

SPARSE-GROUP LOG-SUM PENALIZED GRAPHICAL MODEL LEARNING FOR TIME SERIES

Jitendra K. Tugnait

Department of Electrical & Computer Engineering
Auburn University, Auburn, AL 36849, USA

ABSTRACT

We consider the problem of inferring the conditional independence graph (CIG) of a high-dimensional stationary multivariate Gaussian time series. A sparse-group lasso based frequency-domain formulation of the problem has been considered in the literature where the objective is to estimate the sparse inverse power spectral density (PSD) of the data. The CIG is then inferred from the estimated inverse PSD. In this paper we investigate use of a sparse-group log-sum penalty (LSP) instead of sparse-group lasso penalty. An alternating direction method of multipliers (ADMM) approach for iterative optimization of the non-convex problem is presented. We provide sufficient conditions for local convergence in the Frobenius norm of the inverse PSD estimators to the true value. This results also yields a rate of convergence. We illustrate our approach using numerical examples utilizing both synthetic and real data.

Keywords: Sparse graph learning; graph estimation; time series; undirected graph; inverse spectral density estimation.

1. INTRODUCTION

Graphical models are an important and useful tool for analyzing multivariate data [1]. A central concept is that of conditional independence. Given a collection of random variables, one wishes to assess the relationship between two variables, conditioned on the remaining variables. Consider a graph $\mathcal{G} = (V, \mathcal{E})$ with a set of p vertices (nodes) $V = \{1, 2, \dots, p\} = [p]$, and a corresponding set of (undirected) edges $\mathcal{E} \subseteq [p] \times [p]$. Also consider a stationary (real-valued), zero-mean, p -dimensional multivariate Gaussian time series $\mathbf{x}(t)$, $t = 0, \pm 1, \pm 2, \dots$, with i th component $x_i(t)$, and correlation (covariance) matrix function $\mathbf{R}_{xx}(\tau) = \mathbb{E}\{\mathbf{x}(t + \tau)\mathbf{x}^T(t)\}$, $\tau = 0, \pm 1, \dots$. Given $\{\mathbf{x}(t)\}$, in the corresponding graph \mathcal{G} , each component series $\{x_i(t)\}$ is represented by a node (i in V), and associations between components $\{x_i(t)\}$ and $\{x_j(t)\}$ are represented by edges between nodes i and j of \mathcal{G} . In a conditional independence graph (CIG), there is no edge between nodes i and j if and only if (iff) $x_i(t)$ and $x_j(t)$ are conditionally independent given the remaining $p-2$ scalar series $x_\ell(t)$, $\ell \in [p]$, $\ell \neq i$, $\ell \neq j$.

A key insight in [2] was to transform the series to the frequency domain and express the graph relationships in the frequency domain. Denote the power spectral density (PSD) matrix of $\{\mathbf{x}(t)\}$ by $\mathbf{S}_x(f)$, where $\mathbf{S}_x(f) = \sum_{\tau=-\infty}^{\infty} \mathbf{R}_{xx}(\tau)e^{-j2\pi f\tau}$. In [2] it was shown that conditional independence of two time series components given all other components of the time series, is encoded by zeros in the inverse PSD, that is, $\{i, j\} \notin \mathcal{E}$ iff the (i, j) -th element of $\mathbf{S}_x(f)$, $[\mathbf{S}_x^{-1}(f)]_{ij} = 0$ for every f . Hence one can use estimated inverse PSD of observed time series to infer the associated graph.

This work was supported by NSF Grant ECCS-2040536. Author's email: tugnajk@auburn.edu

Graphical models were originally developed for random vectors [3, p. 234]. Such models have been extensively studied, and found to be useful in a wide variety of applications [4–8]. Graphical modeling of real-valued time-dependent data (stationary time series) originated with [9], followed by [2]. Nonparametric approaches for graphical modeling of real time series in high-dimensional settings (p is large and/or sample size n is of the order of p) have been formulated in frequency-domain in [10, 11] using a neighborhood regression scheme, and in the form of group-lasso penalized log-likelihood in frequency-domain in [12]. Sparse-group lasso penalized log-likelihood approach in frequency-domain has been considered in [13, 14].

In this paper we consider a sparse-group log-sum penalty (SGLSP) instead of sparse-group lasso (SGL) penalty ([13, 14]) to regularize the problem, motivated by [16]. For sparse solutions, ideal penalty is ℓ_0 which is non-convex and the problem is usually impossible to solve. So one relaxes the problem using ℓ_1 (lasso) penalty which is convex. [16] notes that a key difference between the ℓ_1 and ℓ_0 norms is the dependence on magnitude: “larger coefficients are penalized more heavily in the ℓ_1 norm than smaller coefficients, unlike the more democratic penalization of the ℓ_0 norm.” Their solution to rectify this imbalance, is iterative reweighted ℓ_1 minimization, and to construct an analytical framework, [16] suggests the log-sum penalty. We present an ADMM approach for iterative optimization of the non-convex problem. We provide sufficient conditions for consistency of a local estimator of inverse PSD. We illustrate our approach using numerical examples utilizing both synthetic and real data. Synthetic data example shows that our SGLSP approach significantly outperforms the SGL and other approaches in correctly detecting the graph edges.

Notation: $|\mathbf{A}|$ and $\text{tr}(\mathbf{A})$ denote the determinant and the trace of the square matrix \mathbf{A} , respectively. $[\mathbf{B}]_{ij}$ denotes the (i, j) -th element of \mathbf{B} , and so does B_{ij} . \mathbf{I} is the identity matrix. The superscripts $*$ and H denote the complex conjugate and the Hermitian (conjugate transpose) operations, respectively. The notation $\mathbf{x} \sim \mathcal{N}_c(\mathbf{m}, \mathbf{\Sigma})$ denotes a random vector \mathbf{x} that is circularly symmetric (proper) complex Gaussian with mean \mathbf{m} and covariance $\mathbf{\Sigma}$.

2. SPARSE-GROUP LASSO PENALIZED NEGATIVE LOG-LIKELIHOOD

Given $\mathbf{x}(t)$ for $t = 0, 1, 2, \dots, n-1$. Define the (normalized) DFT $\mathbf{d}_x(f_m)$ of $\mathbf{x}(t)$, ($j = \sqrt{-1}$, $f_m = m/n$), over $m = 0, 1, \dots, n-1$ as $\mathbf{d}_x(f_m) = \frac{1}{\sqrt{n}} \sum_{t=0}^{n-1} \mathbf{x}(t) \exp(-j2\pi f_m t)$. It is established in [15] that the set of random vectors $\{\mathbf{d}_x(f_m)\}_{m=0}^{n/2}$ is a sufficient statistic for any inference problem based on dataset $\{\mathbf{x}(t)\}_{t=0}^{n-1}$. Suppose $\mathbf{S}_x(f_k)$ is locally smooth, so that $\mathbf{S}_x(f_k)$ is (approximately) constant over $K = 2m_t + 1$ consecutive frequency points

f_m s. Pick $M = \lfloor (\frac{n}{2} - m_t - 1)/K \rfloor$ and

$$\tilde{f}_k = \frac{(k-1)K + m_t + 1}{n}, \quad k = 1, 2, \dots, M, \quad (1)$$

yielding M equally spaced frequencies \tilde{f}_k in the interval $(0, 0.5)$. By local smoothness

$$\mathbf{S}_x(\tilde{f}_{k,\ell}) = \mathbf{S}_x(\tilde{f}_k) \text{ for } \ell = -m_t, -m_t + 1, \dots, m_t, \quad (2)$$

$$\text{where } \tilde{f}_{k,\ell} = \frac{(k-1)K + m_t + 1 + \ell}{n}. \quad (3)$$

It is known ([17, Theorem 4.4.1]) that asymptotically (as $n \rightarrow \infty$), $\mathbf{d}_x(f_m)$, $m = 1, 2, \dots, (n/2) - 1$, (n even), are independent proper, complex Gaussian $\mathcal{N}_c(\mathbf{0}, \mathbf{S}_x(f_m))$ random vectors, respectively, provided all elements of $\mathbf{R}_{xx}(\tau)$ are absolutely summable. Denote the joint probability density function of $\mathbf{d}_x(f_m)$, $m = 1, 2, \dots, (n/2) - 1$, as $f_D(\mathbf{D})$.

We wish to estimate inverse PSD matrix $\Phi_k := \mathbf{S}_x^{-1}(\tilde{f}_k)$. In terms of Φ_k , we have the log-likelihood [13] (up to some constant)

$$\ln f_D(\mathbf{D}) \propto -G(\{\Phi\}, \{\Phi^*\}) \quad (4)$$

$$:= \sum_{k=1}^M \frac{1}{2} \left[(\ln |\Phi_k| + \ln |\Phi_k^*|) - \text{tr} \left(\hat{\mathbf{S}}_k \Phi_k + \hat{\mathbf{S}}_k^* \Phi_k^* \right) \right] \quad (5)$$

where the PSD estimator using unweighted frequency-domain smoothing is $\hat{\mathbf{S}}_k = (1/K) \sum_{\ell=-m_t}^{m_t} \mathbf{d}_x(\tilde{f}_{k,\ell}) \mathbf{d}_x^H(\tilde{f}_{k,\ell})$. In the high-dimension case of $K < p(p-1)/2$ (# of unknowns in $\mathbf{S}_x^{-1}(\tilde{f}_k)$), one may need to use penalty terms to enforce sparsity and to make the problem well-conditioned. Imposing a sparse-group sparsity constraint (cf. [4, 18, 19]), [13] minimizes a penalized version of negative log-likelihood w.r.t. $\{\Phi\}$

$$L_{SGL}(\{\Phi\}) = G(\{\Phi\}, \{\Phi^*\}) + \bar{P}(\{\Phi\}), \quad (6)$$

$$\bar{P}(\{\Phi\}) = \bar{\lambda}_1 \sum_{k=1}^M \sum_{i \neq j}^p \left| [\Phi_k]_{ij} \right| + \bar{\lambda}_2 \sum_{i \neq j}^p \|\Phi^{(ij)}\| \quad (7)$$

$$\text{where } \Phi^{(ij)} := [[\Phi_1]_{ij} \ [\Phi_2]_{ij} \ \dots \ [\Phi_M]_{ij}]^T \in \mathbb{C}^M \quad (8)$$

and $\bar{\lambda}_1, \bar{\lambda}_2 \geq 0$ are tuning parameters. An analysis of the properties of the minimizer $\{\Phi\}$ of $L_{SGL}(\{\Phi\})$ is given in [14]. In [14], $\bar{\lambda}_1 = \alpha\lambda$ and $\bar{\lambda}_2 = (1 - \alpha)\lambda$ with $\lambda > 0$, and $\alpha \in [0, 1]$ providing a convex combination of lasso and group-lasso penalties [18, 19].

3. PROPOSED SPARSE-GROUP LOG-SUM PENALIZED NEGATIVE LOG-LIKELIHOOD

With $0 < \epsilon \ll 1$, following [16], define the log penalty for $\theta \in \mathbb{R}$,

$$p_\lambda(\theta) = \lambda \ln(1 + |\theta|/\epsilon). \quad (9)$$

Replace $\bar{P}(\{\Phi\})$ in (7) with $P(\{\Phi\})$, defined as

$$P(\{\Phi\}) = \sum_{k=1}^M \sum_{i \neq j}^p p_{\bar{\lambda}_1}(|[\Phi_k]_{ij}|) + \sum_{i \neq j}^p p_{\bar{\lambda}_2}(\|\Phi^{(ij)}\|). \quad (10)$$

Now replace (6) with (11)

$$L_{LSP}(\{\Phi\}) = G(\{\Phi\}, \{\Phi^*\}) + P(\{\Phi\}), \quad (11)$$

to define the sparse-group log-sum penalized log-likelihood function. Unlike $L_{SGL}(\{\Phi\})$, we now have a non-convex function of $\{\Phi\}$ in $L_{LSP}(\{\Phi\})$.

As for the SCAD (smoothly clipped absolute deviation) penalty in [20], we solve the problem $\min_{\Phi_k \succ \mathbf{0}, k=1, \dots, M} L_{LSP}(\{\Phi\})$ iteratively, where in each iteration, the problem is convex. Using $\partial p_\lambda(|\theta|)/\partial|\theta| = \lambda/(|\theta| + \epsilon)$, a local linear approximation to $p_\lambda(|\theta|)$ around θ_0 yields

$$p_\lambda(|\theta|) \approx P_\lambda(|\theta_0|) + \frac{\lambda}{|\theta_0| + \epsilon} (|\theta| - |\theta_0|) \Rightarrow \frac{\lambda}{|\theta_0| + \epsilon} |\theta|, \quad (12)$$

therefore, with θ_0 fixed, we consider only the last term above for optimization w.r.t. θ . Suppose we have a ‘‘good’’ initial solution $\{\Phi\}$ to the problem (from e.g., using $L_{SGL}(\{\Phi\})$ instead of $L_{LSP}(\{\Phi\})$). Then, given $\{\Phi\}$, a local linear approximation to $P(\{\Phi\})$ yields the convex function to be minimized

$$\tilde{L}_{LSP}(\{\Phi\}) = G(\{\Phi\}, \{\Phi^*\}) + \tilde{P}(\{\Phi\}), \quad (13)$$

$$\tilde{P}(\{\Phi\}) = \sum_{k=1}^M \sum_{i \neq j}^p \lambda_{1ij} \left| [\Phi_k]_{ij} \right| + \sum_{i \neq j}^p \lambda_{2ij} \|\Phi^{(ij)}\| \quad (14)$$

$$\lambda_{1ij} = \bar{\lambda}_1 / (|[\Phi_k]_{ij}| + \epsilon), \quad \lambda_{2ij} = \bar{\lambda}_2 / (\|\Phi^{(ij)}\| + \epsilon). \quad (15)$$

This is then quite similar to adaptive lasso [21] or adaptive sparse group lasso [22], except that [21] and [22] have $\epsilon = 0$. If we initialize with $[\Phi_k]_{ij} = 0$ for all $i \neq j$, we obtain an SGL cost.

4. OPTIMIZATION

Minimization of $L_{LSP}(\{\Phi\})$ is done iteratively where in each iteration, we minimize $\tilde{L}_{LSP}(\{\Phi\})$. Initially we take $\lambda_{1ij} = \bar{\lambda}_1$ and $\lambda_{2ij} = \bar{\lambda}_2 \forall (i, j)$. The result of this iteration is then used as $\{\Phi\}$ in (15) for next iteration. To optimize $\tilde{L}_{LSP}(\{\Phi\})$, using variable splitting, we first reformulate as in [13]:

$$\min_{\{\Phi\}, \{\mathbf{W}\}} \left\{ G(\{\Phi\}, \{\Phi^*\}) + P(\{\mathbf{W}\}) \right\} \quad (16)$$

subject to $\mathbf{W}_k = \Phi_k \succ \mathbf{0}$, $k = 1, 2, \dots, M$, where $\{\Phi\} = \{\Phi_k, k = 1, 2, \dots, M\}$ and $\{\mathbf{W}\} = \{\mathbf{W}_k, k = 1, 2, \dots, M\}$. In ADMM, we consider the scaled augmented Lagrangian for this problem [4, 23]

$$L_\rho(\{\Phi\}, \{\mathbf{W}\}, \{\mathbf{U}\}) = G(\{\Phi\}, \{\Phi^*\}) + \tilde{P}(\{\mathbf{W}\}) + \frac{\rho}{2} \sum_{k=1}^M \|\Phi_k - \mathbf{W}_k + \mathbf{U}_k\|_F^2 \quad (17)$$

where $\{\mathbf{U}\} = \{\mathbf{U}_k, k = 1, 2, \dots, M\}$ are dual variables, and $\rho > 0$ is the ‘‘penalty parameter’’ [23].

4.1. ADMM Algorithm

Given the results $\{\Phi^{(m)}\}, \{\mathbf{W}^{(m)}\}, \{\mathbf{U}^{(m)}\}$ of the m th iteration, in the $(m+1)$ st iteration, an ADMM algorithm executes the following three updates:

- $\{\Phi^{(m+1)}\} \leftarrow \arg \min_{\{\Phi\}} L_\rho(\{\Phi\}, \{\mathbf{W}^{(m)}\}, \{\mathbf{U}^{(m)}\})$
- $\{\mathbf{W}^{(m+1)}\} \leftarrow \arg \min_{\{\mathbf{W}\}} L_\rho(\{\Phi^{(m+1)}\}, \{\mathbf{W}\}, \{\mathbf{U}^{(m)}\})$
- $\{\mathbf{U}^{(m+1)}\} \leftarrow \{\mathbf{U}^{(m)}\} + \left(\{\Phi^{(m+1)}\} - \{\mathbf{W}^{(m+1)}\} \right)$

Update (a): Notice that up to some terms not dependent upon Φ_k 's, $L_\rho(\{\Phi\}, \{\mathbf{W}^{(m)}\}, \{\mathbf{U}^{(m)}\}) = \sum_{k=1}^M \frac{1}{2} L_{\rho k}(\Phi_k, \mathbf{W}_k^{(m)}, \mathbf{U}_k^{(m)})$ (i.e., it is separable in k), where

$$L_{\rho k}(\Phi_k, \mathbf{W}_k^{(m)}, \mathbf{U}_k^{(m)}) := -\ln |\Phi_k| - \ln |\Phi_k^*| + \text{tr} \left(\hat{\mathbf{S}}_k \Phi_k + \hat{\mathbf{S}}_k^* \Phi_k^* \right) + \rho \|\Phi_k - \mathbf{W}_k^{(m)} + \mathbf{U}_k^{(m)}\|_F^2. \quad (18)$$

The solution to $\arg \min_{\Phi_k} L_{\rho k}(\Phi_k, \mathbf{W}_k^{(m)}, \mathbf{U}_k^{(m)})$ is as follows [23, Sec. 6.5], [4, 13]. Let $\mathbf{V}\mathbf{D}\mathbf{V}^H$ denote the eigen-decomposition of the matrix $(\hat{\mathbf{S}}_k - \rho\mathbf{W}_k^{(m)} + \rho\mathbf{U}_k^{(m)})$. Then $\Phi_k^{(m+1)} = \mathbf{V}\tilde{\mathbf{D}}\mathbf{V}^H$ where $\tilde{\mathbf{D}}$ is the diagonal matrix with ℓ th diagonal element

$$\tilde{D}_{\ell\ell} = \frac{1}{2\rho} \left(-D_{\ell\ell} + \sqrt{|D_{\ell\ell}|^2 + 4\rho} \right) \quad (19)$$

By construction $\tilde{D}_{\ell\ell} > 0$ for any $\rho > 0$, hence, $\Phi_k^{(m+1)} = \mathbf{V}\tilde{\mathbf{D}}\mathbf{V}^H > \mathbf{0}$.

Update (b): Update $\{\mathbf{W}_k^{(m+1)}\}_{k=1}^M$ as the minimizer of

$$\frac{\rho}{2} \sum_{k=1}^M \|\mathbf{W}_k - (\Phi_k^{(m+1)} + \mathbf{U}_k^{(m)})\|_F^2 + \tilde{P}(\{\mathbf{W}\}) \quad (20)$$

w.r.t. $\{\mathbf{W}\}_{k=1}^M$. The solution follows from [13] (which follows real-valued results of [18]). Define $(b)_+ := \max(0, b)$, soft-thresholding operator $S(b, \beta) := (1 - \beta/|b|)_+ b$, and vector operator $[S(\mathbf{a}, \beta)]_j = S(a_j, \beta)$, $a_j = [\mathbf{a}]_j$. Let $\mathbf{A}_k = \Phi_k^{(m+1)} + \mathbf{U}_k^{(m)}$. The solution to minimization of (20) is

$$[\hat{\mathbf{W}}_k]_{ij} = \begin{cases} [\mathbf{A}_k]_{ij}, & \text{if } i = j \\ S([\mathbf{A}_k]_{ij}, \frac{\lambda_{1ij}}{\rho}) \left(1 - \frac{\lambda_{2ij}}{\rho \|S(\mathbf{A}_k^{(ij)}, \lambda_{1ij}/\rho)\|} \right)_+, & \text{if } i \neq j \end{cases}$$

Update (c): For the scaled Lagrangian formulation of ADMM [23], for $k = 1, 2, \dots, M$, update $\mathbf{U}_k^{(m+1)} = \mathbf{U}_k^{(m)} + (\Phi_k^{(m+1)} - \mathbf{W}_k^{(m+1)})$.

Edge Selection: Denote the converged estimates as $\hat{\Phi}_k$, $k = 1, \dots, M$. If $\|\hat{\Phi}^{(ij)}\| > 0$ then $\{i, j\} \in \mathcal{E}$, else $\{i, j\} \notin \mathcal{E}$, $i \neq j$.

4.2. BIC for selection of λ and α

Given n , K and M , the Bayesian information criterion (BIC) is given by $\text{BIC}(\lambda, \alpha) = 2K \sum_{k=1}^M \left(-\ln |\hat{\Phi}_k| + \text{tr}(\hat{\mathbf{S}}_k \hat{\Phi}_k) \right) + \ln(2KM) \sum_{k=1}^M (\# \text{ of nonzero elements in } \hat{\Phi}_k)$ where $2KM$ are total number of real-valued measurements in frequency-domain and $2K$ are number of real-valued measurements per frequency point, with total M frequencies in $(0, \pi)$. Pick α and λ to minimize BIC. We use BIC to first select λ over a grid of values with fixed α , and then with selected λ , we search over α values in $[0, 0.3]$. We search over λ in the range $[\lambda_\ell, \lambda_u]$ selected via a heuristic as in [22]. For $\alpha = \alpha_0 (=0.1)$, find the smallest λ , labeled λ_{sm} , for which we get a no-edge model; then we set $\lambda_u = \lambda_{sm}/2$ and $\lambda_\ell = \lambda_u/10$.

5. CONSISTENCY

Define $p \times (pM)$ matrix Ω as

$$\Omega = [\Phi_1 \ \Phi_2 \ \dots \ \Phi_M], \quad (21)$$

and denote $L_{LSP}(\Omega) = L_{LSP}(\{\Phi\})$, given by (15). Set $\bar{\lambda}_1 = \alpha\lambda$ and $\bar{\lambda}_2 = (1 - \alpha)\lambda$, with $0 \leq \alpha \leq 1$. We now allow p , M , K (see (1)), and λ to be functions of sample size n , denoted as p_n , M_n , K_n and λ_n , respectively. Note that $K_n M_n \approx n/2$. Pick $K_n = a_1 n^\gamma$ and $M_n = a_2 n^{1-\gamma}$ for some $0.5 < \gamma < 1$ so that $M_n/K_n \rightarrow 0$ as $n \rightarrow \infty$. Assume

(A1) The p -dimensional time series $\{\mathbf{x}(t)\}_{t=-\infty}^\infty$ is zero-mean stationary, Gaussian, satisfying $\sum_{\tau=-\infty}^\infty \|[\mathbf{R}_{xx}(\tau)]_{k\ell}\| < \infty$ for every $k, \ell \in V$.

(A2) Denote the true edge set of the graph by \mathcal{E}_0 , implying that $\mathcal{E}_0 = \{\{i, j\} : [\mathbf{S}_0^{-1}(f)]_{ij} \neq 0, i \neq j, 0 \leq f \leq 0.5\}$ where $\mathbf{S}_0(f)$ denotes the true PSD of $\mathbf{x}(t)$. (We also use Φ_{0k} for $\mathbf{S}_0^{-1}(\tilde{f}_k)$ where \tilde{f}_k is as in (1), and use Ω_0 to denote the true value of Ω). Assume that $\text{card}(\mathcal{E}_0) = |(\mathcal{E}_0)| \leq s_{n0}$.

(A3) The minimum and maximum eigenvalues of $p_n \times p_n$ PSD $\mathbf{S}_0(f) > \mathbf{0}$ satisfy $0 < \beta_{\min} \leq \min_{f \in [0, 0.5]} \phi_{\min}(\mathbf{S}_0(f))$ and $\max_{f \in [0, 0.5]} \phi_{\max}(\mathbf{S}_0(f)) \leq \beta_{\max} < \infty$. Here β_{\min} and β_{\max} are not functions of n (or p_n).

Let $\hat{\Omega}_\lambda = \arg \min_{\Omega : \Phi_k > \mathbf{0}} L_{LSP}(\Omega)$. Theorem 1 whose proof is omitted for lack of space, establishes local consistency of $\hat{\Omega}_\lambda$.

Theorem 1 (Consistency). For $\tau > 2$, let

$$C_0 = 80 \max_{\ell, f} ([\mathbf{S}_0(f)]_{\ell\ell}) \sqrt{\frac{2 \ln(16p_n^\tau M_n)}{\ln(p_n)}}. \quad (22)$$

Given any real numbers $\delta_1 \in (0, 1)$, $\delta_2 > 0$ and $C_1 > 1$, let

$$R = C_2 C_0 / \beta_{\min}^2, \quad C_2 = 2(2 + C_1 + \delta_2)(1 + \delta_1)^2, \quad (23)$$

$$r_n = \sqrt{\frac{M_n(p_n + s_{n0}) \ln(p_n)}{K_n}}, \quad C_2 r_n = o(1), \quad (24)$$

$$N_1 = \arg \min \{n : K_n > 2 \ln(16p_n^\tau M_n)\}, \quad (25)$$

$$N_2 = \arg \min \left\{ n : r_n \leq \max \left(\frac{\delta_1}{\beta_{\min} R}, \frac{\epsilon(C_1 - 1)}{R} \right) \right\}. \quad (26)$$

Suppose the regularization parameter λ_n and $\alpha \in [0, 1]$ satisfy

$$\begin{aligned} C_0 C_1 \sqrt{\frac{\ln(p_n)}{K_n}} &\leq \frac{\lambda_n}{\epsilon \sqrt{M_n}} \\ &\leq \frac{C_1 C_0}{1 + \alpha(\sqrt{M_n} - 1)} \sqrt{\left(1 + \frac{p_n}{s_{n0}} \right) \frac{\ln(p_n)}{K_n}}. \end{aligned} \quad (27)$$

Under assumptions (A1)-(A3), for sample size $n > \max\{N_1, N_2\}$, there exists a local minimizer $\hat{\Omega}_\lambda$ such that

$$\|\hat{\Omega}_\lambda - \Omega_0\|_F \leq R r_n \quad (28)$$

with probability greater than $1 - 1/p_n^{\tau-2}$. A sufficient condition for the lower bound in (27) to be less than the upper bound for every $\alpha \in [0, 1]$ is $C_1 = 2(1 + \alpha(\sqrt{M_n} - 1))$. In terms of rate of convergence, $\|\hat{\Omega}_\lambda - \Omega_0\|_F = \mathcal{O}_P(C_1 r_n)$.

Remark 1. Proof of Theorem 1 is patterned after [14, Theorem 1] pertaining to the SGL approach; in turn, [14] exploits [24, 25]. Convergence in [14, Theorem 1] is global while here it is only local since LSP is non-convex. Assumption (A1) is needed to invoke [17, Theorem 4.4.1] to establish asymptotic statistics of DFT $\mathbf{d}_x(f_m)$'s; it is needed in [14, Theorem 1] also where it was not explicitly stated. Assumptions (A2) and (A3) here are assumptions (A1) and (A2), respectively, in [14]. We use $C_1 > 1$ to bound $\ln(1 + x) \geq x/C_1$ for $0 \leq x \leq C_1 - 1$. \square

6. NUMERICAL EXAMPLES

All ADMM approaches used variable penalty parameter ρ as in [23, Sec. 3.4.1], and the stopping (convergence) criterion following [23, Sec. 3.3.1]. We picked $\epsilon = 0.0001$ in LSP.

Synthetic Data: Consider $p = 128$, 16 clusters (communities) of 8 nodes each, where nodes within a community are not connected to any nodes in other communities. Within any community of 8

nodes, the data are generated using a vector autoregressive (VAR) model of order 3. Consider community q , $q = 1, 2, \dots, 16$. Then $\mathbf{x}^{(q)}(t) \in \mathbb{R}^8$ is generated as

$$\mathbf{x}^{(q)}(t) = \sum_{i=1}^3 \mathbf{A}_i^{(q)} \mathbf{x}^{(q)}(t-i) + \mathbf{w}^{(q)}(t)$$

with $\mathbf{w}^{(q)}(t)$ as i.i.d. zero-mean Gaussian with identity covariance matrix. Only 10% of entries of $\mathbf{A}_i^{(q)}$'s are nonzero and the nonzero elements are independently and uniformly distributed over $[-0.8, 0.8]$. We then check if the VAR(3) model is stable with all eigenvalues of the companion matrix ≤ 0.95 in magnitude; if not, we re-draw randomly till this condition is fulfilled. The overall data $\mathbf{x}(t)$ is given by $\mathbf{x}(t) = [\mathbf{x}^{(1)\top}(t) \dots \mathbf{x}^{(16)\top}(t)]^\top \in \mathbb{R}^p$. First 100 samples are discarded to eliminate transients. This set-up leads to approximately 3.5% connected edges.

Simulation results based on 100 runs are shown in Fig. 1. We used $M = 4$ for all samples sizes $n = 128, 256, 512, 1024, 2048$ with corresponding $K = 15, 31, 63, 127, 255$, respectively. The performance measure is F_1 -score for efficacy in edge detection. The F_1 -score is defined as $F_1 = 2 \times \text{precision} \times \text{recall} / (\text{precision} + \text{recall})$ where precision = $|\hat{\mathcal{E}} \cap \mathcal{E}_0| / |\hat{\mathcal{E}}|$, recall = $|\hat{\mathcal{E}} \cap \mathcal{E}_0| / |\mathcal{E}_0|$, and \mathcal{E}_0 and $\hat{\mathcal{E}}$ denote the true and estimated edge sets, respectively. Six approaches were tested: (i) Proposed sparse-group log-sum penalty based approach, labeled ‘‘SGLSP’’ in Fig. 1. (ii) Sparse-group lasso (labeled ‘‘SGL’’) of [13, 14] optimized using ADMM as in this paper (this approach also initializes SGLSP). (iii) Sparse-group SCAD-penalized method (labeled ‘‘SGSCAD’’) optimized using ADMM, where instead of LSP we use the non-convex SCAD penalty (see [20] for SCAD). (iv) An i.i.d. modeling approach that exploits only the sample covariance $\frac{1}{n} \sum_{t=0}^{n-1} \mathbf{x}(t)\mathbf{x}^\top(t)$ (labeled ‘‘IID’’), implemented via the ADMM (adaptive) lasso approach ([23, Sec. 6.4]). In this approach edge $\{i, j\}$ exists in the CIG iff $\Omega_{ij} \neq 0$ where precision matrix $\Omega = \mathbf{R}_{xx}^{-1}(0)$. (v) The ADMM approach of [12], labeled ‘‘GMS’’ (graphical model selection), which was applied with $F = 4$ (four frequency points, corresponds to our $M = 4$) and all other default settings of [12] to compute the PSDs. The tuning parameters, (α, λ) for SGLSP, SGL and SGSCAD, and lasso parameter λ for IID and GMS, were selected via an exhaustive search over a grid of values to maximize the F_1 -score (which requires knowledge of the true edge-set). The results shown in Fig. 1 are based on these optimized tuning parameters. (vi) The sixth approach labeled ‘‘SGLSP with BIC’’ in Fig. 1, denotes our proposed ADMM approach for SGLSP where (α, λ) were selected via BIC, not requiring knowledge of the ground-truth.

It is seen that with F_1 -score as the performance metric, our proposed SGLSP method significantly outperforms other approaches. The BIC-based SGLSP method yields performance that is close to that based on optimized parameter selection. SGLSCAD offers very little improvement over SGL, in contrast with proposed SGLSP.

Real data: Financial Time Series: We consider daily share prices (at close of the day) of 97 stocks in S&P 100 index from Jan. 1, 2013 through Jan. 1, 2018, yielding 1259 samples. If $y_m(t)$ is share price of m th stock on day t , we consider (as is conventional in such studies) $x_m(t) = \ln(y_m(t)/y_m(t-1))$ as the time series to analyze, yielding $n = 1258$ and $p = 97$. These 97 stocks are classified into 11 sectors and we order the nodes to group them as information technology (nodes 1-12), health care (13-27), financials (28-44), real estate (45-46), consumer discretionary (47-56), industrials (57-68), communication services (69-76), consumer staples (77-87), energy (88-92), materials (93), utilities (94-97). The weighted adjacency

matrices resulting from the IID modeling approach (estimated $|\Omega_{ij}|$ is the edge weight), and the proposed SGLSP (with BIC) approach with $M = 4$, ($\sqrt{\sum_{k=1}^M |[\hat{\Phi}_k]_{ij}|^2}$ is the edge weight), are shown in Fig. 2. In both cases we used BIC to determine the tuning parameters. While the ground truth is unknown, the dependent time series based proposed approach yields sparser, more interpretable CIG (321 edges for the proposed approach versus 579 edges for IID modeling) which also conforms better with the sector classification according to the Global Industry Classification Standard.

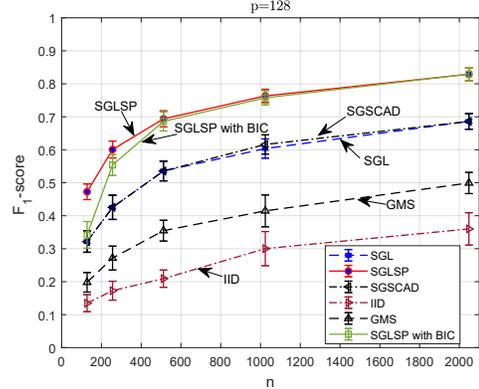
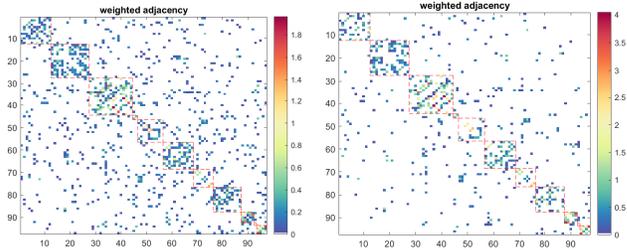


Fig. 1: F_1 -scores based on 100 runs for 6 approaches. SGLSP: proposed sparse-group log-sum penalty approach; SGL: sparse-group lasso [13, 14]; SGLSCAD: sparse-group SCAD-penalized method optimized using ADMM, where LSP is replaced with the non-convex SCAD penalty [20]; IID: exploits only $\mathbf{R}_{xx}(0)$; GMS: method of [12]; SGLSP with BIC: proposed SGLSP with tuning parameters selected via BIC.



(a) Estimated $|\Omega_{ij}|$ as edge weight; 579 edges. (b) Estimated $\sqrt{\sum_{k=1}^M |[\hat{\Phi}_k]_{ij}|^2}$ as edge weight; 321 edges.

Fig. 2: Weighted adjacency matrices for financial time series. The red squares (in dashed lines) show the 11 sectors – they are not part of the adjacency matrices.

7. CONCLUSIONS

Sparse-group lasso penalized log-likelihood approach in frequency-domain has been considered in [13, 14] for graph learning for dependent time series. In this paper we considered a sparse-group log-sum penalty instead of the sparse-group lasso (SGL) penalty to regularize the problem. An ADMM approach for iterative optimization of the non-convex problem was presented. We provided sufficient conditions for consistency of a local estimator of inverse PSD. We illustrated our approach using numerical examples utilizing both synthetic and real data. Synthetic data example showed that our SGLSP approach significantly outperformed the SGL and other approaches in correctly detecting the graph edges with F_1 -score as performance metric.

8. REFERENCES

- [1] S.L. Lauritzen, *Graphical models*. Oxford, UK: Oxford Univ. Press, 1996.
- [2] R. Dahlhaus, "Graphical interaction models for multivariate time series," *Metrika*, vol. 51, pp. 157-172, 2000.
- [3] M. Eichler, "Graphical modelling of multivariate time series," *Probability Theory and Related Fields*, vol. 153, issue 1-2, pp. 233-268, June 2012.
- [4] P. Danaher, P. Wang and D.M. Witten, "The joint graphical lasso for inverse covariance estimation across multiple classes," *J. Royal Statistical Society, Series B (Methodological)*, vol. 76, pp. 373-397, 2014.
- [5] N. Friedman, "Inferring cellular networks using probabilistic graphical models," *Science*, vol 303, pp. 799-805, 2004.
- [6] S.L. Lauritzen and N.A. Sheehan, "Graphical models for genetic analyses," *Statistical Science*, vol. 18, pp. 489-514, 2003.
- [7] N. Meinshausen and P. Bühlmann, "High-dimensional graphs and variable selection with the Lasso," *Ann. Statist.*, vol. 34, no. 3, pp. 1436-1462, 2006.
- [8] K. Mohan, P. London, M. Fazel, D. Witten and S.I. Lee, "Node-based learning of multiple Gaussian graphical models," *J. Machine Learning Research*, vol. 15, pp. 445-488, 2014.
- [9] D.R. Brillinger, "Remarks concerning graphical models of times series and point processes," *Revista de Econometria (Brazilian Rev. Econometr.)*, vol. 16, pp. 1-23, 1996.
- [10] A. Jung, R. Heckel, H. Bölcskei, and F. Hlawatsch, "Compressive nonparametric graphical model selection for time series," in *Proc. IEEE ICASSP-2014*, Florence, Italy, May 2014.
- [11] A. Jung, "Learning the conditional independence structure of stationary time series: A multitask learning approach," *IEEE Trans. Signal Process.*, vol. 63, no. 21, pp. 5677-5690, Nov. 1, 2015.
- [12] A. Jung, G. Hannak and N. Goertz, "Graphical LASSO based model selection for time series," *IEEE Signal Process. Lett.*, vol. 22, no. 10, pp. 1781-1785, Oct. 2015.
- [13] J.K. Tugnait, "Graphical modeling of high-dimensional time series," in *Proc. 52nd Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, Oct. 29 - Oct. 31, 2018, pp. 840-844.
- [14] J.K. Tugnait, "Consistency of sparse-group lasso graphical model selection for time series," in *Proc. 54th Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, Nov. 1-4, 2020, pp. 589-593.
- [15] J.K. Tugnait, "Edge exclusion tests for graphical model selection: Complex Gaussian vectors and time series," *IEEE Trans. Signal Process.*, vol. 67, no. 19, pp. 5062-5077, Oct. 1, 2019.
- [16] E.J. Candès, M.B. Wakin and S.P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J. Fourier Anal. Appl.*, vol. 14, pp. 877-905, 2008.
- [17] D.R. Brillinger, *Time Series: Data Analysis and Theory*, Expanded edition. New York: McGraw Hill, 1981.
- [18] J. Friedman, T. Hastie and R. Tibshirani, "A note on the group lasso and a sparse group lasso," *arXiv:1001.0736v1 [math.ST]*, 5 Jan 2010.
- [19] N. Simon, J. Friedman, T. Hastie and R. Tibshirani, "A sparse-group lasso," *J. Computational Graphical Statistics*, vol. 22, pp. 231-245, 2013.
- [20] C. Lam and J. Fan, "Sparsistency and rates of convergence in large covariance matrix estimation," *Ann. Statist.*, vol. 37, no. 6B, pp. 4254-4278, 2009.
- [21] H. Zou, "The adaptive lasso and its oracle properties," *J. American Statistical Assoc.*, vol. 101, pp. 1418-1429, 2006.
- [22] J.K. Tugnait, "Sparse-group lasso for graph learning from multi-attribute data," *IEEE Trans. Signal Process.*, vol. 69, pp. 1771-1786, 2021. (Corrections: vol. 69, p. 4758, 2021.)
- [23] S. Boyd, N. Parikh, E. Chu, B. Peleato and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1-122, 2010.
- [24] A.J. Rothman, P.J. Bickel, E. Levina and J. Zhu, "Sparse permutation invariant covariance estimation," *Electronic J. Statistics*, vol. 2, pp. 494-515, 2008.
- [25] J.K. Tugnait, "On sparse complex Gaussian graphical model selection," in *Proc. 2019 IEEE Intern. Workshop on Machine Learning for Signal Processing (MLSP 2019)*, Pittsburgh, PA, Oct. 13-16, 2019, pp. 1-6.