High-dimensional structure learning of sparse vector autoregressive models using fractional marginal pseudo-likelihood

Kimmo Suotsalo¹ · Yingying Xu² · Jukka Corander^{3,4,5} · Johan Pensar⁶

Received: 2 November 2020 / Accepted: 22 August 2021 / Published online: 17 September 2021
© The Author(s) 2021

Abstract

Learning vector autoregressive models from multivariate time series is conventionally approached through least squares or maximum likelihood estimation. These methods typically assume a fully connected model which provides no direct insight to the model structure and may lead to highly noisy estimates of the parameters. Because of these limitations, there has been an increasing interest towards methods that produce sparse estimates through penalized regression. However, such methods are computationally intensive and may become prohibitively time-consuming when the number of variables in the model increases. In this paper we adopt an approximate Bayesian approach to the learning problem by combining fractional marginal likelihood and pseudo-likelihood. We propose a novel method, PLVAR, that is both faster and produces more accurate estimates than the state-of-the-art methods based on penalized regression. We prove the consistency of the PLVAR estimator and demonstrate the attractive performance of the method on both simulated and real-world data.

Keywords Vector autoregression · Pseudo-likelihood · Fractional marginal likelihood · Gaussian graphical models · Multivariate time series

1 Introduction

Vector autoregressive (VAR) models (Lütkepohl 2005; Brockwell and Davis 2016; Neusser 2016) have become standard tools in many fields of science and engineering. They are used in, for example, economics (Ang and Piazzesi 2003; Ito and Sato 2008; Zang and Baimbridge 2012), psychology (Wild et al. 2010; Bringmann et al. 2013; Epskamp et al. 2018), and sustainable energy technology (Dowell and Pinson 2016;

Cavalcante et al. 2017; Zhao et al. 2018). In neuroscience, VAR models have been widely applied in brain connectivity analysis with data collected through functional magnetic resonance imaging (Baccalá and Sameshima 2001; Harrison et al. 2003; Roebroeck et al. 2005), magnetoencephalography (Tsiaras et al. 2011; Michalareas et al. 2013; Fukushima et al. 2015), or electroencephalography (EEG) (Supp et al. 2007; Gómez-Herrero et al. 2008; Chiang et al. 2009).

The conventional approach to learning a VAR model is to utilize either multivariate least squares (LS) estimation or maximum likelihood (ML) estimation (Lütkepohl 2005). Both of these methods produce dense model structures where each variable interacts with all the other variables. Some of these interactions may be weak but it is typically impossible to tell which variables are directly related to each other, making the estimation results difficult to interpret. Yet, the interpretation of the model dynamics is a goal of itself in many applications, such as reconstructing genetic networks from gene expression data (Abegaz and Wit 2013), or identifying functional connectivity between brain regions (Chiang et al. 2009).

Estimating a dense VAR model also suffers from another problem, namely, the large number of model parameters may lead to noisy estimates and give rise to unstable predictions

✉ Kimmo Suotsalo
kimmo.suotsalo@aalto.fi

¹ Aalto University, Espoo, Finland

² RIKEN, Center for Advanced Intelligence Project
AIP/iTHEMS, Wako, Japan

³ Department of Biostatistics, Faculty of Medicine, University
of Oslo, Oslo, Norway

⁴ Parasites and Microbes, Wellcome Sanger Institute,
Cambridge, UK

⁵ Helsinki Institute of Information Technology, Department of
Mathematics and Statistics, University of Helsinki, Helsinki,
Finland

⁶ Department of Mathematics, University of Oslo, Oslo,
Norway

(Davis et al. 2016). In a dense VAR model, the number of parameters grows linearly with respect to the lag length and quadratically with respect to the number of variables. The exact number of parameters in such a model equals $kd^2 + d(d+1)/2$ where k denotes the lag length and d is the number of variables.

Due to the reasons mentioned above, methods that provide sparse estimates of VAR model structures have received a considerable amount of attention in the literature (Valdés-Sosa et al. 2005; Arnold et al. 2007; Hsu et al. 2008; Haufe et al. 2010; Bańbura et al. 2010; Melnyk and Banerjee 2016). A particularly popular group of methods is based on penalized regression where a penalty term is included to regularize the sparsity in the estimated model. Different penalty measures have been investigated (Valdés-Sosa et al. 2005; Hsu et al. 2008; Haufe et al. 2010), including the L_1 norm, the L_2 norm, hard-thresholding, smoothly clipped absolute deviation (SCAD) (Fan and Li 2001), and mixture penalty. Of these, using the L_2 norm is also known as ridge regression, and using the L_1 norm as the least absolute shrinkage and selection operator (LASSO) (Tibshirani 1996). While ridge regression penalizes strong connections between the variables, it is incapable of setting them to exactly zero. In contrast, LASSO and SCAD have the ability to force some of the parameter values to exactly zero. Hence, these two methods are suitable for variable selection and, in the context of VAR models, learning a sparse model structure while simultaneously estimating its parameters.

Despite their usefulness in many applications, LASSO and SCAD also have certain shortcomings with respect to VAR model estimation. LASSO is likely to give inaccurate estimates of the model structure when the sample size is small (Arnold et al. 2007; Melnyk and Banerjee 2016). In particular, the structure estimates obtained by LASSO tend to be too dense (Haufe et al. 2010). While SCAD has been reported to perform better (Abegaz and Wit 2013), it still leaves room for improvement in many cases. In addition, both LASSO and SCAD involve one or more hyperparameters whose value may have a significant effect on the results. Suitable hyperparameter values may be found through cross-validation but this is typically a very time-consuming task.

In order to overcome the challenges related to LASSO and SCAD, we propose a novel method for learning VAR model structures. We refer to this method as pseudo-likelihood vector autoregression (PLVAR) because it utilizes a pseudo-likelihood based scoring function. The PLVAR method works in two steps: it first learns the temporal structure of the VAR model and then proceeds to estimate the contemporaneous structure. The PLVAR method is shown to yield highly accurate results even with small sample sizes. Importantly, it is also very fast compared with the other methods.

The contributions of this paper are summarised as follows. (1) We extend the idea of learning the structure of Gaussian

graphical models via pseudo-likelihood and Bayes factors (Leppä-aho et al. 2017; Pensar et al. 2017) from the cross-sectional domain into the time-series domain of VAR models. The resulting PLVAR method is able to infer the complete VAR model structure, including the lag length, in a single run. (2) We show that the PLVAR method produces a consistent structure estimator in the class of VAR models. (3) We present an iterative maximum likelihood procedure for inferring the model parameters for a given sparse VAR structure, providing a complete VAR model learning pipeline. (4) We demonstrate through experiments that the proposed PLVAR method is more accurate and much faster than the current state-of-the-art methods.

The organization of this paper is as follows. The first section provides an introduction to VAR models and their role in the recent literature. The second section describes the PLVAR method, a consistency result for the method, and finally a procedure for estimating the model parameters. The third section presents the experiments and results, and the fourth section discusses the findings and makes some concluding remarks. Appendices A and B provide a search algorithm and a consistency proof, respectively, for the PLVAR estimator.

1.1 VAR models

In its simplest form, a VAR model can be written as

$$\mathbf{y}_t = \sum_{m=1}^k \mathbf{A}_m \mathbf{y}_{t-m} + \boldsymbol{\epsilon}_t \quad (1)$$

where $\mathbf{y}_t$ is the vector of current observations, $\mathbf{A}_m$ are the lag matrices, $\mathbf{y}_{t-m}$ are the past observations, and $\boldsymbol{\epsilon}_t$ is a random error. The observations are assumed to be made in d dimensions. The lag length k determines the number of past observations influencing the distribution of the current time step in the model. In principle, the random error can be assumed to follow any distribution, but it is often modeled using a multivariate normal distribution with no dependencies between the time steps. In this paper we always assume the error terms to be independent and identically distributed (i.i.d.) and that $\boldsymbol{\epsilon}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma})$.

It is also possible to express a VAR model as

$$\mathbf{y}_t = \sum_{m=1}^k \mathbf{A}_m \mathbf{y}_{t-m} + \mathbf{b}t + \mathbf{c} + \boldsymbol{\epsilon}_t \quad (2)$$

where $\mathbf{b}t$ accounts for a linear time trend and $\mathbf{c}$ is a constant. However, for time series that are (at least locally) stationary, one can always start by fitting a linear model to the data. This gives estimates of $\mathbf{b}$ and $\mathbf{c}$, denoted as $\hat{\mathbf{b}}$ and $\hat{\mathbf{c}}$, respectively. The value $\hat{\mathbf{b}}t + \hat{\mathbf{c}}$ can then be subtracted from each $\mathbf{y}_t$, thereby reducing the model back to (1).

A distinctive characteristic of VAR models is their linearity. The fact that the models are linear brings both advantages and drawbacks. One of the biggest advantages is the simplicity of the models that allows closed-form analytical derivations, especially when the error terms are i.i.d. and Gaussian. The simplicity of the models also makes them easy to interpret: it is straightforward to see which past observations have effect on the current observations. A clear drawback of VAR models is that linear mappings may fail to fully capture the complex relationships between the variables and the observations. In practice, however, VAR models have shown to yield useful results in many cases: some illustrative examples can be found in Harrison et al. (2003), Gómez-Herrero et al. (2008), Wild et al. (2010), and Michalareas et al. (2013).

Relationship with state space models VAR models are also related to linear state space models. A VAR(k) model of the form (1) can always be written (Lütkepohl 2005) as an equivalent VAR(1) model

$$\mathbf{Y}_t = \mathbf{D}\mathbf{Y}_{t-1} + \mathbf{E}_t \quad (3)$$

where

$$\mathbf{Y}_t = \begin{bmatrix} \mathbf{y}_t \\ \mathbf{y}_{t-1} \\ \vdots \\ \mathbf{y}_{t-k+1} \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \cdots & \mathbf{A}_{k-1} & \mathbf{A}_k \\ \mathbf{I} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & & \mathbf{0} & \mathbf{0} \\ \vdots & & \ddots & & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{I} & \mathbf{0} \end{bmatrix},$$

$$\mathbf{E}_t = \begin{bmatrix} \boldsymbol{\epsilon}_t \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}. \quad (4)$$

As Lütkepohl (2005) points out, (3) can be seen as the transition equation of the state space model

$$\mathbf{x}_t = \mathbf{G}_t \mathbf{x}_{t-1} + \mathbf{u}_t \quad (5)$$

$$\mathbf{y}_t = \mathbf{H}_t \mathbf{x}_t + \mathbf{v}_t \quad (6)$$

where $\mathbf{H}_t = \mathbf{I}$ and $\mathbf{v}_t = \mathbf{0}$. This can be interpreted as the system state being equal to the observed values, or equivalently, as a case where one does not include a hidden state in the model. While a hidden state is a natural component in certain applications, its estimation requires computationally expensive methods such as Kalman filters. Also, in some cases, the transition equation may not be known or is not backed up with a solid theory. In such cases VAR modeling allows the estimation and the analysis of the relationships between the current and the past observations. Although the VAR modeling approach concentrates on the observations rather than the underlying system that generates them, it can still pro-

vide highly useful results such as implications of Granger causality or accurate predictions of future observations.

1.2 Graphical VAR modeling

Similar to Dahlhaus and Eichler (2003), this paper utilizes a time series chain (TSC) graph representation for modeling the structure of the VAR process. Let $Y = \{Y_t^i : t \in \mathbb{Z}, i \in V\}$, where $V = \{1, \dots, d\}$, represent a stationary d -dimensional process. For any $S \subseteq V$, let Y^S denote the subprocess given by the indices in S and let $\bar{Y}_t^S = \{Y_u^S : u < t\}$ denote the history of the subprocess at time t . In the TSC graph representation by Dahlhaus and Eichler (2003), the variable Y_t^i at a specific time step is represented by a separate vertex in the graph. More specifically, the TSC graph is denoted by $G_{TS} = (V_{TS}, E_{TS})$, where $V_{TS} = V \times \mathbb{Z}$ is a time-step-specific node set, and E_{TS} is a set of directed and undirected edges such that

$$(a, t-u) \rightarrow (b, t) \notin E_{TS} \Leftrightarrow u \leq 0 \text{ or } Y_{t-u}^a \perp Y_t^b \mid \bar{Y}_t \setminus \{Y_{t-u}^a\} \quad (7)$$

and

$$(a, t-u) \leftrightarrow (b, t) \notin E_{TS} \Leftrightarrow u \neq 0 \text{ or } Y_t^a \perp Y_t^b \mid \bar{Y}_t \cup \{Y_t^{V \setminus \{a,b\}}\}. \quad (8)$$

As noted in Dahlhaus and Eichler (2003), conditions (7) and (8) imply that the process adheres to the pairwise AMP Markov property for G_{TS} (Andersson et al. 2001).

In terms of a VAR process of form (1) with $\boldsymbol{\epsilon}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma})$, the conditional independence relations in (7) and (8) correspond to simple restrictions on the lag matrices $\mathbf{A}_m$, $m = 1, \dots, k$, and the precision matrix $\boldsymbol{\Omega} = \boldsymbol{\Sigma}^{-1}$ (Dahlhaus and Eichler 2003). More specifically, we have that

$$(a, t-u) \rightarrow (b, t) \in E_{TS} \Leftrightarrow u \in \{1, \dots, k\} \text{ and } \mathbf{A}_u(b, a) \neq 0$$

and

$$(a, t) \leftrightarrow (b, t) \in E_{TS} \Leftrightarrow \boldsymbol{\Omega}(a, b) \neq 0.$$

Thus, finding the directed and undirected edges in G_{TS} is equivalent to finding the non-zero elements of the lag matrices $\mathbf{A}_m$ and the non-zero off-diagonal elements of the precision matrix $\boldsymbol{\Omega}$, respectively.

As a concrete example of the above relationships, consider the TSC-graph in Fig. 1a which represents the dependence structure of the following sparse VAR(2) process:

$$\mathbf{A}_1 = \begin{bmatrix} 0.3 & 0 & 0 & 0 \\ -0.2 & 0.2 & 0 & 0 \\ 0 & 0 & -0.3 & 0 \\ 0 & 0 & 0.2 & -0.2 \end{bmatrix}, \quad \mathbf{A}_2 = \begin{bmatrix} 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.1 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

Fig. 1 **a** TSC graph and **b** corresponding GVAR structure of a sparse VAR(2) process

$$\Omega = \begin{bmatrix} 1 & 0 & 0.2 & 0 \\ 0 & 1 & 0 & 0 \\ 0.2 & 0 & 1 & 0.2 \\ 0 & 0 & 0.2 & 1 \end{bmatrix}.$$

To compactly encode G_{TS} , which essentially is a graph over all time steps, we will use the subgraphs $G_d \triangleq (V_d, E_d)$ and $G_u \triangleq (V_u, E_u)$. The temporal graph G_d , with respect to the reference time step t , contains the directed edges from the lagged node sets $V_{t-1}, \dots, V_{t-k}$ to V_t . The contemporaneous graph G_u contains the undirected edges between the nodes in the same time step; $V_u = V_t$. Finally, following the graphical VAR modeling framework of Epskamp et al. (2018), we refer to the pair $G \triangleq (G_d, G_u)$ as a graphical VAR (GVAR) structure. For an example of a GVAR structure, see Fig. 1b.

2 Methods

The PLVAR method proposed in this paper utilizes a two-step approach for learning the GVAR model structure. Each step uses a score-based approach, where a candidate structure is scored according to how well it explains the dependence structure observed in the data. The proposed scoring function is based on a Bayesian framework for model selection. More specifically, PLVAR is based on the concept of objective

Bayes factors for Gaussian graphical models, an approach that has been used successfully in the cross-sectional setting for learning the structure of decomposable undirected models (Carvalho and Scott 2009), directed acyclic graph (DAG) models (Consonni and Rocca 2012), and general undirected models (Leppä-aho et al. 2017). More recently, this technique has also been used for learning the contemporaneous structure G_u of a graphical VAR model under the assumption that the graph is chordal (decomposable) (Paci and Consonni 2020). By utilizing the Bayesian scoring function, we here develop an algorithm that infers the complete model structure (G_d, G_u) , where G_u is allowed to be non-chordal. The method requires very little input from the user and is proven to be consistent in the large sample limit.

Given a scoring function, we still need an efficient search algorithm for finding the optimal structure. To enable scalability to high-dimensional systems, we propose a search algorithm that exploits the structure of the considered scoring function via a divide-and-conquer type approach. More specifically, the initial global structure learning problem is split up into multiple node-wise sub-problems whose solutions are combined into a final global structure. Since the number of possible structures in each sub-problem is still too vast for finding the global optimum for even moderate values of d and k , a greedy search algorithm is used to find a local maximum in a reasonable time.

2.1 Bayesian model selection

The scoring function of the PLVAR method is derived using a Bayesian approach. Ideally, we are interested in finding a GVAR structure G that is close to optimal in the sense of maximum a posteriori (MAP) estimation. By the Bayes' rule,

$$p(G|\mathbf{Y}) = \frac{p(\mathbf{Y}|G)p(G)}{\sum p(\mathbf{Y}|G)p(G)} \quad (9)$$

where $\mathbf{Y} \triangleq [\mathbf{y}_1 \dots \mathbf{y}_N]^\top$ is the observation matrix, and the sum in the denominator is taken over all the possible GVAR structures. Thus, the MAP estimate is given by

$$\hat{G}_{MAP} = \arg \max p(\mathbf{Y}|G)p(G). \quad (10)$$

A key part of the above score is the marginal likelihood $p(\mathbf{Y}|G)$ which is obtained by integrating out the model parameters θ in the likelihood $p(\mathbf{Y}|G, \theta)$ under some prior on the parameters.

Objective Bayesian model selection poses a problem in the specification of the parameter prior, as uninformative priors are typically improper. A solution to this is given by the fractional Bayes factor (FBF) approach (O'Hagan 1995), where a fraction of the likelihood is used for updating the improper uninformative prior into a proper fractional prior. The proper

fractional prior is then used with the remaining part of likelihood to give the FBF. It has been shown that the FBF provides better robustness against misspecified priors compared to the use of proper priors in a situation where only vague information is available about the modeled system (O'Hagan 1995).

In this work, we make use of the objective FBF approach for Gaussian graphical models (Carvalho and Scott 2009; Consonni and Rocca 2012; Leppä-aho et al. 2017; Paci and Consonni 2020). In particular, our procedure is inspired by the fractional marginal pseudo-likelihood (FMPL) (Leppä-aho et al. 2017), which allows for non-chordal undirected graphs when inferring the contemporaneous part of the network.

2.2 Fractional marginal pseudo-likelihood

In terms of Gaussian graphical models, the FBF was introduced for decomposable undirected models by Carvalho and Scott (2009), and extended to DAG models by Consonni and Rocca (2012), who referred to the resulting score as the fractional marginal likelihood (FML). Finally, Leppä-aho et al. (2017) extended the applicability of the FBF approach to the general class of non-decomposable undirected models by combining it with the pseudo-likelihood approximation (Besag 1975). Our procedure is inspired by the FMPL and can be thought of as an extension of the approach to time-series data.

As stated in Leppä-aho et al. (2017), the FMPL of a Gaussian graphical model is given by

$$pl(\mathbf{Y}|\mathbf{G}) = \prod_{i=1}^d p(\mathbf{Y}_i|\mathbf{Y}_{mb(i)}) \\ = \prod_{i=1}^d \pi^{-\frac{N-1}{2}} \frac{\Gamma\left(\frac{N+p_i}{2}\right)}{\Gamma\left(\frac{p_i+1}{2}\right)} N^{-\frac{2p_i+1}{2}} \left(\frac{|\mathbf{S}_{fa}(i)|}{|\mathbf{S}_{mb}(i)|}\right)^{-\frac{N-1}{2}}, \quad (11)$$

where $mb(i)$ denotes the Markov blanket of i , p_i denotes the number of nodes in $mb(i)$, and $fa(i) = mb(i) \cup i$ is the family of i . Moreover, N denotes the number of observations, $\mathbf{S} = \mathbf{Y}^\top \mathbf{Y}$, and $\mathbf{S}_{fa}(i)$ and $\mathbf{S}_{mb}(i)$ denote the submatrices of $\mathbf{S}$ restricted to the nodes that belong to the family and Markov blanket, respectively, of node i . The FMPL is well-defined if $N-1 \geq \max_i(p_i)$, since it ensures (with probability one) positive definiteness of $\mathbf{S}_{fa}(i)$ for each $i = 1, \dots, d$. Thus, we assume that this (rather reasonable) sparsity condition holds in the following.

Whereas Leppä-aho et al. (2017) is solely concerned with the learning of undirected Gaussian graphical models, we are interested in learning a GVAR structure which is made up of both a directed temporal part G_d and an undirected contemporaneous part G_u . Rather than inferring these simul-

taneously, we break up the problem into two consecutive learning problems that both can be approached using the FMPL. First, in terms of G_d , we apply a modified version of (11) where the Markov blankets are built from nodes in earlier time steps. In terms of G_u , we use the inferred directed structure $\hat{G}_d$ to account for the temporal dependence between the observations. As this reduces the problem of inferring G_u to the cross-sectional setting considered in Leppä-aho et al. (2017), the score in (11) can be applied as such.

2.3 Learning the temporal network structure

In order to learn the temporal part of the GVAR structure, we begin by creating a lagged data matrix

$$\mathbf{Z} \triangleq [\mathbf{Y}_0 \mathbf{Y}_{-1} \cdots \mathbf{Y}_{-k}] \quad (12)$$

where $\mathbf{Y}_{-m} = [\mathbf{y}_{k-m+1} \cdots \mathbf{y}_{N-m}]^\top$. The lagged data matrix thus contains $(k+1)d$ columns, each corresponding to a time-specific variable; more specifically, the variable of the m :th time series at lag l corresponds to column $ld+m$. As a result of the missing lag information in the beginning of the time-series, the effective number of observations (rows) is reduced from N to $N-k$; this, however, poses no problems as long as N is clearly larger than k .

Now, the problem of inferring the temporal structure can be thought of as follows. For each node i , we want to identify, among the lagged variables, a minimal set of nodes that will shield node i from the remaining lagged variables, that is, a type of temporal Markov blanket. As seen in (7), the temporal Markov blanket will equal the set of lagged variables from which there is a directed edge to node i , commonly referred to as the parents of i . Assuming a structure prior that factorizes over the Markov blankets and a given lag length k , we can frame this problem as the following optimization problem based on the lagged data matrix and a modified version of (11):

$$\arg \max_{mb(i)} \prod_{i=1}^d p(\mathbf{Z}_i|\mathbf{Z}_{mb(i)}, k) p(mb(i)|k) \quad (13)$$

subject to $mb(i) \subseteq \{d+1, \dots, (k+1)d\}$,

where the so-called local FMPL is here defined as

$$p(\mathbf{Z}_i|\mathbf{Z}_{mb(i)}, k) \triangleq \pi^{-\frac{N-k-1}{2}} \frac{\Gamma\left(\frac{N-k+p_i}{2}\right)}{\Gamma\left(\frac{p_i+1}{2}\right)} (N-k)^{-\frac{2p_i+1}{2}} \left(\frac{|\mathbf{S}_{fa}(i)|}{|\mathbf{S}_{mb}(i)|}\right)^{-\frac{N-k-1}{2}}. \quad (14)$$

The unscaled covariance matrix in (14) is now obtained from the lagged data matrix $\mathbf{Z}$, that is, $\mathbf{S} = \mathbf{Z}^\top \mathbf{Z}$.

The only thing left to specify in (13) is the structure prior. To further promote sparsity in the estimated network, we draw inspiration from the extended Bayesian information criterion (eBIC; see, for example, (Barber and Drton 2015)) and define our prior as

$$p(G_d|k) = \prod_{i=1}^d p(\text{mb}(i)|k) = \prod_{i=1}^d (kd)^{-\gamma p_i}, \quad (15)$$

where p_i is the number of nodes in the Markov blanket of node i and γ is a tuning parameter adjusting the strength of the sparsity promoting effect. As the default value for the tuning parameter, we use $\gamma = 0.5$. As seen later, this prior will be essential for selecting the lag length.

From a computational point of view, problem (13) is very convenient in the sense that it is composed of d separate sub-problems that can be solved independently of each other. To this end, we apply the same greedy search algorithm as in Pensar et al. (2017), Leppä-aho et al. (2017).

2.4 Search algorithm

The PLVAR method uses a greedy search algorithm to find the set $\text{mb}(i)$ (with $p_i \leq N - 1$) that maximizes the local FMPL for each $i \in V$. The search algorithm used by the PLVAR method has been published as Algorithm 1 in Leppä-aho et al. (2017) and originally introduced as Algorithm 1 in Pensar et al. (2017), based on Tsamardinos et al. (2006). This algorithm works in two interleaving steps, deterministically adding and removing individual nodes in the set $\text{mb}(i)$. At each step, the node $j \notin \text{mb}(i)$ that yields the largest increase in the local FMPL is added to or removed from $\text{mb}(i)$. This is a hybrid approach aimed for low computational cost without losing much of the accuracy of an exhaustive search (Pensar et al. 2017). Pseudo-code for the algorithm is given in “Appendix A” section.

2.5 Selecting the lag length

As noted in Sect. 2.3, the FMPL of a graphical VAR model depends on the lag length k . Unfortunately, the true value of k is usually unknown in real-world problems. For the PLVAR method, we adopt the approach of setting an upper bound K and learning the temporal structure separately for each $k = 1, \dots, K$. The value of K should preferably be based on knowledge about the problem domain, e.g. by considering the time difference between the observations and the maximum time a variable can influence itself or the other variables.

Once the upper bound K has been set, a score is calculated for each k . Typical scores used for this purpose include the Bayesian information criterion and the Akaike information criterion, both of which are to be minimized. For the

PLVAR method, we eliminate the need for external scores by selecting the k that maximizes our FMPL-based objective function in (13). To ensure comparable data sets for different values on k , the time steps included in the lagged data matrices are determined by the maximum lag K . Note that our lag selection criterion is made possible by our choice of structure prior (15) which depends on k . As there is no restriction that a temporal Markov blanket should include nodes of all considered lags, the FMPL part of the objective function does not explicitly penalize overly long lag lengths.

2.6 Learning the contemporaneous network structure

Once the temporal network structure has been estimated, we proceed to learn the contemporaneous structure in three steps. First, the non-zero parameters of the lag matrices $\mathbf{A}_m$ are estimated via the ordinary least squares method (which is applicable under the sparsity condition $p_i \leq N - 1$): this is done separately for each node given its temporal Markov blanket (or parents) in the temporal network. Next, the fully estimated lag matrices $\hat{\mathbf{A}}_m$ are used to calculate the residuals $\hat{\epsilon}_t$ between the true and the predicted observations:

$$\hat{\epsilon}_t = \mathbf{y}_t - \sum_{m=1}^k \hat{\mathbf{A}}_m \mathbf{y}_{t-m}. \quad (16)$$

If the estimates $\hat{\mathbf{A}}_m$ are accurate, the residuals $\hat{\epsilon}_t$ will form a set of approximately cross-sectional data, generated by the contemporaneous part of the model, and the problem is reduced to the setting of the original FMPL method. Thus, as our objective function, we use the standard FMPL (11) in combination with our eBIC type prior:

$$p(G_u) = \prod_{i=1}^d p(\text{mb}(i)) = \prod_{i=1}^d (d-1)^{-\gamma p_i}, \quad (17)$$

The only difference to (15) is the number of possible Markov blanket members which is now $d - 1$. Again, as the default value for the tuning parameter, we use $\gamma = 0.5$.

Although the Markov blankets are now connected through the symmetry constraint $i \in \text{mb}(j) \Leftrightarrow j \in \text{mb}(i)$, we use the same divide-and-conquer search approach as was used for the temporal part of the network. However, since the search may then result in asymmetries in the Markov blankets, we use the OR-criterion as described in Leppä-aho et al. (2017) when constructing the final graph: if $i \in \text{mb}(j)$ or $j \in \text{mb}(i)$, then $(i, j) \in E_u$.

2.7 Consistency of PLVAR

A particularly desirable characteristic for any statistical estimator is consistency. This naturally also applies when learning the GVAR structure. Importantly, it turns out that the proposed PLVAR method enjoys consistency. More specifically, as the sample size tends to infinity, PLVAR will recover the true GVAR structure, including the true lag length k^* , if the maximum considered lag length K is larger than or equal to k^* . A brief outline of the rationale is provided in this subsection, and a formal proof of the key steps is provided in “Appendix B” section.

The authors of Leppä-aho et al. (2017) show that their FMPL score in combination with the greedy search algorithm is consistent for finding the Markov blankets in Gaussian graphical models. Assuming a correct temporal structure, the lag matrices can be consistently estimated by the LS method (Lütkepohl 2005), and the residuals will thus tend towards i.i.d. samples from $N(\mathbf{0}, \Sigma)$, where the sparsity pattern of $\Omega = \Sigma^{-1}$ corresponds to the contemporaneous structure. This is the setting considered in Leppä-aho et al. (2017), for which consistency has been proven.

What remains to be shown is that the first step of PLVAR is consistent in terms of recovering the temporal structure. Since PLVAR uses a variant of FMPL where the set $mb(i)$ is selected from lagged variables, it can be proven, using a similar approach as in Leppä-aho et al. (2017), that PLVAR will recover the true temporal Markov blankets (or parents) in the large sample limit if $K \geq k^*$ (see “Appendix B” section for more details). Finally, since the structure prior (15) is designed to favor shorter lag lengths, it will in combination with the FMPL ensure that the estimated lag length will equal k^* in the large sample limit.

As opposed to LASSO and SCAD, PLVAR needs no conditions on the measurement matrix and relation to true signal strength for model selection consistency (Zhao and Yu 2006; Bühlmann 2013; Huang and Xie 2007). The reason for this lies in the difference between the approaches. In LASSO and SCAD, the regularization term introduces bias to the cost function, thereby deviating the estimation from the true signal (Zhao and Yu 2006; Bühlmann 2013). On the other hand, PLVAR estimates the graph structure using the FMPL score: this is a Bayesian approach and thus has no bias in asymptotic behavior (Consonni and Rocca 2012; Leppä-aho et al. 2017).

2.8 Estimating the remaining parameters

Once the PLVAR method has been used to infer the model structure, the remaining parameters can be estimated as described in this subsection. This combination provides a complete pipeline for the learning of graphical VAR models.

The parameters that still remain unknown at this point are the non-zero elements of the lag matrices $\mathbf{A}_m$, $m = 1, \dots, k$ and the non-zero elements of the precision matrix Ω . We calculate ML estimates for these parameters iteratively until a convergence criterion is met.

In each iteration, we first estimate the remaining elements of $\mathbf{A}_m$ given the current precision matrix. As initial value for the precision matrix, we use the identity matrix. In order to enforce the sparsity pattern implied by the temporal network, we use the ML estimation procedure described in Lütkepohl (2005). Next, we calculate the residuals between the true and predicted observations. We then estimate the remaining elements of Ω from the residuals while enforcing the sparsity pattern learned earlier for the contemporaneous network: this is done by applying the ML estimation approach outlined in Hastie et al. (2009).

The iterations are repeated until the absolute difference between the current and the previous maximized log-likelihood of Ω is smaller than a threshold δ .

3 Experiments and results

In order to assess the performance of the PLVAR method, we made a number of experiments on both synthetic and real-world data. For the synthetic data, we used a single node in the Puhti supercomputer of CSC – IT Center for Science, Finland, equipped with 4 GB RAM and an Intel® Xeon® Gold 6230 processor with 20 physical cores at 2.10 GHz base frequency. For the real-world data, we used a laptop workstation equipped with 8 GB RAM and an Intel® Core™ i7-6700HQ CPU with 4 physical cores at 2.60 GHz base frequency. The algorithms were implemented in the R language. For the sparsity-constrained ML estimation of Ω , we used the implementation in the R package mixggm (Fop et al. 2019). The threshold δ for determining the convergence of the parameter estimation was set to 10^{-6} .

3.1 Synthetic data

The performance of the PLVAR method was first evaluated in a controlled setting on data generated from known models with a ground truth structure. Having access to the true network structure, the quality of the inferred network structures were assessed by *precision*: the proportion of true edges among the inferred edges, and *recall*: the proportion of detected edges out of all true edges. Precision and recall were calculated separately for the temporal and the contemporaneous part of the network.

The R package `SparseTSCGM`¹ Abegaz and Wit (2013) was used to generate GVAR models with a random network structure and lag length $k = 2$. We performed two experiments. In the first experiment, we fixed the degree of sparsity and varied the number of variables, $d \in \{20, 40, 80\}$. In the second experiment, we fixed the number of variables to $d = 40$ and varied the degree of sparsity. The degree of sparsity in the generated networks is controlled by an edge inclusion probability parameter, called `prob0`. In order to have similar adjacency patterns for different number of variables, we specified the edge inclusion parameter as $q/(kd)$, meaning that the average indegree in the temporal structure is q and the average number of neighbors in the contemporaneous structure is $q/2$. In the first experiment, we set $q = 3$, resulting in rather sparse structures, and in the second experiment, we let $q \in \{5, 7, 9\}$, resulting in increasingly denser structures. For each considered configuration of d and q , we generated 20 stable GVAR models. Finally, a single time series over 800 time points was generated from each model, and the N first observations, with N varying between 50 and 800, were used to form the final collection of data sets.

For comparison, LASSO and SCAD were also applied to the generated data, using the implementations available in `SparseTSCGM`. While LASSO and SCAD were given the correct lag length $k = 2$, PLVAR was given a maximum allowed lag length $K = 5$ and thereby tasked with the additional challenge of estimating the correct lag length. In terms of the tuning parameters, the γ prior parameter in PLVAR was set to 0.5, and the remaining parameters of LASSO and SCAD were set according to the example provided in the `SparseTSCGM` manual. More specifically, the maximum number of outer and inner loop iterations were set to 10 and 100, respectively. The default grid of candidate values for the tuning parameters (`lam1` and `lam2`) were used, and the optimal values were selected using the modified Bayesian information criterion (`bic_mod`).

The precision, recall and running times of the considered methods are summarized in Fig. 2 for the first experiment and in Fig. 3 for the second experiment. For some of the most challenging settings: $N = 50$ for $d \in \{40, 80\}$ and $N = 100$ for $d = 80$, LASSO and SCAD exited with an error and the results are therefore missing. In terms of precision (Figs. 2a and 3a), PLVAR is clearly the most accurate of the methods, both for the temporal and the contemporaneous part of the network structure. In terms of recall (Figs. 2b and 3b), all methods perform very well for the larger samples, recovering close to all edges. For smaller sample sizes, PLVAR has a slight edge on both LASSO and SCAD in the sparser settings, $q = 3$ and $q = 5$, while it is overall quite even in the denser setting $q = 7$, and the situation is almost reversed for $q = 9$.

¹ More specifically, we used the function `sim.data()`; see Abegaz and Wit (2013) for more details regarding the simulator.

As expected, increasing the number of variables or the number of edges makes the learning problem harder, leading to a reduced accuracy. However, there is no drastic reduction in accuracy for any of the methods. The most drastic change between the settings is perhaps the computation time of LASSO and SCAD when the number of variables is increased (Fig. 2c). Overall, PLVAR was orders of magnitude faster than LASSO and SCAD, both of which use cross-validation to select the regularization parameters. However, it must be noted that PLVAR solves a simpler problem in that it focuses on learning the model structure (yet without a given lag length), while the likelihood-based methods also estimate the model parameters along with the structure. Still, the problem of estimating the model (parameters) is greatly facilitated once the model structure has been inferred.

Finally, when it comes to selecting the correct lag length ($k = 2$) from the candidate lag lengths, PLVAR was highly accurate and estimated the lag length to be higher than 2 for only a few cases in both the sparse and the denser settings (Fig. 4). It is also worth noting that while the maximum lag length supported by the PLVAR method is only limited by the available computing resources, the LASSO and SCAD implementations in `SparseTSCGM` currently only support lag lengths 1 and 2.

3.2 Electroencephalography data

Since VAR models have found extensive use in EEG analysis, the experiments made on synthetic data were complemented with experiments on real-world EEG data. The dataset chosen for this purpose is a subset of the publicly available CHB-MIT Scalp EEG Database (CHBMIT) (Shoeb 2009; Goldberger et al. 2000).

The CHBMIT database contains a total of 24 cases of EEG recordings, collected from 23 pediatric subjects at the Children's Hospital Boston. Each case comprises multiple recordings stored in European Data Format (.EDF) files. However, only the first recording of each unique subject was used in order to speed up the experiments. Case 24 was excluded since it has been added afterwards and is missing gender and age information. The recordings included in the experiments are listed in Table 1.

The recordings in the CHBMIT database have been collected using the international 10-20 system for the electrode placements. The channels derived from the electrodes vary between files and some channels only exist in a subset of the recordings. Hence, the experiments were made using only those 21 channels that can be found in all the recordings, implying $d = 21$ in the VAR models. Table 2 lists the channels included in the experiments.

All the recordings have a sampling rate of 256 Hz. The first 60 s were skipped in each recording because the signal-to-noise ratio in the beginning of some recordings is very

Fig. 2 Results of the simulation study with a fixed degree of sparsity ($q = 3$) and an increasing number of variables: **a** Precision and **b** recall (y-axis) for the temporal and contemporaneous part of the network structure for various sample sizes (x-axis) and model size, d . **c** Recorded computation times (y-axis in log scale) for learning the com-

plete network structure. Part of the results are missing for LASSO and SCAD since they exited with an error for some of the settings. The statistics shown in the plots have been calculated from a sample of 20 random GVAR models

Fig. 3 Results of the simulation study with a fixed number of variables ($d = 40$) and a decreasing degree of sparsity: **a** Precision and **b** recall (y-axis) for the temporal and contemporaneous part of the network structure for various sample sizes (x-axis) and degrees of sparsity, as controlled by q . **c** Recorded computation times (y-axis in log scale) for

learning the complete network structure. Part of the results are missing for LASSO and SCAD since they exited with an error when $N = 50$. The statistics shown in the plots have been calculated from a sample of 20 random GVAR models

Fig. 4 Estimated lag length by PLVAR for the different sample sizes in the simulation studies: **a** $q = 3$ and **b** $d = 40$. The x-axis represents the candidate lag lengths and the y-axis represents the percentage of times a specific lag length was chosen. The correct lag length is 2 in all cases

Table 1 Recordings included in the EEG experiments

Recording filename	Subject gender	Subject age (years)
chb01_01.edf	Female	11
chb02_01.edf	Male	11
chb03_01.edf	Female	14
chb04_01.edf	Male	22
chb05_01.edf	Female	7
chb06_01.edf	Female	1.5
chb07_01.edf	Female	14.5
chb08_02.edf	Male	3.5
chb09_01.edf	Female	10
chb10_01.edf	Male	3
chb11_01.edf	Female	12
chb12_06.edf	Female	2
chb13_02.edf	Female	3
chb14_01.edf	Female	9
chb15_02.edf	Male	16
chb16_01.edf	Female	7
chb17a_03.edf	Female	12
chb18_02.edf	Female	18
chb19_02.edf	Female	19
chb20_01.edf	Female	6
chb22_01.edf	Female	9
chb23_06.edf	Female	6

Table 2 Channels included in the EEG experiments

FP1-F7	F7-T7	T7-P7	P7-O1
FP1-F3	F3-C3	C3-P3	P3-O1
FP2-F4	F4-C4	C4-P4	P4-O2
FZ-CZ	CZ-PZ	T7-FT9	
FT9-FT10	FT10-T8		

low. The next 256 samples were used as training data and presented to each method for learning the model structure and the model parameters. The following 512 samples were used as test data to evaluate the performance of the methods.

Each experiment was made using 6 different algorithms. In addition to PLVAR, LASSO, and SCAD, conventional LS estimation was included for comparison. As noted before, the SparseTSCGM implementations of LASSO and SCAD only allow lag lengths 1 and 2. Therefore, the final set of algorithms comprised PLVAR, LS, PLVAR2, LS2, LASSO2, and SCAD2. In this set, algorithm names ending with 2 refer to the respective methods provided with lag length 2 for straightforward comparison. The names PLVAR and LS refer to respective methods with no a priori lag length information.

The performance of each algorithm was evaluated using the following metrics: the number of edges in the estimated temporal network (n_t), the number of edges in the esti-

Fig. 5 Results on EEG data. n_t , n_c , and MSE denote the number of edges in the temporal network, the number of edges in the contemporaneous network, and the mean squared error of 1-step predictions on the test data, respectively. Note the logarithmic y axis in both subplots

mated contemporaneous network (n_c), the mean squared error (MSE) of 1-step predictions on the test data, and the estimation time in seconds (t). The results are summarized in Fig. 5. An example of the estimated temporal and contemporaneous networks is given in Fig. 6 where the networks are visualized through their adjacency matrices.

Figure 5 shows that PLVAR and PLVAR2 produced the sparsest estimates for both the temporal network and the contemporaneous network. The dense estimates produced by LS and LS2 had an order of magnitude more edges than the estimates obtained through PLVAR and PLVAR2. The difference between the likelihood-based methods and the PLVAR methods is smaller but still evident. The MSEs on the test data were somewhat equal for all the methods except for LS and LS2. The MSEs obtained through LS2 were slightly higher than the MSEs obtained through the likelihood-based and the PLVAR methods. The MSEs produced by LS were about 100 times as high.

The computation times were minimal for LS and LS2 but this makes hardly a difference if the goal is to learn a sparse VAR model. The computation time required by PLVAR2 was an order of magnitude shorter than the time required by LASSO2 and SCAD2. The PLVAR method, tasked with the additional challenge of estimating the lag length, was slower than PLVAR2 but still faster than LASSO2 and SCAD2.

The example in Fig. 6 highlights the difference in sparsity of the model estimates obtained through PLVAR2 and the likelihood-based methods. In this particular case, it is evident that PLVAR2 produced the sparsest precision matrix and lag matrices, corresponding to the contemporaneous network and temporal network, respectively. One can find some com-

mon structures in all the precision matrices but the structures in the likelihood-based estimates are much denser.

4 Discussion

In this paper we have extended the idea of structure learning of Gaussian graphical models via pseudo-likelihood and Bayes factors from the cross-sectional domain into the time-series domain of VAR models. We have presented a novel method, PLVAR, that is able to infer the complete VAR model structure, including the lag length, in a single run. We have shown that the PLVAR method produces a consistent structure estimator in the class of VAR models. We have also presented an iterative maximum likelihood procedure for inferring the model parameters for a given sparse VAR structure, providing a complete VAR model learning pipeline.

The experimental results on synthetic data suggest that the PLVAR method is better suited for recovering sparse VAR structures than the available likelihood-based methods, a result that is in line with previous research (Leppä-aho et al. 2017; Pensar et al. 2017). The results on EEG data provide further evidence on the feasibility of the PLVAR method. Even though LASSO and SCAD reached quite similar MSEs, PLVAR was able to do so while producing much sparser models and using significantly less computation time. The conventional LS and LS2 methods were even faster but they are unable to learn sparse VAR models. Also, the MSEs produced by LS and LS2 were the highest among all the methods under consideration. All in all, the results show that PLVAR is a highly competitive method for high-dimensional VAR

Fig. 6 Adjacency matrix estimates, obtained from the EEG recording chb06_01.edf using the PLVAR2, the LASSO2, and the SCAD2 methods. The black and white pixels correspond to non-zero and zero elements in the matrices, respectively. Ω refers to the precision matrix, A_1 to the first lag matrix, and A_2 to the second lag matrix. Note the sparsity of the PLVAR2 estimates as compared to the other methods

structure learning, from both an accuracy point of view and a purely computational point of view.

It should be noted that the PLVAR implementation used in the experiments is purely sequential. It would be relatively straightforward to make the search algorithm run in parallel, thereby reducing the computation time up to $1/n_{\text{cpus}}$ where n_{cpus} is the number of CPU cores available on the system. While it is common for modern desktop and laptop workstations to be equipped with 4 to 8 physical cores, parallelization can be taken much further by utilizing cloud computing services. Several cloud vendors offer high-power virtual machines with tens of CPUs and the possibility of organizing them into computing clusters. This makes PLVAR

highly competitive against methods that cannot be easily parallelized, especially when the number of variables in the model is large.

In summary, PLVAR is an accurate and fast method for learning VAR models. It is able to produce sparse estimates of the model structure and thereby allows the model to be interpreted to gain additional knowledge in the problem domain.

In future work, we plan to apply the PLVAR method to domains that make extensive use of VAR modeling and where improvements in estimation accuracy and speed have the potential to enhance the end results significantly. An example of these domains is brain research where brain connectivity is measured in terms of Granger causality index, directed

transfer function, and partial directed coherence. Since these measures are derived from VAR models estimated from the data, one can expect more accurate estimates to result in more accurate measures of connectivity. Moreover, the enhanced possibilities in VAR model fitting introduced here allow more complex models to be studied. Further research with applications to other high-dimensional problems is also encouraged.

Supplementary material

The R code written for the experiments is publicly available in the repository <https://github.com/kimmo-suotsalo/plvar>. The EEG data can be freely obtained from PhysioNet (Goldberger et al. 2000) at <https://www.physionet.org/content/chbmit/1.0.0/>.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11222-021-10049-z>.

Acknowledgements This work was supported by RIKEN Special Postdoctoral Researcher Program (Yingying Xu). The authors wish to acknowledge CSC – IT Center for Science, Finland, for computational resources.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendices

A Pseudo-code for the search algorithm

The pseudo-code for the greedy search procedure is presented in Algorithm 1. The algorithm has been published previously in Leppä-aho et al. (2017) and was originally introduced in Pensar et al. (2017), based on Tsamardinos et al. (2006).

B Consistency proof

The below result states that the modified FMPL given in (13) is consistent for inferring the temporal Markov blankets in the considered model class. The proof is adapted from Leppä-

Algorithm 1 Greedy search for finding the temporal Markov blanket of node i

```

1: function FIND_MB( $i, \mathbf{Z}, k$ )
2:    $\widehat{mb}(i), \widehat{mb}(i) \leftarrow \emptyset$ 
3:   while  $\widehat{mb}(i)$  has changed AND  $|\widehat{mb}(i)| < N - 1$  do
4:      $C \leftarrow V \setminus \widehat{fa}(i)$ 
5:      $\triangleright$  above,  $V$  contains lagged variables from  $t - 1$  to  $t - k$ 
6:      $\widehat{mb}(i) \leftarrow \widehat{mb}(i)$ 
7:     for  $j \in C$  do
8:       if  $p(\mathbf{Z}_i | \mathbf{Z}_{\widehat{mb}(i) \cup \{j\}}, k) > p(\mathbf{Z}_i | \widehat{mb}(i), k)$  then
9:          $\widehat{mb}(i) \leftarrow \widehat{mb}(i) \cup \{j\}$ 
10:      end if
11:    end for
12:    while  $\widehat{mb}(i)$  has changed do
13:       $\widehat{mb}(i) \leftarrow \widehat{mb}(i)$ 
14:      for  $j \in \widehat{mb}(i)$  do
15:        if  $p(\mathbf{Z}_i | \mathbf{Z}_{\widehat{mb}(i) \setminus \{j\}}, k) > p(\mathbf{Z}_i | \widehat{mb}(i), k)$  then
16:           $\widehat{mb}(i) \leftarrow \widehat{mb}(i) \setminus \{j\}$ 
17:        end if
18:      end for
19:    end while
20:  end while
21:  return  $\widehat{mb}(i)$ 
22: end function

```

aho et al. (2017) where the authors use a similar reasoning for Gaussian graphical models with cross-sectional data.

Theorem 1 Define $\mathbf{Z}$ and $p(\mathbf{Z}_i | \mathbf{Z}_{mb(i)}, k)$ as in Eqs. (12) and (14), respectively. Let $S = \{d + 1, \dots, (k + 1)d\}$ contain the true temporal Markov blanket $mb(i)^*$ or, equivalently, let $k \geq k^*$ where k^* is the true lag length. Then,

$$\text{plim}_{N \rightarrow \infty} \arg \max_{mb(i) \subseteq S} p(\mathbf{Z}_i | \mathbf{Z}_{mb(i)}, k) = mb(i)^*. \quad (18)$$

Proof Under the current assumption of i.i.d. $N(\mathbf{0}, \Sigma)$ error terms, it can be shown that $\mathbf{y}_t$ is a Gaussian process, meaning that the subcollections $\mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-k}$ have a multivariate normal distribution (Lütkepohl 2005). Thus, $\mathbf{Z}$ can be considered a data matrix generated from a multivariate normal distribution over the variables $Z_1, \dots, Z_{(k+1)d}$ where the dependence structure of the model is determined by the GVAR structure. In particular, following from the directed Markov property in (7), the dependence structure will be such that

$$\mathbf{Z}_i \perp \mathbf{Z}_{mb(i)^*} \mid \mathbf{Z}_{S \setminus mb(i)^*},$$

where $i \in \{1, \dots, d\}$. Following a similar line of reasoning as in Leppä-aho et al. (2017), we can now show through Lemmas 1 and 2 that (18) holds. More specifically, Lemma 1 guarantees that as $N \rightarrow \infty$,

$$p(\mathbf{Z}_i | \mathbf{Z}_{mb(i)}, k) < p(\mathbf{Z}_i | \mathbf{Z}_{mb(i)^*}, k)$$

for any $\text{mb}(i)$ that does not contain all the nodes in $\text{mb}(i)^*$. Similarly, Lemma 2 guarantees that as $N \rightarrow \infty$,

$$p(\mathbf{Z}_i | \mathbf{Z}_{\text{mb}(i)}, k) < p(\mathbf{Z}_i | \mathbf{Z}_{\text{mb}(i)^*}, k)$$

for any $\text{mb}(i)$ that contains one or more nodes not in $\text{mb}(i)^*$. $\square$

The proofs of Lemmas 1 and 2 follow closely those of Leppä-aho et al. (2017). However, for completeness, we give here the proofs in detail.

Lemma 1 Let $\text{mb}(i) \subset \text{mb}(i)^*$ and $A \subset V \setminus \text{fa}(i)^*$. Then

$$\text{plim}_{N \rightarrow \infty} \log \frac{p(\mathbf{Z}_i | \mathbf{Z}_{\text{mb}(i)^* \cup A}, k)}{p(\mathbf{Z}_i | \mathbf{Z}_{\text{mb}(i) \cup A}, k)} = \infty. \quad (19)$$

Proof By defining $p_i \triangleq |\text{mb}(i)^* \cup A|$ and $a \triangleq |\text{mb}(i) \cup A| - p_i$, we have that

$$\begin{aligned} \log \frac{p(\mathbf{Z}_i | \mathbf{Z}_{\text{mb}(i)^* \cup A}, k)}{p(\mathbf{Z}_i | \mathbf{Z}_{\text{mb}(i) \cup A}, k)} &= \log \frac{\Gamma\left(\frac{N-k+p_i}{2}\right)}{\Gamma\left(\frac{N-k+p_i+a}{2}\right)} \\ &+ \log \frac{\Gamma\left(\frac{p_i+a+1}{2}\right)}{\Gamma\left(\frac{p_i+1}{2}\right)} + a \log(N-k) \\ &- \frac{N-k-1}{2} \log \frac{|\mathbf{S}_{\text{fa}(i)^*} \cup A| |\mathbf{S}_{\text{mb}(i)} \cup A|}{|\mathbf{S}_{\text{mb}(i)^*} \cup A| |\mathbf{S}_{\text{fa}(i)} \cup A|}. \end{aligned} \quad (20)$$

The first term on the right-hand side equals

$$\log \Gamma(M) - \log \Gamma\left(M + \frac{a}{2}\right) \quad (21)$$

where $M = (N - k + p_i)/2$. As $N \rightarrow \infty$, $M \rightarrow \infty$ and by Stirlings's formula

$$\begin{aligned} \log \Gamma(M) - \log \Gamma\left(M + \frac{a}{2}\right) &= \left(M - \frac{1}{2}\right) \log M - M \\ &- \left[\left(M + \frac{a}{2} - \frac{1}{2}\right) \log \left(M + \frac{a}{2}\right) - \left(M + \frac{a}{2}\right)\right] + O(1). \end{aligned} \quad (22)$$

Here, the right-hand side can be written as

$$M \log \frac{M}{M + \frac{a}{2}} + \frac{1}{2} \log \frac{M + \frac{a}{2}}{M} - \frac{a}{2} \log \left(M + \frac{a}{2}\right) + O(1) \quad (23)$$

where we have included the term $a/2$ in $O(1)$. Since

$$\begin{aligned} M \log \frac{M}{M + \frac{a}{2}} &= \frac{1}{2} \log \left(\frac{1}{1 + \frac{a}{2M}}\right)^{2M} \xrightarrow{M \rightarrow \infty} \\ \frac{1}{2} \log e^{-a} &= -\frac{a}{2} = O(1) \end{aligned} \quad (24)$$

and

$$\frac{1}{2} \log \frac{M + \frac{a}{2}}{M} = \frac{1}{2} \log \left(1 + \frac{a}{2M}\right) \xrightarrow{M \rightarrow \infty} 0, \quad (25)$$

we have that

$$\begin{aligned} \log \frac{\Gamma\left(\frac{N-k+p_i}{2}\right)}{\Gamma\left(\frac{N-k+p_i+a}{2}\right)} &\xrightarrow{N \rightarrow \infty} -\frac{a}{2} \log \left(\frac{N-k+p_i+a}{2}\right) \\ &+ O(1). \end{aligned} \quad (26)$$

The second term on the right-hand side of Eq. (20) can be included in $O(1)$ since it is constant with respect to N . The fourth term

$$-\frac{N-k-1}{2} \log \frac{|\mathbf{S}_{\text{fa}(i)^*} \cup A| |\mathbf{S}_{\text{mb}(i)} \cup A|}{|\mathbf{S}_{\text{mb}(i)^*} \cup A| |\mathbf{S}_{\text{fa}(i)} \cup A|} \quad (27)$$

can be expressed using the covariance Σ : since $\mathbf{S} = \mathbf{Z}^\top \mathbf{Z} = N \hat{\Sigma}_{\text{ML}}$, it holds that

$$\frac{|\mathbf{S}_{\text{fa}(i)^*} \cup A| |\mathbf{S}_{\text{mb}(i)} \cup A|}{|\mathbf{S}_{\text{mb}(i)^*} \cup A| |\mathbf{S}_{\text{fa}(i)} \cup A|} \xrightarrow{N \rightarrow \infty} \frac{|\Sigma_{\text{fa}(i)^*} \cup A| |\Sigma_{\text{mb}(i)} \cup A|}{|\Sigma_{\text{mb}(i)^*} \cup A| |\Sigma_{\text{fa}(i)} \cup A|} \quad (28)$$

in probability. Let us now partition $\Sigma_{\text{fa}(i)^*} \cup A$ as

$$\Sigma_{\text{fa}(i)^*} \cup A = \begin{bmatrix} \sigma^2(z_i) & \sigma(z_i, \mathbf{z}_{\text{mb}(i)^*} \cup A) \\ \sigma(z_i, \mathbf{z}_{\text{mb}(i)^*} \cup A)^\top & \Sigma_{\text{mb}(i)^*} \cup A \end{bmatrix} \quad (29)$$

where $\sigma(z_i, \mathbf{z}_{\text{mb}(i)^*} \cup A)$ denotes a row vector whose elements are the covariances between the variable z_i and each of the variables in $\text{mb}(i)^* \cup A$. The determinant of this partitioned matrix is

$$\begin{aligned} &|\Sigma_{\text{mb}(i)^*} \cup A| \cdot \\ &\left[\sigma^2(z_i) - \sigma(z_i, \mathbf{z}_{\text{mb}(i)^*} \cup A) (\Sigma_{\text{mb}(i)^*} \cup A)^{-1} \right. \\ &\left. \sigma(z_i, \mathbf{z}_{\text{mb}(i)^*} \cup A)^\top \right] \end{aligned} \quad (30)$$

where the term

$$\sigma(z_i, \mathbf{z}_{\text{mb}(i)^*} \cup A) (\Sigma_{\text{mb}(i)^*} \cup A)^{-1} \sigma(z_i, \mathbf{z}_{\text{mb}(i)^*} \cup A)^\top \triangleq \sigma(\hat{z}_i | \mathbf{z}_{\text{mb}(i)^*} \cup A) \quad (31)$$

is the variance of the linear LS predictor of z_i from $\mathbf{z}_{\text{mb}(i)^*} \cup A$. By definition, the residual variance of z_i after subtracting the variance of the linear LS predictor is the partial variance $\sigma(z_i | \mathbf{z}_{\text{mb}(i)^*} \cup A)$; hence,

$$|\Sigma_{\text{fa}(i)^*} \cup A| = |\Sigma_{\text{mb}(i)^*} \cup A| \sigma(z_i | \mathbf{z}_{\text{mb}(i)^*} \cup A) \quad (32)$$

which implies that

$$\frac{|\Sigma_{fa}(i)^* \cup A|}{|\Sigma_{mb}(i) \cup A|} \frac{|\Sigma_{mb}(i)^* \cup A|}{|\Sigma_{fa}(i) \cup A|} = \frac{\sigma(z_i | \Sigma_{mb}(i)^* \cup A)}{\sigma(z_i | \Sigma_{mb}(i) \cup A)}. \quad (33)$$

Combining the results so far, Eq. (20) can be rewritten as

$$\log \frac{p(\mathbf{Z}_i | \Sigma_{mb}(i)^* \cup A, k)}{p(\mathbf{Z}_i | \Sigma_{mb}(i) \cup A, k)} = -\frac{a}{2} \log \left(\frac{N-k+p_i+a}{2} \right) + a \log(N-k) - \frac{N-k-1}{2} \log \frac{\sigma(z_i | \Sigma_{mb}(i)^* \cup A)}{\sigma(z_i | \Sigma_{mb}(i) \cup A)} + O(1). \quad (34)$$

The second term on the right hand side equals

$$\frac{a}{2} \log(N-k)^2 = \frac{a}{2} \log(N-k) + \frac{a}{2} \log(N-k) \quad (35)$$

and therefore

$$\log \frac{p(\mathbf{Z}_i | \Sigma_{mb}(i)^* \cup A, k)}{p(\mathbf{Z}_i | \Sigma_{mb}(i) \cup A, k)} = \frac{a}{2} \log(N-k) - \frac{N}{2} \log \frac{\sigma(z_i | \Sigma_{mb}(i)^* \cup A)}{\sigma(z_i | \Sigma_{mb}(i) \cup A)} + O(1). \quad (36)$$

Now, consider the argument of the last logarithm term. Assume first that $\text{mb}(i) \neq \emptyset$ and let $B \triangleq \text{mb}(i)^* \setminus \text{mb}(i)$. Then

$$\begin{aligned} \sigma(\hat{z}_i[\Sigma_{mb}(i)^* \cup A]) &= \sigma(\hat{z}_i[\Sigma_{mb}(i) \cup A \cup B]) \\ &= \sigma(\hat{z}_i[\Sigma_{mb}(i) \cup A]) + \sigma(\hat{z}_i[\mathbf{Z}_B - \hat{\mathbf{Z}}_B[\Sigma_{mb}(i) \cup A]]) \end{aligned} \quad (37)$$

where the last term $\sigma(\hat{z}_i[\mathbf{Z}_B - \hat{\mathbf{Z}}_B[\Sigma_{mb}(i) \cup A]]) > 0$ since $B \subset \text{mb}(i)^*$. Hence,

$$\sigma(\hat{z}_i[\Sigma_{mb}(i)^* \cup A]) > \sigma(\hat{z}_i[\Sigma_{mb}(i) \cup A]) \quad (38)$$

which is equivalent to

$$\sigma^2(z_i) - \sigma(\hat{z}_i[\Sigma_{mb}(i)^* \cup A]) < \sigma^2(z_i) - \sigma(\hat{z}_i[\Sigma_{mb}(i) \cup A]). \quad (39)$$

This implies that

$$\sigma(z_i | \Sigma_{mb}(i)^* \cup A) < \sigma(z_i | \Sigma_{mb}(i) \cup A) \quad (40)$$

and therefore the last logarithm term

$$-\frac{N}{2} \log \frac{\sigma(z_i | \Sigma_{mb}(i)^* \cup A)}{\sigma(z_i | \Sigma_{mb}(i) \cup A)} \xrightarrow{N \rightarrow \infty} \infty. \quad (41)$$

Assume then that $\text{mb}(i) = \emptyset$. This implies that

$$\sigma(\hat{z}_i[\mathbf{Z}_A]) + \sigma(\hat{z}_i[\Sigma_{mb}(i)^* - \hat{\Sigma}_{mb}(i)^*[\mathbf{Z}_A]]) > \sigma(\hat{z}_i[\mathbf{Z}_A]) \quad (42)$$

which is equivalent to

$$\sigma(\hat{z}_i[\Sigma_{mb}(i)^* \cup A]) > \sigma(\hat{z}_i[\mathbf{Z}_A]) \quad (43)$$

and further to Inequation (38). Hence, the last logarithm term behaves exactly as for $\text{mb}(i) \neq \emptyset$.

Since $\text{mb}(i) \subset \text{mb}(i)^*$, $a < 0$, and the second-last logarithm term

$$\frac{a}{2} \log(N-k) \xrightarrow{N \rightarrow \infty} -\infty. \quad (44)$$

Combining all the results, we have that

$$\begin{aligned} \log \frac{p(\mathbf{Z}_i | \Sigma_{mb}(i)^* \cup A, k)}{p(\mathbf{Z}_i | \Sigma_{mb}(i) \cup A, k)} &= \frac{a}{2} \log(N-k) - \frac{N}{2} \log \frac{\sigma(z_i | \Sigma_{mb}(i)^* \cup A)}{\sigma(z_i | \Sigma_{mb}(i) \cup A)} + O(1) \end{aligned} \quad (45)$$

where the first and the second term on the right-hand side tend to $-\infty$ and ∞ , respectively. However, since $-\log N$ decreases slower than N increases, the right-hand side of the equation tends to ∞ . This proves the statement of Lemma 1. $\square$

Lemma 2 Let $\text{mb}(i)^* \subset \text{mb}(i)$. Then

$$\text{plim}_{N \rightarrow \infty} \log \frac{p(\mathbf{Z}_i | \Sigma_{mb}(i)^*, k)}{p(\mathbf{Z}_i | \Sigma_{mb}(i), k)} = \infty. \quad (46)$$

Proof By defining $p_i \triangleq |\text{mb}(i)^*|$ and $a \triangleq |\text{mb}(i)| - p_i$, we have that

$$\begin{aligned} \log \frac{p(\mathbf{Z}_i | \Sigma_{mb}(i)^*, k)}{p(\mathbf{Z}_i | \Sigma_{mb}(i), k)} &= \log \frac{\Gamma\left(\frac{N-k+p_i}{2}\right)}{\Gamma\left(\frac{N-k+p_i+a}{2}\right)} \\ &\quad + \log \frac{\Gamma\left(\frac{p_i+a+1}{2}\right)}{\Gamma\left(\frac{p_i+1}{2}\right)} + a \log(N-k) \\ &\quad - \frac{N-k-1}{2} \log \frac{|\mathbf{S}_{fa}(i)^*| |\mathbf{S}_{mb}(i)|}{|\mathbf{S}_{mb}(i)^*| |\mathbf{S}_{fa}(i)|}. \end{aligned} \quad (47)$$

This equation is identical to (20) except for the last term on the right-hand side. As $N \rightarrow \infty$, this term tends to

$$-N \log \frac{|\mathbf{S}_{fa}(i)^*| |\mathbf{S}_{mb}(i)|}{|\mathbf{S}_{mb}(i)^*| |\mathbf{S}_{fa}(i)|} \cdot \frac{1}{2} + O(1). \quad (48)$$

Now select $B \subset \text{mb}(i)$ so that $\text{mb}(i) = \text{mb}(i)^* \cup B$. Then

$$-N \log \frac{|\mathbf{S}_{\text{fa}}(i)^*| |\mathbf{S}_{\text{mb}}(i)|}{|\mathbf{S}_{\text{mb}}(i)^*| |\mathbf{S}_{\text{fa}}(i)|} = \text{dev}(z_i \perp \mathbf{z}_B \mid \mathbf{z}_{\text{mb}(i)^*}) \quad (49)$$

where, asymptotically, the deviance $\text{dev}(z_i \perp \mathbf{z}_B \mid \mathbf{z}_{\text{mb}(i)^*})$ follows (Whittaker 1990) a χ^2 distribution. This implies that the deviance is bounded in probability. Consequently,

$$-\frac{N-k-1}{2} \log \frac{|\mathbf{S}_{\text{fa}}(i)^*| |\mathbf{S}_{\text{mb}}(i)|}{|\mathbf{S}_{\text{mb}}(i)^*| |\mathbf{S}_{\text{fa}}(i)|} = O(1). \quad (50)$$

Now, applying the same reasoning as for Lemma 1, we arrive at

$$\log \frac{p(\mathbf{Z}_i \mid \mathbf{Z}_{\text{mb}(i)^*}, k)}{p(\mathbf{Z}_i \mid \mathbf{Z}_{\text{mb}(i)}, k)} = \frac{a}{2} \log(N-k) + O(1). \quad (51)$$

Since $\text{mb}(i)^* \subset \text{mb}(i)$, $a > 0$, and the right-hand side of the equation tends to ∞ . This proves the statement of Lemma 2. $\square$

References

- Abegaz, F., Wit, E.: Sparse time series chain graphical models for reconstructing genetic networks. *Biostatistics* **14**(3), 586–599 (2013)
- Andersson, S.A., Madigan, D., Perlman, M.D.: Alternative Markov properties for chain graphs. *Scand. J. Stat.* **28**(1), 33–85 (2001)
- Ang, A., Piazzesi, M.: A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *J. Monet. Econ.* **50**(4), 745–787 (2003)
- Arnold, A., Liu, Y., Abe, N.: Temporal causal modeling with graphical Granger methods. In: *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 66–75 (2007)
- Baccalá, L.A., Sameshima, K.: Partial directed coherence: a new concept in neural structure determination. *Biol. Cybern.* **84**(6), 463–474 (2001)
- Bañura, M., Giannone, D., Reichlin, L.: Large Bayesian vector autoregressions. *J. Appl. Econ.* **25**(1), 71–92 (2010)
- Barber, R.F., Drton, M.: High-dimensional Ising model selection with Bayesian information criteria. *Electron. J. Stat.* **9**(1), 567–607 (2015)
- Besag, J.: Statistical analysis of non-lattice data. *J. R. Stat. Soc.: Ser. D (Stat.)* **24**(3), 179–195 (1975)
- Bringmann, L.F., Vissers, N., Wichers, M., Geschwind, N., Kuppens, P., Peeters, F., Borsboom, D., Tuerlinckx, F.: A network approach to psychopathology: new insights into clinical longitudinal data. *PLoS ONE* **8**(4), e60188 (2013)
- Brockwell, P.J., Davis, R.A.: *Introduction to Time Series and Forecasting*. Springer, Berlin (2016)
- Bühlmann, P.: Statistical significance in high-dimensional linear models. *Bernoulli* **19**(4), 1212–1242 (2013)
- Carvalho, C.M., Scott, J.G.: Objective Bayesian model selection in Gaussian graphical models. *Biometrika* **96**(3), 497–512 (2009)
- Cavalcante, L., Bessa, R.J., Reis, M., Browell, J.: LASSO vector autoregression structures for very short-term wind power forecasting. *Wind Energy* **20**(4), 657–675 (2017)
- Chiang, J., Wang, Z.J., McKeown, M.J.: Sparse multivariate autoregressive (mAR)-based partial directed coherence (PDC) for electroencephalogram (EEG) analysis. In: *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 457–460. IEEE (2009)
- Consonni, G., Rocca, L.L.: Objective Bayes factors for Gaussian directed acyclic graphical models. *Scand. J. Stat.* **39**(4), 743–756 (2012)
- Dahlhaus, R., Eichler, M.: Causality and graphical models in time series analysis. In: Green, P.J., Hjort, N.L., Richardson, S. (eds.) *Highly Structured Stochastic Systems*, Chapter 4. Oxford University Press, Oxford (2003)
- Davis, R.A., Zang, P., Zheng, T.: Sparse vector autoregressive modeling. *J. Comput. Graph. Stat.* **25**(4), 1077–1096 (2016)
- Dowell, J., Pinson, P.: Very-short-term probabilistic wind power forecasts by sparse vector autoregression. *IEEE Trans. Smart Grid* **7**(2), 763–770 (2016)
- Epskamp, S., Waldorp, L.J., Möttus, R., Borsboom, D.: The Gaussian graphical model in cross-sectional and time-series data. *Multivar. Behav. Res.* **53**(4), 453–480 (2018)
- Fan, J., Li, R.: Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Am. Stat. Assoc.* **96**(456), 1348–1360 (2001)
- Fop, M., Murphy, T.B., Scrucca, L.: Model-based clustering with sparse covariance matrices. *Stat. Comput.* **29**(4), 791–819 (2019)
- Fukushima, M., Yamashita, O., Knösche, T.R., Sato, M.A.: MEG source reconstruction based on identification of directed source interactions on whole-brain anatomical networks. *NeuroImage* **105**, 408–427 (2015)
- Goldberger, A.L., Amaral, L.A.N., Glass, L., Hausdorff, J.M., Ivanov, P.C., Mark, R.G., Mietus, J.E., Moody, G.B., Peng, C.K., Stanley, H.E.: PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation* **101**(23), e215–e220 (2000)
- Gómez-Herrero, G., Atienza, M., Egiazarian, K., Cantero, J.L.: Measuring directional coupling between EEG sources. *NeuroImage* **43**(3), 497–508 (2008)
- Harrison, L., Penny, W.D., Friston, K.: Multivariate autoregressive modeling of fMRI time series. *NeuroImage* **19**(4), 1477–1491 (2003)
- Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd edn. Springer Science & Business Media, Berlin (2009)
- Haufe, S., Müller, K.R., Nolte, G., Krämer, N.: Sparse causal discovery in multivariate time series. *Proc. Mach. Learn. Res.* **6**, 97–106 (2010)
- Hsu, N.J., Hung, H.L., Chang, Y.M.: Subset selection for vector autoregressive processes using Lasso. *Comput. Stat. Data Anal.* **52**(7), 3645–3657 (2008)
- Huang, J., Xie, H.: Asymptotic oracle properties of SCAD-penalized least squares estimators. In: *Asymptotics: Particles, Processes and Inverse Problems*, pp. 149–166. Institute of Mathematical Statistics (2007)
- Ito, T., Sato, K.: Exchange rate changes and inflation in post-crisis Asian economies: vector autoregression analysis of the exchange rate pass-through. *J. Money, Credit, Bank* **40**(7), 1407–1438 (2008)
- Leppä-aho, J., Pensar, J., Roos, T., Corander, J.: Learning Gaussian graphical models with fractional marginal pseudo-likelihood. *Int. J. Approx. Reason.* **83**, 21–42 (2017)
- Lütkepohl, H.: *New Introduction to Multiple Time Series Analysis*. Springer Science & Business Media, Berlin (2005)
- Melnyk, I., Banerjee, A.: Estimating structured vector autoregressive models. *Proc. Mach. Learn. Res.* **48**, 830–839 (2016)
- Michalareas, G., Schoffelen, J.M., Paterson, G., Gross, J.: Investigating causality between interacting brain areas with multivariate autoregressive models of MEG sensor data. *Hum. Brain Mapp.* **34**(4), 890–913 (2013)
- Neusser, K.: *Time Series Econometrics*. Springer, Berlin (2016)
- O’Hagan, A.: Fractional Bayes factors for model comparison. *J. R. Stat. Soc.: Ser. B (Methodol.)* **57**(1), 99–118 (1995)

- Paci, L., Consonni, G.: Structural learning of contemporaneous dependencies in graphical var models. *Comput. Stat. Data Anal.* **144**, 106880 (2020)
- Pensar, J., Nyman, H., Niiranen, J., Corander, J.: Marginal pseudo-likelihood learning of discrete Markov network structures. *Bayesian Anal.* **12**(4), 1195–1215 (2017)
- Roebroeck, A., Formisano, E., Goebel, R.: Mapping directed influence over the brain using Granger causality and fMRI. *NeuroImage* **25**(1), 230–242 (2005)
- Shoeb, A.H.: Application of machine learning to epileptic seizure onset detection and treatment. Ph.D. thesis, Massachusetts Institute of Technology (2009)
- Supp, G.G., Schlögl, A., Trujillo-Barreto, N., Müller, M.M., Gruber, T.: Directed cortical information flow during human object recognition: analyzing induced EEG gamma-band responses in brain's source space. *PLoS ONE* **2**(8), e684 (2007)
- Tibshirani, R.: Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc.: Ser. B (Methodol.)* **58**(1), 267–288 (1996)
- Tsamardinos, I., Brown, L.E., Aliferis, C.F.: The max–min hill-climbing Bayesian network structure learning algorithm. *Mach. Learn.* **65**(1), 31–78 (2006)
- Tsiaras, V., Simos, P.G., Rezaie, R., Sheth, B.R., Garyfallidis, E., Castillo, E.M., Papanicolaou, A.C.: Extracting biomarkers of autism from MEG resting-state functional connectivity networks. *Comput. Biol. Med.* **41**(12), 1166–1177 (2011)
- Valdés-Sosa, P.A., Sánchez-Bornot, J.M., Lage-Castellanos, A., Vega-Hernández, M., Bosch-Bayard, J., Melie-García, L., Canales-Rodríguez, E.: Estimating brain functional connectivity with sparse multivariate autoregression. *Philos. Trans. R. Soc. B: Biol. Sci.* **360**(1457), 969–981 (2005)
- Whittaker, J.: *Graphical Models in Applied Multivariate Statistics*. Wiley, Hoboken (1990)
- Wild, B., Eichler, M., Friederich, H.C., Hartmann, M., Zipfel, S., Herzog, W.: A graphical vector autoregressive modelling approach to the analysis of electronic diary data. *BMC Med. Res. Methodol.* **10**, 1–13 (2010)
- Zang, W., Baimbridge, M.: Exports, imports and economic growth in South Korea and Japan: a tale of two economies. *Appl. Econ.* **44**(3), 361–372 (2012)
- Zhao, P., Yu, B.: On model selection consistency of Lasso. *J. Mach. Learn. Res.* **7**, 2541–2563 (2006)
- Zhao, Y., Ye, L., Pinson, P., Tang, Y., Lu, P.: Correlation-constrained and sparsity-controlled vector autoregressive model for spatio-temporal wind power forecasting. *IEEE Trans. Power Syst.* **33**(5), 5029–5040 (2018)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.