

Enhancing Graphical Lasso: A Robust Scheme for Non-Stationary Mean Data

S. Rey[†], E. Curbelo[‡], L. Martino^{*}, F. Llorente[◇], A. G. Marques[†]

[†] Dept. of Signal Theory and Communications, Universidad Rey Juan Carlos, Madrid, Spain

[‡] Dept. of Statistics, Universidad Carlos III de Madrid, Madrid, Spain.

^{*} Dept. of Economics and Business, Università degli Studi di Catania, Catania, Italy.

[◇] Computing and Data Science Directorate, Brookhaven National Laboratory, Upton, USA.

Abstract—This work addresses the problem of graph learning from data following a Gaussian Graphical Model (GGM) with a time-varying mean. Graphical Lasso (GL), the standard method for estimating sparse precision matrices, assumes that the observed data follows a zero-mean Gaussian distribution. However, this assumption is often violated in real-world scenarios where the mean evolves over time due to external influences, trends, or regime shifts. When the mean is not properly accounted for, applying GL directly can lead to estimating a biased precision matrix, hence hindering the graph learning task. To overcome this limitation, we propose *Graphical Lasso with Adaptive Targeted Adaptive Importance Sampling (GL-ATAIS)*, an iterative method that jointly estimates the time-varying mean and the precision matrix. Our approach integrates Bayesian inference with frequentist estimation, leveraging importance sampling to obtain an estimate of the mean while using a regularized maximum likelihood estimator to infer the precision matrix. By iteratively refining both estimates, GL-ATAIS mitigates the bias introduced by time-varying means, leading to more accurate graph recovery. Our numerical evaluation demonstrates the impact of properly accounting for time-dependent means and highlights the advantages of GL-ATAIS over standard GL in recovering the true graph structure.

Index Terms—Network topology inference, graph learning, graphical models, Bayesian inversion, importance sampling.

I. INTRODUCTION

A fundamental problem in data analysis is the inference of pairwise interactions between variables, which can be naturally framed as learning the topology of a graph. Graphs serve as powerful tools for modeling structured data in diverse fields such as neuroscience, genomics, and finance [1]–[3]. Among the many approaches for inferring graph structure from nodal observations [4]–[6], Graphical Lasso (GL) stands out as a prominent and widely used method [7]–[10]. GL operates under the assumption that the data follows a Gaussian Graphical Model (GGM), where the graph structure is encoded in the precision matrix, i.e., the inverse of the covariance matrix. By leveraging this assumption, GL provides a regularized

maximum likelihood estimator for inferring the underlying graph.

Over the years, numerous extensions of GL have been developed to enhance its robustness and flexibility. These include adaptations for scenarios with partial observability [11]–[13], more sophisticated structural priors to encode advanced graph information [10], [14]–[16], and methods for jointly estimating multiple graphs [17]–[19]. Despite these advancements, a key underlying assumption remains: the observed data follows a *zero-mean* multivariate Gaussian distribution. While in some cases the sample mean can be removed as a preprocessing step, this approach is inadequate in many real-world settings. For instance, time series data frequently exhibit time-dependent mean structures due to external influences, trends, or regime shifts. If the mean is not properly estimated before applying GL, the inferred precision matrix may be biased, leading to incorrect graph topology estimation. Thus, accounting for temporal variability in the mean is crucial for ensuring reliable graph learning with GL.

Several works have explored graph learning in dynamic environments. The works in [20]–[23] consider scenarios where the observations follow a GGM with a time-evolving precision matrix, which implies that the graph changes over time. Building on this premise, later approaches have developed efficient online algorithms for learning dynamic graph structures from streaming data, hence adapting to changes in graph connectivity [24]–[26]. Despite their importance, these methods generally assume that the mean of the observed data is zero, which limits their applicability in settings where the mean is time-dependent. More related to our work is [27], which focuses on estimating the time-dependent mean of a GGM, with the covariance matrix obtained as a byproduct. However, this approach does not explicitly address the problem of graph topology inference, and the only available information about the precision matrix comes from the inverse of the sample covariance, which is known to be a poor estimate in high-dimensional regimes.

This work addresses the aforementioned limitations of GL by considering the problem of graph learning when observations are drawn from a GGM with a time-dependent mean. Accurately estimating the mean is essential to prevent bias in

This work was partially supported by the Spanish AEI Grants PID2022-136887NB-I00, PID2023-149457OB-I00, the Community of Madrid (via grants CAM-URJC F1180 (CP2301), TEC-2024/COM-89, and Madrid EL-LIS Unit), by the "PIA no di inCentivi per la Ricerca di Ateneo 2026" (UPB 28722052159) and by PIACERI Starting Grant BA-GRAPH (UPB 28722052144) of the University of Catania, and the Office of Naval Research (award 558 no. N62909-24-1-2095).

the learned precision matrix. To this end, we propose a novel approach, termed Graphical Lasso with Adaptive Targeted Adaptive Importance Sampling (GL-ATAIS), which integrates Bayesian inference with frequentist estimation to overcome the limitations of GL in the presence of time-dependent means. GL-ATAIS is an alternating optimization algorithm where we first leverage importance sampling to estimate the time-dependent mean, and then harness this estimate to center the observations and apply GL to learn the precision matrix. By iteratively refining both the mean and the precision matrix estimates, GL-ATAIS provides a principled approach for recovering the true graph structure under time-varying mean conditions.

II. FUNDAMENTALS OF GAUSSIAN GRAPHICAL MODELS

In this section, we provide the necessary definitions and background for our problem of interest.

Graphs and graph signals. Let $\mathcal{G} := (\mathcal{V}, \mathcal{E})$ be an undirected and weighted graph, where $\mathcal{V}$ denotes the set of nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ represents the set of edges. The graph consists of $|\mathcal{V}| = N$ nodes, and for any pair $(i, j) \in \mathcal{V} \times \mathcal{V}$, an edge $(i, j) \in \mathcal{E}$ exists if and only if nodes i and j are connected. The connectivity of $\mathcal{G}$ is encoded in the weighted adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, where $A_{ij} \neq 0$ if and only if $(i, j) \in \mathcal{E}$. Together with the graph $\mathcal{G}$, we consider signals defined on its node set $\mathcal{V}$. Formally, a *graph signal* is a function from the node set to the real field, $x : \mathcal{V} \rightarrow \mathbb{R}$, or equivalently, an N -dimensional vector $\mathbf{x} \in \mathbb{R}^N$, where x_i denotes the signal value at node i .

Gaussian Graphical Models (GGMs). GGMs describe the conditional dependence structure of multivariate Gaussian distributions as a graph. Consider a random graph signal $\mathbf{x} \in \mathbb{R}^N$ drawn from a Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma} \succ 0$, i.e., $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The inverse covariance matrix $\boldsymbol{\Theta} = \boldsymbol{\Sigma}^{-1}$, known as the *precision matrix*, fully characterizes the conditional dependence structure of $\mathbf{x}$. Specifically, for any distinct pair $(i, j) \in \mathcal{V} \times \mathcal{V}$, the variables x_i and x_j are conditionally independent if and only if $\Theta_{ij} = 0$ [7], [28]. This Markovian property naturally induces a graph $\mathcal{G}$, where the nonzero off-diagonal entries of $\boldsymbol{\Theta}$ define the weighted edges representing conditional dependencies between variables.

Graphical Lasso (GL). Let the matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_R] \in \mathbb{R}^{N \times R}$ collect R independent samples drawn from $\mathcal{N}(\mathbf{0}, \boldsymbol{\Theta}^{-1})$. A widely used method to estimate the graph structure encoded in $\boldsymbol{\Theta}$ from $\mathbf{X}$ is GL [8], which solves the following convex optimization problem

$$\begin{aligned} \hat{\boldsymbol{\Theta}} &= \underset{\boldsymbol{\Theta} \succeq 0}{\operatorname{argmin}} \frac{1}{2R} \sum_{r=1}^R \mathbf{x}_r^\top \boldsymbol{\Theta} \mathbf{x}_r - \frac{1}{2} \log \det(\boldsymbol{\Theta}) + \lambda \|\boldsymbol{\Theta}\|_1, \\ &= \underset{\boldsymbol{\Theta} \succeq 0}{\operatorname{argmin}} \frac{1}{2} \operatorname{tr}(\hat{\boldsymbol{\Sigma}}_{\text{emp}} \boldsymbol{\Theta}) - \frac{1}{2} \log \det(\boldsymbol{\Theta}) + \lambda \|\boldsymbol{\Theta}\|_1. \end{aligned} \quad (1)$$

Here, $\hat{\boldsymbol{\Sigma}}_{\text{emp}} = \frac{1}{R} \mathbf{X} \mathbf{X}^\top$ is the empirical covariance matrix, the constrain $\boldsymbol{\Theta} \succeq 0$ represents the set of positive semidefinite matrices, and $\|\boldsymbol{\Theta}\|_1$ is the elementwise ℓ_1 -norm of $\boldsymbol{\Theta}$ excluding

its diagonal entries to promote sparsity. Indeed, the objective can be interpreted as a negative log-likelihood function, which we use in the next section to define the Bayesian model.

The GL formulation in (1) provides a regularized maximum likelihood estimator of $\boldsymbol{\Theta}$, relying on the assumption that the mean of the distribution is zero, i.e., $\boldsymbol{\mu} = \mathbb{E}[\mathbf{x}] = \mathbf{0}$. When $\boldsymbol{\mu} \neq \mathbf{0}$, the standard approach is to use the sample mean estimate $\hat{\boldsymbol{\mu}} = \frac{1}{R} \sum_{r=1}^R \mathbf{x}_r$ to center the data before applying GL. However, when $\boldsymbol{\mu}$ varies over time, $\hat{\boldsymbol{\mu}}$ becomes a poor estimate, leading to inaccurate recovery of the graph topology encoded in $\boldsymbol{\Theta}$. To overcome this limitation, we next propose an iterative procedure where we first estimate the time-varying mean $\boldsymbol{\mu}$ and then infer $\boldsymbol{\Theta}$ using the debiased observations. This strategy ensures a more accurate estimation of the underlying graph structure in scenarios where the mean exhibits temporal variations.

III. GGMs WITH TIME-VARYING MEAN

This work focuses on the general setting where the observations $\mathbf{x}_r$ follow a random time series modeled by a GGM with a time-varying mean, i.e., $\mathbf{x}_r \sim \mathcal{N}(\boldsymbol{\mu}_r, \boldsymbol{\Theta}^{-1})$, with r denoting a discrete time index. Furthermore, we assume that the mean $\boldsymbol{\mu}_r$ can be accurately represented by a (possibly non-linear) parametric function $\mathbf{f}_r(\boldsymbol{\phi})$, where $\boldsymbol{\phi} \in \mathbb{R}^M$ denotes the unknown parameters, and the subscript r emphasizes the time dependency. Intuitively, this is equivalent to the following signal model

$$\mathbf{x}_r = \mathbf{f}_r(\boldsymbol{\phi}) + \mathbf{v}_r, \quad (2)$$

where $\mathbf{v}_r \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Theta}^{-1})$ represents zero-mean fluctuations modeled by a standard GGM. We note that $\mathbf{f}_r$ is a vector valued-function (meaning that the time-varying function associated with each node can be different). With this notation at hand, we formally state our graph learning problem below.

Problem 1 *Let the matrix $\mathbf{X} \in \mathbb{R}^{N \times R}$ collect R observations of a time series following (2). Our goal is to estimate the graph topology encoded in the precision matrix $\boldsymbol{\Theta} \in \mathbb{R}^{N \times N}$, assuming that the function $\mathbf{f}_r(\boldsymbol{\phi})$ with unknown parameters $\boldsymbol{\phi} \in \mathbb{R}^M$, can be evaluated.*

While our primary interest lies in learning the graph structure, accurately estimating the parameters $\boldsymbol{\phi}$ is crucial, as it enables an improved estimation of the mean of the underlying process. Notably, we do not assume prior knowledge of the parametric function $\mathbf{f}_r(\boldsymbol{\phi})$, nor do we impose any specific structure, such as convexity or differentiability, thereby allowing for a broad range of possible mean functions. With these considerations in mind, we now introduce an iterative approach designed to address this challenge and solve Problem 1.

A. GL with MAP estimation of time-varying mean

The cornerstone of our method is accurately estimating the mean of the GGM, which amounts to estimating the parameters of the function $\mathbf{f}_r(\boldsymbol{\phi})$ in (2). A natural approach is to formulate a two-step process where one first estimates the parameters $\boldsymbol{\phi}$, uses them to center the observations collected

in $\mathbf{X}$, and then solves GL as in (1). A key observation is that any method capable of accurately estimating the mean (i.e., the parameters ϕ) can be employed for this purpose [27], [29]

In this work, we advocate for a more involved method that integrates Bayesian inference with frequentist estimation. Let us assume some priors $g_\phi(\phi)$ and $g_\Theta(\Theta) \propto \exp(-\lambda\|\Theta\|_1)$, and a likelihood function $\ell(\mathbf{X}|\Theta, \phi) = \prod_r \mathcal{N}(\mathbf{x}_r|\mathbf{f}_r(\phi), \Theta^{-1})$, so that the posterior is given by $\pi(\phi, \Theta|\mathbf{X}) \propto \ell(\mathbf{X}|\Theta, \phi)g_\phi(\phi)g_\Theta(\Theta)$. Specifically, we aim to compute the maximum a posteriori (MAP) estimate of ϕ and the sparsity-regularized maximum likelihood estimator of Θ , respectively given by

$$\hat{\phi}_{\text{MAP}} = \underset{\phi}{\operatorname{argmax}} \log \pi(\phi|\mathbf{X}, \hat{\Theta}_{\text{GL}}), \quad (3)$$

$$\hat{\Theta}_{\text{GL}} = \underset{\Theta}{\operatorname{argmax}} \log \ell(\mathbf{X}|\Theta, \hat{\phi}_{\text{MAP}}) - \lambda\|\Theta\|_1. \quad (4)$$

Note that the maximizer $\hat{\Theta}_{\text{GL}}$ coincides with the one in (1) when we assume $\mathbf{f}_r(\phi) = \mathbf{0}$ for all r . Here, $\pi(\phi|\mathbf{X}, \hat{\Theta}_{\text{GL}})$ denotes the conditional posterior of ϕ and $\ell(\mathbf{X}|\Theta, \hat{\phi}_{\text{MAP}})$ is the likelihood of Θ given $\hat{\phi}_{\text{MAP}}$. Following [27], we estimate both $\hat{\phi}_{\text{MAP}}$ and $\hat{\Theta}_{\text{GL}}$ via an iterative algorithm, which we term *Graphical Lasso with Adaptive Target Adaptive Importance Sampling (GL-ATAIS)*. GL-ATAIS is an alternating optimization algorithm where the following steps are sequentially performed for a maximum of K iterations.

Step 1 - Mean estimation. The first step is to estimate $\hat{\phi}^{(k+1)}$ by maximizing the log-posterior $\pi(\phi|\mathbf{X}, \hat{\Theta}_{\text{GL}}^{(k)})$, where $\hat{\Theta}_{\text{GL}}^{(k)}$ is the estimate from the previous iteration. This amounts to solving the following optimization problem

$$\min_{\phi} \left[\frac{1}{2} \sum_{r=1}^R (\mathbf{x}_r - \mathbf{f}_r(\phi))^\top \hat{\Theta}_{\text{GL}}^{(k)} (\mathbf{x}_r - \mathbf{f}_r(\phi)) - \log g_\phi(\phi) \right], \quad (5)$$

where $\mathbf{x}_r$ are the columns of $\mathbf{X}$ and $g_\phi(\phi)$ represents a prior on the distribution of ϕ .

Since $\mathbf{f}_r(\phi)$ is not restricted to be convex or differentiable, we rely on an importance sampling (IS) procedure to find an approximate solution to (5). More precisely, given a *proposal* density $q(\phi|\tilde{\mu}^{(k)}, \tilde{\Sigma}^{(k)})$ with mean $\tilde{\mu}^{(k)}$ and covariance $\tilde{\Sigma}^{(k)}$, we draw P samples $\{\phi_k^{(p)}\}_{p=1}^P$ and compute $\hat{\phi}^{(k+1)}$ as

$$\hat{\phi}^{(k+1)} = \phi_k^{(p_k^*)}, \text{ with } p_k^* := \underset{p \in \{1, \dots, P\}}{\operatorname{argmax}} \pi(\phi_k^{(p)}|\mathbf{X}, \hat{\Theta}_{\text{GL}}^{(k)}). \quad (6)$$

The impact of the IS procedure becomes evident in the third step, where the samples $\{\phi_k^{(p)}\}_{p=1}^P$ are used to compute the importance weights and update $\tilde{\mu}^{(k)}$ and $\tilde{\Sigma}^{(k)}$.

Step 2 - Precision estimation. After estimating $\hat{\phi}^{(k+1)}$, the mean of the GGM underlying our data $\mathbf{X}$ is given by $\mathbf{f}_r(\hat{\phi}^{(k+1)})$. Thus, we center the observations by removing the estimated mean and compute the sample covariance as

$$\hat{\Sigma}^{(k+1)} = \frac{1}{R} \sum_{r=1}^R (\mathbf{x}_r - \mathbf{f}_r(\hat{\phi}^{(k+1)}))(\mathbf{x}_r - \mathbf{f}_r(\hat{\phi}^{(k+1)}))^\top, \quad (7)$$

and then estimate $\hat{\Theta}^{(k+1)}$ via GL by solving

$$\hat{\Theta}^{(k+1)} = \underset{\Theta \succeq 0}{\operatorname{argmin}} \frac{1}{2} \operatorname{tr}(\hat{\Sigma}^{(k+1)} \Theta) - \frac{1}{2} \log \det(\Theta) + \lambda\|\Theta\|_1. \quad (8)$$

Clearly, this is equivalent to solving the original GL formulation in (1), but replacing the sample covariance $\hat{\Sigma}_{\text{emp}}$ with its centered estimate $\hat{\Sigma}^{(k+1)}$. Since (8) is strictly convex, a variety of efficient algorithms can be employed to obtain the global minimum [25], [30], [31].

Step 3 - Adaptation. The final step updates the estimates $\hat{\phi}_{\text{MAP}}^{(k+1)}$ and $\hat{\Theta}_{\text{GL}}^{(k+1)}$, as well as the mean and covariance of the proposal density function $q(\phi|\tilde{\mu}^{(k)}, \tilde{\Sigma}^{(k)})$. With the conditional log-posterior of ϕ given by

$$\begin{aligned} \log \pi(\phi|\mathbf{X}, \Theta) &= -\frac{1}{2} \sum_{r=1}^R (\mathbf{x}_r - \mathbf{f}_r(\phi))^\top \Theta (\mathbf{x}_r - \mathbf{f}_r(\phi)) \\ &\quad + \frac{1}{2} \log \det(\Theta) + \log g_\phi(\phi) - \lambda\|\Theta\|_1, \end{aligned} \quad (9)$$

we update the estimates as follows. We set

$$\hat{\phi}_{\text{MAP}}^{(k+1)} = \hat{\phi}^{(k+1)}, \quad \hat{\Theta}_{\text{GL}}^{(k+1)} = \hat{\Theta}^{(k+1)}, \quad (10)$$

if $\log \pi(\hat{\phi}^{(k+1)}|\mathbf{X}, \hat{\Theta}^{(k+1)}) > \log \pi(\hat{\phi}_{\text{MAP}}^{(k)}|\mathbf{X}, \hat{\Theta}_{\text{GL}}^{(k)})$. Otherwise, we retain the previous estimates, i.e., $\hat{\phi}_{\text{MAP}}^{(k+1)} = \hat{\phi}_{\text{MAP}}^{(k)}$ and $\hat{\Theta}_{\text{GL}}^{(k+1)} = \hat{\Theta}_{\text{GL}}^{(k)}$.

Next, using the samples $\{\phi_k^{(p)}\}_{p=1}^P$ from Step 1, we compute the importance weights

$$w_k^{(p)} = \frac{\pi(\phi_k^{(p)}|\mathbf{X}, \hat{\Theta}^{(k)})}{q(\phi_k^{(p)}|\tilde{\mu}^{(k)}, \tilde{\Sigma}^{(k)})}. \quad (11)$$

The proposal distribution parameters are then updated as

$$\tilde{\mu}^{(k+1)} = \hat{\phi}_{\text{MAP}}^{(k+1)}, \quad (12)$$

$$\tilde{\Sigma}^{(k+1)} = \sum_{p=1}^P \bar{w}_k^{(p)} (\phi_k^{(p)} - \bar{\phi}_k)(\phi_k^{(p)} - \bar{\phi}_k)^\top + \delta_k \mathbf{I}, \quad (13)$$

where $\bar{w}_k^{(p)} = \frac{w_k^{(p)}}{\sum_{p=1}^P w_k^{(p)}}$, $\bar{\phi}_k = \sum_{p=1}^P \bar{w}_k^{(p)} \phi_k^{(p)}$, and $\delta_k > 0$. Updating the proposal parameters following (12) and (13) improves the performance and robustness of IS, mitigating catastrophic scenarios and leading to more stable results, particularly as the problem dimension increases [32], [33]. Furthermore, updating $\hat{\phi}_{\text{MAP}}^{(k+1)}$ and $\hat{\Theta}_{\text{GL}}^{(k+1)}$ only when the evaluation of the regularized log-posterior improves helps avoid convergence to suboptimal stationary points. After running steps 1–3 for K iterations, GL-ATAIS provides a frequentist estimate of the precision matrix $\hat{\Theta}_{\text{GL}} = \hat{\Theta}_{\text{GL}}^{(K)}$ and, consequently, the graph topology. Moreover, the Bayesian estimation of $\hat{\phi}_{\text{MAP}} = \hat{\phi}_{\text{MAP}}^{(K)}$ not only yields a point estimate but also enables uncertainty quantification of the estimator. Developing a Bayesian inference method to learn Θ and assess the uncertainty of the estimated edges remains an intriguing direction for future research.

B. Accelerating the convergence

The convergence of the global optimization problem can be accelerated by implementing minor modifications in (5) in the first iterations. Namely, in order to find a good region of the space for the initial point of the alternating optimization procedure, we can have a number of iterations $K_0 < K$ of the algorithm considering

$$\phi_{\text{MAP}}^{(k)} = \arg \min_{\phi} \left[\frac{1}{2} \sum_{r=1}^R \|\mathbf{x}_r - \mathbf{f}_r(\phi)\|^2 - \log g_{\phi}(\phi) \right], \quad (14)$$

with $k = 1, \dots, K_0$, instead of finding the minimum of the optimization problem in (5). This is equivalent to set $\hat{\Theta}_{\text{GL}}^{(k)} = \mathbf{I}_N$ for $k = 0, \dots, K_0 - 1$ in (5), where $\mathbf{I}_N$ is the $N \times N$ identity matrix. Thus, in the first K_0 iterations, we focus only on finding a good point $\phi_{\text{MAP}}^{(K_0)}$. Indeed, note that if there exists a point ϕ^* such that $\sum_{r=1}^R \|\mathbf{x}_r - \mathbf{f}_r(\phi^*)\|^2 = 0$, then this point ϕ^* is also a root for $\sum_{r=1}^R (\mathbf{x}_r - \mathbf{f}_r(\phi^*))^\top \hat{\Theta} (\mathbf{x}_r - \mathbf{f}_r(\phi^*)) = 0$ for any possible precision matrix $\hat{\Theta}$.

IV. NUMERICAL EVALUATION

This section evaluates how the time-varying mean of the data impacts the performance of GL, illustrating the importance of estimating its mean and the benefits of GL-ATAIS.

Experiment setup. We generate Erdős-Rényi (ER) graphs with $N = 10$ nodes and an edge probability of $p = 0.1$. The true precision matrix is then constructed as $\Theta = \mathbf{A} + \epsilon \mathbf{I}$, where $\epsilon > 0$ is a small constant ensuring that Θ is positive semidefinite. We sample R graph signals from a Gaussian distribution $\mathcal{N}(\mathbf{f}_r(\phi), \Theta^{-1})$, with $\phi \in \mathbb{R}^4$ and a time-dependent mean function $\mathbf{f}_r(\phi)$ given by

$$\begin{aligned} f_{r,[1]}(\phi) &= -\phi_4 \tau_r + 5\phi_1^2 \\ f_{r,[2]}(\phi) &= 2\phi_3 \sin(-\phi_2 \tau_r) \\ f_{r,[3]}(\phi) &= \phi_1 - \phi_3 + \phi_1 \cos(2\tau_r) \\ f_{r,[4]}(\phi) &= 3\phi_4 + 3\phi_2 + \phi_1 \exp(0.1\tau_r) \\ f_{r,[5]}(\phi) &= \phi_3^2 - 2\phi_1 + 3\phi_2 - \exp(0.8\phi_3) \exp(1 - \tau_r) \\ f_{r,[6]}(\phi) &= 5(\phi_4 + \phi_3) - \phi_2 \log(1 + 2\tau_r) \\ f_{r,[7]}(\phi) &= 3\phi_2 - 0.2\tau \sin(\phi_3) \\ f_{r,[8]}(\phi) &= 3\phi_1 + 5\phi_3 - 20 \sin(\phi_4) \cos(2\tau_r + \frac{\pi}{4}) \\ f_{r,[9]}(\phi) &= \phi_2 + 4\phi_4 + 5 \exp\left(\frac{1}{1 + \phi_3}\right) \tau_r \\ f_{r,[10]}(\phi) &= 5\phi_1 + 10\phi_3 - 5\phi_4 \sin(\tau_r), \end{aligned} \quad (15)$$

Here, $f_{r,[i]}$ denotes the value of the function $\mathbf{f}_r(\phi)$ at the node i , and τ_r specifies the timestamp of the r -th sample. More precisely, the values of τ_r correspond to R equally spaced timestamps in the interval $[0, 4]$.

We compare the following methods: i) Standard GL: the classical GL formulation (1), where the observed data is centered using the sample mean estimate; ii) GL with true mean: an oracle method where the true μ_r is used to center the observations before applying GL, hence serving as an upper bound for the performance of other methods; iii) GL-ATAIS: our proposed algorithm, which jointly estimates the mean and the precision matrix; and iv) ATAIS [27]: a method that estimates Θ as the inverse of the inferred covariance matrix.




Fig. 1: Average F1-score over 100 repetitions for four algorithms, when the number of observations varies in $\{50, 100, 150, 200\}$. The ATAIS repetitions are run for $T=30$ iterations, with $N=30000$ particles.




Fig. 2: Average F1-score over 100 repetitions for four algorithms versus the number of particles in ATAIS.

Finally, we evaluate the performance using the average F-score of the estimated graph over 100 independent realizations.

Test case 1 - Impact of the number of observations. In this experiment, we analyze how the number of observed signals R influences the quality of the estimated graph. Fig. 1 depicts the F-score of different methods as R increases from 50 to 200. As expected, performance improves with larger R , except for standard GL, which remains consistently poor. In contrast, GL with the true mean outperforms all alternatives, reinforcing our claim that standard GL cannot recover the graph topology (i.e., the support of Θ), when the mean varies over time. This result underscores the importance of properly estimating the mean before applying GL. While relying on the true mean is unrealistic in practical scenarios, it serves as an upper bound for graph recovery performance. Interestingly, GL-ATAIS achieves an F-score comparable to this oracle case, highlighting the effectiveness of our method. Moreover, GL-ATAIS significantly outperforms ATAIS, particularly when R is small, demonstrating the advantages of using GL for precision matrix estimation instead of directly inverting the estimated covariance.

Test case 2 - Impact of the number of particles. Next, we investigate how the number of particles P used to estimate the mean parameters ϕ affects the recovery of Θ . Fig. 2 presents the results, where P ranges from 30 to 30000. Note that the number of particles only affects the performance of ATAIS and GL-ATAIS since they are the only methods based on IS. Consistent with the previous experiment, we observe that standard GL fails to recover the graph when using the sample mean, while it performs best when the true mean is available. For large P , GL-ATAIS outperforms ATAIS and approaches the performance of GL with the true mean, validating the importance of jointly estimating the mean and precision matrix. However, when P is small, the estimation of ϕ deteriorates, leading to a poor graph recovery. This highlights the trade-off between computational cost and estimation accuracy when selecting the number of particles in GL-ATAIS.

V. CONCLUSION

This work addressed the problem of learning graphs from nodal signals when such signals are drawn from a GGM with a time-varying mean. To tackle this challenge, we proposed GL-ATAIS, an iterative method that combines Bayesian inference with frequentist estimation to jointly learn the time-varying mean and the precision matrix. Assuming that the mean is represented by a parametric function $f_r(\phi)$, GL-ATAIS employs an importance sampling strategy to refine the estimate of the parameters ϕ , leverages this estimate to center the observed data, and then applies GL to obtain the maximum likelihood estimator Θ . Notably, this approach enables uncertainty quantification for the parameters ϕ , and developing a Bayesian estimator for Θ remains an interesting direction for future research. Finally, our numerical experiments demonstrated that the zero-mean assumption of GL severely hinders its ability to recover the true graph topology, while GL-ATAIS effectively mitigates the bias introduced by time-varying means, resulting in significantly improved graph recovery.

REFERENCES

- [1] E. Belilovsky, G. Varoquaux, and M. B. Blaschko, "Testing for differences in Gaussian graphical models: Applications to brain connectivity," in *Conf. Neural Inform. Process. Syst.*, 2016, vol. 29.
- [2] S. M. Smith, K. L. Miller, G. Salimi-Khorshidi, M. Webster, C. F. Beckmann, T. E. Nichols, J. D. Ramsey, and M. W. Woolrich, "Network modelling methods for FMRI," *NeuroImage*, vol. 54, no. 2, pp. 875–891, 2011.
- [3] O. Stegle, S. A. Teichmann, and J. C. Marioni, "Computational and analytical challenges in single-cell transcriptomics," *Nature Reviews Genetics*, vol. 16, no. 3, pp. 133–145, 2015.
- [4] V. Kalofolias, "How to learn a graph from smooth signals," in *Intl. Conf. Artif. Intel. Statist. (AISTATS)*. J. Mach. Learning Res., 2016, pp. 920–929.
- [5] S. Segarra, A. G. Marques, G. Mateos, and A. Ribeiro, "Network topology inference from spectral templates," *IEEE Trans. Signal, Inform. Process. Networks*, vol. 3, no. 3, pp. 467–483, Sep. 2017.
- [6] 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.
- [7] S. L. Lauritzen, *Graphical Models*, vol. 17, Clarendon Press, 1996.
- [8] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.
- [9] 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.
- [10] A. Buciulea, J. Ying, A. G. Marques, and D. P. Palomar, "Polynomial graphical lasso: Learning edges from gaussian graph-stationary signals," 2025.
- [11] V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky, "Latent variable graphical model selection via convex optimization," *Annu. Allerton Conf. Commun., Control, Comput.*, vol. 40, no. 4, pp. 1935–1967, 2012.
- [12] Z. Meng, B. Eriksson, and A. Hero, "Learning latent variable gaussian graphical models," in *Int. Conf. Mach. Learning*. PMLR, 2014, pp. 1269–1277.
- [13] A. Buciulea, S. Rey, and A. G. Marques, "Learning graphs from smooth and graph-stationary signals with hidden variables," vol. 8, pp. 273–287, 2022.
- [14] S. Kumar, J. Ying, J. Cardoso, and D. Palomar, "Structured graph learning via Laplacian spectral constraints," *Conf. Neural Inform. Process. Syst.*, vol. 32, 2019.
- [15] M. Navarro and S. Segarra, "Joint network topology inference via a shared graphon model," vol. 70, pp. 5549–5563, 2022.
- [16] S. Rey, T. M. Roddenberry, S. Segarra, and A. G. Marques, "Enhanced graph-learning schemes driven by similar distributions of motifs," vol. 71, pp. 3014–3027, 2023.
- [17] L. Jacob, G. Obozinski, and J.-P. Vert, "Group lasso with overlap and graph lasso," in *Int. Conf. Mach. Learning*, 2009, pp. 433–440.
- [18] P. Danaher, P. Wang, and D. M. Witten, "The joint graphical lasso for inverse covariance estimation across multiple classes," *J. Royal Statistical Soc. Series B: Statistical Methodology*, vol. 76, no. 2, pp. 373–397, 2014.
- [19] S. Rey, M. Navarro, A. Buciulea, S. Segarra, and A. G. Marques, "Joint graph learning from gaussian observations in the presence of hidden nodes," in *Asilomar Conf. Signals, Systems and Comput.* IEEE, 2022, pp. 53–57.
- [20] D. Hallac, Y. Park, S. Boyd, and J. Leskovec, "Network inference via the time-varying graphical lasso," in *ACM SIGKDD Int. Conf. Knowledge Discovery & Data Mining*, 2017, pp. 205–213.
- [21] K. Greenewald, S. Park, S. Zhou, and A. Giessing, "Time-dependent spatially varying graphical models, with application to brain fmri data analysis," *Conf. Neural Inform. Process. Syst.*, vol. 30, 2017.
- [22] B. Cai, G. Zhang, A. Zhang, J. M. Stephen, T. W. Wilson, V. D. Calhoun, and Y.-P. Wang, "Capturing dynamic connectivity from resting state fmri using time-varying graphical lasso," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 7, pp. 1852–1862, 2018.
- [23] J. Yang and J. Peng, "Estimating time-varying graphical models," *J. Comput. Graphical Statist.*, vol. 29, no. 1, pp. 191–202, 2020.
- [24] H. Yang, Z. Xu, I. King, and M. R. Lyu, "Online learning for group lasso," in *Int. Conf. Mach. Learning*, 2010, pp. 1191–1198.
- [25] A. Natali, E. Isufi, M. Coutino, and G. Leus, "Learning time-varying graphs from online data," *IEEE Open J. Signal Signal Process.*, vol. 3, pp. 212–228, 2022.
- [26] S. Rey, B. Das, and E. Isufi, "Online learning of expanding graphs," *IEEE Open J. Signal Signal Process.*, 2025.
- [27] E. Curbelo, L. Martino, F. Llorente, and D. Delgado-Gomez, "Adaptive posterior distributions for uncertainty analysis of covariance matrices in bayesian inversion problems for multioutput signals," *J. Franklin Institute*, vol. 362, no. 2, pp. 107441, 2025.
- [28] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, MIT press, 2009.
- [29] L. Martino, F. Llorente, E. Curbelo, J. Lopez-Santiago, and J. Míguez, "Automatic tempered posterior distributions for bayesian inversion problems," *Mathematics*, vol. 9, no. 7, pp. 784, 2021.
- [30] D. M. Witten, J. H. Friedman, and N. Simon, "New insights and faster computations for the graphical lasso," *J. Comput. Graphical Statist.*, vol. 20, no. 4, pp. 892–900, 2011.
- [31] M. Navarro, S. Rey, A. Buciulea, A. G. Marques, and S. Segarra, "Fair glasso: Estimating fair graphical models with unbiased statistical behavior," *Conf. Neural Inform. Process. Syst.*, vol. 37, pp. 139589–139620, 2025.
- [32] Y. El-Laham, V. Elvira, and M. F. Bugallo, "Robust covariance adaptation in adaptive importance sampling," *IEEE Signal Process. Lett.*, vol. 25, no. 7, pp. 1049–1053, 2018.
- [33] F. Llorente, L. Martino, D. Delgado, and J. Lopez-Santiago, "Marginal likelihood computation for model selection and hypothesis testing: an extensive review," *SIAM review*, vol. 65, no. 1, pp. 3–58, 2023.