, and M. Palviainen. Estimation of causal effects using linear non-gaussian causal models with hidden variables. *International Journal of Approximate Reasoning*, 49(2):362–378, 2008.
- A. Hyvärinen, K. Zhang, S. Shimizu, and P. O. Hoyer. Estimation of a structural vector autoregression model using non-gaussianity. *Journal of Machine Learning Research*, 11 (56):1709–1731, 2010.
- A. Hyvärinen and S. M. Smith. Pairwise likelihood ratios for estimation of non-gaussian structural equation models. *Journal of Machine Learning Research (JMLR)*, 14, 2013.
- M. Kalisch. *pcalg: Methods for Graphical Models and Causal Inference*, 2022. URL https: //cran.r-project.org/web/packages/pcalg/. R package version 2.7-7.

- M. J. Katz and M. Sharir. Optimal slope selection via expanders. Information Processing Letters, 47(3):115–122, 1993. URL https://doi.org/10.1016/0020-0190(93) 90234-Z.
- C. Kleiber and A. Zeileis. *Applied Econometrics with R (AER)*. Springer-Verlag, 2008. URL https://CRAN.R-project.org/package=AER.
- J. Matoušek, D. M. Mount, and N. S. Netanyahu. Efficient randomized algorithms for the repeated median line estimator. In *Proceedings of the Fourth ACM-SIAM Annual* Symposium on Discrete Algorithms, pages 74–82, 1993.
- J. Pearl. *Causality: Models, Reasoning and Inference.* Cambridge University Press, 2 edition, 2009.
- H. Peng, S. Wang, and X. Wang. Consistency and asymptotic distribution of the theil–sen estimator. Journal of Statistical Planning and Inference, 138(6):1836–1850, 2008.
- J. Peters, D. Janzing, and B. Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- J. Raymaekers. robslopes: Fast Algorithms for Robust Slopes, 2022. URL https://CRAN. R-project.org/package=robslopes. R package version 1.1.2.
- J. Raymaekers. The r journal: robslopes: Efficient computation of the (repeated) median slope. *The R Journal*, 14:38–49, 2023. ISSN 2073-4859.
- P. J. Rousseeuw and A. M. Leroy. Robust regression and outlier detection. John wiley & sons, 2005.
- P. K. Sen. Estimates of the regression coefficient based on kendall's tau. Journal of the American Statistical Association, 63(324):1379–1389, 1968.
- S. Shimizu, P. O. Hoyer, A. Hyvarinen, and A. Kerminen. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7:2003–2030, 10 2006.

- S. Shimizu, A. Hyvarinen, Y. Kawahara, and T. Washio. A direct method for estimating a causal ordering in a linear non-gaussian acyclic model. UAI, 2009. URL https: //arxiv.org/ftp/arxiv/papers/1408/1408.2038.pdf.
- S. Shimizu, T. Inazumi, Y. Sogawa, A. Hyvarinen, Y. Kawahara, T. Washio, P. O. Hoyer, and K. Bollen. Directlingam: A direct method for learning a linear non-gaussian structural equation model. *Journal of Machine Learning Research*, 12:1225–1248, 04 2011.
- A. F. Siegel. Robust regression using repeated medians. *Biometrika*, 69(1):242–244, 1982.
- S. M. Smith, K. L. Miller, G. Salimi-Khorshidi, M. Webster, C. F. Beckmann, T. E. Nichols, J. D. Ramsey, and M. W. Woolrich. Network modelling methods for fmri. *NeuroImage*, 54:875–891, 2011.
- A. Stein and M. Werman. Finding the repeated median regression line. In SODA '92, 1992.
- G. J. Székely, M. L. Rizzo, and N. K. Bakirov. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35(6):2769 – 2794, 2007. doi: 10.1214/ 009053607000000505.
- T. Tashiro, S. Shimizu, A. Hyvärinen, and T. Washio. ParceLiNGAM: A Causal Ordering Method Robust Against Latent Confounders. *Neural Computation*, 26(1):57–83, 2014.
- H. Theil. A rank-invariant method of linear and polynomial regression analysis. In Proceedings of Koninklijke Nederlandse Akademie van Wetenschappen, 1950.
- W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S*, fourth edition, 2002. URL https://www.stats.ox.ac.uk/pub/MASS4/.
- R. R. Wilcox. Simulation results on extensions of the theil-sen regression estimator. Communications in Statistics - Simulation and Computation, 27(4):1117–1126, 1998.
- H. Zou. The adaptive lasso and its oracle properties. *Journal of the American statistical* association, 101(476):1418–1429, 2006.

A Proofs

A.1 Proof of Lemma 1

Proof.

1) Assume X_j is exogenous:

If X_j is exogenous we have $X_j = e_j$. From X = Ae we have $X_i = a_{ij}X_j + \sum_{h \neq j} a_{ih}e_h$ $(i \neq j)$. Here X_j is independent of $\sum_{h \neq j} a_{ih}e_h$ since $X_j = e_j$ and all the e_i are mutually independent. Then:

$$r_i^{(j)} = X_i - T(X_j, X_i)X_j$$

= $a_{ij}X_j + \sum_{h \neq j} a_{ih}e_h - T(X_j, a_{ij}X_j + \sum_{h \neq j} a_{ih}e_h)X_j$
= $a_{ij}X_j + \sum_{h \neq j} a_{ih}e_h - (a_{ij} + T(X_j, \sum_{h \neq j} a_{ih}e_h))X_j$
= $a_{ij}X_j + \sum_{h \neq j} a_{ih}e_h - (a_{ij} + 0)X_j$
= $\sum_{h \neq j} a_{ih}e_h$

where in the second equality we substituted X_i , in the third equality we used regression equivariance and in the fourth equality we used the independence property. We obtain that X_j is independent of $r_i^{(j)}$ for all $i \neq j$ as X_j is independent of $\sum_{h \neq j} a_{ih} e_h$.

2) Assume X_j is not exogenous:

If X_j is endogenous, then there exists an exogenous variable $X_h = e_h$ such that X_h has a

directed path to X_j :

$$\begin{aligned} r_h^{(j)} &= X_h - T(X_j, X_h) X_j \\ &= e_h - T(X_j, X_h) \left(\sum_{k(t) < k(j)} a_{jt} e_t + e_j \right) \\ &= (1 - T(X_j, X_h) a_{jh}) e_h - T(X_j, X_h) \sum_{k(t) < k(j), t \neq h} a_{jt} e_t - T(X_j, X_h) e_j \end{aligned}$$

and

$$X_{j} = \sum_{k(t) < k(j)} a_{jt}e_{t} + e_{j} = \sum_{k(t) < k(j), t \neq h} a_{jt}e_{t} + a_{jh}e_{h} + e_{j}$$

If we now proof that $T(X_j, X_h)$ is nonzero, then the Darmois-Skitovitch theorem gives us that $r_h^{(j)}$ and X_j are dependent as all the e_k are independent and non-Gaussian. For this, we proceed as follows. First we have that

$$T(X_h, X_j) = T(e_h, \sum_{k(t) < k(j), t \neq h} a_{jt}e_t + a_{jh}e_h + e_j) = a_{jh}$$

using independence and regression equivariance. Second one has that

$$\operatorname{cov}(X_h, X_j) = \operatorname{cov}(e_h, \sum_{k(t) < k(j), t \neq h} a_{jt}e_t + a_{jh}e_h + e_j) = a_{jh}\operatorname{var}(e_h)$$
$$\implies a_{jh} = \operatorname{cov}(X_h, X_j)/\operatorname{var}(e_h) \text{ as } \operatorname{var}(e_h) > 0$$

Here $cov(X_h, X_j)$ cannot be zero under the correlation-faithfulness assumption as we have a directed path connecting them. Hence $T(X_h, X_j) = a_{jh}$ is nonzero. If we now use the third assumption in (6) we have that $T(X_j, X_h)$ is also nonzero and we are done.

For the Theil-Sen slope assumptions 1 and 2 from (6) are immediately satisfied by using the stronger property of Fisher consistency. For the third assumption we proceed as follows. First we note that the Theil-Sen slope uses medians to estimate the regression slope between variables Y and X. For these we have that the median of $\frac{Y-Y'}{X-X'}$ is nonzero if and only if the median of $\frac{X-X'}{Y-Y'}$ is nonzero. To see this: suppose without loss of generality that the median of the slopes $\frac{Y-Y'}{X-X'}$ is larger than zero. Then, as all slopes larger than zero yield an inverse slope larger than zero, and likewise all slopes smaller than zero yield an inverse slope smaller than zero, the median of $\frac{X-X'}{Y-Y'}$ is also larger than zero. Hence the median preserves the sign when swapping numerator and denominator and as we assume continuous variables, division by zero when swapping occurs with a negligible probability of zero. Hence $T(X, Y) \neq 0 \implies T(Y, X) \neq 0$ and we are done.

A.2 Proof of Lemma 2

Proof.

Assume that the mixing matrix A in X = Ae has already been permuted to lower triangularity with ones on the diagonal and assume without loss of generality that $X_j = X_1 = e_1$. Since X_1 is exogenous, we have that the a_{i1} for $i \neq 1$ are the slope coefficients when X_i is regressed on X_1 using T:

$$T(X_1, X_i) = T(e_1, \sum_{t \le i} a_{it}e_t)$$

= $a_{i1} + T(e_1, \sum_{t \le i, t \ne 1} a_{it}e_t)$
= $a_{i1} + 0 = a_{i1}$

where we used regression equivariance in the second equality and independence in the third equality. Hence when we remove the effect of X_1 from X_i by switching to the residuals $r_i^{(1)} = X_i - T(X_1, X_i)X_1 = X_i - a_{i1}e_1$, the first column of A becomes a zero vector. As $r^{(1)}$ is independent of X_1 , we get for $r^{(1)}$ a new $(p-1) \times (p-1)$ dimensional lower triangular matrix $A^{(1)} = [A]_{2:p,2:p}$ with ones on the diagonal: $r^{(1)} = A^{(1)}e^{(1)}$ with $e^{(1)} = [e]_{2:p}$. Therefore a LiNGAM holds for the residual vector $r^{(1)}$.

Also, when switching to $r^{(1)}$, the corresponding matrix $A^{(1)}$ is the lower triangular submatrix formed by removing the first row and the first column of A. Hereby the causal order is not altered and hence switching to the residuals preserves the causal order.

A.3 Proof of Theorem 1

Proof.

We have that:

$$Y - Y' = \beta X + \varepsilon - \beta X' - \varepsilon' = \beta (X - X') + \varepsilon - \varepsilon', \quad \text{with } \varepsilon \stackrel{d}{=} \varepsilon'$$
$$\implies \frac{Y - Y'}{X - X'} = \beta + \frac{\varepsilon - \varepsilon'}{X - X'}$$

Hence the Theil-Sen slope is Fisher-consistent:

$$\iff \operatorname{med}\left(\frac{\varepsilon - \varepsilon'}{X - X'}\right) = 0$$
$$\iff \mathbb{P}\left(\frac{\varepsilon - \varepsilon'}{X - X'} \le 0\right) = 0.5$$

Now for $\frac{\varepsilon - \varepsilon'}{X - X'}$ holds that $\varepsilon - \varepsilon'$ and X - X' are symmetric about zero. Therefore $\frac{\varepsilon - \varepsilon'}{X - X'}$ is also symmetric about zero as numerator and denominator are independent and symmetric about zero, and thus we obtain that $\mathbb{P}\left(\frac{\varepsilon - \varepsilon'}{X - X'} \le 0\right) = 0.5$.

A.4 Proof of Theorem 2

Proof.

For $X \stackrel{d}{=} X'$, $\varepsilon \stackrel{d}{=} \varepsilon'$ independent, continuous random variables, we have that

$$\begin{split} Y - Y' &= \beta X + \varepsilon - \beta X' - \varepsilon' = \beta (X - X') + \varepsilon - \varepsilon' \\ \Longrightarrow \ \frac{Y - Y'}{X - X'} &= \beta + \frac{\varepsilon - \varepsilon'}{X - X'} \end{split}$$

Hence, for Fisher consistency of the repeated median, it is needed that:

$$\underset{X,\varepsilon}{\operatorname{med}} \left(\underset{X',\varepsilon'}{\operatorname{med}} \left(\frac{\varepsilon - \varepsilon'}{X - X'} \right) \right) = 0$$

$$\iff \mathbb{P}_{X,\varepsilon} \left(\underset{X',\varepsilon'}{\operatorname{med}} \left(\frac{\varepsilon - \varepsilon'}{X - X'} \right) \le 0 \right) = 0.5$$

$$\iff \mathbb{P}_{X,\varepsilon} \left(x, e : \left[\underset{X',\varepsilon'}{\operatorname{med}} \left(\frac{e - \varepsilon'}{x - X'} \right) \le 0 \right] \right) = 0.5$$

We compute:

$$\operatorname{med}_{X',\varepsilon'} \left(\frac{e - \varepsilon'}{x - X'} \right) \leq 0 \iff \mathbb{P}_{X',\varepsilon'} \left(\frac{e - \varepsilon'}{x - X'} \leq 0 \right) \geq 0.5$$

$$\Leftrightarrow \mathbb{P}(e - \varepsilon' \leq 0 \ \cap \ x - X' \geq 0) + \mathbb{P}(e - \varepsilon' \geq 0 \ \cap \ x - X' \leq 0) \geq 0.5$$

$$\Leftrightarrow \mathbb{P}(e - \varepsilon' \leq 0) \cdot \mathbb{P}(x - X' \geq 0) + \mathbb{P}(e - \varepsilon' \geq 0) \cdot \mathbb{P}(x - X' \leq 0) \geq 0.5 \quad [\text{independence}]$$

$$\Leftrightarrow (1 - F_{\varepsilon}(e)) \cdot F_X(x) + F_{\varepsilon}(e) \cdot (1 - F_X(x)) \geq 0.5 \quad [\text{continuous r.v.}]$$

This implies:

$$\mathbb{P}_{X,\varepsilon}\left(x,e:\left[\inf_{X',\varepsilon'}\left(\frac{e-\varepsilon'}{x-X'}\right)\le 0\right]\right)$$

$$=\mathbb{P}_{X,\varepsilon}\left(x,e:\left(1-F_{\varepsilon}(e)\right)\cdot F_{X}(x)+F_{\varepsilon}(e)\cdot\left(1-F_{X}(x)\right)\ge 0.5\right)$$

$$=\mathbb{P}\left(\left(1-F_{\varepsilon}(\varepsilon)\right)\cdot F_{X}(X)+F_{\varepsilon}(\varepsilon)\cdot\left(1-F_{X}(X)\right)\ge 0.5\right)$$

$$=\mathbb{P}\left(\left(1-U'\right)U+U'(1-U)\ge 0.5\right) \quad \text{with } U\coloneqq F_{X}(X), \ U'\coloneqq F_{\varepsilon}(\varepsilon) \text{ uniform and independent}$$

$$\left[\left(1-U'\right)U+U'(1-U)\ge 0.5 \iff U\ge 0.5\cap U'\le 0.5 \text{ or } U\le 0.5\cap U'\ge 0.5\right]$$

$$=\int_{0.5}^{1}\int_{0}^{0.5}dsdt+\int_{0}^{0.5}\int_{0.5}^{1}dsdt=\frac{1}{2}\cdot\frac{1}{2}+\frac{1}{2}\cdot\frac{1}{2}=0.5$$

Hence the Repeated-Median is Fisher-consistent for continuous, random errors. \Box

B Additional simulation results

Table 4 and 5 present the simulation results for p = 2 and p = 5 variables respectively. It is clear that TSLiNGAM performs best overall. In particular, it outperforms DirectLiNGAM on heavy-tailed and skewed distributions, and the outperformance is more pronounced as the tails gets heavier. For lighter tails, the performance becomes similar to DirectLiNGAM. Using the repeated median works well, but provides no improvement over TSLiNGAM. The dcorr independence measure performs somewhat worse than the kernel-based independence measure, but becomes competitive at larger sample sizes. EASE is not doing very well, and needs larger sample sizes to become competitive.

The computation times for the simulation with p = 2 and p = 5 variables with t_5 distributions are given in Tables 6 and 7 respectively. It is clear that TSLiNGAM and

sample size	5	10	25	50	100	5	10	25	50	100
			t_1			Pareto				
TSLiNGAM	609	753	924	982	999	643	790	949	979	996
DirectLiNGAM	565	680	817	896	933	589	743	854	919	956
EASE		568	739	879	962		471	490	451	459
Repeated Median & KBI	582	758	905	974	996	609	740	904	956	991
Theil-Sen & dcorr	571	744	920	977	994	608	778	943	977	995
Repeated Median & dcorr	540	725	890	967	992	590	733	902	957	991
			t_2			lognormal				
TSLiNGAM	568	639	795	908	960	624	761	925	971	996
DirectLiNGAM	545	624	771	871	943	569	708	895	935	975
EASE		540	654	771	864		508	481	473	458
Repeated Median & KBI	542	629	776	904	953	587	724	898	960	988
Theil-Sen & dcorr	547	616	772	892	962	540	745	912	965	994
Repeated Median & dcorr	519	625	744	885	954	545	700	882	951	984
			t_5				ex	ponent	ial	
TSLiNGAM	538	522	618	698	789	604	703	884	961	974
DirectLiNGAM	517	513	616	692	796	572	676	861	926	980
EASE		535	557	610	696		474	412	406	379
Repeated Median & KBI	522	527	617	677	796	555	674	849	935	973
Theil-Sen & dcorr	533	518	596	655	758	567	696	888	959	980
Repeated Median & dcorr	505	525	580	644	718	547	670	860	938	972

Table 4: Number of correct causal orders out of 1000 trials for p = 2.

DirectLiNGAM have similar computational costs. Using dcorr as independence measure decreases the computational costs by a factor of roughly 3 for p = 2 and a factor of 10 for p = 5. EASE is again by far the fastest method.

sample size	30	50	100	200	30	50	100	200
	t ₁ Pareto				reto			
TSLiNGAM	709	893	947	992	742	887	979	999
DirectLiNGAM	504	620	740	794	589	713	796	926
EASE	183	399	630	826	88	119	146	206
Repeated Median & KBI	663	857	943	991	645	827	943	996
Theil-Sen & dcorr	567	769	895	960	677	815	959	995
Repeated Median & dcorr	528	738	872	951	557	743	919	986
		t	2		lognormal			
TSLiNGAM	405	551	747	853	688	847	954	990
DirectLiNGAM	390	513	709	809	558	748	873	947
EASE	128	171	297	495	79	84	112	169
Repeated Median & KBI	366	520	743	850	591	772	930	982
Theil-Sen & dcorr	283	435	688	892	600	779	929	986
Repeated Median & dcorr	264	396	652	859	487	679	885	964
		t	5			expon	ential	
TSLiNGAM	157	220	327	407	514	716	875	962
DirectLiNGAM	158	225	333	429	502	679	846	950
EASE	82	77	122	161	58	60	62	88
Repeated Median & KBI	160	217	322	394	458	647	852	940
Theil-Sen & dcorr	71	102	185	267	466	691	879	970
Repeated Median & dcorr	72	91	151	218	375	592	818	936

Table 5: Number of correct causal orders out of 1000 trials for p = 5.

C Additional figures

sample size	5	10	25	50	100
TSLiNGAM	3.42	4.09	4.99	5.84	7.32
DirectLiNGAM	3.55	4.13	4.98	5.70	6.71
EASE		0.23	0.21	0.22	0.25
Repeated Median & KBI	3.47	4.19	5.16	6.25	8.02
Theil-Sen & dcorr	1.11	1.07	1.13	1.21	1.26
Repeated Median & dcorr	1.05	1.03	1.23	1.37	1.73

Table 6: Computational time in seconds for 1000 runs for p = 2 for a Student-*t* distribution with 5 degrees of freedom.

sample size	30	50	100	200
TSLiNGAM	1.60	1.80	2.19	2.88
DirectLiNGAM	1.57	1.78	2.12	2.81
EASE	0.01	0.01	0.02	0.01
Repeated Median & KBI	1.66	1.93	2.40	3.18
Theil-Sen & dcorr	0.10	0.13	0.21	0.35
Repeated Median & dcorr	0.13	0.20	0.36	0.55

Table 7: Computational time in minutes for 1000 runs for p = 5 for a Student-*t* distribution with 5 degrees of freedom.



Figure 10: Causal graph found by Theil-Sen and distance correlation with correct arrows in black, wrongly directed arrows in red, unverifiable arrows in purple and redundant arrows in blue.