

Inferring Brain Networks through Graphical Models with Hidden Variables

Justin Dauwels*, Hang Yu*, Xueou Wang⁺, Francois Vialatte[†], Charles Latchoumane⁺⁺, Jaeseung Jeong^{††}, and Andrzej Cichocki[‡]

School of Electrical and Electronic Engineering*
School of Physical and Mathematical Sciences⁺
Nanyang Technological University, Singapore, 639798

ESPCI ParisTech[†]
10 Rue Vauquelin 75005 Paris, France

Center for Neural Science, KIST⁺⁺
Seongbuk-gu, Seoul 136-791, Republic of Korea

Brain Dynamics Laboratory, KAIST^{††}
Daejeong 305-701, Republic of Korea

RIKEN Brain Science Institute[‡]
2-1 Hirosawa, Wako City, Saitama 351-0198 Japan

Abstract. Inferring the interactions between different brain areas is an important step towards understanding brain activity. Most often, signals can only be measured from some specific brain areas (e.g., cortex in the case of scalp electroencephalograms). However, those signals may be affected by brain areas from which no measurements are available (e.g., deeper areas such as hippocampus). In this paper, the latter are described as hidden variables in a graphical model; such model quantifies the statistical structure in the neural recordings, conditioned on hidden variables, which are inferred in an automated fashion from the data. As an illustration, electroencephalograms (EEG) of Alzheimer’s disease patients are considered. It is shown that the number of hidden variables in AD EEG is not significantly different from healthy EEG. However, there are fewer interactions between the brain areas, conditioned on those hidden variables. Explanations for these observations are suggested.

1 Introduction

Sparse graphical models (see, e.g., [1]) provide an effective way to capture statistical structure in high-dimensional data such as multi-electrode brain recordings. A sparse graph displays the most significant interactions between variables (e.g., brain areas), and may help to interpret the data.

In practice, it is quite common that data is unavailable for some relevant variables. Most often, signals can only be measured from some specific brain areas (e.g., cortex in the case of scalp electroencephalograms). However, those signals may be affected by brain areas from which no measurements are available (e.g., deeper areas such as hippocampus). The latter may then be treated as hidden variables in a statistical model, providing a simple explanation for the statistical relations between the observed brain regions.

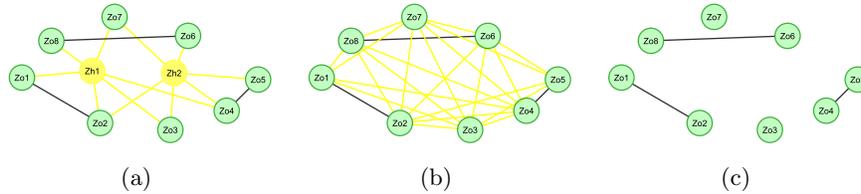


Fig. 1. Hidden variable graphical model (a) joint graphical model. (b) marginal graphical model. (c) conditional graphical model.

Therefore, sparse graphical models with hidden variables are quite powerful models for a large variety of neural datasets. When the observed variables Z_o (green nodes in Fig. 1(a)) and hidden variables Z_h (yellow nodes in Fig. 1(a)) are jointly Gaussian distributed, the structure of the graphical model can be defined by the precision matrix (inverse covariance matrix) of the observed and hidden variables. Recently, Chandrasekaran et al. [2] decomposed the (marginal) precision matrix of Z_o into a sparse matrix K_o (conditional precision matrix) and a low-rank matrix L , which describes the coupling between the observed and hidden variables. The conditional graphical model K_o and the number of hidden variables (rank of L) are inferred by solving a convex regularized maximum-likelihood problem. The conditional precision matrix K_o is represented as a graph (Fig. 1(c)), where nodes i and j are connected by an edge iff the corresponding element (i, j) in K_o is non-zero. That graph visualizes the dependence among the observed variables, conditioned on the hidden variables. In other words, such graph shows those interactions between observed variables that cannot be explained by hidden variables alone.

To infer sparse graphical models with hidden variables from neural recordings, two issues need to be addressed:

1. Neural data is often non-Gaussian. Therefore, we apply copula Gaussian graphical models with hidden variables, which we developed in recent work to model non-Gaussian data [3].
2. So far, such models have only been applied to i.i.d. samples. On the other hand, neural signals are time series, and therefore, we will follow the theory for graphical models of multivariate time series (“conditional independence graphs”) [4–6]; we infer the graphical models (with hidden variables) in frequency domain.

This paper is organized as follows. In Section 2, we first review several relevant graphical models. In Section 3, we briefly explain how the graphical models can be applied to time series, such as neural signals. In Section 5, we apply the copula Gaussian graphical model with hidden variables to electroencephalograms (EEG) of mild Alzheimer’s patients. In Section 6, we offer concluding remarks.

2 Graphical Models

In the following, we briefly describe the copula Gaussian graphical model, the Gaussian hidden variable graphical model (GHVGM), and the copula Gaussian hidden variable graphical model (CGHVGM) that we proposed in recent work [3].

2.1 Copula Gaussian Graphical Model

We denote the observed non-Gaussian variables and hidden Gaussian variables as Y_1, \dots, Y_p and Z_1, \dots, Z_p respectively. A Gaussian copula graphical model is defined as [7, 8]:

$$Z \sim \mathcal{N}(0, K^{-1}) \quad (1)$$

$$Y_k = F_k^{-1}(\Phi(Z_k)), \quad (2)$$

where K is the precision matrix whose inverse (the covariance matrix) has normalized diagonal, Φ is the CDF (cumulative distribution function) of the standard Gaussian distribution, and F_k is the CDF of Y_k . The latter is often approximated by the empirical distributions \hat{F}_k . Note that F_k^{-1} is the pseudo-inverse of F_k , which is defined as:

$$F^{-1}(y) = \inf_{x \in \mathcal{X}} \{F(x) \geq y\}. \quad (3)$$

2.2 Gaussian Hidden Variable Graphical Model

Suppose we have Gaussian distributed observed variables Z_o and hidden variables Z_h . The joint precision matrix $K_{(oh)}$ (Fig. 1(a)) associated with these variables is given by:

$$K_{(oh)} = \begin{bmatrix} K_o & K_{o,h} \\ K_{h,o} & K_h \end{bmatrix}. \quad (4)$$

According to Schur complement, the marginalized precision matrix \tilde{K}_o of Z_o (Fig. 1(b)) can be written as:

$$\tilde{K}_o = K_o - K_{o,h}K_h^{-1}K_{h,o} = K_o - L, \quad (5)$$

with product matrix $L = K_{o,h}K_h^{-1}K_{h,o}$. Those two components have specific properties [2]: K_o is the supposedly sparse conditional precision matrix of Z_o , conditioned on Z_h (Fig. 1(c)); the product matrix L summarizes the effect of marginalization over the hidden variables. The rank of that matrix (equal to the number of hidden variables Z_h) is low, since the number of hidden variables is supposed to be small.

Given i.i.d. samples of Z_o , our objective is to estimate K_o and L ; we are especially interested in the rank of L , since it equals the number of hidden variables Z_h . Those matrices may be recovered by solving the convex relaxation [2]:

$$(\hat{K}_o, \hat{L}) = \underset{K_o, L}{\operatorname{argmin}} \operatorname{trace}((K_o - L)\Sigma_o) - \log \det(K_o - L) + \lambda(\gamma\|K_o\|_1 + \operatorname{trace}(L)), \quad (6)$$

where \hat{K}_o and \hat{L} are the estimates of K_o and L respectively, and Σ_o is the empirical marginal covariance of Z_o . The convex problem (6) can be solved efficiently by the Newton-CG primal proximal point algorithm [9].

2.3 Copula Gaussian Hidden Variable Graphical Model

The observed (continuous) variables Y (blue nodes in Fig. 2) are non-Gaussian, and each of them is associated with a Gaussian distributed hidden variable Z_o (green nodes in Fig. 2), as in the copula Gaussian model. However, besides the hidden variables Z_o , there exist several hidden variables Z_h (yellow nodes in Fig. 2) that are not associated with observed variables. In the corresponding graphical model, the nodes Z_h are only connected to hidden variables; they are not connected to observed variables Y .

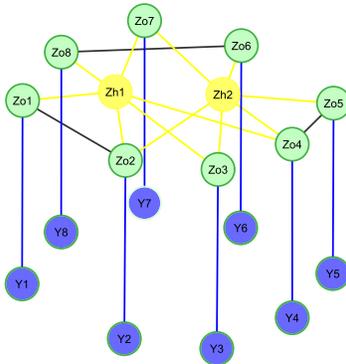


Fig. 2. Copula Gaussian hidden variable graphical model (blue nodes denote observed non-Gaussian variables, green nodes denote latent Gaussian variables and yellow nodes denote hidden variables in the latent layer)

In other words, the variables Y and Z_o constitute a Gaussian copula graphical model, while the variables Z_o and Z_h form a Gaussian hidden variable graphical model. Together, the variables (Y, Z_o, Z_h) form a copula Gaussian hidden variable graphical model with associated conditional precision matrix K_o and product matrix L (cf. (5)).

Given i.i.d. samples of the non-Gaussian variables Y , we wish to infer the conditional precision matrix K_o of Z_o (conditioned on Z_h), and the product matrix L , associated with the CGHVGGM.

As a first step, we transform the non-Gaussian observed variables Y into Gaussian distributed hidden variables Z_o (associated with the observed variables Y):

$$Z_{ok} = \Phi^{-1}(\hat{F}_k(Y_k)), \quad (7)$$

where Φ is the CDF of the standard Gaussian distribution and \hat{F}_k is the empirical CDF of Y_k . As a result, we are dealing with Gaussian variables Z_o which together with Z_h constitute a GHVGGM.

In the second step, we follow the procedure of GHVGGM to infer the sparse conditional precision matrix K_o of Z_o and the low-rank product matrix L .

3 Graphical Models for Multivariate Time Series

A graphical model $G = (V, E)$ may also represent the conditional independence among Gaussian stationary processes $X(t) = (X_1(t), \dots, X_p(t))$ with $t \in \mathbb{Z}$, where each node V_i represent a time series $X_i(t)$. The absence of an edge between two nodes V_i and V_j denotes the conditional independence between the two processes $X_i(t)$ and $X_j(t)$. Specifically, the stationary Gaussian processes X_i and X_j are considered conditionally independent if (see, e.g., [4])

$$\text{cov}\{\epsilon_i(t), \epsilon_j(t + \tau)\} = 0 \quad \forall \tau, \quad (8)$$

where $\tau \in \mathbb{Z}$ is a time lag and $\epsilon_i(t) = X_i(t) - \hat{X}_i(t)_{|i,j}$. The quantity $\hat{X}_i(t)_{|i,j}$ stands for the best linear predictor of $X_i(t)$, obtained from all the values (at all time points) of all Gaussian processes $X(t) = (X_1(t), \dots, X_p(t))$ except $X_i(t)$ and $X_j(t)$. In other words, conditional independence for stationary Gaussian processes means that the correlation between the residuals of the best linear predictors is zero for all possible lags [4].

Assuming that the cross-covariance function of $X(t)$ is summable, i.e.,

$$\sum_{\tau=-\infty}^{\infty} |\text{cov}\{X_i(t), X_j(t + \tau)\}| < \infty, \quad \forall i, j, \quad (9)$$

we can define the spectral density matrix $S(\omega)$ of X as:

$$S_{i,j}(\omega) = \mathcal{F}\{\text{cov}\{X_i, X_j\}\}, \quad (10)$$

where \mathcal{F} represents the Fourier transform. Conditional independence can also be formulated in terms of the spectral density matrix $S(\omega)$ of X . Specifically, the stationary Gaussian processes X_i and X_j are conditionally independent iff [5]

$$\{S(\omega)^{-1}\}_{i,j} = 0 \quad \forall \omega. \quad (11)$$

Consequently, the inverse spectral matrix S^{-1} plays a similar role for Gaussian random processes as the precision matrix K does for Gaussian random variables (cf. (1)).

The condition (11) relies on $S(\omega)$, which can be easily estimated from tapered periodograms [5]. In contrast, inferring the conditional cross-covariance in (8) is more computationally complex, since it requires all $2\binom{p}{2}$ linear predictors $\hat{X}_i(t)_{|i,j}$.

We infer a copula Gaussian hidden variable graphical model from time series as follows. First we transform the time series into a Gaussian time series (assuming ergodicity), similarly as (7), and compute the spectral density matrix $S(\omega)$ from the resulting Gaussian time series [5]. We then solve (6), where Σ_o is replaced by $S(\omega)$. We will provide more details in a future publication.

4 EEG Data Set

Our EEG data set consists of 24 healthy control subjects (age: 69.4 ± 11.5 years old; 10 males) and 17 patients with mild AD (age: 77.6 ± 10.0 years old; 9 males) [10]. The patient group underwent full battery of cognitive tests (Mini Mental State Examination, Rey Auditory Verbal Learning Test, Benton Visual Retention Test, and memory recall tests). The EEG time series were recorded using 19 electrodes positioned according to Maudsley system, similar to the 10-20 international system, at a sampling frequency of 128 Hz. All recording sessions were conducted with the subjects in an awake but resting state with eyes closed, and the length of the EEG recording was about 5 minutes, for each subject. Only those subjects were retained in the analysis whose EEG recordings contained at least 20s of artifact-free data.

5 Numerical Results and Discussion

It is well known that the EEG of AD patients is generally less synchronous (less correlated) [11]. Therefore, we would expect the (functional) brain networks to be sparser in AD patients. Indeed, we found that the conditional precision matrix K_0 in CGHVGM is sparser for AD EEG than healthy EEG ($p = 0.0121$ for Mann-Whitney test; see Fig. 3). Based on sparsity only as input feature, we can classify the two populations with a success rate of about 80% (linear discriminant analysis with leave-one-out crossvalidation). In other words, conditioned on the hidden factors, the AD EEG signals are substantially less mutually dependent than healthy EEG signals. The correlations between the AD EEG signals are to a larger extent due to hidden factors. Interestingly, the number of hidden variables is not significantly different in healthy subjects and AD patients (typically between 3 and 6). No significant differences were observed for GHVGM, which seems to suggest that modeling the non-Gaussian nature of the data with a copula is crucial. Likewise, no significant differences (see Fig. 4) between the AD patients and healthy subjects were found for graphical models without hidden variables, i.e., Gaussian graphical models (glasso; [1]) and copula Gaussian

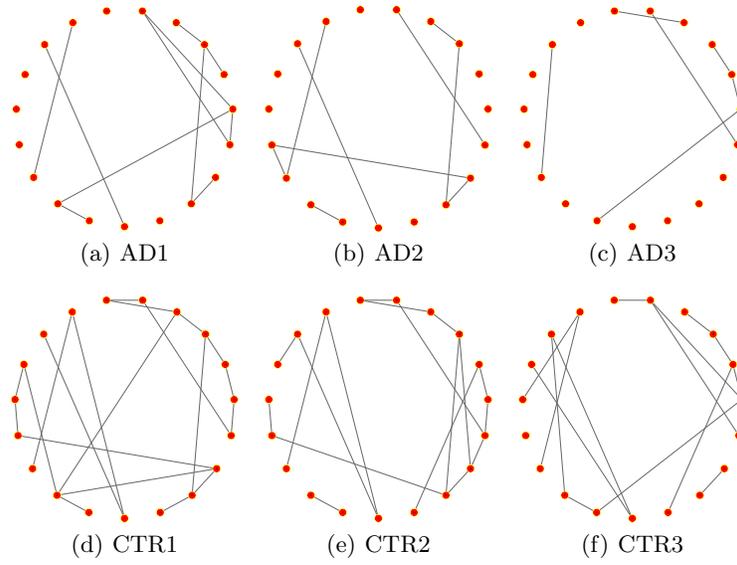


Fig. 3. Conditional precision matrix K_0 (CGHVGM) of first three AD patients and control subjects.

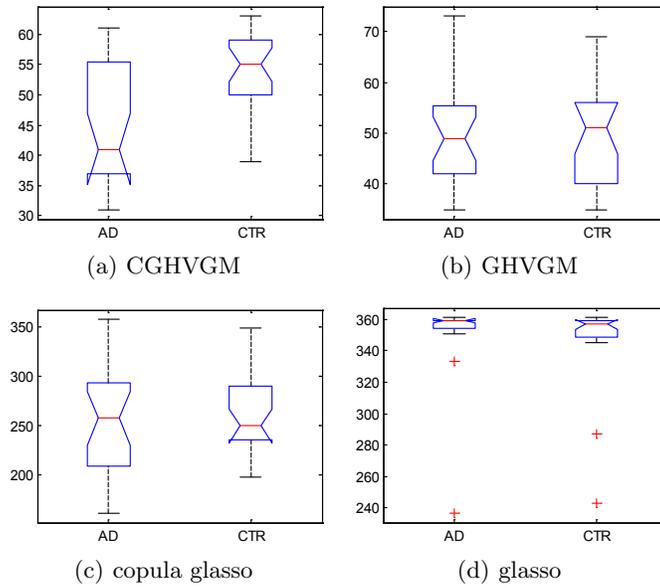


Fig. 4. Boxplots for the sparsity measure of CGHVGM and GHVGM (The corresponding p -value are 0.012, 0.90, 0.99 and 0.16)

graphical models (copula glasso; [8]). Those models are also much denser (and hence more complicated) than the models with hidden variables, as can be seen from Fig. 4.

Note that all parameters in the considered graphical models are inferred automatically from data; none of the parameters have been optimized, to minimize the risk of overfitting and false positives.

6 Conclusions

We have developed a method to infer brain networks, conditioned on latent factors. Such models are quite relevant, since oftentimes neural recordings are often affected by brain areas from which it is impractical to record. Such brain areas are viewed as hidden variables in the proposed approach. The method was illustrated on an EEG data set of mild Alzheimer’s patients.

References

1. J. Friedman, T. Hastie, and R. Tibshirani, “Sparse inverse covariance estimation with the graphical lasso,” *Biostatistics* vol. 9, pp. 432–441, 2008.
2. V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky, “Latent Variable Graphical Model Selection via Convex Optimization,” *Technical report, Massachusetts Institute of Technology*, 2010.
3. J. Dauwels, H. Yu and X. O. Wang, “Copula Gaussian Graphical Models with Hidden Variables,” *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2012*, Kyoto, Japan, March 25–30, 2012.
4. D.R. Billinger, “Remarks concerning graphical models for time series and point processes”, *Revista de Econometria*, vol. 16, pp. 1–23, 1996.
5. R. Dahlhaus, “Graphical interaction models for multivariate time series,” *Metrika*, vol.5, pp. 157–172, 2000.
6. F. R. Bach, M. I. Jordan, “Learning graphical models for stationary time series,” *IEEE Transactions on Signal Processing*, vol. 52, pp. 2189–2199, 2004.
7. A. Dobra and A. Lenkoski, “Copula Gaussian graphical models and their application to modeling functional disability data,” *Annals of Applied Statistics*, vol. 5, pp. 969–993, 2011.
8. H. Liu, J. Lafferty, and L. Wasserman, “The Nonparanormal: Semiparametric Estimation of High Dimensional Undirected Graphs,” *Journal of Machine Learning Research*, vol. 10, pp. 2295–2328, 2009.
9. C. Wang, D. Sun, and K. C. Toh, “Solving log-determinant optimization problems by a Newton-CG primal proximal point algorithm,” *Society for Industrial and Applied Mathematics*, vol. 20, pp. 2994–3013, 2009.
10. G. Henderson, E. Ifeakor, N. Hudson, C. Goh, N. Outram, S. Wimalaratna, C. Del Percio, and F. Vecchio, “Development and assessment of methods for detecting dementia using the human electroencephalogram,” *IEEE Transactions on Biomedical Engineering*, vol. 53, pp. 1557–1568, 2006.
11. J. Dauwels, F. Vialatte, and A. Cichocki, “Diagnosis of alzheimer’s disease from EEG signals: Where are we standing?” *Current Alzheimer Research*, vol. 7, pp. 487–505, 2010.