

Non-negative DAG Learning from Time-Series Data

Samuel Rey* and Gonzalo Mateos†

*Dept. of Signal Theory and Communications, Rey Juan Carlos University, Madrid, Spain

Email: samuel.rey.escudero@urjc.es.

†Dept. of Electrical and Computer Engineering, University of Rochester, NY, USA.

Email: gmateosb@ur.rochester.edu

Abstract—This work aims to learn the directed acyclic graph (DAG) that captures the instantaneous dependencies underlying a multivariate time series. The observed data follow a linear structural vector autoregressive model (SVARM) with both instantaneous and time-lagged dependencies, where the instantaneous structure is modeled by a DAG to reflect potential causal relationships. While recent continuous relaxation approaches impose acyclicity through smooth constraint functions involving powers of the adjacency matrix, they lead to non-convex optimization problems that are challenging to solve. In contrast, we assume that the underlying DAG has only non-negative edge weights, and leverage this additional structure to impose acyclicity via a convex constraint. This enables us to cast the problem of non-negative DAG recovery from multivariate time-series data as a convex optimization problem in abstract form, which we solve using the method of multipliers. Crucially, the convex formulation guarantees global optimality of the solution. Finally, we assess the performance of the proposed method on synthetic time-series data, where it outperforms existing alternatives.

Index Terms—DAG structure learning, dynamic Bayesian networks, causal discovery, graph learning, graph signal processing

I. INTRODUCTION

Understanding dependencies in multivariate data requires representations that go beyond simple correlation matrices. Directed acyclic graphs (DAGs) have emerged as the canonical representation for modeling complex systems where directionality plays a key role [1], enabling probabilistic reasoning in Bayesian networks and causal inference tasks [2]–[5]. Consequently, DAGs are ubiquitous in applications requiring the identification of driving forces, ranging from machine learning [3], [6], [7] and signal processing [8], [9], to biology [10] and genetics [11]. Despite their widespread adoption, the underlying DAG structure is rarely known *a priori*. This necessitates the development of structure learning methods capable of recovering the graph from observational data.

Inferring the graph topology from nodal observations constitutes a central problem, connecting the statistical properties of data to the underlying graph structure [12], [13]. For undirected graphs, noteworthy approaches include smoothness priors [14]–[16], Gaussian graphical models [17]–[19], and graph stationary models [20]–[24], among others. However, when the focus shifts to DAGs, structural equation models

(SEMs) become the prevailing formalism [25]–[29]. Strictly enforcing the acyclicity constraint renders DAG structure learning a combinatorial, NP-hard endeavor [30], [31]. This complexity is further exacerbated in time-series settings, as temporal lags drastically increase the dimensionality of the search space, rendering traditional discrete search methods computationally prohibitive [32], [33].

Recent advancements have revolutionized the field of DAG structure learning, with impact to causal discovery. The seminal work in [25] reformulated the combinatorial problem of learning the DAG structure from nodal observations into a continuous optimization task, enabling efficient exploration of the DAG space. This breakthrough sparked a wave of variants exploring alternative acyclicity constraints and score functions [26], [27], [29], [34]. In the time-series domain, dynamic Bayesian networks (DBNs) have long served as the standard for modeling temporal processes, particularly over discrete state spaces [33], [35]. Building on structural vector autoregressive models (SVARMs) [13], [36], recent approaches like DYNOTEARS [32], [37] extended the continuous acyclicity constraint to the dynamic setting. Further extensions incorporating sparsity-promoting score functions have also been proposed [38]. However, a fundamental limitation persists: the resulting optimization landscape is inherently non-convex. Consequently, state-of-the-art methods often rely on heuristics or specific initializations [27], frequently converging to local minima rather than the global optimum [39].

In this work, we surmount these optimization challenges by harnessing recent advances in non-negative DAG structure learning [40]. Specifically, we tackle the problem of recovering DAGs with non-negative edge weights from time-series data governed by linear SVARMs. Crucially, the non-negativity constraint enables characterizing acyclicity via a *convex* function. We formulate the inference task as an *abstract convex optimization problem* and propose an iterative algorithm based on the method of multipliers to solve it. Unlike previous approaches, our formulation warrants the recovery of the global minimizer, eliminating the sensitivity to initialization. The merits of this approach are demonstrated in our numerical evaluation, where the proposed method outperforms state-of-the-art alternatives on synthetic datasets.

II. DAG STRUCTURE LEARNING FROM TIME SERIES

Let $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ denote a directed graph with a set of N nodes \mathcal{V} and edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. The graph \mathcal{D} is a DAG if

This work was supported in part by the Spanish AEI under Grants PID2019-105032GB-I00, TED2021-130347B-I00 and PID2022-136887NB-I00 funded by MCIN/AEI/10.13039/501100011033, and by the Community of Madrid via the Ellis Madrid Unit and grants URJC/CAM F1180 and TEC-2024/COM-89.

it contains no directed cycles. The connectivity of the DAG is captured by the weighted adjacency matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$, where $W_{ij} \neq 0$ if and only if $(i, j) \in \mathcal{E}$. Hence, the entry W_{ij} denotes the weight of the directed link from node i to node j . Finally, we define a graph signal on \mathcal{D} as a vector $\mathbf{x} \in \mathbb{R}^N$, where the i -th entry x_i denotes the signal value at node i .

We consider the problem of learning the causal structure from a time series of graph signals $\{\mathbf{x}_t\}_{t=1}^T$, where $\mathbf{x}_t \in \mathbb{R}^N$ represents the graph signal at time index t . We assume the data generation process adheres to a SVARM of order P [36]. Formally, for each time instant t , the signal \mathbf{x}_t is modeled as

$$\mathbf{x}_t = \mathbf{W}^\top \mathbf{x}_t + \sum_{p=1}^P \mathbf{A}_p^\top \mathbf{x}_{t-p} + \mathbf{z}_t, \quad (1)$$

where the matrix \mathbf{W} encodes the instantaneous dependencies, and the set of matrices $\{\mathbf{A}_p\}_{p=1}^P$, with $\mathbf{A}_p \in \mathbb{R}^{N \times N}$, captures the time-lagged influences from the p -th previous step. The term \mathbf{z}_t represents a vector of exogenous noises, assumed to be an i.i.d. random vector with covariance matrix $\Sigma_{\mathbf{z}} = \sigma^2 \mathbf{I}$.

We can derive a compact matrix formulation for the SVARM defined in (1). To that end, we collect the observed signals in the matrix $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_T] \in \mathbb{R}^{N \times T}$. Similarly, let $\mathbf{Y}_p \in \mathbb{R}^{N \times T}$ denote the time-lagged data matrix for lag p , constructed by shifting the columns of \mathbf{X} appropriately. By stacking the lagged matrices vertically as $\mathbf{Y} := [\mathbf{Y}_1^\top, \dots, \mathbf{Y}_P^\top]^\top \in \mathbb{R}^{NP \times T}$ and the corresponding coefficients as $\mathbf{A} := [\mathbf{A}_1^\top, \dots, \mathbf{A}_P^\top]^\top \in \mathbb{R}^{NP \times N}$, the signal model takes the form

$$\mathbf{X} = \mathbf{W}^\top \mathbf{X} + \mathbf{A}^\top \mathbf{Y} + \mathbf{Z}, \quad (2)$$

where $\mathbf{Z} \in \mathbb{R}^{N \times T}$ collects the noise vectors. Note that the standard linear SEM is recovered as a particular case of (2) when the autoregressive order is $P = 0$ (i.e., $\mathbf{A} \equiv \mathbf{0}$). As stated in [32], the SVARM is identifiable under relatively mild assumptions. In particular, \mathbf{W} is identifiable whenever the exogenous noise vectors \mathbf{z}_t are either non-Gaussian or standard Gaussian $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. Identifiability of \mathbf{A} follows from standard results on vector autoregressive models.

With the previous definitions in place, the joint inference of the instantaneous DAG structure \mathbf{W} and the time-lagged dependencies \mathbf{A} from the observed time series \mathbf{X} can be cast as the following optimization problem

$$\min_{\mathbf{W}, \mathbf{A}} F(\mathbf{W}, \mathbf{A}; \mathbf{X}) \quad \text{subject to} \quad \mathbf{W} \in \mathbb{D}, \quad (3)$$

where $F(\cdot)$ denotes a data-dependent score function quantifying the goodness of fit, and \mathbb{D} represents the set of adjacency matrices corresponding to a DAG. For a valid causal interpretation in the absence of latent variables, \mathbf{W} must be a DAG. Conversely, the lagged coefficient matrices collected in \mathbf{A} are not required to be acyclic, as they encapsulate how past system states influence the current state at time t .

The optimization problem in (3) is challenging to solve due to the non-convex and combinatorial nature of the constraint $\mathbf{W} \in \mathbb{D}$. Recent advances propose addressing this issue by replacing the discrete set \mathbb{D} with a continuous *acyclicity*

constraint of the form $h(\mathbf{W}) = 0$, where $h : \mathbb{R}^{N \times N} \mapsto \mathbb{R}$ is a differentiable function whose zero level set coincides with the space of DAGs \mathbb{D} . This approach was pioneered in [25] and subsequently followed by several works proposing alternative continuous acyclicity functions [26], [27]. Replacing the combinatorial constraint with $h(\mathbf{W}) = 0$ constitutes a significant advancement, enabling the use of standard continuous optimization methods. However, existing acyclicity functions $h(\mathbf{W})$ are generally non-convex. Consequently, the resulting optimization landscape remains fraught with local minima, hindering the recovery of the true DAG structure. To overcome this limitation, we henceforth assume \mathbf{W} has *non-negative weights*. This allows us to leverage a convex acyclicity function and develop an algorithm to obtain the *global minimizer* of an optimization problem equivalent to (3).

III. LEARNING NON-NEGATIVE DAGS FROM SVAR DATA

We approach the DAG structure learning problem by enforcing a non-negativity constraint on the entries of \mathbf{W} . This assumption is practically relevant, extending to binary DAGs and other pragmatic settings [8]. Crucially, this additional structure simplifies the optimization landscape of (3), enabling the use of a convex acyclicity function to recover the global minimizer of the score function. In the sequel, we detail the proposed optimization problem and the algorithmic solution.

Central to modern DAG learning methods is the use of smooth constraints to enforce the absence of cycles. As discussed in [25], an effective acyclicity function $h(\mathbf{W})$ must be differentiable, computationally efficient, and satisfy $h(\mathbf{W}) = 0$ if and only if $\mathbf{W} \in \mathbb{D}$. In this work, we leverage recent advancements introduced in [40], which demonstrate that imposing non-negativity on \mathbf{W} allows for a *convex* characterization of acyclicity. The details about this acyclicity function are provided in the following proposition.

Proposition 1 ([40]) *Let $\mathbf{W} \in \mathbb{R}_+^{N \times N}$ be a non-negative matrix with spectral radius bounded by $\rho(\mathbf{W}) < s$ for some $s > 0$. Consider the function*

$$h(\mathbf{W}) := N \log(s) - \log \det(s\mathbf{I} - \mathbf{W}). \quad (4)$$

Then, $h(\mathbf{W})$ is convex over its domain, differentiable with gradient $\nabla h(\mathbf{W}) = (s\mathbf{I} - \mathbf{W})^{-\top}$, and satisfies $h(\mathbf{W}) \geq 0$. Furthermore, $h(\mathbf{W}) = 0$ if and only if $\mathbf{W} \in \mathbb{D}$.

When \mathbf{X} adheres to a SVARM, a standard choice for the score function is the least-squares data-fidelity term augmented with ℓ_1 -norm regularization to promote sparsity on both \mathbf{W} and \mathbf{A} . Leveraging this convex score function, the continuous acyclicity constraint, and enforcing entry-wise non-negativity on both \mathbf{W} and \mathbf{A} , the DAG structure learning problem can be alternatively formulated as

$$\begin{aligned} \{\hat{\mathbf{W}}, \hat{\mathbf{A}}\} = & \arg \min_{\mathbf{W}, \mathbf{A}} \frac{1}{2T} \|\mathbf{X} - \mathbf{W}^\top \mathbf{X} - \mathbf{A}^\top \mathbf{Y}\|_F^2 \\ & + \lambda_{\mathbf{W}} \sum_{i,j} W_{ij} + \lambda_{\mathbf{A}} \sum_{k,l} A_{kl} \\ \text{subject to} & \mathbf{W} \geq 0, \mathbf{A} \geq 0, h(\mathbf{W}) = 0. \end{aligned} \quad (5)$$

Here, the regularization parameters $\lambda_{\mathbf{W}}, \lambda_{\mathbf{A}} > 0$ control the trade-off between data fidelity and sparsity for the instantaneous and time-lagged coefficients, respectively. The term $h(\mathbf{W})$ denotes the convex acyclicity function defined in (4). Note that, due to the non-negativity constraints, the standard ℓ_1 penalties $\|\mathbf{W}\|_1$ and $\|\mathbf{A}\|_1$ simplify to the linear sums of the entries appearing in the objective function. It is worth noting that while we assume both \mathbf{W} and \mathbf{A} to be non-negative for consistency, only the non-negativity of \mathbf{W} is required for the convexity of the acyclicity function. Therefore, the proposed framework can readily accommodate signed lagged coefficients by simply removing the constraint $\mathbf{A} \geq 0$ and replacing the linear sum penalty with the standard ℓ_1 -norm.

Ensuring the acyclicity of \mathbf{W} via a convex function such as h confers significant advantages. Foremost, the convexity of h renders (5) an *abstract convex optimization problem* [41], thereby guaranteeing the recovery of the global minimizer. To see this, recall from Proposition 1 that $h(\mathbf{W}) \geq 0$, thus the equality constraint $h(\mathbf{W}) = 0$ is mathematically equivalent to the inequality constraint $h(\mathbf{W}) \leq 0$. Since the sublevel sets of a convex function form a convex set, the feasible region of (5) is convex. In turn, the ability to recover the global minimum, combined with the identifiability of the SVARM in (2) under mild assumptions, paves the way for establishing strong statistical guarantees for the estimated DAG, a promising avenue for future research.

A. DAG structure learning via the method of multipliers

Estimating the instantaneous and time-lagged dependencies encoded in \mathbf{W} and \mathbf{A} requires solving the constrained optimization problem introduced in (5). To that end, we rely on the method of multipliers [42, Ch. 4.2], an iterative approach based on the augmented Lagrangian to tackle equality-constrained problems, which offers convergence guarantees.

The augmented Lagrangian of (5) is given by

$$L_c(\mathbf{W}, \mathbf{A}, \alpha) = \frac{1}{2T} \|\mathbf{X} - \mathbf{W}^\top \mathbf{X} - \mathbf{A}^\top \mathbf{Y}\|_F^2 + \lambda_{\mathbf{W}} \sum_{i,j} W_{ij} + \lambda_{\mathbf{A}} \sum_{k,l} A_{kl} + \alpha h(\mathbf{W}) + \frac{c}{2} h_{\text{det}}^2(\mathbf{W}), \quad (6)$$

where $\alpha \in \mathbb{R}_+$ is the Lagrange multiplier associated with the acyclicity constraint, and $c \in \mathbb{R}_+$ serves as the quadratic penalty parameter. Observe that the augmented Lagrangian L_c is convex with respect to the primal variables \mathbf{W} and \mathbf{A} . This follows from Proposition 1, which guarantees that $h(\mathbf{W}) \geq 0$, and hence, the term $h^2(\mathbf{W})$ represents the composition of a convex function with a convex, non-decreasing function [41]. Finally, note that the non-negativity constraints $\mathbf{W} \geq 0$ and $\mathbf{A} \geq 0$ are not explicitly included in L_c since they can be enforced through a simple projection.

Given the augmented Lagrangian formulation, to solve (5) we adopt an iterative approach where we perform the following sequence of steps for $k = 1, \dots, K$ iterations.

Step 1. We update the primal variables $\mathbf{W}^{(k+1)}$ and $\mathbf{A}^{(k+1)}$ by minimizing

$$\{\mathbf{W}^{(k+1)}, \mathbf{A}^{(k+1)}\} = \arg \min_{\mathbf{W} \geq 0, \mathbf{A} \geq 0} L_{c^{(k)}}(\mathbf{W}, \mathbf{A}, \alpha^{(k)}). \quad (7)$$

Due to the convexity of $L_{c^{(k)}}$, we can recover the global minimizer using standard convex optimization methods, such as (accelerated) projected gradient descent.

Step 2. The update of the Lagrange multiplier $\alpha^{(k+1)}$ depends on the current constraint violation, given by

$$\alpha^{(k+1)} = \alpha^{(k)} + c^{(k)} h(\mathbf{W}^{(k+1)}). \quad (8)$$

This update is equivalent to a gradient ascent step, as the constraint violation corresponds to the gradient of $L_{c^{(k)}}$ with respect to α .

Step 3. The penalty parameter $c^{(k)}$ is progressively increased to ensure that the constraint $h(\mathbf{W}) = 0$ is satisfied as $K \rightarrow \infty$. A typical update scheme is given by

$$c^{(k+1)} = \begin{cases} \beta c^{(k)} & \text{if } h(\mathbf{W}^{(k+1)}) > \gamma h(\mathbf{W}^{(k)}) \\ c^{(k)} & \text{otherwise,} \end{cases} \quad (9)$$

where $0 < \gamma < 1$ and $\beta > 1$ are user-defined constants. Intuitively, $c^{(k)}$ is increased only if the constraint violation is not decreased by a factor of γ .

Upon completion of the sequence of iterations, the estimated DAG structure and time-lagged dependencies are given by $\hat{\mathbf{W}} = \mathbf{W}^{(K)}$ and $\hat{\mathbf{A}} = \mathbf{A}^{(K)}$, respectively. The convexity of the augmented Lagrangian guarantees that the tuple $\{\mathbf{W}^{(k)}, \mathbf{A}^{(k)}\}$ corresponds to the global minimum of the subproblem in (7) at each iteration k . Therefore, as established in [42, Prop. 4.2.1], every limit point of the sequence $\{\mathbf{W}^{(k)}, \mathbf{A}^{(k)}\}$ constitutes a global minimum of the constrained problem in (5).

While the method of multipliers was also employed in [25], subsequent work identified a significant drawback in their formulation [26]. Specifically, every DAG constitutes a stationary point of the acyclicity constraint used in [25], implying that $\nabla h_{\text{notears}}(\mathbf{W}) = \mathbf{0}$ for any $\mathbf{W} \in \mathbb{D}$. This implies that a feasible solution can satisfy the KKT optimality conditions only if it is also a stationary point of the unconstrained objective, posing algorithmic challenges for finding the true constrained minimizer. In contrast, observe that DAGs are not stationary points of our acyclicity function (i.e., $\nabla h(\mathbf{W}) \neq \mathbf{0}$). Consequently, our method avoids these algorithmic pitfalls.

IV. NUMERICAL EXPERIMENTS

Here, we evaluate the performance of the proposed method across various scenarios. The code for the algorithm and all implementation details is publicly available on GitHub¹.

We assess the estimation accuracy using two metrics. We measure the normalized Frobenius error (NFE) calculated as

$$\text{NFE}(\hat{\mathbf{W}}, \mathbf{W}^*) = \frac{\|\mathbf{W}^* - \hat{\mathbf{W}}\|_F^2}{\|\mathbf{W}^*\|_F^2}, \quad (10)$$

¹https://github.com/reysam93/cvx_dyn_dag




Fig. 1: Performance evaluation of the proposed method (CVX-DYN) compared to DYNOTEARS (we use DYNOTEARS in the legends for clarity) across different scenarios. (a) NFE of the estimated weights $\hat{\mathbf{W}}$ and $\hat{\mathbf{A}}$ as a function of the time-series length T . (b) F1-score of $\hat{\mathbf{W}}$ and $\hat{\mathbf{A}}$ as the number of nodes N increases. (c) NFE for varying autoregressive orders P .

as well as the F1-score, defined as the harmonic mean of precision and recall computed on the support of the adjacency matrices, to evaluate structure recovery. We report these metrics for both the instantaneous effects $\hat{\mathbf{W}}$ and the lagged effects $\hat{\mathbf{A}}$. In all experiments, we report the median values along with the 25th and 75th percentiles obtained from 50 independent realizations. We compare our approach with DYNOTEARS [32], a closely related alternative that also assumes a SVARM and estimates \mathbf{W} and \mathbf{A} simultaneously, but leverages a continuous non-convex acyclicity constraint.

Unless otherwise stated, we simulate Erdős-Rényi (ER) DAGs with $N = 50$ nodes and an average degree of 4, generated by sampling lower-triangular adjacency matrices and randomly permuting their rows and columns. The $P = 2$ time-lagged matrices $\{\mathbf{A}_p\}_{p=1}^P$ are sampled as ER graphs with an average degree of 1. Edge weights for \mathbf{W} and \mathbf{A} are sampled uniformly from the interval $[0.1, 0.5]$. Following [32], we also apply an exponential decay $e^{-1.5p}$ to \mathbf{A}_p to promote the stability of the SVARM process. Finally, time series data is generated according to (2) with length $T = 5000$, driven by i.i.d. noise \mathbf{Z} sampled from a standard Gaussian distribution.

Test case 1 - Number of samples. The first experiment evaluates the estimation accuracy of the proposed method (CVX-DYN) relative to DYNOTEARS as the sample size T increases. As illustrated in Fig. 1 (a), leveraging a convex acyclicity constraint consistently yields superior performance compared to the baseline. Notably, while the error for DYNOTEARS saturates, the error associated with our method exhibits a continuous decay, converging towards zero as the number of samples increases. This empirical behavior supports the intuition that, since the SVARM is identifiable and our convex formulation recovers the global minimum, the method should be capable of recovering the true underlying structure given sufficient data. Additionally, we observe that the error for the lagged coefficients $\hat{\mathbf{A}}$ is consistently higher than for the instantaneous weights $\hat{\mathbf{W}}$. This discrepancy can be attributed to the data generation process, where the exponential decay

applied to \mathbf{A}_p reduces the influence of lagged dependencies, making them inherently more challenging to estimate.

Test case 2 - Number of nodes. Next, we test the scalability of the methods by analyzing the F1-score for $\hat{\mathbf{W}}$ and $\hat{\mathbf{A}}$ as the network size N increases (see Fig. 1 (b)). The results demonstrate that CVX-DYN consistently outperforms the baseline, further validating the advantages of our convex approach. Remarkably, our method achieves a perfect F1-score of 1 for the instantaneous matrix $\hat{\mathbf{W}}$, indicating robust support recovery even in high-dimensional regimes where the number of nodes ($N = 1000$) becomes comparable to the sample size ($T = 5000$). Regarding the lagged dependencies $\hat{\mathbf{A}}$, the performance follows the trend observed in the previous test case. While our method still outperforms the baseline, the recovery is inherently more challenging due to the weaker influence of the lagged components.

Test case 3 - Time-lagged dependencies. Finally, in Fig. 1 (c) we assess how the autoregressive order P impacts the performance of the DAG learning alternatives. The results illustrate that the error of the estimated $\hat{\mathbf{W}}$ remains stable for both methods (CVX-DYN and DYNOTEARS), irrespective of the value of P , which is likely partially due to the exponential decay strategy followed to generate the data. In contrast, we observe that the error for $\hat{\mathbf{A}}$ increases monotonically with the number of lags, underscoring that accounting for longer temporal dependencies results in a more challenging inference problem. In addition, Fig. 1 (c) demonstrates that CVX-DYN attains a consistently low estimation error for both $\hat{\mathbf{W}}$ and $\hat{\mathbf{A}}$, which is approximately one order of magnitude smaller than the error attained by DYNOTEARS, emphasizing the merits of our proposed convex approach.

V. CONCLUSIONS

This paper addressed the prominent task of jointly learning the DAG structure, which models instantaneous dependencies, and the matrix \mathbf{A} , which encodes time-lagged dependencies, from nodal observations governed by a SVARM. By focusing

on the estimation of a non-negative matrix $\hat{\mathbf{W}}$, we formulated the DAG learning task as a continuous optimization problem in a convex abstract form. Harnessing the specific structure of our formulation, we proposed an iterative algorithm based on the method of multipliers. To the best of our knowledge, this is the first method for learning DAGs from time series that guarantees the recovery of the global minimizer. We argue that our approach avoids algorithmic issues present in competing DAG learning schemes also based on augmented Lagrangian methods. Finally, we corroborate the algorithmic advantages of our proposal through reproducible numerical experiments on synthetic data, where it consistently outperforms state-of-the-art methods.

REFERENCES

- [1] A. G. Marques, S. Segarra, and G. Mateos, "Signal processing on directed graphs: The role of edge directionality when processing and learning from network data," *IEEE Signal Process. Mag.*, vol. 37, no. 6, pp. 99–116, 2020.
- [2] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search*. MIT Press, 2001.
- [3] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.
- [4] J. Peters, D. Janzing, and B. Schölkopf, *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017.
- [5] L. Yao, Z. Chu, S. Li, Y. Li, J. Gao, and A. Zhang, "A survey on causal inference," *ACM Trans. Knowl. Discovery Data*, vol. 15, no. 5, pp. 1–46, 2021.
- [6] Y. Yu, J. Chen, T. Gao, and M. Yu, "DAG-GNN: DAG structure learning with graph neural networks," in *Intl. Conf. Mach. Learn. (ICML)*. PMLR, 2019, pp. 7154–7163.
- [7] S. Rey, H. Ajourlou, and G. Mateos, "Convolutional learning on directed acyclic graphs," *arXiv preprint arXiv:2405.03056*, 2024.
- [8] B. Seifert, C. Wendler, and M. Püschel, "Causal Fourier analysis on directed acyclic graphs and posets," *IEEE Trans. Signal Process.*, vol. 71, pp. 3805–3820, 2023.
- [9] P. Misiakos, V. Mihal, and M. Püschel, "Learning signals and graphs from time-series graph data with few causes," in *IEEE Intl. Conf. Acoustics, Speech Signal Process. (ICASSP)*, 2024, pp. 9681–9685.
- [10] P. J. F. Lucas, L. C. Van der Gaag, and A. Abu-Hanna, "Bayesian networks in biomedicine and health-care," *Artif Intell Med*, vol. 30, no. 3, pp. 201–214, 2004.
- [11] B. Zhang, C. Gaiteri, L.-G. Bodea, Z. Wang, J. McElwee, A. A. Podtelezhnikov, C. Zhang, T. Xie, L. Tran, R. Dobrin, et al., "Integrated systems approach identifies genetic nodes and networks in late-onset Alzheimer's disease," *Cell*, vol. 153, no. 3, pp. 707–720, 2013.
- [12] G. Mateos, S. Segarra, A. G. Marques, and A. Ribeiro, "Connecting the dots: Identifying network structure via graph signal processing," *IEEE Signal Process. Mag.*, vol. 36, no. 3, pp. 16–43, 2019.
- [13] G. B. Giannakis, Y. Shen, and G. V. Karanikolas, "Topology identification and learning over graphs: Accounting for nonlinearities and dynamics," *Proc. IEEE*, vol. 106, no. 5, pp. 787–807, 2018.
- [14] V. Kalofolias, "How to learn a graph from smooth signals," in *Intl. Conf. Artif. Intel. Statist. (AISTATS)*, 2016, pp. 920–929.
- [15] X. Dong, D. Thanou, P. Frossard, and P. Vandergheynst, "Learning Laplacian matrix in smooth graph signal representations," *IEEE Trans. Signal Process.*, vol. 64, no. 23, pp. 6160–6173, 2016.
- [16] S. S. Saboksayr and G. Mateos, "Accelerated graph learning from smooth signals," *IEEE Signal Process. Lett.*, vol. 28, pp. 2192–2196, 2021.
- [17] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.
- [18] H. E. Egilmez, E. Pavez, and A. Ortega, "Graph learning from data under Laplacian and structural constraints," *IEEE J. Select. Topics Signal Process.*, vol. 11, no. 6, pp. 825–841, 2017.
- [19] S. Rey, T. M. Roddenberry, S. Segarra, and A. G. Marques, "Enhanced graph-learning schemes driven by similar distributions of motifs," *IEEE Trans. Signal Process.*, vol. 71, pp. 3014–3027, 2023.
- [20] S. Segarra, A. G. Marques, G. Mateos, and A. Ribeiro, "Network topology inference from spectral templates," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 3, pp. 467–483, Sep. 2017.
- [21] R. Shafipour and G. Mateos, "Online topology inference from streaming stationary graph signals with partial connectivity information," *Algorithms*, vol. 13, no. 9, pp. 228, 2020.
- [22] T. M. Roddenberry, M. Navarro, and S. Segarra, "Network topology inference with graphon spectral penalties," in *IEEE Intl. Conf. Acoustics, Speech Signal Process. (ICASSP)*, 2021, pp. 5390–5394.
- [23] A. Buciuiea, S. Rey, and A. G. Marques, "Learning graphs from smooth and graph-stationary signals with hidden variables," *IEEE Trans. Signal, Inform. Process. Networks*, vol. 8, pp. 273–287, 2022.
- [24] M. Navarro, S. Rey, A. Buciuiea, A. G. Marques, and S. Segarra, "Joint network topology inference in the presence of hidden nodes," *IEEE Trans. Signal Process.*, vol. 72, pp. 2710–2725, 2024.
- [25] X. Zheng, B. Aragam, P. K. Ravikumar, and E. P. Xing, "DAGs with NO TEARS: Continuous optimization for structure learning," *Advances Neural Inf. Process. Syst.*, vol. 31, 2018.
- [26] D. Wei, T. Gao, and Y. Yu, "DAGs with No Fears: A closer look at continuous optimization for learning Bayesian networks," *Advances Neural Inf. Process. Syst.*, vol. 33, pp. 3895–3906, 2020.
- [27] K. Bello, B. Aragam, and P. Ravikumar, "DAGMA: Learning DAGs via M-matrices and a log-determinant acyclicity characterization," *Advances Neural Inf. Process. Syst.*, vol. 35, pp. 8226–8239, 2022.
- [28] C. Peng and U. Mitra, "Causal graph identification under soft intervention," in *IEEE Intl. Symp. Information Theory (ISIT)*, 2025, pp. 1–6.
- [29] S. S. Saboksayr, G. Mateos, and M. Tepper, "CoLiDE: Concomitant linear DAG estimation," *Intl. Conf. Learn. Repr. (ICLR)*, 2024.
- [30] D. M. Chickering and D. Heckerman, "Efficient approximations for the marginal likelihood of Bayesian networks with hidden variables," *Mach. Learn.*, vol. 29, pp. 181–212, 1997.
- [31] D. M. Chickering, D. Heckerman, and C. Meek, "Large-sample learning of Bayesian networks is NP-hard," *J. Mach. Learn. Res.*, vol. 5, pp. 1287–1330, 2004.
- [32] R. Pamfil, N. Sriwattanaworachai, S. Desai, P. Pilgerstorfer, Konstantinos G., Paul B., and B. Aragam, "DYNOTEARS: Structure learning from time-series data," in *Intl. Conf. Artif. Intel. Statist. (AISTATS)*, 2020, pp. 1595–1605.
- [33] K. P. Murphy, *Dynamic bayesian networks: representation, inference and learning*. University of California, Berkeley, 2002.
- [34] I. Ng, A. Ghassami, and K. Zhang, "On the role of sparsity and DAG constraints for learning linear DAGs," *Advances Neural Inf. Process. Syst.*, vol. 33, pp. 17943–17954, 2020.
- [35] T. Dean and K. Kanazawa, "A model for reasoning about persistence and causation," *Comput. Intell.*, vol. 5, no. 2, pp. 142–150, 1989.
- [36] N. R. Swanson and C. W. J. Granger, "Impulse response functions based on a causal approach to residual orthogonalization in vector autoregressions," *J. Amer. Statist. Assoc.*, vol. 92, no. 437, pp. 357–367, 1997.
- [37] T. Gao, D. Bhattacharjya, E. Nelson, M. Liu, and Y. Yu, "IDYNO: Learning nonparametric DAGs from interventional dynamic data," in *Intl. Conf. Learn. Repr. (ICLR)*. Pmlr, 2022, pp. 6988–7001.
- [38] P. Misiakos and M. Püschel, "Learning dags and root causes from time-series data," 2025.
- [39] C. Deng, K. Bello, P. K. Ravikumar, and B. Aragam, "Global optimality in bivariate gradient-based DAG learning," in *Intl. Conf. Mach. Learn. (ICML)*, 2023.
- [40] S. Rey, S. S. Saboksayr, and G. Mateos, "Non-negative weighted DAG structure learning," in *IEEE Intl. Conf. Acoustics, Speech Signal Process. (ICASSP)*. IEEE, 2025, pp. 1–5.
- [41] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [42] D. P. Bertsekas, *Dynamic Programming*, vol. 4. Athena Scientific, 1997.