

# Structure Learning with Continuous Optimization: A Sober Look and Beyond

**Ignavier Ng**

*Carnegie Mellon University*

IGNAVIERNG@CMU.EDU

**Biwei Huang**

*University of California San Diego*

BIH007@UCSD.EDU

**Kun Zhang**

*Carnegie Mellon University*

*Mohamed bin Zayed University of Artificial Intelligence*

KUNZ1@CMU.EDU

**Editors:** Francesco Locatello and Vanessa Didelez

## Abstract

This paper investigates in which cases continuous optimization for directed acyclic graph (DAG) structure learning can and cannot perform well and why this happens, and suggests possible directions to make the search procedure more reliable. Reisach et al. [47] suggested that the remarkable performance of several continuous structure learning approaches is primarily driven by a high agreement between the order of increasing marginal variances and the topological order, and demonstrated that these approaches do not perform well after data standardization. We analyze this phenomenon for continuous approaches assuming equal and non-equal noise variances, and show that the statement may not hold in either case by providing counterexamples, justifications, and possible alternative explanations. We further demonstrate that nonconvexity may be a main concern especially for the non-equal noise variances formulation, while recent advances in continuous structure learning fail to achieve improvement in this case. Our findings suggest that future works should take into account the non-equal noise variances formulation to handle more general settings and for a more comprehensive empirical evaluation. Lastly, we provide insights into other aspects of the search procedure, including thresholding and sparsity, and show that they play an important role in the final solutions.

**Keywords:** structure learning, continuous optimization, directed acyclic graphs

## 1. Introduction

Bayesian networks are a class of probabilistic graphical models that encode probabilistic distributions in a compact way [42, 30]. Learning their graphical structures from data, represented by directed acyclic graphs (DAGs), has been applied in various fields, including genetics [44] and education [22]. Classical approaches for structure learning typically involve discrete procedures, such as constraint-based methods using conditional independence tests [54, 55] and score-based methods that search for a high-scoring structure [29, 53, 14, 67, 12, 43]. Greedy search is often employed in score-based methods because of the large space of possible structures [11, 13], such as GES [12] and GDS [43].

Recently, Zheng et al. [72] proposed a smooth characterization of acyclicity and transformed the structure learning problem of discrete nature into a continuous, nonconvex optimization problem, thus enabling the application of efficient gradient-based optimization. This formulation has been extended and applied to a wide range of settings, including nonlinear cases [65, 32, 73, 39, 27], interventional

data [8, 19], unobserved confounding [6, 4], incomplete data [60, 20], time series [41, 57], multi-task learning [10], multi-domain data [68], federated learning [36, 21], and representation learning [64].

Given the growing interest in continuous structure learning [58], various theoretical and empirical aspects of these approaches have gained considerable attention. In particular, Wei et al. [61], Ng et al. [38] studied the optimality conditions and convergence property of continuous constrained approaches [72], while Deng et al. [15] showed that a proper optimization scheme converges to the global minimum of the least squares objective in the bivariate case. Zhang et al. [70], Bello et al. [3] demonstrated that existing DAG constraints [72, 65] may encounter gradient vanishing issues in practice and proposed improved variants. Reisach et al. [47] suggested that the remarkable performance of several continuous structure learning approaches [72, 37] is primarily driven by a high agreement between the order of increasing marginal variances and the topological order, and demonstrated that these approaches do not perform well after data standardization. Similar phenomenon was also observed by Kaiser and Sipos [26]. These empirical findings regarding data standardization provide insights into the performance of continuous structure learning, and may at first appear surprising. One of our goals is to provide further analysis of this phenomenon.

In this work, we investigate in which cases continuous structure learning approaches can and cannot perform well and why this happens, and suggest possible directions to make the search procedure more reliable. Our main contributions are:

- We analyze the statements and observations by Reisach et al. [47] in the linear case with equal (Section 3.1) and non-equal noise variances (Section 3.2), and show that the statements may not hold in either case by providing counterexamples and justifications. We also provide possible alternative explanations for the observations by Reisach et al. [47] that continuous structure learning approaches do not perform well after data standardization.
- We show that nonconvexity may be a main concern especially for the non-equal noise variances formulation, while recent advances in continuous structure learning, including search strategies, DAG constraints, and nonlinear approaches, fail to achieve improvement in this case (Section 4). Our findings suggest that future works should take into account the non-equal noise variances formulation to handle more general settings and for a more comprehensive empirical evaluation.
- We provide insights into other aspects of the search procedure, including thresholding (Section 5.1) and sparsity (Section 5.2), and show that they play a crucial role in the final solutions.

## 2. Preliminaries

In this section, we briefly describe several existing continuous structure learning approaches that this work focuses on, and explain the notion of varsortability proposed by Reisach et al. [47].

### 2.1. Structure Learning with Continuous Optimization

**Setup.** Let  $\mathcal{G}$  be a DAG with node set  $\{X_1, \dots, X_d\}$ . In a Bayesian network, each node  $X_i$  of  $\mathcal{G}$  corresponds to a random variable, and the joint distribution of random variables  $X = (X_1, \dots, X_d)$ , denoted as  $P(X)$ , is Markov w.r.t. DAG  $\mathcal{G}$ . We consider the linear case in which variables  $X$  follow a linear structural equation model (SEM)  $X = B^T X + N$ , where  $B \in \mathbb{R}^{d \times d}$  is an acyclic weighted adjacency matrix whose nonzero entries represent the edges in DAG  $\mathcal{G}$ , and  $N = (N_1, \dots, N_d)$  consists of the independent noise variables with covariance matrix  $\Omega := \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$ . If  $\sigma_1^2 = \dots = \sigma_d^2$ , we refer to it as the equal noise variances (EV) case, and otherwise as the non-equal

noise variances (NV) case. Structure learning aims at estimating DAG  $\mathcal{G}$  or its Markov equivalence class (MEC) given the data matrix  $\mathbf{X} \in \mathbb{R}^{n \times d}$  consisting of  $n$  i.i.d. samples from distribution  $P(X)$ .

**NOTEARS.** Zheng et al. [72] proposed to solve the constrained optimization problem

$$\min_{B \in \mathbb{R}^{d \times d}} \ell(B; \mathbf{X}) := \frac{1}{2n} \|\mathbf{X} - \mathbf{X}B\|_F^2 \quad \text{subject to } h(B) = 0,$$

where  $\ell(B; \mathbf{X})$  is the least squares objective and  $h(B) := \text{tr}(e^{B \odot B}) - d$  is the DAG constraint term. An  $\ell_1$  penalty term  $\lambda \|B\|_1$  is also incorporated into the objective function, where  $\|\cdot\|_1$  is defined element-wise and  $\lambda$  is a hyperparameter. Since the above formulation focuses on the linear case with equal noise variances [43, 33], we refer to it as NOTEARS-EV throughout this paper.

**GOLEM.** Ng et al. [37] solve an unconstrained optimization problem:

$$\min_{B \in \mathbb{R}^{d \times d}} \mathcal{L}(B; \mathbf{X}) - \log |\det(I - B)| + \lambda_1 \|B\|_1 + \lambda_2 h(B), \quad (1)$$

where  $\mathcal{L}(B; \mathbf{X})$  is defined as

$$\mathcal{L}_{\text{EV}}(B; \mathbf{X}) = \frac{d}{2} \log \|\mathbf{X} - \mathbf{X}B\|_F^2 \quad \text{and} \quad \mathcal{L}_{\text{NV}}(B; \mathbf{X}) = \frac{1}{2} \sum_{i=1}^d \log \|\mathbf{X}_{\cdot, i} - \mathbf{X}B_{\cdot, i}\|_2^2$$

in the linear Gaussian case with equal and non-equal noise variances, respectively. Here,  $\lambda_1$  and  $\lambda_2$  are hyperparameters. The above two variants are denoted as GOLEM-EV and GOLEM-NV, respectively.

## 2.2. Varsortability

Recently, Reisach et al. [47] introduced *varsortability*, denoted as  $\nu$ , as a measure of agreement between increasing marginal variances and topological order. It is defined as the proportion of directed paths which start from a variable with lower marginal variance than the variable they end in [47]. When all variables have higher marginal variances than their ancestors, we have  $\nu = 1$ .

Reisach et al. [47] provided insights into the performance of several continuous structure learning approaches [72, 37], and suggested that their remarkable performance is primarily driven by high varsortability. Here we list some excerpts from Reisach et al. [47]: “Our experiments demonstrate that varsortability dominates the optimization and helps achieve state-of-the-art performance provided the ground-truth data scale” and “we focus on the first optimization steps to explain a) why continuous structure learning algorithms that assume equal noise variance work remarkably well in the presence of high varsortability”. For clarity and ease of further analysis, we give a partial formulation of the statements (on which this work focuses) in the following. We provide a further discussion of how these statements are formulated and possible alternative formulations in Appendix A.

**Statement 1 (Equal Noise Variances Formulation)** *Continuous structure learning approaches that assume equal noise variances, specifically NOTEARS-EV and GOLEM-EV, perform well in the presence of high varsortability.*

**Statement 2 (Non-Equal Noise Variances Formulation)** *Continuous structure learning approaches that assume non-equal noise variances, specifically GOLEM-NV, perform well in the presence of high varsortability.*

Reisach et al. [47] showed that the synthetic data used by Zheng et al. [72], Ng et al. [37] to benchmark their methods exhibits high varsortability, e.g., higher than 0.94 on average, and that these approaches do not perform well after data standardization which removes such patterns in the marginal variances.

### 3. Varsortability and Data Standardization

In this section, we analyze the statements and observations by Reisach et al. [47] in the linear case with equal and non-equal noise variances formulations, and show that the statements may not hold in either case. That is, continuous structure learning approaches often do not perform well even in the presence of high varsortability. We also provide possible alternative explanations for the observations by Reisach et al. [47] that continuous approaches do not perform well after data standardization.

#### 3.1. With Equal Noise Variances Formulation

**Varsortability and continuous structure learning.** We first take a closer look at Statement 1 and consider the least squares score [33, 72]. For bivariate case, say  $X = (X_1, X_2)$  with ground truth  $X_1 \rightarrow X_2$ , Reisach et al. [47] showed that estimations based on varsortability and least squares are consistent in the sense that  $\nu = 1$ , i.e.,  $\text{Var}(X_1) < \text{Var}(X_2)$  if and only if the least squares score computed with  $X_1 \rightarrow X_2$  is smaller than that with  $X_2 \rightarrow X_1$ ; see Reisach et al. [47, Appendix A] for a proof. Thus, with high varsortability in the bivariate case, i.e.,  $\nu = 1$ , NOTEARS-EV, loosely speaking, is able to asymptotically output the correct structure assuming that the global minimizer can be found. Therefore, Statement 1 holds in the bivariate case, at least for NOTEARS-EV.

However, the property in the bivariate case above cannot be extended to general cases with more than two variables; see the following example with three variables.

**Example 1** Consider the linear SEM over  $X = (X_1, X_2, X_3)$  with the true weighted adjacency matrix  $\tilde{B}$  and noise covariance matrix  $\tilde{\Omega}$ , as well as an alternative weighted adjacency matrix  $\hat{B}$ :

$$\tilde{B} = \begin{bmatrix} 0 & 0 & 0 \\ \frac{1}{2} & 0 & 0 \\ 1 & \frac{1}{2} & 0 \end{bmatrix}, \quad \tilde{\Omega} = \begin{bmatrix} \frac{1}{2} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{1}{2} \end{bmatrix}, \quad \hat{B} = \begin{bmatrix} 0 & \frac{2}{3} & 0 \\ 0 & 0 & 0 \\ \frac{5}{4} & -\frac{1}{3} & 0 \end{bmatrix}.$$

In the large sample limit, we have

$$\text{Var}(X_1) > \text{Var}(X_2) > \text{Var}(X_3) \quad \text{and} \quad \ell(\hat{B}; \mathbf{X}) < \ell(\tilde{B}; \mathbf{X}).$$

That is, even in the presence of high varsortability, i.e.,  $\nu = 1$ , the estimation with least squares will be the structure indicated by  $\hat{B}$ , instead of that indicated by  $\tilde{B}$ , and thus Statement 1 does not hold.

The weighted adjacency matrices  $\tilde{B}$  and  $\hat{B}$  represent the triangle structures in Figures 1a and 1b, respectively. Here a discrete exhaustive DAG search with least squares score cannot return the true structure, and similarly for NOTEARS-EV assuming that the global minimizer can be found. This is also the case for GOLEM-EV because  $\mathcal{L}_{\text{EV}}(\hat{B}; \mathbf{X}) < \mathcal{L}_{\text{EV}}(\tilde{B}; \mathbf{X})$ . To verify it, we conduct 100 simulations using  $\tilde{B}$  and  $\tilde{\Omega}$  defined in Example 1, and generate  $10^6$  random samples in each simulation. In all simulations, we have  $\text{Var}(X_1) > \text{Var}(X_2) > \text{Var}(X_3)$ , and observe that NOTEARS-EV (with  $\lambda = 0$  and threshold of 0.1), GOLEM-EV (with  $\lambda_1 = 0.02$ ,  $\lambda_2 = 5$ , and threshold of 0.1), and exhaustive search with least squares score return the structure in Figure 1b. In the example above, even in the presence of high varsortability, i.e.,  $\nu = 1$ , NOTEARS-EV and GOLEM-EV return a completely incorrect structure, indicating that Statement 1 does not hold in general.

One may wonder whether the parameters of  $\tilde{B}$  and  $\tilde{\Omega}$  in Example 1 have to be exactly “tuned” to obtain such an outcome, which is analogous to the violation of faithfulness assumption [55] that occurs with Lebesgue measure zero. We show that this is generally not the case via the following proposition; specifically, such counterexample exists for a set of parameters with nonzero Lebesgue measure. The proof is provided in Appendix B.1.




Figure 1: Examples of triangle structures.




Figure 2: Noise ratio after data standardization.

**Proposition 1** Consider the parameters  $(B, \Omega)$  of the linear SEMs over variables  $X = (X_1, \dots, X_d)$ ,  $d \geq 3$ , where the noise variables follow Gaussian distributions. In the large sample limit, the set of parameters such that varsortability equals one and that the true DAG does not have the lowest least squares score has a nonzero Lebesgue measure.

On the other hand, the proposition below shows that there are cases where the varsortability is low but the true DAG yields the lowest least squares score, with a proof given in Appendix B.2. To illustrate, if  $\mathcal{G}$  is a fully connected DAG or a chain, then we have  $v_{\mathcal{G}}^{\text{source}} = \frac{2^{d-1}-1}{2^d-d-1}$  or  $v_{\mathcal{G}}^{\text{source}} = \frac{2}{d}$ , respectively. Note that, as  $d \rightarrow \infty$ , their limits equal 0.5 or 0, respectively, indicating a low varsortability. However, loosely speaking, methods that adopt least squares score, such as NOTEARS-EV and discrete search, are able to asymptotically find the correct structure assuming that the global minimizer can be found.

**Proposition 2** Consider the parameters  $(B, \Omega)$  of the linear SEMs over variables  $X = (X_1, \dots, X_d)$  induced by DAG  $\mathcal{G}$ , where the noise variables follow Gaussian distributions. In the large sample limit, the set of parameters such that varsortability equals  $v_{\mathcal{G}}^{\text{source}}$  and that the true weighted adjacency matrix  $B$  yields the lowest least squares score has a nonzero Lebesgue measure, where

$$v_{\mathcal{G}}^{\text{source}} := \frac{\text{Number of distinct paths from the source nodes in DAG } \mathcal{G}}{\text{Number of distinct paths in DAG } \mathcal{G}}.$$

**An alternative explanation by noise ratio.** The above argument indicates that Statement 1 generally does not hold, and so may not explain why NOTEARS-EV does not perform well after data standardization. Here, we provide a possible alternative explanation of this phenomenon, i.e., the theoretical guarantee of least squares used by NOTEARS-EV does not accommodate standardization in general; thus, such phenomenon may not be surprising. Specifically, define the noise ratio  $r = \frac{\max(\sigma_1^2, \dots, \sigma_d^2)}{\min(\sigma_1^2, \dots, \sigma_d^2)}$ , which may intuitively be viewed as a measure of how far the SEM is from having equal noise variances. Loh and Bühlmann [33, Theorem 9] showed that if  $r < 1 + \frac{\xi}{d}$ , where  $\xi$  is defined as the difference between the score of the true DAG and the next best DAG, then, in the large sample limit, minimizing the least squares in the space of acyclic weighted adjacency matrices  $B$  returns the true structure for linear SEMs. This allows a certain degree of misspecification of the noise variances—if  $\xi$  is larger, then the least squares score will be more robust to such misspecification. The result of equal noise variances [33, Theorem 7] can be viewed as a special case with  $r = 1 < 1 + \frac{\xi}{d}$ .

Thus, assuming that the global minimizer can be found, NOTEARS-EV, loosely speaking, is able to return the true DAG in the large sample limit under the assumption specified above, i.e.,  $r < 1 + \frac{\xi}{d}$ . However, after data standardization, this assumption may no longer hold, i.e., the noise ratio becomes  $r' = \frac{\max(\sigma_1^2 / \text{Var}(X_1), \dots, \sigma_d^2 / \text{Var}(X_d))}{\min(\sigma_1^2 / \text{Var}(X_1), \dots, \sigma_d^2 / \text{Var}(X_d))}$ , and there is no guarantee that  $r' < 1 + \frac{\xi'}{d}$  holds. In this case, The-

orem 9 of Loh and Bühlmann [33] does not apply and we are left with no guarantee that NOTEARS-EV or least-squares-based methods can find the true structure. Therefore, they may not perform well after data standardization, and the observation by Reisach et al. [47] may not be surprising.

In fact, after data standardization, the data may be far from having equal noise variances, i.e., the noise ratio  $r'$  may be very large. An example is provided in Figure 2, in which we conduct 1000 random simulations on linear Gaussian-EV model and Erdős–Rényi [17] graphs with varying number of variables and degrees, and further compute the noise ratio in the large sample limit after standardizing the data. One observes that a larger number of variables and degree leads to a higher noise ratio  $r'$ ; in this setting considered, the noise ratio appears to increase approximately exponentially in the degree. For example, for 50-node graphs with degree of 4, i.e., the largest number of variables and degrees considered by Reisach et al. [47], the noise ratio  $r'$  is  $2079.30 \pm 189.6$ , which is far from the equal noise variances case.

With such a large noise ratio  $r'$  in the scenarios above, the guarantee [33, Theorem 9] of least squares may not apply after standardization. In other words, least-squares-based search methods, whether it be continuous or discrete ones, are not expected to perform well after standardizing the data.

**Empirical studies.** To validate the argument above, we compare the performance of various structure learning methods that assume equal noise variances on different noise ratios  $r$ . Following previous work [72, 37], we simulate Erdős–Rényi [17] graphs with  $kd$  edges, denoted as  $ERk$ , where  $k$  represents the average degree. Unless otherwise stated, each edge weight in the true weighted adjacency matrix  $B$  is sampled uniformly at random from  $[-2, -0.5] \cup [0.5, 2]$ . We consider  $ER1$  graphs with  $d \in \{15, 50\}$  variables and  $n \in \{100, 10^6\}$  samples. We use Gaussian noises where the variances of two randomly chosen noise variables are set to 1 and  $r$ , respectively, while the variances of the remaining noise variables are sampled uniformly at random from  $[1, r]$ . The noise ratios considered are  $r \in \{1, 2, 4, 8, 16, 32, 64\}$ . Apart from continuous approaches, i.e., NOTEARS-EV and GOLEM-EV, we also consider discrete ones, including a greedy approach GDS [43] and an exact approach  $A^*$  [67]. For both discrete approaches, we adopt least squares score. We report the structural Hamming distance (SHD), F1 score, and recall over 30 simulations, as well as their standard errors. Further details about the approaches and evaluation metrics are provided in Appendix E.1.

The experiment results for  $d = 15$  and  $n = 10^6$  are shown in Figure 3, while complete results are available in Figure 15 in Appendix F.1. As the noise ratio increases, the performance of all four methods deteriorates in all cases, especially for continuous approaches, i.e., NOTEARS-EV and GOLEM-EV. This demonstrates that, even with a large sample size, these methods that assume equal noise variances are not expected to perform well when the noise ratio  $r$  is large. As shown in Figure 2, the noise ratio after data standardization  $r'$  is very large, e.g., could be as large as  $10^6$ , which is much larger than that considered in our experiment here. Thus, it may not be surprising that the structure learning methods considered do not perform well.

### 3.2. With Non-Equal Noise Variances Formulation

We next turn to the general linear Gaussian case, i.e., the non-equal noise variances formulation. We provide two arguments to explain why Statement 2 may not hold.

**Argument 1.** *GOLEM-NV often does not perform well in the presence of high varsortability.* We demonstrate that GOLEM-NV with the standard initialization scheme (i.e., zero matrix) does not perform well in the presence of high varsortability, by comparing its estimated completed partially




Figure 3: Linear Gaussian-EV formulation without standardization.




Figure 4: Linear Gaussian-NV formulation without standardization.




Figure 5: Linear Gaussian-NV formulation with standardization.

DAGs (CPDAGs) to standard methods such as PC [54] and FGES [45]. Note that GOLEM-NV outputs a DAG, and thus an additional step is needed to convert it into a CPDAG that represents the MEC. We consider linear Gaussian models with 50 variables and ER1 graphs, where the noise ratio  $r$  is 2. The sample size is set to  $n = 10^6$  to reduce finite-sample errors. Here, the varsortability is  $0.97 \pm 0.003$ , which is consistent with the high varsortability reported by Reisach et al. [47, Section 3.3].

The results in Table 1 show that PC and FGES can nearly recover the true CPDAG with a high F1 score, while GOLEM-NV has a low F1 score and high SHD. Note that the performance of GOLEM-NV is worse than an empty graph whose SHD is 50. Therefore, GOLEM-NV does not perform well in the presence of high varsortability, and it seems that high varsortability does not help GOLEM-NV achieve a remarkable performance. The question is then why Ng et al. [37] observed a remarkable performance for GOLEM-NV on synthetic data as noticed by Reisach et al. [47], which we address next.

Table 1: Empirical results.

	SHD of CPDAG	F1 of skeleton
GOLEM-NV	$120.27 \pm 2.91$	$0.50 \pm 0.00$
PC	$1.57 \pm 0.35$	$0.99 \pm 0.00$
FGES	$0.63 \pm 0.47$	$1.00 \pm 0.00$

**Argument 2.** *The quality of the final solution estimated by GOLEM-NV depends largely on the initial solution possibly owing to nonconvexity.* Ng et al. [37, Section 4.1] adopted an initialization scheme for GOLEM-NV which uses the solution returned by GOLEM-EV as the initial solution, because the authors noticed that the optimization procedure of GOLEM-NV is prone to local solutions. Here, we experiment with different initialization schemes and find that they play an important role in the quality of the final solutions, possibly because of the nonconvex optimization formulation.

**Empirical studies.** We consider 50-variable linear Gaussian model with ER1 graphs and noise ratios  $r \in \{1, 2, 4, 8, 16, 32, 64\}$ . To focus on the aspect of nonconvex optimization, we use a large sample size, i.e.,  $n = 10^6$ , to reduce finite-sample errors. We compare various initialization schemes, including GOLEM-NV initialized with zero matrix, with the solution of GOLEM-EV, and with the solution of FGES<sup>1</sup>. We also provide empirical studies for GOLEM-NV with random initializations in Appendix C.1, and for GOLEM-EV with random perturbations in Appendix C.2

The SHDs before and after data standardization are available in Figures 4 and 5, where the other metrics are available in Figure 16 in Appendix F.2. Consistent with the results in Table 1, GOLEM-NV initialized with zero matrix does not perform well across all settings. Moreover, before standardization, the performance of GOLEM-EV and GOLEM-NV initialized with GOLEM-EV exhibit similar

1. Since FGES outputs a CPDAG, we generate a DAG consistent with it following the procedure developed by Dor and Tarsi [16], and compute the least squares coefficients for this specific DAG.

trend, both of which degrade as the noise ratio increases, and their gap with PC and FGES is also enlarged. A possible reason is that GOLEM-NV may be susceptible to suboptimal local solutions owing to nonconvexity, and thus its performance largely depends on the quality of solution of GOLEM-EV. As the noise ratio increases, the quality of solution returned by GOLEM-EV deteriorates, as discussed in Section 3.1; thus, GOLEM-NV does not have an initial solution of decent quality. Similar explanation may also apply to its results after standardization, i.e., GOLEM-NV initialized with GOLEM-EV has a poor performance after standardization (also observed by Reisach et al. [47]) because GOLEM-EV does not perform well owing to the large noise ratio  $r'$ , as discussed in Section 3.1. By contrast, GOLEM-NV initialized with the solution of FGES performs well in all cases, where the results are close to FGES and PC. This is because FGES returns a solution with high quality across different noise ratio  $r$ , which serves as a good initial starting point for optimization of GOLEM-NV.

Note that the formulation of GOLEM-NV in Eq. (1) involves only matrix  $B$  because it profiles out the parameter  $\Omega$  (i.e., the noise variances) [37, Appendix C.1]. In Appendix C.3, we further consider the likelihood function without profiling out such parameter and observe that it also leads to a poor performance when initialized with zero matrix, similar to the original version of GOLEM-NV.

Instead of high varsortability, the observations above suggest that the performance of GOLEM-NV largely depends on the initial solution. A possible reason is that the optimization problem of GOLEM-NV may be highly nonconvex and contain many suboptimal local solutions, because (1) a large sample size  $n = 10^6$  is used which leads to small finite-sample errors, and (2) GOLEM-NV has been shown to be consistent under mild conditions, assuming that the global minimizer can be found and  $\ell_0$  penalty is used (note that Ng et al. [37] did not provide guarantee for  $\ell_1$  penalty).<sup>2</sup> This is in contrast with the observation by Zheng et al. [72, Section 5.3] for the equal noise variances formulation, which finds that the solution of NOTEARS-EV is in practice close to the global minimizer. We provide further analysis on the nonconvexity of GOLEM-NV in Appendix C.4. Moreover, the above empirical studies may explain the remarkable performance of GOLEM-NV in the synthetic data experiments conducted by Ng et al. [37], as noticed by Reisach et al. [47], which simulate linear Gaussian model with a relatively small noise ratio  $r \leq 4$ . In this case, GOLEM-EV achieves a decent performance, and using its solution to initialize GOLEM-NV leads to a remarkable performance.

#### 4. Nonconvexity of Non-Equal Noise Variances Formulation

The empirical studies in Section 3.2 suggest that continuous structure learning approach that assumes non-equal noise variances, specifically GOLEM-NV, is prone to suboptimal local solutions, possibly owing to nonconvexity. On the other hand, many existing works [72, 66, 9, 3] considered the equal noise variances formulation (by adopting least squares objective) and observed remarkable performance. This indicates that the optimization procedure of non-equal noise variances formulation may in practice be more susceptible to suboptimal local solutions than the equal noise variances formulation, and thus empirically suggests that the nonconvexity issue might be more severe for the former formulation. In this section, we demonstrate that the recent advances for the equal noise variances formulation fail to achieve improvement for the non-equal noise variances formulation, thus suggesting that nonconvexity may be a key concern for the latter formulation. In particular, we consider recently developed search strategies and DAG constraints in Sections 4.1 and 4.2, respectively, and other cases in Section 4.3, namely linear non-Gaussian and nonlinear cases.

---

2. With global minimizer and  $\ell_0$  penalty, the theoretical guarantee of GOLEM-NV [37, Theorem 2] holds in the large sample limit under the corresponding conditions, regardless of the initial solution and whether the data is standardized.






Figure 6: Search strategies.




Figure 7: DAG constraints.




Figure 8: Linear non-Gaussian case.

#### 4.1. Search Strategies

We investigate whether different search strategies, namely NOTEARS [72], DAGMA [3], NOCURL [66], DPDAG [9], and GOLEM [37], perform well for the non-equal noise variances formulation. These methods represent different strategies to traverse the search space for estimating a DAG with continuous optimization; a summary of them is provided in Appendix E.2. For these methods except GOLEM, we replace their least squares objective  $\ell(B; \mathbf{X})$  with  $\mathcal{L}_{\text{NV}}(B; \mathbf{X})$ , which corresponds to the likelihood of linear Gaussian DAGs assuming non-equal noise variances [37, Appendix C.1] (note that the term  $-\log |\det(I - B)|$  is omitted). We denote the resulting methods by NOTEARS-NV, DAGMA-NV, NOCURL-NV, and DPDAG-NV, respectively. We consider linear Gaussian model with 50-node ER1 graphs and noise ratio  $r$  of 16, as well as sample sizes  $n \in \{10^2, 10^3, 10^4, 10^5, 10^6\}$ .

The SHDs computed over CPDAGs are available in Figure 6, where the other metrics are available in Figure 17 in Appendix F.3. It is observed that none of the continuous structure learning approaches perform well, compared to PC and FGES. For PC and FGES, as the sample size increases, PC and FGES improve with a larger sample size and eventually reach SHDs of  $5.97 \pm 0.43$  and  $5.13 \pm 0.51$ , respectively. By contrast, the quality of solution estimated by continuous structure learning approaches does not improve much with a larger sample size and reaches a performance plateau after  $n = 10^3$ . Specifically, the SHD of these continuous approaches is larger than 70 even when the sample size is  $n = 10^6$ . Note that an empty graph has an SHD of 50. A possible reason is that the optimization problem of these continuous approaches may be highly nonconvex and contain many suboptimal local solutions, thus leading to estimated CPDAGs that are far from the true ones.

#### 4.2. DAG Constraints

We consider the DAG constraints commonly adopted by continuous structure learning approaches based on matrix exponential [72] and binomial [65], as well as more recent DAG constraints based on log-determinant [3] and truncated geometric series [70] that have been demonstrated to be less sensitive to the gradient vanishing issue encountered by the former two constraints. We apply these four constraints to GOLEM-NV, denoted as GOLEM-NV-Exp, GOLEM-NV-Bin, GOLEM-NV-LogDet, and GOLEM-NV-TMPI, respectively, and similarly for NOTEARS-NV.

Here we use the same setup of simulated data as that of Section 4.1. The SHDs of GOLEM-NV are depicted in Figure 7, while complete results including those for NOTEARS-NV are available in Figure 18 in Appendix F.4. Similar to the observation in Section 4.1, one observes that GOLEM-NV and NOTEARS-NV equipped with these DAG constraints quickly reach a performance plateau and do not improve much with increasing sample sizes. Specifically, the SHDs of GOLEM-NV and NOTEARS-NV with these DAG constraints are larger than 120 and 80, respectively. Note that Bello

et al. [3], Zhang et al. [70] observed that their proposed DAG constraints lead to a huge improvement over those existing ones [72, 65] for the equal noise variances formulation. Clearly, this experiment demonstrates that such an improvement cannot be translated into the non-equal noise variances formulation, at least in the settings considered here. Furthermore, this suggests that the optimization problem of the non-equal noise variances formulation might be more challenging than the equal noise variances formulation as the former may contain more suboptimal local solutions.

### 4.3. Other Settings

We now investigate whether the observations in previous linear Gaussian experiments generalize to the other settings, namely linear non-Gaussian [51] and nonlinear [23] cases.

**Linear non-Gaussian case.** Zheng [71] developed an extension of NOTEARS to handle linear non-Gaussian case, called NOTEARS-ICA. It is worth noting that Zheng [71, Section 4.3.2] provided a general formulation in the non-equal noise variances case; as the same time, their experiments, as discussed by Zheng [71, Section 4.5], seem to adopt the equal noise variances formulation. Here, we consider both equal and non-equal noise variances formulation for NOTEARS-ICA, denoted as NOTEARS-ICA-EV and NOTEARS-ICA-NV, respectively. We simulate 15-node ER1 graphs and linear SEMs with standard Laplace noises. The SHDs are reported in Figure 8, while complete results are shown in Figure 19 in Appendix F.5. We observe that NOTEARS-ICA-EV performs similarly to DirectLiNGAM [52, 24], both of which improve with more samples; specifically, their SHDs are very close to zero when the sample size is large. However, the performance of NOTEARS-ICA-NV does not improve much even when the sample size is large. For example, when  $n = 10^6$ , NOTEARS-ICA-EV and DirectLiNGAM have SHDs of  $1.93 \pm 0.59$  and  $1.00 \pm 0.19$ , respectively, while NOTEARS-ICA-NV has SHD of  $26.40 \pm 0.80$ . Note that an empty graph leads to an SHD of 15. This suggests that the optimization problem of NOTEARS-ICA-NV may be more challenging than that of NOTEARS-ICA-EV and contain more suboptimal local solutions.

**Nonlinear case.** There are several nonlinear extensions of NOTEARS, some of which rely on various pre/post-processing steps such as pruning [32, 39], making it difficult to investigate the performance of continuous optimization procedure. We consider NOTEARS-MLP [73] that does not rely on such pre/post-processing steps, and extend it to the non-equal noise variances formulation by replacing its least squares objective with  $\mathcal{L}_{\text{NV}}(B, \mathbf{X})$ , in which the residuals are computed with multilayer perceptrons (MLPs) instead of linear regressions. We consider the data generating procedure used by Zheng et al. [73] with MLPs and standard Gaussian noises. The experiment results are available in Figure 20 in Appendix F.5. Similar to the observation in the linear non-Gaussian case, NOTEARS-MLP-EV achieves a much lower SHD when the sample size is large, while NOTEARS-MLP-NV does not improve much and its SHD remains at  $10.07 \pm 0.57$  even when the sample size is as large as  $10^5$ .

## 5. Thresholding and Sparsity Penalty

Apart from the possible nonconvexity issue discussed in Section 4, we show that the performance of continuous approaches is sensitive to other aspects of the search procedure, i.e., thresholding and sparsity. These technical issues should be made transparent, as they play an important role in the final solutions. In concurrent work, Xu et al. [63] also studied the aspects of thresholding and sparsity in NOTEARS-EV. Here, our analysis of thresholding includes GOLEM-EV and discrete approaches such as A\* and GDS, as well as both small ( $n = 100$ ) and large ( $n = 10^6$ ) sample sizes; therefore,

the analysis and observations are different. For sparsity, Xu et al. [63] adopted the adaptive  $\ell_1$  penalty [75] for NOTEARS-EV, while we consider the smoothly clipped absolute deviation (SCAD) penalty [18] and minimax concave penalty (MCP) [69] for NOTEARS-EV and GOLEM-EV.

### 5.1. Thresholding

We investigate the impact of threshold specification on continuous structure learning, and show that it plays a crucial role for the final solutions. Existing continuous approaches [72, 37, 66, 3] typically use edge weights sampled uniformly at random from  $[-2, -0.5] \cup [0.5, 2]$  for the true weighted adjacency matrix and apply a threshold of 0.3 to identify the final structure. Here, we consider 15-node ER1 graphs with different ranges of edge weights, i.e.,  $[-2\alpha, -0.5\alpha] \cup [0.5\alpha, 2\alpha]$  where  $\alpha \in \{0.25, 0.5, 0.75, 1, 1.25\}$ , and sample sizes  $n \in \{100, 10^6\}$ . Apart from continuous approaches, i.e., NOTEARS-EV and GOLEM-EV, we also report results for discrete approaches A\* and GDS. The thresholds considered are 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, and 0.6. Note that A\* and GDS output a DAG, and thus we apply thresholding on its least squares coefficients. The SHDs are provided in Figure 9, while complete results for F1 score and recall are given in Figure 21 in Appendix F.6. The observations are described below.

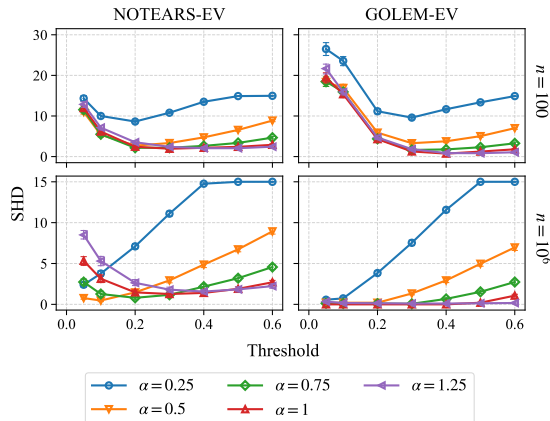


Figure 9: SHDs of different thresholds and 15 variables under different weight scales.

**Observation 1.** When  $n = 100$ , using a relatively large threshold (e.g., 0.3) is beneficial for all methods considered. When the sample size is small, the estimated DAGs by both continuous and discrete approaches may contain many false discoveries owing to finite-sample errors, and using a relatively large threshold helps remove some of them. This appears to increase the precision at the cost of decreasing the recall.

**Observation 2.** When  $n = 10^6$ , using a relatively large threshold (e.g., 0.3) is harmful in many cases. In this case, using a threshold of 0.3 increases the SHDs of all methods for  $\alpha = 0.25, 0.5$ , since the recall is decreased. A possible reason is that the threshold applied is larger than many edge weights in the true weighted adjacency matrix. Therefore, the fixed threshold of 0.3 used in existing works may be undesirable in many cases, as it may wrongly remove many edges and lead to a low recall. To avoid such a case, we adopt a relatively small threshold (e.g., 0.1) in the other experiments.

**Observation 3.** When  $\alpha = 1$ , the optimal threshold for NOTEARS-EV is 0.3 for  $n = 100, 10^6$ . Specifically, its SHDs are  $1.93 \pm 0.36$  and  $1.27 \pm 0.24$  for  $n = 100, 10^6$ , respectively. Moreover, the optimal thresholds for NOTEARS-EV are different across all settings. Since the empirical studies conducted by Zheng et al. [72] and several follow-up works adopt  $\alpha = 1$ , this indicates that their comparison with existing baselines, particularly discrete approaches, may not be completely fair as it may implicitly assume prior knowledge about the range of edge weights in the data generating procedure.

**Observation 4.** When  $n = 10^6$ , NOTEARS-EV requires a relatively large threshold (e.g., 0.3) to perform well for  $\alpha = 0.75, 1, 1.25$ , while a small threshold (i.e., 0.05) is sufficient for GOLEM-EV for

all  $\alpha$  considered. Specifically, with a threshold of 0.05, GOLEM-EV achieves SHDs of  $0.57 \pm 0.18$ ,  $0.20 \pm 0.10$ ,  $0.13 \pm 0.10$ ,  $0 \pm 0$ , and  $0.33 \pm 0.18$  for  $\alpha = 0.25, 0.5, 0.75, 1, 1.25$ , respectively, indicating that the estimated DAGs by GOLEM-EV are very close to ground truths. Since the sample size  $n = 10^6$  used here is large and leads to considerably small finite-sample errors, this appears to suggest that the nonconvex landscape of GOLEM-EV may contain fewer suboptimal local solutions in this setting; thus, when initialized with zero matrix, gradient-based optimization procedure reaches a solution that is close to the ground truth. By contrast, this is not the case of NOTEARS-EV, which, with a threshold of 0.05, achieves SHDs of  $14.33 \pm 0.70$ ,  $11.00 \pm 0.61$ ,  $11.53 \pm 0.72$ ,  $11.97 \pm 0.68$ , and  $12.87 \pm 0.69$  for  $\alpha = 0.25, 0.5, 0.75, 1, 1.25$ , respectively. For these values of  $\alpha$ , a threshold of 0.3 leads to better performance (roughly SHD of 2) for NOTEARS-EV. This indicates that it is more likely for NOTEARS-EV to return suboptimal local solutions as compared to GOLEM-EV, at least in this setting considered where the sample size is large. This observation also suggests that, in addition to reducing false discoveries resulting from finite-sample errors, *thresholding may help reduce false discoveries caused by nonconvexity*, especially for NOTEARS-EV.

It is clear from Observations 1, 2, 3, and 4 that different approaches have different optimal thresholds in different settings. Therefore, a possible direction is to develop a general procedure (e.g., adaptive thresholding) that can be reliably applied in different settings to identify edges from the solutions of continuous structure learning approaches. As discussed above, this may also help reduce false discoveries resulting from both finite-sample errors and nonconvexity (see Observation 4).

### 5.2. Sparsity Penalty

Apart from nonconvexity, another key factor that differentiates continuous structure learning approaches from discrete search approaches (e.g., A\* [67], dynamic programming [53], GES, [12], and GDS [43]) is that the latter use  $\ell_0$  penalty, while the former adopt  $\ell_1$  penalty, which is known to be biased and negatively affect the performance [18, 7]. The reason is that it penalizes all coefficients with the same intensity, including those with small values. Thus, we consider alternative forms of sparsity penalty that help overcome such an issue, namely the SCAD penalty [18] and MCP [69], where large coefficients are penalized less than small ones. Moreover, SCAD penalty and MCP do not require the incoherence condition for support recovery [34] that is needed by  $\ell_1$  penalty in several statistical problems [59, 46], which might be a rather restrictive assumption in practice. This is also the case for structure learning, i.e., Aragam et al. [1] established high dimensional structure consistency of least squares in the space of acyclic weighted adjacency matrices (corresponding to the formulation of NOTEARS), for which the incoherence condition is not needed when using MCP.

We compare the performance of NOTEARS-EV and GOLEM-EV under specification of the sparsity penalties discussed above. We consider ER1, ER2, and ER4 graphs with graph sizes

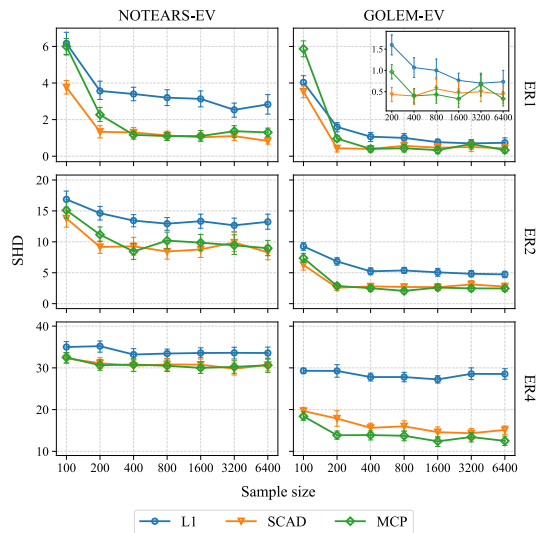


Figure 10: SHDs of different sparsity penalties and 15 variables under different sample sizes.

$d \in \{15, 50\}$  and sample sizes  $n \in \{100, 200, 400, 800, 1600, 3200, 6400\}$ . The SHDs of  $d = 15$  are reported in Figure 10, and complete results are provided in Figure 22 in Appendix F.7. In most settings, we observe that SCAD penalty and MCP achieve significant improvement over  $\ell_1$  penalty for NOTEARS-EV and GOLEM-EV, especially when the degree is large. For 15-node graphs, SCAD penalty and MCP lead to similar performance, while SCAD penalty slightly outperforms MCP on 50-node graphs. A further comparison between SCAD and  $\ell_0$  penalties is provided in Appendix D.

## 6. Conclusion and Discussion

We investigate in which cases continuous structure learning approaches can and cannot perform well and why this happens. We focus on several specifications of the data and search procedure, including varsortability, data standardization, nonconvexity, thresholding, and sparsity. Despite the simplicity of continuous structure learning approaches, we demonstrate that they may suffer from various technical issues discussed below. Our goal here is not to resolve all these issues, but rather to analyze and make them transparent, as well as suggesting possible directions. We hope that our studies could stimulate future works on developing more reliable continuous structure learning approaches, e.g., to overcome the nonconvexity and thresholding issues. Detailed discussions are provided below.

**Varsortability and data standardization.** We show that the statements by Reisach et al. [47] may not hold for formulations with equal and non-equal noise variances, and provide possible alternative explanations for the observations that continuous approaches do not perform well after data standardization [47]. For the equal noise variances formulation, we provide counterexamples to the statement, and demonstrate that the performance of continuous approaches assuming equal noise variances degrade as the noise ratio increases, which may explain why they do not perform well after standardization. For the non-equal noise variances formulation, we show that nonconvexity may be the reason why continuous approaches do not perform well both before and after data standardization.

**Nonconvexity.** Many existing works in continuous structure learning considered only the equal noise variances formulation by adopting least squares and observed remarkable performance. We demonstrate that these recent advances for the equal noise variances formulation fail to achieve improvement for the non-equal noise variances one. This implies that nonconvexity may be a main concern especially for the latter formulation. Our findings suggest that future works should take into account the non-equal noise variances formulation to handle more general settings and for a more comprehensive empirical evaluation, and develop reliable approaches that mitigate the nonconvexity issue.

**Thresholding.** Our experiments indicate that the choice of threshold plays an important role in the final solutions, and that the optimal thresholds may differ across settings. Specifically, (1) using a relatively large threshold of 0.3 as in existing works may be harmful and remove many true positives, especially when the true edge weights are small. (2) Thresholding may be beneficial in certain cases as it not only helps reduce false discoveries resulting from finite-sample errors, but may also help reduce those caused by nonconvexity, especially for NOTEARS. Thus, the choice of threshold should be treated with care. A possible future direction is to develop a general procedure (e.g., adaptive thresholding) that can be applied in different settings to identify edges from the solutions of continuous structure learning approaches, while also reducing false discoveries arising from nonconvexity.

**Sparsity penalty.** We demonstrate that sparsity penalty plays a crucial role for continuous structure learning approaches, e.g., GOLEM-EV and NOTEARS-EV. Specifically,  $\ell_1$  penalty may not perform well possibly due to its bias, while other penalties such as SCAD and MCP could help remedy it.

## Acknowledgments

The authors would like to thank the anonymous reviewers for helpful comments and suggestions. This project is partially supported by NSF Grant 2229881, the National Institutes of Health (NIH) under Contract R01HL159805, and grants from Apple Inc., KDDI Research Inc., Quris AI, and Infinite Brain Technology.

## References

- [1] B. Aragam, A. Amini, and Q. Zhou. Globally optimal score-based learning of directed acyclic graphs in high-dimensions. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [2] J. Bang-Jensen and G. Gutin. *Digraphs. Theory, Algorithms and Applications*. Springer Monographs in Mathematics, 2009.
- [3] K. Bello, B. Aragam, and P. Ravikumar. DAGMA: Learning DAGs via M-matrices and a log-determinant acyclicity characterization. In *Advances in Neural Information Processing Systems*, 2022.
- [4] A. Bellot and M. van der Schaar. Deconfounded score method: Scoring DAGs with dense unobserved confounding. *arXiv preprint arXiv:2103.15106*, 2021.
- [5] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2nd edition, 1999.
- [6] R. Bhattacharya, T. Nagarajan, D. Malinsky, and I. Shpitser. Differentiable causal discovery under unmeasured confounding. In *International Conference on Artificial Intelligence and Statistics*, 2021.
- [7] P. Breheny and J. Huang. Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *The Annals of Applied Statistics*, 5(1):232–253, 2011.
- [8] P. Brouillard, S. Lachapelle, A. Lacoste, S. Lacoste-Julien, and A. Drouin. Differentiable causal discovery from interventional data. In *Advances in Neural Information Processing Systems*, 2020.
- [9] B. Charpentier, S. Kibler, and S. Günnemann. Differentiable DAG sampling. In *International Conference on Learning Representations*, 2022.
- [10] X. Chen, H. Sun, C. Ellington, E. Xing, and L. Song. Multi-task learning of order-consistent causal graphs. In *Advances in Neural Information Processing Systems*, 2021.
- [11] D. M. Chickering. Learning Bayesian networks is NP-complete. In *Learning from Data: Artificial Intelligence and Statistics V*. Springer, 1996.
- [12] D. M. Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3(Nov):507–554, 2002.

- [13] D. M. Chickering, D. Heckerman, and C. Meek. Large-sample learning of Bayesian networks is NP-hard. *Journal of Machine Learning Research*, 5, 2004.
- [14] J. Cussens. Bayesian network learning with cutting planes. In *Conference on Uncertainty in Artificial Intelligence*, 2011.
- [15] C. Deng, K. Bello, P. K. Ravikumar, and B. Aragam. Global optimality in bivariate gradient-based DAG learning. In *Advances in Neural Information Processing Systems*, 2023.
- [16] D. Dor and M. Tarsi. A simple algorithm to construct a consistent extension of a partially oriented graph. *Technical Report R-185, Cognitive Systems Laboratory, UCLA*, 1992.
- [17] P. Erdős and A. Rényi. On random graphs I. *Publicationes Mathematicae*, 6:290–297, 1959.
- [18] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- [19] G. R. A. Faria, A. Martins, and M. A. T. Figueiredo. Differentiable causal discovery under latent interventions. In *Conference on Causal Learning and Reasoning*, 2022.
- [20] E. Gao, I. Ng, M. Gong, L. Shen, W. Huang, T. Liu, K. Zhang, and H. Bondell. MissDAG: Causal discovery in the presence of missing data with continuous additive noise models. In *Advances in Neural Information Processing Systems*, 2022.
- [21] E. Gao, J. Chen, L. Shen, T. Liu, M. Gong, and H. Bondell. FedDAG: Federated DAG structure learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.
- [22] W. Gong, D. Smith, Z. Wang, C. Barton, S. Woodhead, N. Pawlowski, J. Jennings, and C. Zhang. NeurIPS competition instructions and guide: Causal insights for learning paths in education. *arXiv preprint arXiv:2208.12610*, 2022.
- [23] P. Hoyer, D. Janzing, J. M. Mooij, J. Peters, and B. Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems*, 2009.
- [24] T. Ikeuchi, M. Ide, Y. Zeng, T. N. Maeda, and S. Shimizu. Python package for causal discovery based on LiNGAM. *Journal of Machine Learning Research*, 24(14):1–8, 2023.
- [25] X. Jiang, L.-H. Lim, Y. Yao, and Y. Ye. Statistical ranking and combinatorial Hodge theory. *Mathematical Programming*, 127, 11 2011.
- [26] M. Kaiser and M. Sipos. Unsuitability of NOTEARS for causal graph discovery when dealing with dimensional quantities. *Neural Processing Letters*, 54:1–9, 06 2022.
- [27] D. Kalainathan, O. Goudet, I. Guyon, D. Lopez-Paz, and M. Sebag. Structural agnostic modeling: Adversarial learning of causal graphs. *Journal of Machine Learning Research*, 23 (219):1–62, 2022.
- [28] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2014.

- [29] M. Koivisto and K. Sood. Exact Bayesian structure discovery in Bayesian networks. *Journal of Machine Learning Research*, 5(Dec):549–573, 2004.
- [30] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, MA, 2009.
- [31] W. Kool, H. Van Hoof, and M. Welling. Stochastic beams and where to find them: The Gumbel-top-k trick for sampling sequences without replacement. In *International Conference on Machine Learning*, 2019.
- [32] S. Lachapelle, P. Brouillard, T. Deleu, and S. Lacoste-Julien. Gradient-based neural DAG learning. In *International Conference on Learning Representations*, 2020.
- [33] P.-L. Loh and P. Bühlmann. High-dimensional learning of linear causal networks via inverse covariance estimation. *Journal of Machine Learning Research*, 15(88):3065–3105, 2014.
- [34] P.-L. Loh and M. J. Wainwright. Support recovery without incoherence: A case for nonconvex regularization. *The Annals of Statistics*, 45(6):2455–2482, 2017.
- [35] G. Mena, D. Belanger, S. Linderman, and J. Snoek. Learning latent permutations with Gumbel-Sinkhorn networks. In *International Conference on Learning Representations*, 2018.
- [36] I. Ng and K. Zhang. Towards federated Bayesian network structure learning with continuous optimization. In *International Conference on Artificial Intelligence and Statistics*, 2022.
- [37] I. Ng, A. Ghassami, and K. Zhang. On the role of sparsity and DAG constraints for learning linear DAGs. In *Advances in Neural Information Processing Systems*, 2020.
- [38] I. Ng, S. Lachapelle, N. R. Ke, S. Lacoste-Julien, and K. Zhang. On the convergence of continuous constrained optimization for structure learning. In *International Conference on Artificial Intelligence and Statistics*, 2022.
- [39] I. Ng, S. Zhu, Z. Fang, H. Li, Z. Chen, and J. Wang. Masked gradient-based causal structure learning. In *SIAM International Conference on Data Mining*, 2022.
- [40] J. Nocedal and S. J. Wright. *Numerical optimization*. Springer series in operations research and financial engineering. Springer, 2nd edition, 2006.
- [41] R. Pamfil, N. Sriwattanaworachai, S. Desai, P. Pilgerstorfer, P. Beaumont, K. Georgatzis, and B. Aragam. DYNOTEARS: Structure learning from time-series data. In *International Conference on Artificial Intelligence and Statistics*, 2020.
- [42] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [43] J. Peters and P. Bühlmann. Identifiability of Gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2013.
- [44] J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference - Foundations and Learning Algorithms*. MIT Press, 2017.



- [45] J. Ramsey, M. Glymour, R. Sanchez-Romero, and C. Glymour. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International Journal of Data Science and Analytics*, 3(2):121–129, 2017.
- [46] P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu. High-dimensional covariance estimation by minimizing  $\ell_1$ -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.
- [47] A. Reisach, C. Seiler, and S. Weichwald. Beware of the simulated DAG! Causal discovery benchmarks may be easy to game. In *Advances in Neural Information Processing Systems*, 2021.
- [48] M. Schatzman and J. Taylor. *Numerical Analysis A Mathematical Introduction*. Oxford University Press, 2002.
- [49] R. Scheines, P. Spirtes, C. Glymour, C. Meek, and T. Richardson. The TETRAD project: Constraint based aids to causal model specification. *Multivariate Behavioral Research*, 33: 65–117, 1998.
- [50] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.
- [51] S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(Oct):2003–2030, 2006.
- [52] S. Shimizu, T. Inazumi, Y. Sogawa, A. Hyvärinen, 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(Apr):1225–1248, 2011.
- [53] A. P. Singh and A. W. Moore. Finding optimal Bayesian networks by dynamic programming. Technical report, Carnegie Mellon University, 2005.
- [54] P. Spirtes and C. Glymour. An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9:62–72, 1991.
- [55] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2nd edition, 2001.
- [56] G. W. Stewart. On the continuity of the generalized inverse. *SIAM Journal on Applied Mathematics*, 17(1):33–45, 1969. ISSN 00361399.
- [57] X. Sun, G. Liu, P. Poupart, and O. Schulte. NTS-NOTEARS: Learning nonparametric temporal DAGs with time-series data and prior knowledge. *arXiv preprint arXiv:2109.04286*, 2021.
- [58] M. J. Vowels, N. C. Camgoz, and R. Bowden. D’ya like DAGs? A survey on structure learning and causal discovery. *ACM Computing Surveys*, 55(4), nov 2022. ISSN 0360-0300.
- [59] M. J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using  $\ell_1$ -constrained quadratic programming (Lasso). *IEEE Transactions on Information Theory*, 55(5): 2183–2202, 2009.

- [60] Y. Wang, V. Menkovski, H. Wang, X. Du, and M. Pechenizkiy. Causal discovery from incomplete data: A deep learning approach. *arXiv preprint arXiv:2001.05343*, 2020.
- [61] D. Wei, T. Gao, and Y. Yu. DAGs with no fears: A closer look at continuous optimization for learning Bayesian networks. In *Advances in Neural Information Processing Systems*, 2020.
- [62] J. H. Wilkinson. *Rounding Errors in Algebraic Processes*. The Harlan D. Mills Collection, 1965.
- [63] D. Xu, E. Gao, W. Huang, A. Song, and M. Gong. On the sparse DAG structure learning based on adaptive Lasso. *arXiv preprint arXiv:2209.02946*, 2022.
- [64] M. Yang, F. Liu, Z. Chen, X. Shen, J. Hao, and J. Wang. CausalVAE: Disentangled representation learning via neural structural causal models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [65] Y. Yu, J. Chen, T. Gao, and M. Yu. DAG-GNN: DAG structure learning with graph neural networks. In *International Conference on Machine Learning*, 2019.
- [66] Y. Yu, T. Gao, N. Yin, and Q. Ji. DAGs with no curl: An efficient DAG structure learning approach. In *International Conference on Machine Learning*, 2021.
- [67] C. Yuan and B. Malone. Learning optimal Bayesian networks: A shortest path perspective. *Journal of Artificial Intelligence Research*, 48(1):23–65, 2013.
- [68] Y. Zeng, S. Shimizu, R. Cai, F. Xie, M. Yamamoto, and Z. Hao. Causal discovery with multi-domain LiNGAM for latent factors. In *International Joint Conference on Artificial Intelligence*, 2021.
- [69] C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010.
- [70] Z. Zhang, I. Ng, D. Gong, Y. Liu, E. M. Abbasnejad, M. Gong, K. Zhang, and J. Q. Shi. Truncated matrix power iteration for differentiable DAG learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [71] X. Zheng. *Learning DAGs with Continuous Optimization*. PhD thesis, Carnegie Mellon University, 2020.
- [72] X. Zheng, B. Aragam, P. Ravikumar, and E. P. Xing. DAGs with NO TEARS: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems*, 2018.
- [73] X. Zheng, C. Dan, B. Aragam, P. Ravikumar, and E. P. Xing. Learning sparse nonparametric DAGs. In *International Conference on Artificial Intelligence and Statistics*, 2020.
- [74] Y. Zheng, B. Huang, W. Chen, J. Ramsey, M. Gong, R. Cai, S. Shimizu, P. Spirtes, and K. Zhang. Causal-learn: Causal discovery in Python. *arXiv preprint arXiv:2307.16405*, 2023.
- [75] H. Zou. The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476):1418–1429, 2006.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Preliminaries</b>	<b>2</b>
2.1	Structure Learning with Continuous Optimization . . . . .	2
2.2	Varsortability . . . . .	3
<b>3</b>	<b>Varsortability and Data Standardization</b>	<b>4</b>
3.1	With Equal Noise Variances Formulation . . . . .	4
3.2	With Non-Equal Noise Variances Formulation . . . . .	6
<b>4</b>	<b>Nonconvexity of Non-Equal Noise Variances Formulation</b>	<b>8</b>
4.1	Search Strategies . . . . .	9
4.2	DAG Constraints . . . . .	9
4.3	Other Settings . . . . .	10
<b>5</b>	<b>Thresholding and Sparsity Penalty</b>	<b>10</b>
5.1	Thresholding . . . . .	11
5.2	Sparsity Penalty . . . . .	12
<b>6</b>	<b>Conclusion and Discussion</b>	<b>13</b>
<b>A</b>	<b>Discussion of Statements by Reisach et al. [47]</b>	<b>20</b>
<b>B</b>	<b>Proofs</b>	<b>20</b>
B.1	Proof of Proposition 1 . . . . .	21
B.2	Proof of Proposition 2 . . . . .	23
<b>C</b>	<b>Analysis of Non-Equal Noise Variances Formulation</b>	<b>25</b>
C.1	Random Initializations . . . . .	25
C.2	Equal-Noise Variances Formulation with Random Perturbations . . . . .	25
C.3	Alternative Form of Likelihood Function . . . . .	26
C.4	Analysis of Estimated Score and SHD . . . . .	27
<b>D</b>	<b>Comparison between SCAD and <math>\ell_0</math> Penalties</b>	<b>28</b>
<b>E</b>	<b>Supplementary Experiment Details</b>	<b>28</b>
E.1	Implementation Details and Evaluation Metrics . . . . .	28
E.2	Discussion of Search Strategies . . . . .	29
<b>F</b>	<b>Supplementary Experiment Results</b>	<b>29</b>
F.1	With Equal Noise Variances Formulation . . . . .	30
F.2	With Non-Equal Noise Variances Formulation . . . . .	31
F.3	Search Strategies . . . . .	31
F.4	DAG Constraints . . . . .	32
F.5	Other Settings . . . . .	33
F.6	Thresholding . . . . .	34
F.7	Sparsity Penalty . . . . .	35

## Appendix A. Discussion of Statements by Reisach et al. [47]

We first list a few excerpts from Reisach et al. [47] that involve the relationship between varsortability and the performance of continuous structure learning approaches:

- Abstract: “the remarkable performance of some continuous structure learning algorithms can be explained by high varsortability”.
- Section 1: “Our experiments demonstrate that varsortability dominates the optimization and helps achieve state-of-the-art performance provided the ground-truth data scale”.
- Section 3.4: “We explain how varsortability may dominate the performance of continuous structure learning algorithms”.
- Section 3.4: “For this reason we focus on the first optimization steps to explain a) why continuous structure learning algorithms that assume equal noise variance work remarkably well in the presence of high varsortability”.
- Section 4.3: “the evidence corroborates our claim that the remarkable performance on raw data and the overall behavior upon standardization of the continuous structure learning algorithms may be driven primarily by high varsortability”.

For clarity and ease of further analysis, in Section 2.2 we attempt to provide a partial formulation of the statements (on which this work focuses), separated into two different cases. Specifically, we formulate Statements 1 and 2 by following the wordings of the fourth excerpt above because it has a relatively clear technical interpretation. Note that our analysis in Section 3 also applies if the statements are formulated based on some different excerpts. For instance, we provide alternative formulations of the statements (corresponding to Statements 1 and 2, respectively) that follow the wordings of the second excerpt above.

**Statement 3 (Equal Noise Variances Formulation)** *Varsortability dominates the optimization of continuous structure learning approaches that assume equal noise variances, specifically NOTEARS-EV and GOLEM-EV, and helps them achieve state-of-the-art performance provided the ground-truth data scale.*

**Statement 4 (Non-Equal Noise Variances Formulation)** *Varsortability dominates the optimization of continuous structure learning approaches that assume non-equal noise variances, specifically GOLEM-NV, and helps them achieve state-of-the-art performance provided the ground-truth data scale.*

In particular, the examples in Section 3.1, including Example 1 and Proposition 1, as well as the corresponding empirical studies, demonstrate that a high varsortability does not help NOTEARS-EV and GOLEM-EV achieve a good performance, even when the ground-truth data scale is used. This is also the case for GOLEM-NV, as demonstrated in Section 3.2. Furthermore, these studies suggest that varsortability may not dominate the optimization of these continuous structure learning approaches, at least in the settings considered in this work.

## Appendix B. Proofs

In this section, we provide the proofs of Propositions 1 and 2 given in Section 3.1.

### B.1. Proof of Proposition 1

Let  $(\tilde{B}, \tilde{\Omega})$  be parameters associated with the linear SEMs over variables  $X = (X_1, \dots, X_d)$ ,  $d \geq 3$ , where the noise variables follow Gaussian distributions.<sup>3</sup> By definition,  $\tilde{\Omega}$  is a positive diagonal matrix, while the weighted adjacency matrix  $\tilde{B}$  corresponds to a DAG.

The maximum number of free parameters in matrix  $\tilde{B}$  is  $\frac{d(d-1)}{2}$ , which corresponds to fully connected DAGs with various topological orders. Since there is only a finite number of topological orders with  $d$  nodes, it suffices to consider fully connected DAGs with a specific topological order, e.g., where  $X_{i+1}$  precedes  $X_i$  for  $i = 1, \dots, d-1$ . We denote by  $\tilde{\mathcal{G}}$  the fully connected DAG with this specific topological order. In this case, the corresponding weighted adjacency matrix  $\tilde{B}$  is a strictly lower triangular matrix with  $\frac{d(d-1)}{2}$  free parameters.

Note that, in the large sample limit, the covariance matrix of  $X$  is

$$\tilde{\Sigma} = (I - \tilde{B})^{-T} \tilde{\Omega} (I - \tilde{B})^{-1},$$

and the least squares loss of matrix  $B$  is

$$\ell(B; \tilde{\Sigma}) = \frac{1}{2} \text{tr}((I - B)^T \tilde{\Sigma} (I - B)).$$

Clearly, the varsortability equals one if

$$\text{Var}(X_i) > \text{Var}(X_{i+1}) \iff \tilde{\Sigma}_{i,i} - \tilde{\Sigma}_{i+1,i+1} > 0, \quad \text{for } i = 1, \dots, d-1. \quad (2)$$

It is straightforward to show that the least squares score of  $\tilde{\mathcal{G}}$  is  $\ell(\tilde{B}; \tilde{\Sigma})$  [33, Lemma 6], given by

$$\ell(\tilde{B}; \tilde{\Sigma}) = \frac{1}{2} \text{tr}((I - \tilde{B})^T \tilde{\Sigma} (I - \tilde{B})) = \frac{1}{2} \text{tr}(\tilde{\Omega}). \quad (3)$$

Since matrices  $I - \tilde{B}$  and  $\tilde{\Omega}$  are invertible,  $\tilde{\Sigma}$  is symmetric positive definite. Let  $\tilde{L}$  be the (unique) Cholesky factor with positive diagonal entries, i.e.,

$$\tilde{L} \tilde{L}^T = \tilde{\Sigma}. \quad (4)$$

We now define  $(\hat{B}, \hat{\Omega})$  as

$$\hat{\Omega} := \text{diag}\left((\tilde{L}_{1,1})^2, \dots, (\tilde{L}_{d,d})^2\right) \quad \text{and} \quad \hat{B} := I - \tilde{L}^{-T} \hat{\Omega}^{\frac{1}{2}}, \quad (5)$$

which implies

$$(I - \hat{B})^{-T} \hat{\Omega}^{\frac{1}{2}} = \tilde{L} \quad (6)$$

and that  $\hat{B}$  is strictly upper triangular. By Eqs. (4) and (6), we have

$$(I - \hat{B})^T \tilde{\Sigma} (I - \hat{B}) = \hat{\Omega}. \quad (7)$$

Because  $\hat{B}$  is strictly upper triangular, the DAG defined by  $\hat{B}$  is different from  $\tilde{\mathcal{G}}$  (i.e., the DAG defined by  $\tilde{B}$ ). By Eqs. (5) and (7), the least squares loss of  $\hat{B}$  in the large sample limit is

$$\ell(\hat{B}; \tilde{\Sigma}) = \frac{1}{2} \text{tr}((I - \hat{B})^T \tilde{\Sigma} (I - \hat{B})) = \frac{1}{2} \text{tr}(\hat{\Omega}) = \frac{1}{2} \text{tr}(\tilde{L} \odot \tilde{L}). \quad (8)$$

3. For clarity in this proof, we use  $(\tilde{B}, \tilde{\Omega})$  to denote the parameters instead of  $(B, \Omega)$  as in the proposition statement.

For the alternative weighted adjacency matrix  $\hat{B}$  to yield a lower least squares loss than the true one  $\tilde{B}$ , i.e.,  $\ell(\hat{B}; \tilde{\Sigma}) < \ell(\tilde{B}; \tilde{\Sigma})$ , we have the following inequality by substituting Eqs. (3) and (8):

$$\text{tr}(\tilde{\Omega}) - \text{tr}(\tilde{L} \odot \tilde{L}) > 0. \quad (9)$$

The following lemma concludes the proof of the proposition.

**Lemma 3** Consider strictly lower triangular matrix  $\tilde{B}$  and positive diagonal matrix  $\tilde{\Omega}$ . The set of parameters  $(\tilde{B}, \tilde{\Omega})$  that satisfy Inequalities (2) and (9) has a nonzero Lebesgue measure.

**Proof of Lemma 3** Here we give a proof for the cases of  $d \geq 4$ ; such a proof holds similarly for  $d = 3$ . First note that the LHS of Inequalities (2) and (9) are continuous functions of (the entries of)  $\tilde{B}$  and  $\tilde{\Omega}$  on the given domain, because, given that  $I - \tilde{B}$  is invertible and that  $\tilde{\Sigma}$  is symmetric positive definite, the mappings of matrix inversion [62, 56] and Cholesky factor [48, Chapter 12] are continuous. In this case, it suffices to find a specific point  $(\tilde{B}, \tilde{\Omega})$  that satisfies Inequalities (2) and (9), because these inequalities are defined by continuous functions and thus they hold in an open neighborhood of this point.

According to the given domain, consider the following parameters:

$$\tilde{B} = \begin{bmatrix} 0 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & \cdots & 0 & \frac{1}{2} & 0 & 0 \\ 0 & \cdots & 0 & \frac{3}{4} & \frac{1}{4} & 0 \end{bmatrix} \quad \text{and} \quad \tilde{\Omega} = \text{diag} \left( d-2, d-3, d-4, \dots, 4, 3, 2, \frac{1}{10}, 1, 1 \right),$$

both of which are  $d \times d$  matrices. Note that  $\tilde{B}$  is a strictly lower triangular matrix. We then have

$$\tilde{\Sigma} = (I - \tilde{B})^{-T} \tilde{\Omega} (I - \tilde{B})^{-1} = \begin{bmatrix} d-2 & 0 & \cdots & 0 & 0 & 0 & 0 & 0 \\ 0 & d-3 & \cdots & 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & 3 & 0 & 0 & 0 & 0 \\ 0 & 0 & \cdots & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 & \frac{357}{320} & \frac{23}{32} & \frac{7}{8} \\ 0 & 0 & \cdots & 0 & 0 & \frac{23}{32} & \frac{17}{16} & \frac{1}{4} \\ 0 & 0 & \cdots & 0 & 0 & \frac{7}{8} & \frac{1}{4} & 1 \end{bmatrix}.$$

Clearly, Inequality (2) is satisfied. By simple calculation, the Cholesky factor  $\tilde{L}$  of  $\tilde{\Sigma}$  is

$$\tilde{L} = \begin{bmatrix} \sqrt{d-2} & 0 & \cdots & 0 & 0 & 0 & 0 & 0 \\ 0 & \sqrt{d-1} & \cdots & 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \sqrt{3} & 0 & 0 & 0 & 0 \\ 0 & 0 & \cdots & 0 & \sqrt{2} & 0 & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 & \frac{\sqrt{1785}}{40} & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 & \frac{23\sqrt{1785}}{1428} & \frac{\sqrt{76398}}{357} & 0 \\ 0 & 0 & \cdots & 0 & 0 & \frac{\sqrt{1785}}{51} & \frac{-8\sqrt{76398}}{5457} & \frac{4\sqrt{107}}{107} \end{bmatrix},$$

which indicates that Inequality (9) also holds. ■

## B.2. Proof of Proposition 2

Let  $(\tilde{B}, \tilde{\Omega})$  be parameters associated with the linear SEMs over variables  $X = (X_1, \dots, X_d)$  induced by DAG  $\tilde{\mathcal{G}}$ , where the noise variables follow Gaussian distributions.<sup>4</sup> By definition,  $\tilde{\Omega}$  is a positive diagonal matrix, while the weighted adjacency matrix  $\tilde{B}$  corresponds to a DAG.

We use  $\text{SN}_{\tilde{\mathcal{G}}}$  and  $\text{NSN}_{\tilde{\mathcal{G}}}$  to denote the set of source nodes and non-source nodes in DAG  $\tilde{\mathcal{G}}$ , respectively. We also denote by  $\text{PA}_{\tilde{\mathcal{G}}}(X_i)$ ,  $\text{CH}_{\tilde{\mathcal{G}}}(X_i)$ , and  $\text{DE}_{\tilde{\mathcal{G}}}(X_i)$  the set of parents, children, and descendants of  $X_i$  in  $\tilde{\mathcal{G}}$ , respectively. For any two DAGs  $\mathcal{G}_1$  and  $\mathcal{G}_2$ , we write  $\mathcal{G}_1 \subseteq \mathcal{G}_2$  if all edges in  $\mathcal{G}_1$  are present in  $\mathcal{G}_2$ . Furthermore, let  $\mathbb{D}^d$  be the set of  $d$ -node DAGs, and, for DAG  $\mathcal{G}$ , define  $\mathbb{R}_{\mathcal{G}}$  as

$$\mathbb{R}_{\mathcal{G}} := \{B \in \mathbb{R}^{d \times d} : B_{i,j} = 0 \text{ when } X_i \rightarrow X_j \text{ is not an edge in } \mathcal{G}\}.$$

Note that, in the large sample limit, the covariance matrix of  $X$  is

$$\tilde{\Sigma} = (I - \tilde{B})^{-T} \tilde{\Omega} (I - \tilde{B})^{-1}.$$

The varsortability is equal to  $v_{\tilde{\mathcal{G}}}^{\text{source}}$ , where  $v_{\tilde{\mathcal{G}}}^{\text{source}}$  is defined as in the proposition statement, if the following inequalities hold:

$$\text{Var}(X_i) < \text{Var}(X_j) \iff \tilde{\Sigma}_{i,i} - \tilde{\Sigma}_{j,j} < 0, \quad \text{for } X_i \in \text{SN}_{\tilde{\mathcal{G}}}, X_j \in \text{DE}_{\tilde{\mathcal{G}}}(X_i), \quad (10)$$

and

$$\text{Var}(X_i) > \text{Var}(X_j) \iff \tilde{\Sigma}_{i,i} - \tilde{\Sigma}_{j,j} > 0, \quad \text{for } X_i \in \text{NSN}_{\tilde{\mathcal{G}}}, X_j \in \text{CH}_{\tilde{\mathcal{G}}}(X_i). \quad (11)$$

With a slight abuse of notation, let  $\ell(\mathcal{G}; \tilde{\Sigma})$  denote the least squares score of DAG  $\mathcal{G}$ , and  $\ell(B; \tilde{\Sigma})$  denote the least squares loss of matrix  $B$ . Specifically, we have

$$\ell(\mathcal{G}; \tilde{\Sigma}) := \min_{B \in \mathbb{R}_{\mathcal{G}}} \ell(B; \tilde{\Sigma}). \quad (12)$$

For any DAG  $\hat{\mathcal{G}} \supseteq \tilde{\mathcal{G}}$ , we have  $\ell(\hat{\mathcal{G}}; \tilde{\Sigma}) = \ell(\tilde{\mathcal{G}}; \tilde{\Sigma})$  by Loh and Bühlmann [33, Lemma 6]. Therefore, it suffices to consider the DAG  $\hat{\mathcal{G}}$  where  $\hat{\mathcal{G}} \not\supseteq \tilde{\mathcal{G}}$ . We then consider the following inequality:

$$\ell(\tilde{\mathcal{G}}; \tilde{\Sigma}) - \ell(\hat{\mathcal{G}}; \tilde{\Sigma}) < 0, \quad \text{for } \hat{\mathcal{G}} \in \mathbb{D}^d, \hat{\mathcal{G}} \not\supseteq \tilde{\mathcal{G}}, \quad (13)$$

or equivalently, the true weighted adjacency matrix  $\tilde{B}$  yields the lowest least squares score.

The following lemma concludes the proof of the proposition.

**Lemma 4** *Given DAG  $\tilde{\mathcal{G}}$ , consider matrix  $\tilde{B} \in \mathbb{R}_{\tilde{\mathcal{G}}}$  and positive diagonal matrix  $\tilde{\Omega}$ . The set of parameters  $(\tilde{B}, \tilde{\Omega})$  that satisfy Inequalities (10), (11), and (13) has a nonzero Lebesgue measure.*

**Proof of Lemma 4** First it is clear that the LHS of Inequalities (10) and (11) are continuous functions of (the entries of)  $\tilde{B}$  and  $\tilde{\Omega}$  on the given domain, because  $I - \tilde{B}$  is invertible and the mapping of matrix inversion is continuous [62, 56]. For Inequality (13), the least squares score  $\ell(\hat{\mathcal{G}}; \tilde{\Sigma})$  is computed by solving the optimization problem in Eq. (12); each column, say the  $j$ -th one, of the estimated solution consists of the coefficients obtained from linearly regressing  $X_j$  upon its parents in DAG

4. For clarity in this proof, we use  $(\tilde{B}, \tilde{\Omega})$  to denote the parameters instead of  $(B, \Omega)$  as in the proposition statement. Similarly, we use  $\tilde{\mathcal{G}}$  to denote the DAG instead of  $\mathcal{G}$ .

$\hat{\mathcal{G}}$  [33, Remark 5]. Therefore, it is straightforward to show that the LHS of Inequality (13) is also a continuous function of  $\tilde{B}$  and  $\tilde{\Omega}$  on the given domain, because  $\tilde{\Sigma}$  is symmetric positive definite and so any of its principal submatrix is invertible. In this case, it suffices to find a specific point  $(\tilde{B}, \tilde{\Omega})$  that satisfies Inequalities (10), (11), and (13), because these inequalities are defined by continuous functions and thus they hold in an open neighborhood of this point.

We start with Inequality (13). For the free parameters in diagonal matrix  $\tilde{\Omega}$ , we set them to be equal, i.e.,  $\tilde{\sigma}_1^2 = \dots = \tilde{\sigma}_d^2 = \tilde{\sigma}^2$ . This restriction leads to the equal noise variances case, which, by Loh and Bühlmann [33, Theorem 7], indicates that the true weighted adjacency matrix  $\tilde{B}$  yields the lowest least squares score. Thus, Inequality (13) holds with this restriction.

We now construct parameter  $\tilde{B}$  that such that Inequalities (10) and (11) hold. Consider the variable  $X_j$ . If all parents of  $X_j$  are source nodes in DAG  $\tilde{\mathcal{G}}$ , we set  $\tilde{B}_{i,j}$  to arbitrary (strictly) positive value for each edge  $X_i \rightarrow X_j$  in  $\tilde{\mathcal{G}}$ . If some parents of  $X_j$  are not source nodes in  $\tilde{\mathcal{G}}$ , then we set

$$\tilde{B}_{i,j} := \frac{1}{|\text{PA}_{\tilde{\mathcal{G}}}(X_j)|} \left( \frac{c_j}{\text{Var}(X_i)} \right)^{\frac{1}{2}} \quad (14)$$

for each edge  $X_i \rightarrow X_j$  in  $\tilde{\mathcal{G}}$ , where  $c_j$  is an arbitrary value from the following interval:

$$c_j \in \left( 0, \min_{X_k \in \text{PA}_{\tilde{\mathcal{G}}}(X_j) \cap \text{NSN}_{\tilde{\mathcal{G}}}} \{\text{Var}(X_k)\} - \tilde{\sigma}^2 \right). \quad (15)$$

We construct the above matrix by following the topological order of DAG  $\tilde{\mathcal{G}}$ . Specifically, when we are constructing the  $j$ -th column of  $\tilde{B}$ , every incoming edge into the ancestors of variable  $X_j$  has been assigned a weight. Thus, the variances of the parents of  $X_j$  can be computed (see Eq. (15)).

Note that all entries of the matrix  $\tilde{B}$  constructed above that correspond to the edges in  $\tilde{\mathcal{G}}$  are nonzero. Therefore, for  $X_i \in \text{SN}_{\tilde{\mathcal{G}}}$  and  $X_j \in \text{DE}_{\tilde{\mathcal{G}}}(X_i)$ , we have

$$\text{Var}(X_j) = \text{Var}(\tilde{B}_{\cdot,j}^T X + N_j) = \text{Var}(\tilde{B}_{\cdot,j}^T X) + \text{Var}(N_j) > \text{Var}(N_j) = \tilde{\sigma}^2 = \text{Var}(X_i),$$

which implies that Inequality (10) holds.

Moreover, for  $X_i \in \text{NSN}_{\tilde{\mathcal{G}}}$  and  $X_j \in \text{CH}_{\tilde{\mathcal{G}}}(X_i)$ , we have

$$\begin{aligned} \text{Var}(X_j) &= \text{Var} \left( \sum_{X_l \in \text{PA}_{\tilde{\mathcal{G}}}(X_j)} B_{l,j} X_l + N_j \right) \\ &= \sum_{X_l, X_m \in \text{PA}_{\tilde{\mathcal{G}}}(X_j)} B_{l,j} B_{m,j} \text{Cov}(X_l, X_m) + \text{Var}(N_j) \\ &\leq \sum_{X_l, X_m \in \text{PA}_{\tilde{\mathcal{G}}}(X_j)} B_{l,j} B_{m,j} \text{Var}(X_l)^{\frac{1}{2}} \text{Var}(X_m)^{\frac{1}{2}} + \tilde{\sigma}^2 \quad (\because B_{l,j} B_{m,j} > 0) \\ &= |\text{PA}_{\tilde{\mathcal{G}}}(X_j)|^2 \cdot \frac{c_j}{|\text{PA}_{\tilde{\mathcal{G}}}(X_j)|^2} + \tilde{\sigma}^2 \quad (\text{Substituting Eq. (14)}) \\ &< \min_{X_k \in \text{PA}_{\tilde{\mathcal{G}}}(X_j) \cap \text{NSN}_{\tilde{\mathcal{G}}}} \{\text{Var}(X_k)\} \quad (\text{Substituting Eq. (15)}) \\ &\leq \text{Var}(X_i). \end{aligned}$$

Thus, Inequality (11) also holds. Therefore, we have constructed a specific point  $(\tilde{B}, \tilde{\Omega})$  that satisfies Inequalities (10), (11), and (13).  $\blacksquare$



## Appendix C. Analysis of Non-Equal Noise Variances Formulation

We provide analysis of the non-equal noise variances formulation. We experiment with random initializations in Appendix C.1, and GOLEM-EV with random perturbations in Appendix C.2. We consider alternative form of likelihood function (without profiling out the parameter  $\Omega$ ) in Appendix C.3, and investigate the connection between estimated score and SHD in Appendix C.4.

### C.1. Random Initializations

In Sections 3 and 4, we adopt the initialization scheme with zero matrix for various structure learning methods such as GOLEM-NV and NOTEARS. Note that this initialization scheme has been used by several existing works [72, 37, 3], especially those based on linear SEMs. Here, we conduct experiments for GOLEM-NV with random initializations to examine whether it produces a better solution and helps remedy the possible nonconvexity issue discussed in Sections 3.2 and 4. Specifically, each entry of the initial solution  $B$  is sampled uniformly at random from  $[-\epsilon, \epsilon]$ , where  $\epsilon \in \{0.01, 0.05, 0.1\}$ . Furthermore, we run such random initialization for 10 times, and select the final solution that leads to the best score (see Eq. (1)). Apart from the random initialization scheme, we also report results for GOLEM-NV initialized with the solution of GOLEM-EV and that of FGES.

The empirical results are reported in Table 2. We observe that random initializations perform similarly to the initialization scheme with zero matrix, both of which lead to a much worse performance as compared to GOLEM-NV initialized with GOLEM-EV or with FGES. This suggests that random initializations do not seem to remedy the possible nonconvexity issue, at least in the setting considered here. Moreover, this further validates the findings in Section 3.2, i.e., a proper initialization scheme is crucial to the quality of the final solution for non-equal noise variances formulation.

Table 2: Empirical results of different initialization schemes. The number of variables, sample size, and noise ratio are 50,  $10^6$ , and 2, respectively. The standard errors computed over 30 random repetitions are also reported. Here, “int.” stands for “initialized”.

	SHD of CPDAG	F1 of skeleton	F1 of arrows
GOLEM-NV int. with zero matrix	$120.27 \pm 2.91$	$0.50 \pm 0.00$	$0.04 \pm 0.00$
GOLEM-NV int. with GOLEM-EV	$4.50 \pm 1.11$	$0.98 \pm 0.01$	$0.97 \pm 0.01$
GOLEM-NV int. with FGES	$0.47 \pm 0.47$	$1.00 \pm 0.00$	$0.99 \pm 0.00$
GOLEM-NV int. with Unif $[-0.01, 0.01]$	$123.13 \pm 2.90$	$0.50 \pm 0.01$	$0.04 \pm 0.00$
GOLEM-NV int. with Unif $[-0.05, 0.05]$	$119.20 \pm 2.75$	$0.51 \pm 0.01$	$0.04 \pm 0.01$
GOLEM-NV int. with Unif $[-0.1, 0.1]$	$112.77 \pm 3.11$	$0.51 \pm 0.01$	$0.04 \pm 0.01$

### C.2. Equal-Noise Variances Formulation with Random Perturbations

We investigate if random perturbations on the solution of GOLEM-EV improve the performance in the non-equal noise variances case. Specifically, we randomly perturb the solution of GOLEM-EV by adding a value sampled uniformly at random from  $[-\epsilon, \epsilon]$  to each entry of the estimated matrix. We then use the perturbed solution to initialize GOLEM-EV and run the method again, and repeat this procedure for 10 times. We consider  $\epsilon \in \{0.01, 0.05, 0.1\}$ , and denote the resulting methods as GOLEM-EV-Perturbed-0.01, GOLEM-EV-Perturbed-0.05, and GOLEM-EV-Perturbed-0.1, respectively. Here, we follow the setup (i.e., 50 variables, degree of 2, and  $10^6$  samples) in Section 3.2.

The results of GOLEM-EV-Perturbed, GOLEM-EV, and GOLEM-NV initialized with GOLEM-EV are shown in Figure 11. It is observed that random perturbations on GOLEM-EV does not improve the performance much. Specifically, when the noise ratio  $r \geq 4$ , GOLEM-EV with perturbations, similar to that without perturbations, performs much worse than GOLEM-NV initialized with GOLEM-EV. The reason is that when the noise ratio is high, there is model misspecification for GOLEM-EV because it assumes that the noise variances are equal; therefore, random perturbations on GOLEM-EV do not help. This further validates the importance of non-equal noise variances formulation (e.g., GOLEM-NV) to handle more general settings.

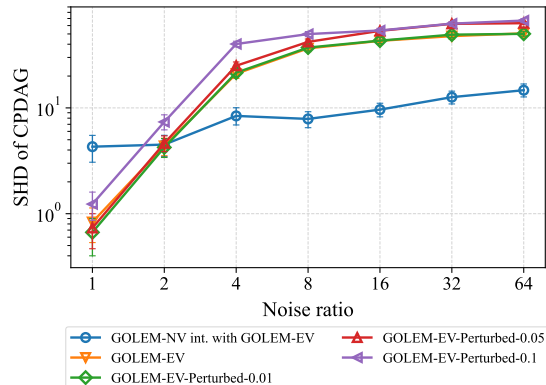


Figure 11: GOLEM-EV with perturbations.

### C.3. Alternative Form of Likelihood Function

In the derivation of GOLEM-NV [37, Appendix C.1], the resulting optimization problem in Eq. (1) involves only the matrix  $B$  because it profiles out the parameter  $\Omega$  that corresponds to the noise variances. Here, we consider the form of likelihood function without profiling out such parameter. Specifically, we consider the optimization problem

$$\min_{\substack{B \in \mathbb{R}^{d \times d} \\ \sigma_1, \dots, \sigma_d > 0}} \mathcal{L}_{\text{NV}}(B, \sigma_1, \dots, \sigma_d; \mathbf{X}) - \log |\det(I - B)| + \lambda_1 \|B\|_1 + \lambda_2 h(B), \quad (16)$$

where

$$\mathcal{L}_{\text{NV}}(B, \sigma_1, \dots, \sigma_d; \mathbf{X}) = \frac{1}{2} \sum_{i=1}^d \left( \log \sigma_i^2 + \frac{\|\mathbf{X}_{:,i} - \mathbf{X}B_{:,i}\|_2^2}{n\sigma_i^2} \right).$$

The derivation of the likelihood function above is available at Ng et al. [37, Appendix C.1]. We follow the setting described in Section 4.1, and report the empirical results in Figure 12. One observes that both forms of likelihood functions with and without profiling out the noise variances lead to a poor performance (when initialized with zero matrix) even when the sample size is large. For instance, with  $10^6$  samples, the SHDs of GOLEM with and without profiling are  $121.4 \pm 2.98$  and  $78.97 \pm 3.34$ , respectively, while PC and FGES have much lower SHDs at  $1.63 \pm 0.32$  and  $0.90 \pm 0.58$ , respectively. This suggests that both forms of likelihood functions may be susceptible to suboptimal local solutions possibly owing to nonconvexity.

Similar to the initialization scheme described in Section 3.2, we also considered using the solution of FGES to initialize the optimization problem (16), by computing the initial solution of weighted matrix  $B$  and noise variances  $\sigma_1, \dots, \sigma_d$ . A similar observation in Figure 4 is obtained (the empirical results are omitted here for brevity), i.e., such an initialization strategy improves the performance of structure learning, which indicates that the alternative formulation of GOLEM-NV in Eq. (16) also depends largely on the initial solution.

It is worth noting that we also conducted the same experiments above with NOTEARS by replacing its least squares objective with  $\mathcal{L}_{\text{NV}}(B, \sigma_1, \dots, \sigma_d; \mathbf{X})$  to handle the non-equal noise variances case, and have the same observations (the empirical results are omitted here for brevity).




Figure 12: Empirical results of different forms of likelihood function under different sample sizes. The number of variables is 50. Error bars represent the standard errors computed over 30 random repetitions.

### C.4. Analysis of Estimated Score and SHD

We investigate the connection between the score function of GOLEM-NV and the quality of the estimated structure. Specifically, we consider GOLEM-NV initialized with zero matrix as well as by solutions of GOLEM-EV and FGES. We also include the comparison with zero matrix and random DAG. For these methods, we compute the SHD of the estimated CPDAGs, and the score of GOLEM-NV (before thresholding) given in Eq. (1). We consider 50-node ER1 graphs with a noise ratio of 16.

The results are reported in Figure 13. One observes that GOLEM-NV initialized with GOLEM-EV and by FGES achieves a low score that is close to the score of the true DAG. At the same time, they also have a low SHD, indicating that their estimated structures are close to the true CPDAGs. However, this is not the case for GOLEM-NV initialized with zero matrix. Specifically, it also achieves a score that is very close to that of the true DAG, but its SHD is  $121.43 \pm 2.98$ . This appears to suggest that (1) a lower score does not necessarily lead to a low SHD-CPDAG, and (2) the nonconvex landscape of GOLEM-NV may contain suboptimal solutions whose scores are very close to the score of the global minimizer, but their corresponding structures may be very far from the true CPDAG. Thus, this demonstrates that nonconvexity may be a severe concern, at least in the setting considered here.

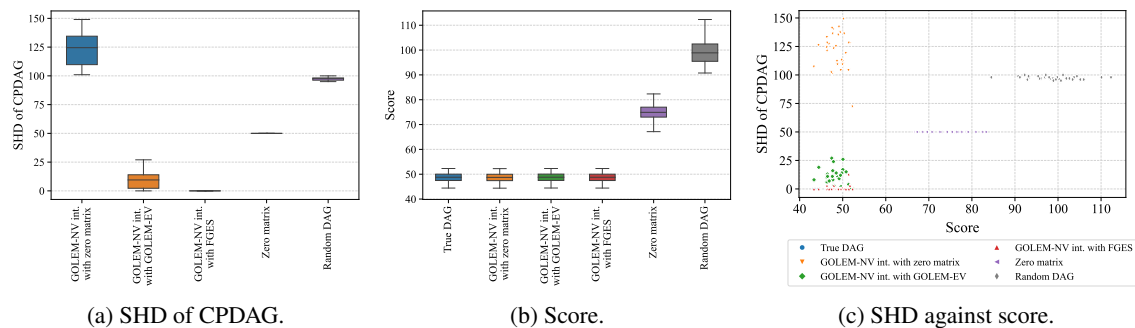


Figure 13: Visualization of estimated score and SHD for different methods over 30 random repetitions. Lower is better for both SHD and score.

## Appendix D. Comparison between SCAD and $\ell_0$ Penalties

Based on the analysis in Section 5.2, one may wonder how different specifications of the sparsity penalty considered above compare to  $\ell_0$  penalty. Since SCAD penalty and MCP perform similarly, we compare GOLEM-EV and NOTEARS-EV equipped with the SCAD penalty to discrete approaches, i.e., A\* and GDS, that employ  $\ell_0$  penalty. We consider 15-node ER2 graphs with varying sample sizes, and report the SHDs in Figure 14. Comparing NOTEARS-EV, A\*, and GDS that all use the least squares objective, one observes that A\*, not surprisingly, has the lowest SHDs, while GDS is on par with it. On the other hand, NOTEARS-EV has a much higher SHD than that of GDS. That is, NOTEARS-EV, even with the SCAD penalty, performs much worse than a greedy search method such as GDS, not to mention exact search method such as A\*. Such a performance gap conveys a cautionary message that continuous structure learning approaches may be inevitably susceptible to the nonconvexity issue that might be severe in practice. This is in contrast with the observation by Zheng et al. [72, Section 5.3] that the estimated solution by NOTEARS-EV is very close to the global minimizer despite its nonconvexity. A possible reason is that we apply a relatively small threshold 0.1, and as discussed in Section 5.1, NOTEARS-EV may require a larger threshold, e.g., 0.3, to achieve a decent performance, possibly owing to nonconvexity.




Figure 14: SCAD and  $\ell_0$  penalties.

## Appendix E. Supplementary Experiment Details

We provide supplementary experiment details for Sections 3.1, 3.2, 4.1, 4.2, 4.3, 5.1, and 5.2.

### E.1. Implementation Details and Evaluation Metrics

**Implementation Details.** For the structure learning methods considered, we use the default hyperparameters and official implementations from the authors, unless otherwise stated. As suggested in Section 5.1, using a relatively large threshold of 0.3 as in existing works may be harmful and remove many true edges, especially when the edge weights in the true weighted adjacency matrix are small. In this work, we consider a smaller threshold of 0.1 for continuous approaches. To ensure a fair comparison with NOTEARS-EV and GOLEM-EV, we also apply the same thresholding on (the least squares coefficients of) the estimated solutions of GDS and A\*. Furthermore, for NOTEARS-NV, DPDAG-NV, DAGMA-NV, and NOCURL-NV, we use the same hyperparameter as GOLEM-NV for sparsity penalty, i.e.,  $\lambda = 0.002$ , and replace their least squares objective  $\ell(B; \mathbf{X})$  with  $\mathcal{L}_{\text{NV}}(B; \mathbf{X})$ , which correspond to the likelihood of linear Gaussian DAGs assuming non-equal noise variances. For sparsity, SCAD penalty and MCP require an additional hyperparameter  $a$ , which we set to 3.7 following Fan and Li [18].

We provide further details for specific methods as follows:

- GOLEM [37]: We use L-BFGS implemented to solve the unconstrained optimization problem as it runs faster than Adam [28] and leads to a similar performance.

- NOTEARS [72], NOTEARS-MLP [73], and NOTEARS-ICA [71]: We use the quadratic penalty method instead of augmented Lagrangian method to solve the constrained optimization problem [5, 40], since Ng et al. [38] demonstrated that their performance is nearly identical, and that the former converges in a fewer number of iterations.
- DPDAG [9]: In the original work [9], DPDAG was combined with variational inference and MLPs, which leads to the VI-DP-DAG method for the nonlinear case. Here, we adapt it to learn linear DAGs (without the parts of variational inference and MLPs).
- A\* [67]: We adapt an implementation from the `causal-learn` package [74] available at the GitHub repository <https://github.com/cmu-phil/causal-learn>, and modify the BIC score [50] to least squares score [33]. We set the coefficient of  $\ell_0$  penalty to 0.01.
- GDS [43]: We use our own implementation in Python. A key difference with the original method [43] is that we use the least squares score [33] instead of likelihood score, both of which assume equal noise variances. We set the coefficient of  $\ell_0$  penalty to 0.01.
- PC [54] and FGES [45]: We use the implementation from the `py-causal` package available at the GitHub repository <https://github.com/bd2kccd/py-causal>, which is a Python wrapper of the TETRAD project [49]. We adopt the Fisher Z test and BIC score [50] for PC and FGES, respectively.

**Evaluation Metrics.** For the linear Gaussian setting with equal noise variances formulations, as well as linear non-Gaussian and nonlinear settings, we report the SHD, F1 score, and recall of the estimated DAG. For the linear Gaussian setting with non-equal noise variances formulations, we report the SHD of estimated CPDAG, as well as the F1 score of estimated skeleton and directed edges. The standard errors over 30 random repetitions are also provided.


## E.2. Discussion of Search Strategies

As discussed in Section 2.1, NOTEARS [72] solves a constrained optimization problem with a hard DAG constraint by using augmented Lagrangian or quadratic penalty method [5, 40], while GOLEM [37] solves an unconstrained optimization with a soft DAG constraint. These essentially represent different strategies to search for a DAG. Similar to NOTEARS, DAGMA [3] solve a constrained optimization problem, but adopts a procedure similar to the barrier method [40] with a log-determinant DAG constraint. Instead of enforcing acyclicity via a constrained optimization problem such as NOTEARS and DAGMA, another line of approaches directly search in the space of DAGs. Specifically, NOCURL [66] developed an algebraic representation of DAGs based on graph Hodge theory [25, 2] that directly outputs weighted adjacency matrix of a DAG, while DPDAG [9] adopts a differentiable DAG sampling procedure that (1) samples a linear ordering of the variables using Gumbel-Sinkhorn [35] or Gumbel-Top-k [31] reparametrizations, and (2) samples a weighted adjacency matrix consistent with the ordering. Note that DPDAG was combined with variational inference and MLPs, which leads to the VI-DP-DAG method for the nonlinear case; here, we adapt it to learn linear DAGs (without the parts of variational inference and MLPs). These several approaches represent different strategies to traverse the search space for estimating a DAG with continuous optimization.


## Appendix F. Supplementary Experiment Results

We provide supplementary experiment results for Sections 3.1, 3.2, 4.1, 4.2, 4.3, 5.1, and 5.2.


**F.1. With Equal Noise Variances Formulation**




(a) 15 variables with 100 samples.



(b) 15 variables with  $10^6$  samples.




(c) 50 variables with 100 samples.




(d) 50 variables with  $10^6$  samples.

Figure 15: Empirical results of structure learning methods assuming equal noise variances under different noise ratios. Error bars represent the standard errors computed over 30 random repetitions.

### F.2. With Non-Equal Noise Variances Formulation



(a) Without data standardization.



(b) With data standardization.

Figure 16: Empirical results of structure learning methods under different noise ratios. The number of variables is 50 and the sample size is  $10^6$ . Error bars represent the standard errors computed over 30 random repetitions. Here, “int.” stands for “initialized”.

### F.3. Search Strategies


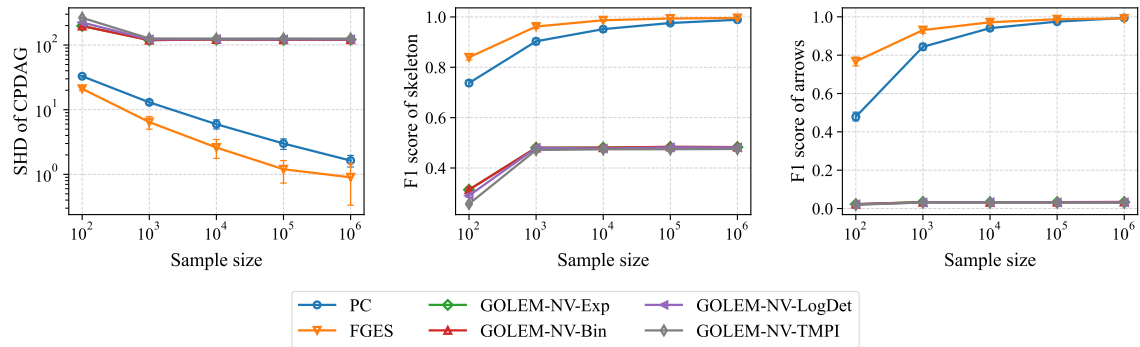
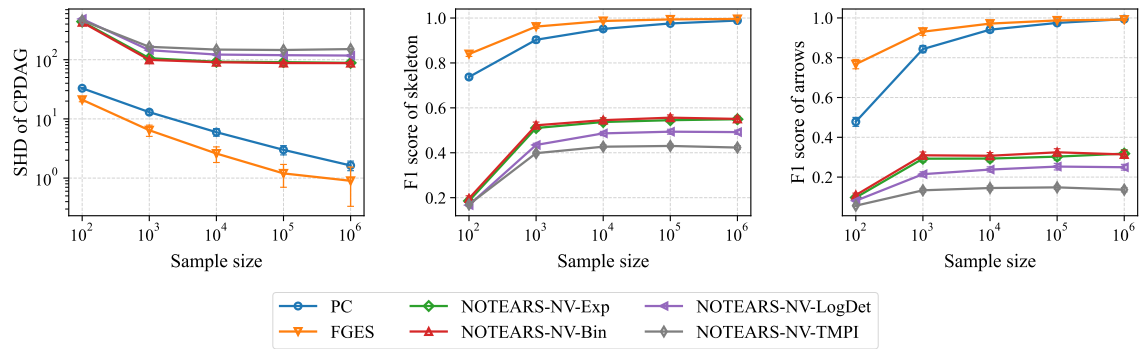


Figure 17: Empirical results of different search strategies under different sample sizes. The number of variables is 50. Error bars represent the standard errors computed over 30 random repetitions.

**F.4. DAG Constraints**



(a) GOLEM-NV with different DAG constraints.



(b) NOTEARS-NV with different DAG constraints.

Figure 18: Empirical results of different DAG constraints under different sample sizes. The number of variables is 50. Error bars represent the standard errors computed over 30 random repetitions.



**F.5. Other Settings**




Figure 19: Empirical results of NOTEARS-ICA under different sample sizes. The number of variables is 15. Error bars represent the standard errors computed over 30 random repetitions.


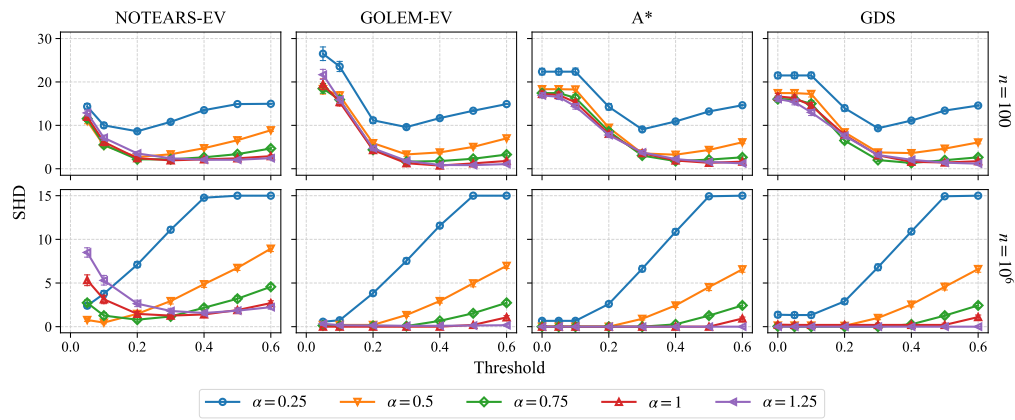
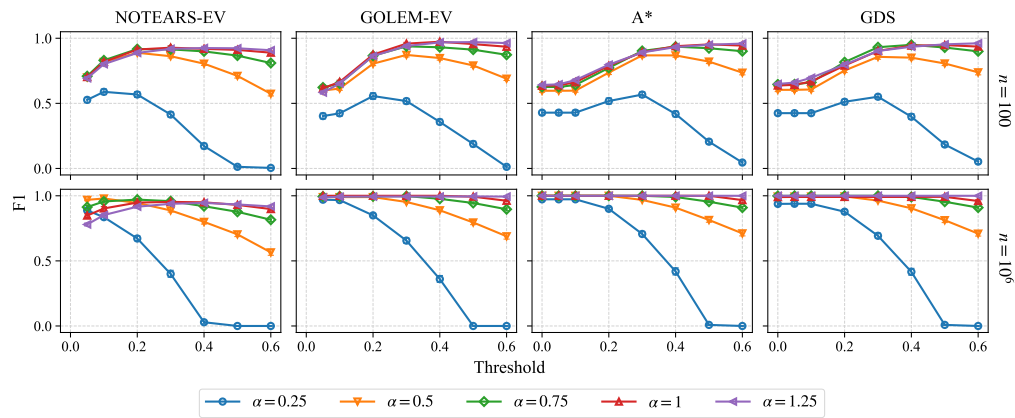


Figure 20: Empirical results of NOTEARS-MLP under different sample sizes. The number of variables is 15. Error bars represent the standard errors computed over 30 random repetitions.

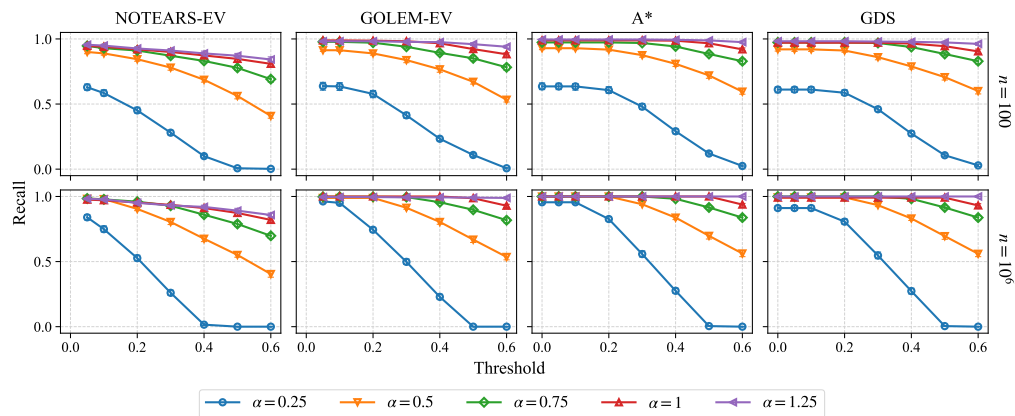
F.6. Thresholding



(a) SHD.



(b) F1 score.



(c) Recall.

Figure 21: Empirical results of different thresholds under different weight scales. The number of variables is 15. Error bars represent the standard errors computed over 30 random repetitions.

E.7. Sparsity Penalty

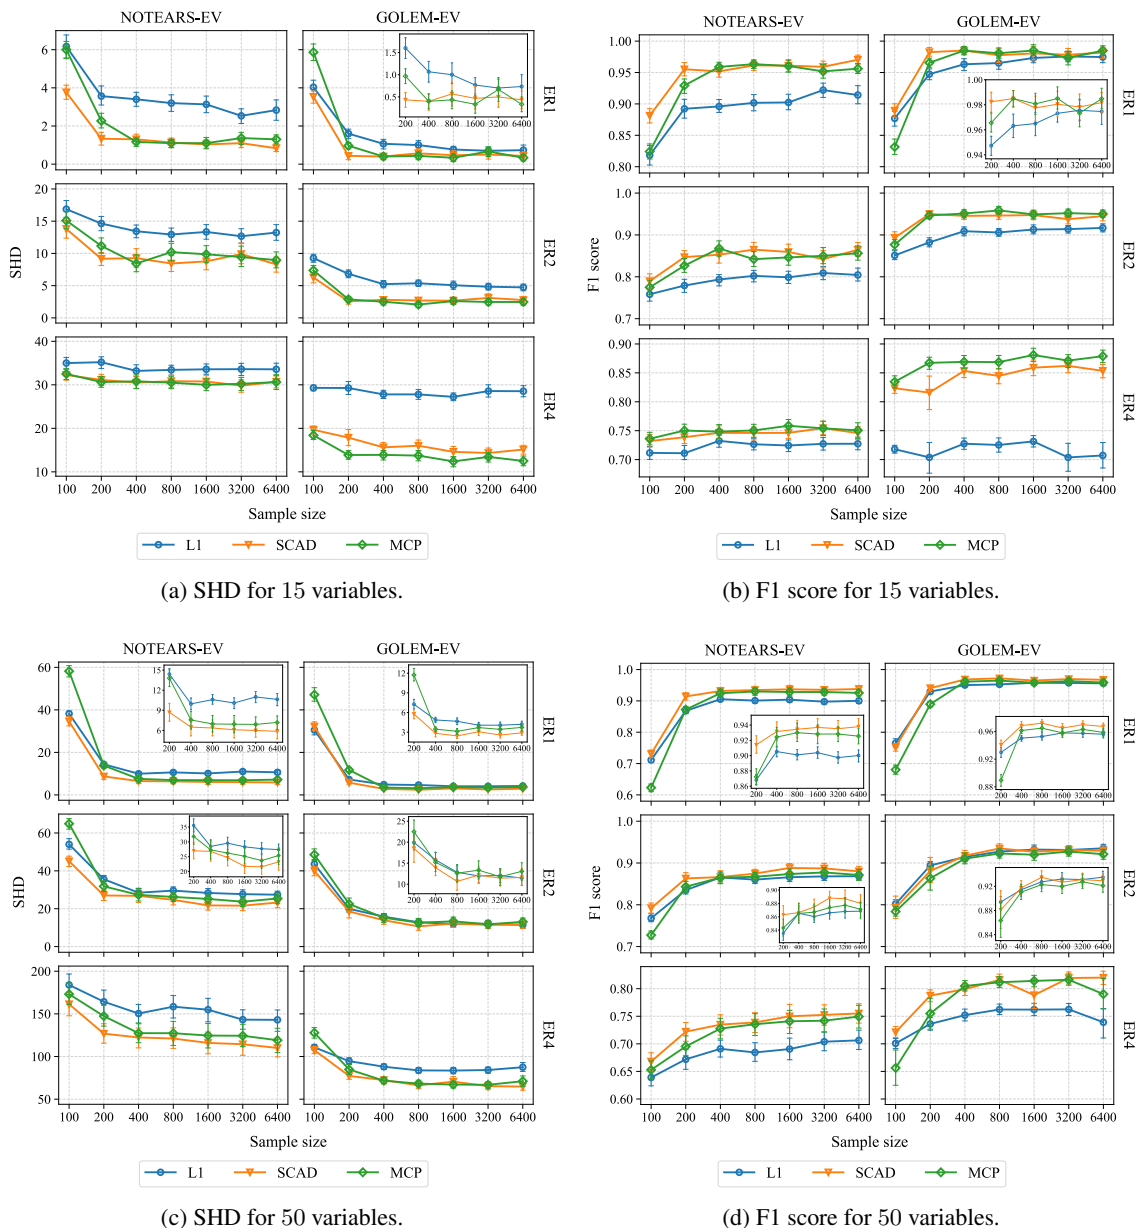


Figure 22: Empirical results of different sparsity penalties under different sample sizes. Error bars represent the standard errors computed over 30 random repetitions.