Technische Universiteit
Eindhoven
University of Technology

Department of Electrical Engineering

Den Dolech 2, 5612 AZ Eindhoven
P.O. Box 513, 5600 MB Eindhoven
The Netherlands
http://w3.ele.tue.nl/nl/

Series title:
Master graduation paper,
Electrical Engineering

Commissioned by Professor:
prof. dr. ir. P.M.J. van den Hof

Group / Chair:
Control systems

Date of final presentation:
August 30,2019

Report number:

Topology Detection in Brain Networks

by

Author: R.J.C. van Esch

Internal supervisors: Shengling Shi MSc
Acknowledgements: This project has received funding from the European Research Council (ERC), Advanced Research Grant SYSDYNET, under the European Union's Horizon 2020 research and innovation programme (grant agreement No 694504).

Topology Detection in Brain Networks

R.J.C. van Esch
Department of Electrical Engineering
Eindhoven University of Technology
Eindhoven, The Netherlands

***Abstract*— Granger-causality analysis [2] is a well-established method for inference of brain network connectivity using fMRI data. In this paper we investigate if Bayesian topology identification [10] can be used as a replacement of Granger-causality analysis for the inference of brain network connectivity. First, we have evaluated the performance of both methods in simulation, where the Bayesian method outperforms the Granger method in inference of the true connectivity of dynamic networks. Furthermore, an extension is developed for the Bayesian method to be able to use it to draw conclusion on differences in connectivity between groups of subjects. In the second part, the influence of Mozart's sonata K448 on brain network connectivity through fMRI data of healthy adults is investigated. And using the developed extension of the Bayesian method, a significant change was found for long listening subjects in the connection between the posterior default mode network and the fronto-parietal right network. Furthermore, a significant change was found for subjects that listened for a relatively shorter time in the connection from the sensori-motor lateral network to the superior temporal gyrus.**

I. Introduction

The study of neural activity and connectivity in the brain gives insight into how the brain processes information. An interesting research question in the study of the brain is the existence of the Mozart effect: several studies have been performed [12] [13] [27] reporting that a Mozart Effect exists, i.e. listening to Mozart music has an effect on the brain and neural activity in the brain. First, in [12] it is reported that listening to Mozart music leads to a reduction in frequency of epileptic seizures in children. Second, in [13] neural activity was measured through spectral analysis of electro encephalogram (EEG) data. After listening to Mozart music an increase in brain wave activity linked to memory, cognition and problem solving was observed in adults. Last, in [27] it is reported that listening to Mozart music can increase spatial reasoning skills in healthy subjects. However, these studies do not focus on inferring interactions between brain networks using time series, so in this work we will focus on methods that can infer, using time series that describe neural activity in brain regions, how these brain regions are connected and interact.

By far the most popular method of inference of brain network connectivity in neuroscience is Granger-causality analysis [2]. While the method originated in the field of economics, it has been used in the past for inference of brain connectivity using EEG [4] and functional magnetic resonance imaging (fMRI) [3] [5] data of the brain. In this paper we will focus on the inference of brain network connectivity from fMRI data.

Besides Granger-causality analysis, several other methods have been developed to infer connectivity of dynamic networks, which could be used to infer the topology of brain networks.

First, there are several methods which use a dynamic network model to infer connectivity of the network [6] [7]. The approach in [6] is specifically designed to model neural activity [26] and interconnectivity using specialized state-space models, for example in [8] excitatory and inhibitory states are used, which is in line with how neurons operate in the brain. While the specialized models are useful to model brain activity more accurately, the estimation of the model is computationally expensive. The method in [7] uses reweighted regularized regression to force a subset of parameters belonging to a connection to zero in an effort to infer an estimate of the connectivity. While this method scales very well with large dynamic networks, it has tuning parameters which must be carefully chosen to achieve a good estimate of the connectivity.

Furthermore, there are approaches which do not use a dynamic network model such as the method detailed in [9]. This method calculates a Wiener filter estimate of the dynamics of one connection to infer if it exists in the dynamic network or not. However, this approach requires a lot of data to compute accurate estimates of the connectivity.

In this paper, we will consider a relatively new method, Bayesian topology identification [10] [11]. It uses Bayesian model selection to estimate the connectivity of a dynamic network. Through the incorporation of prior knowledge this method performs well even when not a lot of data is available. However, it does not scale very well with large networks.

Our first research question is to investigate whether Bayesian topology identification can be used as a replacement for Granger-causality analysis for the inference of brain network connectivity using fMRI data. First, we want to investigate if there is a significant difference in performance between the two methods. Then, since the Bayesian method was originally designed for inference of the brain network connectivity in one person, we want to extend the Bayesian method such that we can draw conclusions on the connectivity of groups of data sets, which is a common research question in neuroscientific studies.

Our second research question is to investigate the influence of Mozart's sonata K448 on brain network connectivity through fMRI data of healthy adults. Furthermore, we want to

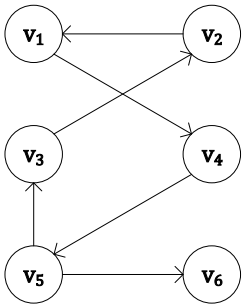

Figure 1: Directed graph structure for a six-node network.

investigate if the amount of time subjects listened to Mozart music leads to different conclusions on brain network connectivity.

This paper is structured as follows: first, in this paper we will detail how we can describe neural activity using fMRI data and how we can use this data to describe brain regions and their neural activity. Second, in section III we detail how we model the interactions between brain regions using a vector autoregressive model. Next, in section IV we describe both the Granger and Bayesian methods and test their performance using simulated dynamic networks. Then, in section V we will explain how we can infer differences in connectivity between two groups of people using Granger-causality. Next, in section VI we develop an extension to the Bayesian method that can test the hypothesis of significant difference between the connectivity of two groups of people. Finally, in section VII we will apply the methods from the previous sections to real data and infer if listening to Mozart music has an effect on the brain network connectivity of the subjects in the study.

II. Data description

To describe brain network connectivity the collected data should be representative of the neural activity in the brain. In this paper we will focus on measurements taken using a method called functional Magnetic Resonance Imaging (fMRI).

A. Functional Magnetic Resonance Imaging

An fMRI measurement indicates the amount of oxygen being used at the measurement location. This is an indirect indication of neural activity called the Blood Oxygenated Level Dependent (BOLD) signal [14]. BOLD signal measurements are acquired for every voxel in the fMRI scan, where a voxel is a unit of space defined by the spatial resolution of the scan, and a trade-off is made between the spatial and temporal resolutions of the scan. The measurements are collected in matrix X, where $x_i(t)$ is the measurement taken at time t and i indicates the index of the voxel where the measurements are taken:

$$X = \begin{pmatrix} x_1(0) & \dots & x_n(0) \\ \vdots & \ddots & \vdots \\ x_1(N-1) & \dots & x_n(N-1) \end{pmatrix}, \tag{1}$$

where each column of $X \in \mathbb{R}^{N \times n}$ is a vector of measurements from one voxel $x_i(t)$ over time.

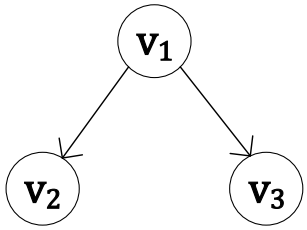

Figure 2: Directed graph structure of a confounding effect.

B. Characterization of brain regions

Typical dimensions of a voxel are $\sim 1\ mm^3$ and as such one distinct region in the brain that processes a certain activity will contain many voxels. To avoid dealing with a very high number of time series for each brain region we will use a single time series that summarizes the activity of that region.

The brain regions can be defined using prior knowledge of their locations. Then the neural activity can be summarized by taking the average over the columns of X of all voxels that are part of one brain region. Alternatively, we can perform a blind source separation procedure, where an algorithm jointly estimates maps of the brain regions and their characteristic time series. The advantage of this approach is that we do not need to rely on prior knowledge of brain regions and the analysis can be performed over all subjects in a study to get comparable maps for all subjects. In this paper we use data processed by a blind source separation procedure called independent component analysis (ICA) [15] [16].

C. Independent component analysis

ICA decomposes data matrix X into spatial maps $\boldsymbol{s}_j$ of each brain region and ICA time series $w_j(t)$:

$$W = \begin{pmatrix} w_1(0) & \dots & w_L(0) \\ \vdots & \ddots & \vdots \\ w_1(N-1) & \dots & w_L(N-1) \end{pmatrix}, \tag{2a}$$

$$\boldsymbol{s}_j = (s_{j,1}, \dots, s_{j,n}), \qquad S = \begin{pmatrix} \boldsymbol{s}_1 \\ \vdots \\ \boldsymbol{s}_L \end{pmatrix}, \tag{2b}$$

$$X = WS, \tag{2c}$$

where the columns of $W \in \mathbb{R}^{N \times L}$ consist of L ICA time series $w_j(t)$ over t and $S \in \mathbb{R}^{L \times n}$ consists of L spatial maps on its rows. Each entry in one spatial map $\boldsymbol{s}_j$ corresponds to one voxel in the scan. Here, L is determined as the amount of non-zero singular values in a probabilistic principal component analysis [17] of data matrix X.

Now, brain region j is defined as the collection of voxels with significantly non-zero map values in $\boldsymbol{s}_j$. As such, each brain region is defined as a group of voxels with 'significantly' correlated neural activity, which is described by ICA time series $w_j(t)$. An example of an ICA time series is shown in Figure 3.

The decomposition into L maps of brain regions in (2) is found by the criteria of statistical independence between the

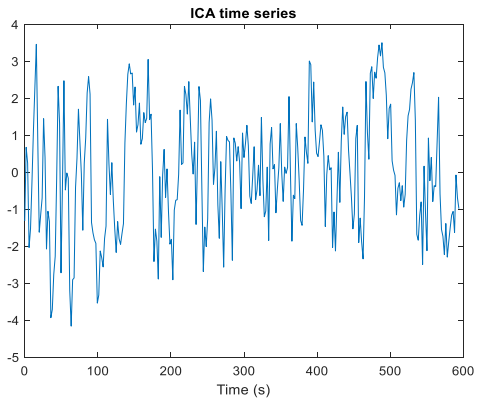

Figure 3: Example of an ICA time series.

maps [16], i.e. the values of one map should give no information on any other map's values.

III. Model Description

To infer the existence of causal interactions between the brain regions defined by the ICA procedure, we will need to choose a modelling framework wherein the interactions between regions can be modelled appropriately.

A. Directed Graph

First, we will define a graph structure, denoted $\mathcal{G} = (V, E)$, that encodes the connectivity between brain regions. In $\mathcal{G}$ each brain region is defined as a node $v_j \in V$ in the graph, where V is the set of all nodes of $\mathcal{G}$. A directed edge from v_i to v_j, denoted $e_{ji} \in E$ or $e_{ji} \in \mathcal{G}$, indicates a directed connection from v_i to v_j, where E is the collection of all edges in $\mathcal{G}$. We denote the set of all incoming edges to v_j as E_j. An example of such a directed graph can be seen in Figure 1.

B. Characterizing Connectivity

In neuroscience we are interested in how the neural activity in one brain region affects the neural activity in other regions. We are thus interested in the causal relationships between brain regions. Therefore, in this paper we will characterize causal connectivity through the conditional Granger causality [2] [18].

The core concept of Granger causality is as follows: For two nodes v_i and v_j we say that v_i Granger-causes v_j if the past of node time series $\{w_i(t-k)|k \geq 1\}$ is useful for predicting $w_j(t)$, given also the past of all other node time series $\{w_{i'}(t-k)|k \geq 1, i' \in I \setminus i\}$. Here $I = \{1, \dots, L\}$ is the index set and $I \setminus i$ indicates the exclusion of index i from set I. Now, graph $\mathcal{G}$ contains all e_{ji} for which v_i Granger-causes v_j over all i, j.

We will now demonstrate with an example why it is important to condition the Granger causality on all other node time series: consider a $\mathcal{G}$ such as in Figure 2, where v_1 granger-causes both v_2 and v_3, i.e. the node time series of v_1 is useful for predicting both of their timeseries. If we do not take v_1 into consideration, it appears that v_2 Granger-causes v_3 and vice-versa, as both contain information from v_1, which is useful for the prediction. This is called a confounding effect of v_1 and ignoring v_1 will not reveal a real causal interaction.

C. Dynamic network model

From literature [19] we learn that the connectivity of brain networks with neural activity described by fMRI measurements can be described by linear models. Furthermore, in general the causal relationship between two regions can be described by an infinite impulse response (IIR) [19]. Now, given brain network connectivity defined by $\mathcal{G}$, we can write the model of the ICA time series $w_j(t)$ as follows:

$$w_j(t) = \sum_{\{i|e_{ji} \in \mathcal{G}\}} G_{ji}(q) w_i(t) + \eta_j(t), \qquad j = 1, \dots, L,$$
$$G_{ji}(q) = \sum_{k=1}^{\infty} \theta_{ji,k} q^{-k}, \tag{3a}$$

where q is the shift operator i.e. $q^{-1} w_j(t) = w_j(t-1)$ and $\eta_j(t)$ is the noise of node v_j. The full model is written as:

$$w(t) = G(q) w(t) + \eta(t) \tag{3b}$$

Where $w(t) = [w_1(t), \dots, w_L(t)]^T$, $\eta(t) = [\eta_1(t), \dots, \eta_L(t)]^T$ and $G(q)$ is a $L \times L$ matrix, where entry (j, i) is $G_{ji}(q)$ if $e_{ji} \in \mathcal{G}$, and is zero everywhere else.

Some assumptions are made on the components of the model in (3b):

- $\left(I - G(q)\right)^{-1}$ is proper and stable.
- $G_{ji}(q)$ is strictly proper and stable for all j, i.
- $\eta_j(t)$ is a white noise process and is independent over j and is Gaussian distributed with unknown variance σ_j^2.

Each IIR $G_{ji}(q)$ can be approximated by a finite order impulse response of order m, which will have no significant effect on the performance of the model if m is chosen large enough. The approximation is written as an order m vector autoregressive model (VAR):

$$w_j(t) = \sum_{i=1}^{L} \left(\sum_{k=1}^{m} w_i(t-k) \theta_{ji,k} \right) + \eta_j(t), \quad j = 1, \dots, L, \tag{3c}$$

and written in matrix form:

$$w_j = \sum_{i \in I} A_{ji} \theta_{ji} + \eta_j, \quad j = 1, \dots, L, \tag{3d}$$

with:

$$A_{ji} = \begin{pmatrix} w_i(-1) & w_i(-2) & \cdots & w_i(-m) \\ w_i(0) & w_i(-1) & \cdots & w_i(-m+1) \\ \vdots & \ddots & \ddots & \vdots \\ w_i(N-1) & w_i(N-2) & \cdots & w_i(N-m) \end{pmatrix}, \tag{3e}$$

where $w_j \in \mathbb{R}^N, \eta_j \in \mathbb{R}^N$ are vectors of time series of v_j and A_{ji} is the matrix of time shifted time series of $w_i(t)$, where $w_i(t) = 0$ for $t < 0$.

Using the VAR model in (3d), we infer the Granger causality of a directed edge e_{ji}, conditional on $\{w_k | k \in I \setminus i\}$, by evaluating if the parameter vector θ_{ji} is non-zero.

IV. Inference methods

Given the ICA time series of one fMRI scan and the model in (3), we would now like to infer an estimate of the brain network graph $\mathcal{G}$, which we will denote $\hat{\mathcal{G}}$. In this section we will detail two methods, which can find such an estimate. First, Granger-causality analysis as an established method of inference. Second, a new method called Bayesian topology identification is proposed as an alternative to Granger-causality analysis.

A. Granger-causality analysis

Granger-causality analysis [18] infers connectivity directly from the Granger-causality of connections. The analysis is based on the idea that if θ_{ji} in (3d) are significantly non-zero, then setting them to 0 will significantly increase the size of the residual. The size of the increase is found by comparing the residuals of two VAR models as in (3d): one where all node time series are included in the regression and one where A_{ji} for a certain connection e_{ji} is removed:

$$w_j = \sum_{k \in I \setminus i} A_{jk} \cdot \theta'_{jk} + \eta'_j, \qquad j = 1, \dots, L \tag{4}$$

The model order m, i.e. the length of vectors θ_{ji}, is selected through the Akaike information criterion (AIC) [20], to sufficiently model the interactions in the data without overfitting, as overfitting can make the Granger-values unreliable [18].

Now, the conditional Granger-causality value is defined as the logarithm of the ratio of the residuals:

$$\mathcal{F}_{e_{ji}} \equiv 2 \ln \left(\frac{\|\eta'_j\|}{\|\eta_j\|} \right). \tag{5}$$

Significant non-zero granger-causality indicates existence of a causal connection e_{ji}. By including all relevant node time series in the regression models in (3d) and (4), we can rule out any confounding effects.

Because $\mathcal{F}_{e_{ji}}$ as $N \to \infty$ asymptotically follows a χ^2 distribution [21], we can perform a statistical test [22] to infer if $\mathcal{F}_{e_{ji}}$ is significantly non-zero. If $\mathcal{F}_{e_{ji}}$ is significantly non-zero, then we add e_{ji} to graph estimate $\hat{\mathcal{G}}$. Now, by performing the statistical test for all e_{ji} independently, we can construct $\hat{\mathcal{G}}$.

B. Bayesian topology identification

Bayesian topology identification [10] [11] is a Bayesian machine learning method that infers from node time series data $D = \{w_1, \dots, w_L\}$ an estimate $\hat{\mathcal{G}}$ of the connectivity. To be precise, a Bayesian model selection [23] approach is used, which compares and selects the optimum posterior odds $P(\mathcal{G}|D)$ among different graphs $\mathcal{G}$. Because we assume that the prior distributions of graphs are uninformative, i.e. $P(\mathcal{G}^1) = P(\mathcal{G}^0)$, the comparison is reduced to a comparison of the marginal likelihood $P(D|\mathcal{G})$.

The marginal likelihood $P(D|\mathcal{G})$ can be found through marginalization over its parametrization θ:

$$P(D|\mathcal{G}) = \int P(D|\theta, \mathcal{G}) P(\theta|\mathcal{G}) d\theta\,, \tag{6a}$$

where $P(D|\theta, \mathcal{G})$ is the likelihood function, given some $\mathcal{G}$ and θ and $P(\theta|\mathcal{G})$ is the prior distribution of θ given some $\mathcal{G}$. $P(D|\theta, \mathcal{G})$ can be computed based on (3) and the Gaussian distributed noise in (3):

$$P(D|\theta, \mathcal{G}) = \prod_{j=1}^{L} \mathcal{N}\left(w_j \middle| \textstyle\sum_i A_{ji}\theta_{ji}, \sigma_j^2 I_N\right), i = \left\{k \middle| e_{jk} \in \mathcal{G}\right\}, \tag{6b}$$

and the parameter prior is chosen as:

$$P(\theta|\mathcal{G}) = \prod_{j=1}^{L} \mathcal{N}\left(0, K_j\right). \tag{6c}$$

Because the parameter vector is modelled as a random vector with a given prior distribution, we can simply choose the model order m in (3d) as some large number, as the parametrization of the prior will determine the relevancy of each parameter. The hyper-parameters K_j and σ_j^2 are estimated using an EM algorithm [10] [23].

The graph $\mathcal{G}^i$ with the maximum $P(D|\mathcal{G})$ is chosen as the estimate $\hat{\mathcal{G}}$. This maximum can be found by comparing all possible graphs, which is computationally expensive when the number of nodes in $\mathcal{G}$ is large. To avoid the combinatorial problem of comparing all possible graphs, a Bayesian greedy search algorithm [24] is employed that efficiently finds $\hat{\mathcal{G}}$.

C. Method performance

To evaluate the potential of the Bayesian method against the Granger-causality analysis, we will investigate the performance of both methods on data generated in simulation. The ICA time series $w_j(t)$ for $j = 1, \dots, L$ are generated from a model as in (3a), where some random directed edges are chosen to form ground truth $\mathcal{G}_0$ and a random Box-Jenkins model [28] is generated for each $G_{ji}(q)$ for all connections in $\mathcal{G}_0$, such that the resulting dynamic network is stable.

The estimated $\hat{\mathcal{G}}$ from both methods are now compared to the generated $\mathcal{G}_0$ and the quality of the estimate is graded by the true positive rate (TPR) and the false positive rate (FPR):

$$TPR = \frac{TP}{P}, \qquad FPR = \frac{FP}{F}, \tag{7}$$

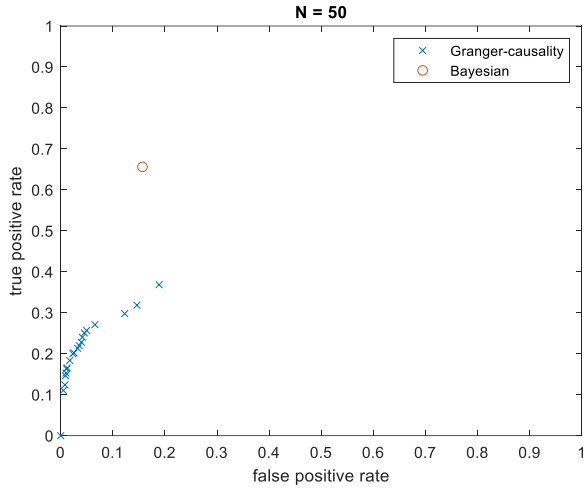

Figure 4: FPR vs TPR of Bayesian and Granger methods for $N = 50$, $L = 6$, $p = \{0.005, 0.05, \ldots, 0.95\}$ from left to right.

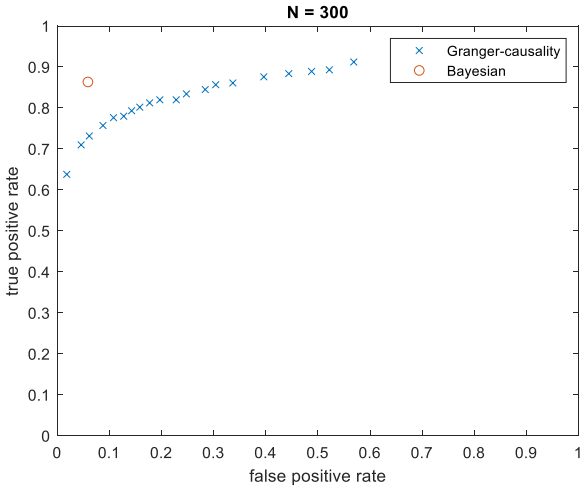

Figure 5: FPR vs TPR of Bayesian and Granger methods for $N = 300$, $L = 6$, $p = \{0.005, 0.05, \ldots, 0.95\}$ from left to right.

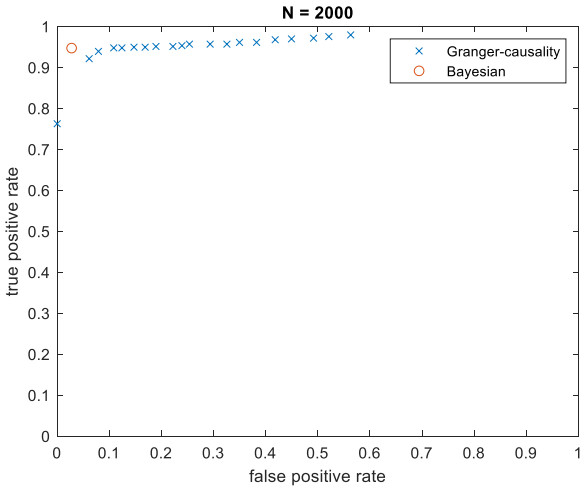

Figure 6: FPR vs TPR of Bayesian and Granger methods for $N = 2000$, $L = 6$, $p = \{0.005, 0.05, \ldots, 0.95\}$ from left to right.

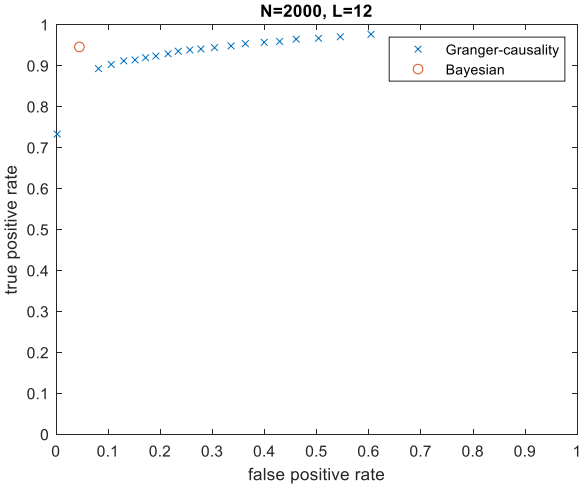

Figure 7: FPR vs TPR of Bayesian and Granger methods for $N = 2000$,$L = 12$, $p = \{0.005, 0.05, \ldots, 0.95\}$ from left to right.

where the TPR is the number of edges in $\hat{\mathcal{G}}$ that also exist in $\mathcal{G}_0$, denoted as TP, over the total number of edges in $\mathcal{G}_0$, denoted as P. The FPR is the number of edges in $\hat{\mathcal{G}}$ that do not exist in $\mathcal{G}_0$, denoted as FP, divided by the total number of edges that do not exist in $\mathcal{G}_0$, denoted as F.

The performance of the methods, given a single data set, does not guarantee overall performance of the method for any given $\mathcal{G}_0$. An estimate of the average expected performance is found by taking the mean of the TPR and FPR from 50 similarly generated dynamic networks, each with a randomly chosen $\mathcal{G}_0$ and random Box-Jenkins model for all connection in $\mathcal{G}_0$.

Furthermore, we evaluate how the performance changes for different data lengths $N = \{50, 300, 2000\}$, for six node networks. Furthermore, we want to infer the performance of twelve node networks for $N = 2000$ and compare to the six node network performance. Finally, for each of the tests we will vary the Granger-causality threshold over a range of thresholds, $p = \{0.005, 0.05, \ldots, 0.95\}$, to infer the performance over these thresholds. The results of the performance evaluation can be found in Figures 4 to 7.

The optimal performance in Figures 4 to 7 is the $(0,1)$ point, because it implies $\hat{\mathcal{G}} = \mathcal{G}_0$. The closer the results are to this point the better the performance of the methods is for that experiment. In Figures 4 to 6 we see that the performance of the Bayesian method moves closer to the $(0,1)$ point in the graph as the data length increases between figures. For the performance of the Granger-causality analysis we can see in Figures 4 to 6 that the performance for only the higher thresholds, which are on the left in the graph move visibly closer to the $(0,1)$ point and we note also that the threshold that is closest to the $(0,1)$ point changes between Figures 4 to 6. Last, we note that the Bayesian method performance is always closer to the optimal point than the performance of Granger-causality analysis for any threshold.

From Figures 4 to 6 we notice that the Bayesian analysis outperforms the Granger-causality analysis for all three data lengths and for all different threshold levels. When comparing Figures 6 and 7 we see that there is barely any change in the performance for the Bayesian algorithm when the number of nodes in the network is increased as the TPR is the same in both

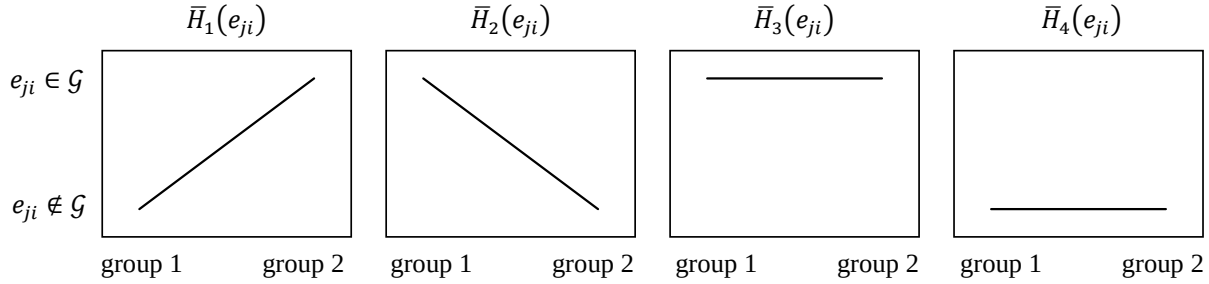

Figure 8: The four between group hypotheses that are tested using Bayesian hypothesis testing.

graphs and the FPR is only slightly larger for the twelve node networks. However, for Granger-causality analysis, the TPR of all threshold levels decreases. Specifically, for the higher thresholds there is a significant difference between six and twelve node networks. Furthermore, we also notice a small increase in the FPR of granger-causality analysis, which becomes even larger for low thresholds. Last, one more important detail is that when we inspect Figures 4 to 6 for the Granger-causality method the optimal performance in each figure does not belong to the same threshold value, instead it varies based on the data length N. But based on the change in performance for each threshold as N becomes large, it seems that the 5% threshold will be the threshold with the optimum performance for large N.

While the Bayesian method outperforms the Granger-causality method in our performance evaluation, there are other advantages and disadvantages to consider. First of all, the Granger-causality is a relatively simple method and the computational speed is fast even for larger amounts of nodes L in the network, while the Bayesian method does not scale well as the number of nodes increases. Furthermore, a disadvantage of the Bayesian method is the reliance on good initial conditions of the EM algorithm to determine K_j and σ_j^2 in (6).

V. Between Group Connectivity Inference

Given the methods we developed in the previous section we now wish to infer how for example a neurological condition might affect the connectivity in the brain of a person, but there exists no ground truth of a 'regular' brain that we can compare to. Therefore, we will need to compare the results of our subject to that of a control subject, a person which does not have this neurological condition.

However, when comparing the brain network connectivity of two subjects we are likely to find many differences in connectivity, as it is also dependent on their 'state of mind' during the fMRI scan. To infer only the differences in connectivity due to the neurological condition, we compare the average connectivity of a group of subjects with a neurological condition to the average connectivity of a group of controls. Now, assuming we have properly chosen our subjects and groups, any significant differences in connectivity can only be due to the neurological condition.

In the previous section, we have detailed methods that infer connectivity of only one subject. In this section, we will define different ways to measure the average connectivity of a group of subjects and how to determine when a difference between the connectivity of two groups is significant. First, we will detail how we can detect significant differences in connectivity between groups using the Granger-causality value. Second, we will formulate a measurement of connectivity by its frequency of occurrence in the estimates of the Bayesian method. And last, we propose an extension to the Bayesian method that allows us to test the hypothesis of significant difference in connectivity between groups.

A. Significant difference using Granger-causality values

Previously, we have defined via Granger-causality that a connection e_{ji} exists if the corresponding parameters θ_{ji} in (3) are significantly non-zero. Furthermore, a significant difference in θ_{ji} between two data sets, leads to a difference in Granger-causality value. A common between group difference statistical test is the two-tailed two-sample t-test [21]:

$$t = \frac{\mathbb{E}_s\left[\mathcal{F}_{e_{ji}}^{1,s}\right] - \mathbb{E}_s\left[\mathcal{F}_{e_{ji}}^{2,s}\right]}{\sqrt{\frac{\sigma_1^2 + \sigma_2^2}{S}}} \tag{8}$$

where $\mathcal{F}_{e_{ji}}^{1,s}$ and $\mathcal{F}_{e_{ji}}^{2,s}$, as defined in (5), indicate the Granger-value of subject s in group 1 and 2, S indicates the total number of subjects across both groups and σ_1^2, σ_2^2 indicate the sample variance of the granger-values of group 1 and 2 respectively.

If this statistic is above the significance threshold, we conclude that there is a significant difference in the connectivity of e_{ji} between group 1 and 2. To determine the threshold, we perform a permutation test [25] of the statistic in (8).

B. Bayesian method selection frequency

The most straightforward way in which we can present the average connectivity of a group of subjects using the Bayesian method is the edge selection frequency: the number of subjects in a group for which $e_{ji} \in \hat{\mathcal{G}}$.

If there is a difference in the selection frequency of some edge e_{ji} between a group of subjects with a neurological condition and controls, this could indicate an overall difference in connectivity caused by the neurological condition. However, there is no way to know when a difference in selection frequency is significantly large enough such that we can draw this conclusion. As such it is clear that just trying to summarize our estimation results from the Bayesian method in the previous section is insufficient and motivates us to develop an alternative method to measure significant differences in connectivity.

VI. Significant change using the Bayesian method

In the Bayesian framework, a difference in connectivity between two subjects due to a neurological condition can only be described by a change in the existence of an edge, i.e. $e_{ji} \in \mathcal{G}$ for one subject and $e_{ji} \notin \mathcal{G}$ for the other. We expand this concept to two groups of subjects, where for example: $e_{ji} \in \mathcal{G}$ for all subjects in group 1 and $e_{ji} \notin \mathcal{G}$ for group 2. In this manner, in the comparison of two groups, there are four different hypotheses on the difference in connectivity of e_{ji} between group 1 and 2, denoted as $\bar{H}_k(e_{ji})$ for $k = 1, \dots, 4$. A visual representation of all $\bar{H}_k(e_{ji})$ is shown in Figure 8.

The hypothesis of subject s in group 1 and 2 are denoted $H_k^{1,s}(e_{ji})$, $H_k^{2,s}(e_{ji})$ and the hypothesis of group 1 is written as $\bar{H}_k^1(e_{ji}) = \{H_k^{1,s}(e_{ji}) | \forall s \in g_1\}$, where g_1 indicates the set of all subjects in group 1. Finally, the between group hypothesis $\bar{H}_k(e_{ji})$ is the collection of the hypotheses of group 1 and 2, i.e. $\bar{H}_k(e_{ji}) = (\bar{H}_k^1(e_{ji}),\ \bar{H}_k^2(e_{ji}))$.

Now, we choose the $\bar{H}_k(e_{ji})$ with the largest posterior probability $P(\bar{H}_k(e_{ji})|\bar{D})$ as the most likely hypothesis, which can be found by comparing all $P(\bar{H}_k(e_{ji})|\bar{D})$, for $k = 1, \dots, 4$, as follows:

$$\frac{P(\bar{H}_1(e_{ji})|\bar{D})}{P(\bar{H}_2(e_{ji})|\bar{D})} \qquad (10a)$$

where $\bar{D} = (\bar{D}^1, \bar{D}^2)$ denotes the data dets of all subjects across both groups, where $\bar{D}^1$ and $\bar{D}^2$ indicate the data of the subjects in group 1 and 2 respectively.

A. Computation of hypotheses

In the Bayesian framework, the hypothesis with the optimum $P(\bar{H}_k(e_{ji})|\bar{D})$ selected using (10a) is the hypothesis we accept as the one most likely to be true. Since the terms in (10a) are not directly computable, the next step in our derivation will be to show how we can connect the hypothesis in (10a) to the model in (3), such that we can compute or approximate $P(\bar{H}_k(e_{ji})|\bar{D})$ for all k.

First, the posterior probability ratio in (10a) can be written as:

$$\frac{P(\bar{H}_1(e_{ji})|\bar{D})}{P(\bar{H}_2(e_{ji})|\bar{D})} = \frac{P\left(\bar{D}\middle|\bar{H}_1(e_{ji})\right) P\left(\bar{H}_1(e_{ji})\right)}{P\left(\bar{D}\middle|\bar{H}_2(e_{ji})\right) P\left(\bar{H}_2(e_{ji})\right)}, \qquad (10b)$$

where $P\left(\bar{D}\middle|\bar{H}_k(e_{ji})\right)$ is the likelihood given $\bar{H}_k(e_{ji})$ and the hypothesis prior $P(\bar{H}_k(e_{ji}))$ is assumed to be uniform. Next using Bayes' rule, we can decompose the right had side of (10b) into:

$$\frac{P(\bar{D}^2|\bar{H}_1^2(e_{ji}), \bar{H}_1^1(e_{ji}), \bar{D}^1) P\left(\bar{D}^1\middle|\bar{H}_1^1(e_{ji}), \bar{H}_1^2(e_{ji})\right)}{P(\bar{D}^2|\bar{H}_2^2(e_{ji}), \bar{H}_2^1(e_{ji}), \bar{D}^1) P\left(\bar{D}^1\middle|\bar{H}_2^1(e_{ji}), \bar{H}_2^2(e_{ji})\right)}. \qquad (10c)$$

Equation (10c) can be further simplified because we assume that, given the hypotheses, the likelihood of $\bar{D}^2$ is not dependent on $\bar{D}^1$ and the likelihood of group 1 is independent of the hypothesis of group 2 and the other way around. Now, (10c) can be written as:

$$\frac{P\left(\bar{D}^2\middle|\bar{H}_1^2(e_{ji})\right) P\left(\bar{D}^1\middle|\bar{H}_1^1(e_{ji})\right)}{P\left(\bar{D}^2\middle|\bar{H}_2^2(e_{ji})\right) P\left(\bar{D}^1\middle|\bar{H}_2^1(e_{ji})\right)}. \qquad (10d)$$

Note now that we can use (10d) to find the optimum hypothesis $\bar{H}_k(e_{ji}) = (\bar{H}_k^1(e_{ji}),\ \bar{H}_k^2(e_{ji}))$ and that $(10d)$ can be separated into the left and right part, which can be used to find the optimum group hypotheses $\bar{H}_k^1(e_{ji})$ and $\bar{H}_k^2(e_{ji})$ separately, which can afterwards be used to decide the optimum $\bar{H}_k(e_{ji})$:

$$\frac{P\left(\bar{D}^g\middle|\bar{H}_1^g(e_{ji})\right)}{P\left(\bar{D}^g\middle|\bar{H}_2^g(e_{ji})\right)}, \qquad g = 1{,}2\,, \qquad (10e)$$

where, because $\bar{H}_k^1(e_{ji})$ and $\bar{H}_k^2(e_{ji})$ are both hypotheses on the existence of e_{ji} in their respective groups, we can use (10e) to select for each group, which one is more likely: $e_{ji} \in \mathcal{G}$ or $e_{ji} \notin \mathcal{G}$ for all subjects in the group. In this way for example we could find $e_{ji} \in \mathcal{G}$ for group 1 and $e_{ji} \notin \mathcal{G}$ for group 2, which corresponds with $\bar{H}^1(e_{ji})$ in Figure 8.

Next, we assume that the $P(D^{g,s}|H_k^{g,s}(e_{ji}))$ for $g = 1{,}2$ are independent over subjects s and that we can write $P(\bar{D}^g|\bar{H}_k^g(e_{ji}))$ in (10d) as:

$$P\left(\bar{D}^g\middle|\bar{H}_k^g(e_{ji})\right) = \prod_s P\left(D^{g,s}\middle|H_k^{g,s}(e_{ji})\right), g = 1{,}2. \qquad (11)$$

Recall that $H_k^{g,s}(e_{ji})$ is an assumption on a single edge in the graph of one subject. Each individual $P\left(D^{g,s}\middle|H_k^{g,s}(e_{ji})\right)$ can now be calculated by marginalization of $P(D^{g,s}|\mathcal{G})$, which we defined in (6), over all $\mathcal{G}$ for which $H_k^{g,s}(e_{ji})$ is true, i.e. $\mathcal{P} = \{\mathcal{G}|H_k^{g,s}(e_{ji}) = 1\}$:

$$P\left(D^{g,s}\middle|H_k^{g,s}(e_{ji})\right) = \sum_{\mathcal{G}\in\mathcal{P}} P(\mathcal{G})\, P(D^{g,s}|\mathcal{G}), \qquad (12)$$

where $P(D^{g,s}|\mathcal{G})$ is determined by the model in (3) and the integral in (6a). Recall that $P(\mathcal{G})$ is uniformly distributed and is thus a simple constant that can be pulled out of the sum. Furthermore, this constant is then present in both the numerator and denominator of (10) and is cancelled out.

In summary, we have defined a framework in which we can test the hypothesis of significant change in (10c). Furthermore, we have detailed how we can calculate the hypothesis in (10c) using the model in (3) and marginalization over $\mathcal{G}$ in (12).

However, graph set G in (12) still contains 2^{L^2-L} graphs and for large L the marginalization in (12) is computationally very

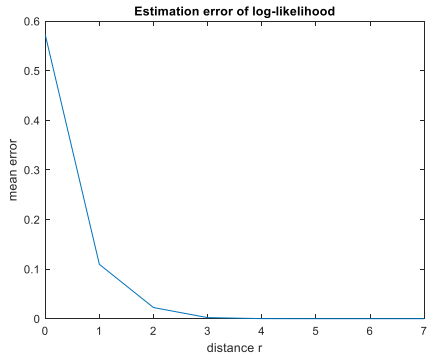

Figure 9: mean error of the estimate of $\log P\left(D\middle|H_k(e_{ji})\right)$*, for increasing size of* $\mathcal{P}_j$ *in (15).*

expensive. Therefore, in the next section, we will develop a method that efficiently approximates the marginalization in (12).

B. Approximation of marginalization of edges

First of all, note that from the definition of our model in (3) and $P(D|\mathcal{G})$ in (6) we see that the marginalization in (12) can be decomposed into L separate marginalization's:

$$P\left(D^{g,s}\middle|H_k^{g,s}(e_{ji})\right)=\prod_{j\in I}\sum_{\mathcal{G}_j\in\mathcal{P}_j}P(D_j^{g,s}|\mathcal{G}_j)P(\mathcal{G}_j) \qquad (13)$$

where $\mathcal{P}_j=\{\mathcal{G}_j|H_k^{g,s}(e_{ji})=1\}$. Only the marginalization in (13) which is affected by the hypothesis needs to be calculated as each other marginalization will result in the same value in both the numerator and denominator of (10) and is cancelled out.

Second, recall the graph estimate of the Bayesian method $\hat{\mathcal{G}}$, where it often appears that the probability mass of $P(D|\mathcal{G})$ is concentrated around graphs close to $P\left(D\middle|\hat{\mathcal{G}}\right)$, where we define the distance between graphs as:

Definition 1: The distance between two graphs $d(\mathcal{G}^1,\mathcal{G}^2)\geq 0$ is defined as the minimum amount of edge modifications required to transform one graph into the other.

This measurement of distance is symmetric, so: $d(\mathcal{G}^i,\mathcal{G}^j)=d(\mathcal{G}^j,\mathcal{G}^i)$ for any i,j. Now, we can define a distance r such that $r\leq L$, which defines a subset of graphs

$$\{\mathcal{G}_j|d(\hat{\mathcal{G}}_j,\mathcal{G}_j)\leq r\}, \qquad (14)$$

where for $r=L$ we recover the set of all possible graphs in (13). If the probability mass is sufficiently concentrated around $P\left(D\middle|\hat{\mathcal{G}}\right)$, then only the terms in the graph set in (14) need to be considered to calculate the marginal in (13), i.e. we define a new G_j that is a subset of (14):

$$\mathcal{P}_j=\{\mathcal{G}_j|d(\hat{\mathcal{G}}_j,\mathcal{G}_j)\leq r,\ H_k^{g,s}(e_{ji})=1\}. \qquad (15)$$

If the cardinality of the set in (15) for some given r is low relative to the full graph set in (12), then it is possible to approximate a computationally efficient estimate of the marginalization using (13) and (15).

C. Performance of marginalization approximation

To test the average performance of the efficient marginalization method, we generate 50 eight node random networks, in the same manner as in section IV and find $\hat{\mathcal{G}}$ using the Bayesian algorithm, for data length $N=300$. We choose a random edge in each graph and calculate the logarithm of the marginal in (13) using $\mathcal{P}_j=\{\mathcal{G}_j|d(\hat{\mathcal{G}}_j,\mathcal{G}_j)\leq r,H_k^{g,s}(e_{ji})=1\}$ for $r=0,\dots,7$, where $r=7$ indicates the complete graph set, given the hypothesis. We measure the performance of the marginalization through the absolute error between the logarithm of the likelihood $\log P(D|H_k(e_{ji}))$ for the complete graph set and the approximations for $r=0,\dots,6$. To obtain the overall performance we take the mean over all 50 networks for each r. The resulting average error is shown in Figure 9.

In Figure 9 we can see that for $r=3$ the approximation is already a very good estimate. Overall, we can see that we can indeed make a good approximation using a low r.

In summary, we want to test hypotheses by selecting the optimum hypothesis with (10a). We have made the test in (10a) computable as shown in (12). But since (12) involves a marginalization over all graphs, to make the computation more efficient, we have proposed an approximation of (12), which as we can see in Figure 9, works well in this case.

VII. Mozart Effect Study

In this section we will apply the methods from section V and the method we developed in section VI to real ICA time series to investigate if we can find evidence of an effect of Mozart music on the brain network connectivity of healthy subjects.

A. Data and Participants

We are provided with a collection of data sets of ICA time series resulting from the following experiment: Sixteen subjects of age 20-65 were part of the study. Four fMRI scans were performed on each subject, each scan one week apart from the other. Each subject was exposed to Mozart music for one week between the second and third scan and the total time of exposure to Mozart music was recorded.

Each fMRI scan has a spatial resolution of 3.5mm^3 and a temporal resolution of 2s. During each scan all subjects were instructed to remain at rest and to keep their eyes open. Each data set of one subject and one scan contains 20 ICA time series, each with 300 data points per time series. The brain regions that correspond to each of the ICA time series are detailed in Appendix A.

B. Model validation

Given the ICA time series of the Mozart data set, we want to verify whether our model assumptions in (3) hold for VAR models of the data. To test these assumptions, we perform a number of tests. First, we calculate a VAR model as in (3) using all 20 ICA time series for some model orders $m=1,\dots,5$ for all 64 data sets in the study, where $m=1$ is typically selected

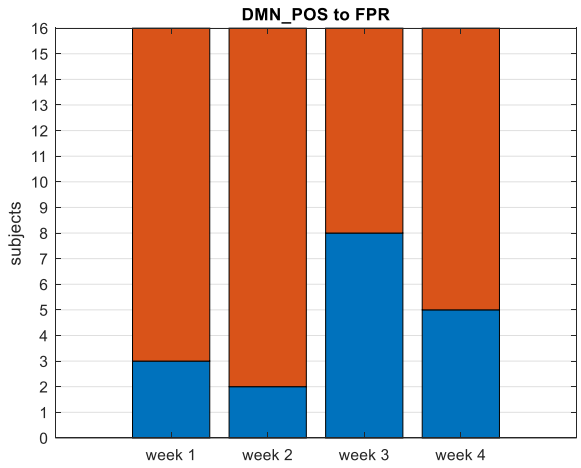

Figure 10: Selection frequency of the connection from default mode, posterior to fronto-parietal right.

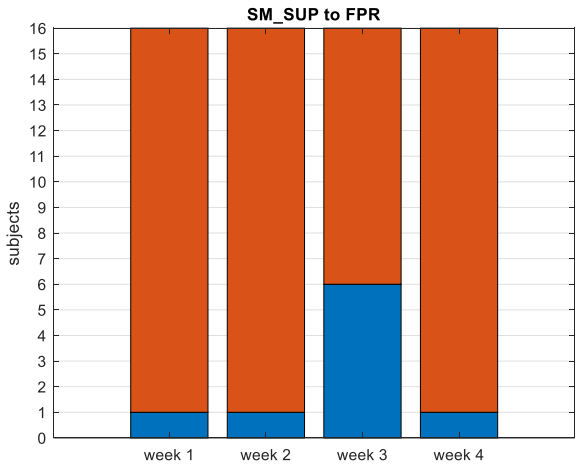

Figure 11: Selection frequency of the connection from sensori-motor superior to fronto-parietal right.

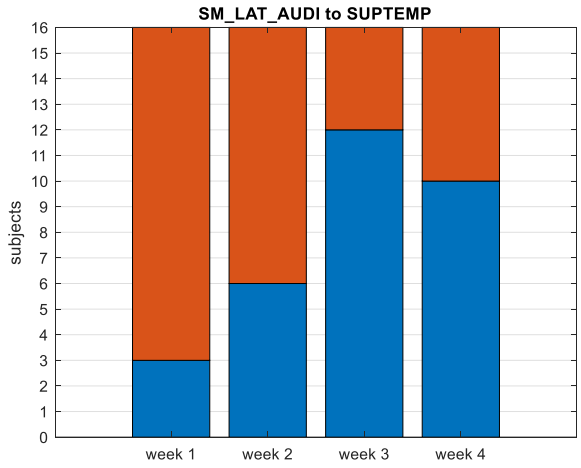

Figure 12: Selection frequency of the connection from sensori-motor, lateral to superior temporal gyrus.

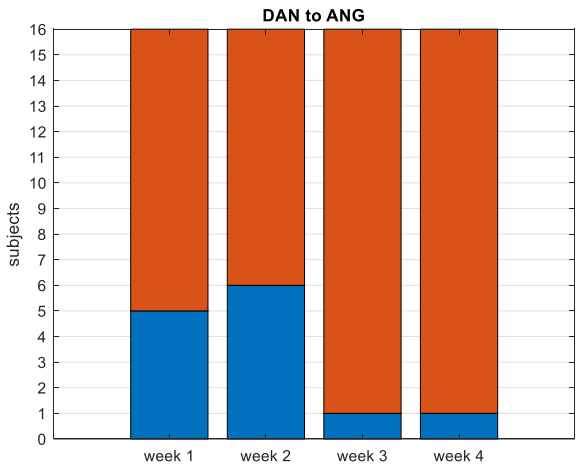

Figure 13: Selection frequency of the connection from dorsal attention network to angular gyrus.

via AIC. The parameters are calculated using ordinary least squares.

Next, we perform a whiteness test of the residuals of each VAR model, to test whether VAR models can sufficiently model the dynamics in the data. If the test is significant, it indicates that there are still unmodeled dynamics in the residual, i.e. it is not white noise. This results in statistics for 20 residuals for each of the 64 data sets, resulting in a total of 1280 statistics. Because the VAR models could have different model orders for which the whiteness test will no longer be significant, we describe in Table 1 the percentage of statistics that is still significant as m increases, thus if this percentage is 0 then all residuals of all VAR models in the analysis are white.

In Table 1 we see that as the model order m increases, an increasingly small percentage of all residuals test significant on the whiteness test. This indicates that the ICA time series can be modeled using (3) in the sense that there exists a m, for which the majority of residuals are white.

Thereafter, we inspect the covariance matrices of the residuals to test our assumption that the noise is independent over j in (3). We observe many correlations between the residual noise of nodes. Therefore, in our search for effects we will be critical of edges that show strong correlations between their residual noise, because correlated noise in our model in (3) can lead to false positives in the inference of connectivity based on these correlations.

TABLE I. PERCENTAGE OF SIGNIFICANT STATISTICS OF THE WHITENESS TEST OVER ALL RESIDUALS.

m	1	2	3	4	5
% significant	68%	16%	3.5%	0.8%	0.6%

C. Mozart data connectivity analysis

We can divide the data sets into four groups based on the weeks in which the scans were taken. Every subject in the study is thus a part of every group and there is no specific control group. Instead, we will use group 1 and 2, correspondent with

week 1 and 2, to infer natural variability in the connectivity of brain networks of the subjects. First, we test if no significant change in connectivity occurs between week 1 and 2. Then, for connections that showed no significant change in week 1 and 2, we can test if there is a significant connectivity change between week 2 and week 3, any changes in connectivity between these weeks is likely to be due to listening to Mozart music in between those scans. Finally, if there was a significant change in connectivity between week 2 and 3, we are interested if the change is lasting after week 3 by testing the hypothesis of no significant change between week 3 and 4.

Based on discussion in other studies [12] [13] [27] we expect to see some changes in connectivity related to motor and auditory brain regions. But mostly we are interested to see if there are any changes in connectivity involving regions that are part of cognitive processing in the brain.

Now that we have specified which effects we are interested in and have defined how we divide the available data into groups, we can apply the methods that we detailed in section V and VI to test for these effects.

D. Bayesian selection frequency

First, given now that we have limited prior knowledge on the possible changes in connectivity in the brain due to Mozart music, we will first search for likely candidates for possible effects using the Bayesian selection frequency.

We will look for edges which have a difference in selection frequency ≥ 5 between week 2 and 3. Then, we discard any edges if the difference in selection frequency between week 1 and 2 is large relative to the difference between week 2 and 3. Using these criteria we report here a selection of the most interesting connections in Figure 10 to 13, the maps of the brain regions involved in these connections are shown in Appendix B.

In Figure 10 and 11 we found two connections to the fronto-parietal right network, which is of interest because this network performs cognitive processing. In Figure 12 we show the selection frequency of the connection between the sensori-motor, lateral and the superior temporal gyrus. Figure 12 is of interest especially because the effect is of a lasting duration, as the connectivity is still very high in week 4. Both networks in this connection are related to auditory processing, which indicates that the increase in connectivity might be attributed simply to the fact that the subjects were listening to music. Last, in Figure 13 we show a connection between the dorsal attention network and the angular gyrus. This edge is of interest, because it shows that a connection could also disappear as a result from listening to Mozart music.

In summary, we have found a small selection of interesting effects in our analysis of the data using the selection frequency. However, while we see some change in connectivity occur between week 2 and 3, the effects are never universal over all subjects and in most cases the edge is not even selected for more than half of the subjects in any week, except in Figure 12. Now, we will now use the methods of section V and VI to test if any changes are significant.

E. Significant change in Granger-values

First, we calculate the Granger-values for each subject over all weeks, for the edges we selected in the previous subsection. Then, we use the two-sample two-tailed t-test in (8), where we calculate the 5% significance threshold with permutation testing [21] of the statistic in (8) and using the Granger-values of all 16 subjects across the two weeks we are comparing.

Next, for all four of the connection in figure 10 to 13, we found no significant difference in Granger-values between week 1 and 2. Then we compare the Granger-values of week 2 and 3 and found no significant evidence of a change in Granger-values. The most significant result we found was in the comparison of week 2 and 3 for the connection in figure 10, from the sensori-motor lateral to the superior temporal gyrus, where in week 2 the average granger value $\mathbb{E}_s[\mathcal{F}^{2,s}] = 0.0105$ with standard deviation $\sigma_2 = 0.0172$ and for week 3, $\mathbb{E}_s[\mathcal{F}^{3,s}] = 0.0205$ with $\sigma_3 = 0.0355$. We can see that the standard deviation of both weeks is larger than the mean of the Granger-values, which is likely why we found no significance using the test in (9). Furthermore, because we found no significant change between week 2 and 3, we do not infer any differences in Granger-values between week 3 and 4, as no effect was detected.

In summary, our analysis of differences in Granger-values yields no outcomes that support our finding in the previous subsection.

F. Bayesian hypothesis test

First, we perform the inference for each subject separately using the Bayesian method in section IV. Then using the graph estimate of each subject over all weeks we calculate an approximation of the marginal in (13), where we have chosen $r = 6$ to build the graph set in (14).

Now, recall the four hypotheses of change we defined in Figure 8 in section VI and the definition of the criteria on the changes in connectivity we want to see, as defined in this section. Now we see that the hypothesis between week 1 and 2 is $\bar{H}_4(e_{ji})$ for the connections in Figures 10 to 12 and $\bar{H}_3(e_{ji})$ for the connection in Figure 13. The hypothesis between week 2 and 3 is then $\bar{H}_1(e_{ji})$ for the connections in Figure 10 to 12 and $\bar{H}_2(e_{ji})$ for the connection in Figure 13. Finally, the hypothesis between week 3 and 4 is $\bar{H}_3(e_{ji})$ for the connections in Figures 10 to 12 and $\bar{H}_4(e_{ji})$ for the connection in Figure 13.

We found that for all considered connections the hypothesis of change between week 2 and 3, $\bar{H}_1(e_{ji})$ and $\bar{H}_2(e_{ji})$ were false and thus no significant change is detected as a result of Mozart music. As no significant changes were detected between week 2 and 3, we do not evaluate between week 1 and 2 and between week 3 and 4 as any results are irrelevant if no change between week 2 and 3 is detected.

In summary, similar to the previous subsection no significant effects were found. Next, we will divide the 16 subjects into two groups of 8 based on their listening time and try again if we can find any significance for subjects that listened longer or shorter than the median value of 19 hours.

G. Inference using listening time

In our final test, the subjects are divided into two groups of 8 subjects. The subjects in group 1 all listened between 19 and 37 hours to Mozart music and the subjects in group 2 all listened between 14 and 19 hours to Mozart music. Both analysis methods were repeated for both of the new groups.

Again, no significant result was found for the test in (9) of the Granger-values for both groups, all comparisons between weeks for both the long and short listeners had no significant change in Granger-values.

Then, using the Bayesian significance test on both groups, and using the same hypotheses as in the inference on the complete group, we have found two significant results. One for group 1 of long listeners in the connection from the default mode, posterior to the fronto-parietal right networks, which was of short duration, and one for group 2 of short listeners in the connection from the sensori-motor lateral to the superior temporal gyrus, which was also of short duration. These finding seem to be mostly in line with the selection frequency plots in Figure 10 and 12, except for in Figure 12 the selection frequency would suggest the effect might be of a longer duration while this is not reflected by the significance test.

To illustrate these findings and how strong the evidence was in favor of the hypotheses, we recall equation (10e) in section VI and calculate the group hypotheses, given the hypotheses between weeks defined above and calculate the $2log$ of the ratio in (10e) for each group hypothesis. Here when evaluating, for example if the hypothesis of a change in connectivity between week 1 and 2 is true, then both log-ratios corresponding to the hypothesis are positive and the hypothesis is optimum among other alternative hypotheses. Furthermore, the log-ratios are an indication of the strength of evidence, defined in [1].

Now for the connection from group 1 of long listeners, the group hypothesis ratios are: for week 1 is 21.37, for week 2 is 56.62, for week 3 is 13.77 and for week 4 is -10.76. And for the connection in group of short listeners, the group hypothesis ratios are: for week 1 is 2.47, for week 2 is 32.66, for week 3 is 91.92 and for week 4 is -15.88.

Since for both connections the group hypotheses of week 1 and 2 are positive, the optimum hypothesis between week 1 and 2 is indeed $\bar{H}_4(e_{ji})$ for both connections. Since the group hypotheses between week 2 and 3 are both positive, the optimum hypothesis between week 2 and 3 is $\bar{H}_1(e_{ji})$ for both connections and last the group hypothesis of week 4 is negative for both connections and so both connections don't exist in week 4, and as such the hypothesis between week 3 and 4, $\bar{H}_3(e_{ji})$ is false and both connection are of short duration.

VIII. Conclusions and Discussion

Bayesian topology identification is evaluated as an alternative to Granger-causality analysis for the inference of brain network connectivity. if we want high accuracy, especially given small N, relative to the number of nodes, we recommend the Bayesian algorithm. In a situation where N and/or L are large, the difference in performance between the methods is small and the optimum threshold appears fixed, Granger-causality analysis becomes a more attractive solution.

While we have focused our approaches on fMRI data inference in this paper, the Bayesian method could also be applied to other types of data that can describe neural activity. An especially interesting application would be in inference of connectivity given EEG data, as it tends to have low spatial resolution, which is where the Bayesian method performs better than Granger-causality analysis.

An extension to the Bayesian method is developed that can test the hypothesis of significant difference of a connection between two groups of subjects. A graph marginalization approximation method is developed in an attempt to reduce the cardinality of the marginalization set. And using numerical results the effectiveness of the method is shown.

However, compared to the group test for Granger-values, the Bayesian hypothesis test can only detect binary changes in connections, i.e. it exists in one group and does not in the other group, while the group test for Granger-values can detect continuous changes in Granger-values and thus in connectivity. In future research we will look for methods to extend the Bayesian hypothesis test beyond binary changes.

Furthermore, while we have limited ourselves to hypotheses on a single connection in the Bayesian hypothesis test, we could extend the method to hypotheses on multiple connections, i.e. local structure in the graph. It would be very interesting to evaluate how these hypotheses on local structure could be used to infer brain network connectivity.

Last, in the application of both Granger and Bayesian methods to inference of the Mozart music study, we have found no significant results for the Granger method. A likely explanation for this is the large standard deviation in Granger-values for each week, which is supported by the large standard deviation of the subjects in week 2 and 3 of the connection from the sensori-motor lateral network to the superior temporal gyrus.

For the Bayesian hypothesis test, we found two significant results, when we divide the subjects into two equal sized groups based on listening time. For the group which listened longer, a significant change was found in the connection from the posterior default mode network to the fronto parietal right network, which was of short duration as the connection disappeared in week 4. This connection is between two networks related to cognitive processing, which could be an indication of a Mozart effect. For the group that listened less long, we found a significant change in connectivity in the connection from the sensori-motor lateral network to the superior temporal gyrus, which was of short duration as there was no change in connectivity between week 3 and 4. The brain regions of this connection are related to auditory processing so the change might not be related to Mozart music but to listening to music in general.

From our analysis we see that a change in brain network connectivity occurred in a connection related to cognitive networks, but only for the group that listened longer. The Mozart effect study would have benefited from a longer time period between scan 2 and 3, where the subjects listened to

Mozart music, because longer exposure might have resulted in more significant results. Also, in the current study the listening time of each subject is variable, which could lead to some subjects to show changes in connectivity, while others do not, due to the exposure time being too short.

While we have focused on the inference of connectivity using ICA time series, it would be interesting to evaluate if there are differences in the inferred connectivity if we use other ways to summarize fMRI data.

REFERENCES

[1] R.E. Kass, A.E. Raftery, "Bayes factors," *Journal of the American Statistical Society*, vol. 90, no. 430, pp. 773-795, 1995.

[2] C. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica 37*, pp. 424-438, 1969.

[3] A. Bernas A.P. Aldenkamp S. Zinger, "Wavelet coherence-based classifier: A resting-state functional MRI study on neurodynamics in adolescents with high-functioning autism," *Methods and Programs in Biomedicine 154*, pp. 143-151, 2018.

[4] A.B. Barrett, M. Murpy, M.A. Bruno, "Granger causality analysis of steady state electroencefalographic signals during propofol-induced aneasthesia.", *PLoS ONE*, 7(1), 2012.

[5] A. Bernas E. Barendse A. Aldenkamp W. Backes P. Hofman, "Brain resting-state networks in adolescents with high-functioning autism: analysis of spatial connectivity and temporal neurodynamics," *Brain and Behaviour*, 2017.

[6] K.J. Friston, L. Harrison and W. Penny, "Dynamic causal modelling," *Neuroimage*, vol. 19, no. 4, pp. 1273-1302, 2003.

[7] Y. Wang M. Sznaier O. Camps, "A super-atomic norm minimization approach to identifying sparse dynamical graphical models," in American Control Conference (ACC), 2016. IEEE, 2016, pp. 1962– 1967.

[8] A. Marreiros S. Kiebel K. Friston, "Dynamic causal modelling for fmri: a two-state model," *Neuroimage*, vol. 39, no. 1, pp. 269-278, 2008.

[9] D. Materassi M. Salapaka, "On the problem of reconstructing an unknown topology via locality properties of the Wiener filter," *IEEE transactions on automatic control*, vol. 57, no. 7, 2012.

[10] S.Shi G. Bottegal P.M.J. van den Hof, "Bayesian topology identification of linear dynamic networks," *ECC*, 2019.

[11] A. Chiuso, G. Pillonetto, "A Bayesian approach to sparse dynamic network identification," *Automatica*, vol. 48, no. 8, pp. 1553-1565, 2012.

[12] G. Coppola, A. Toro, F.F. Operto, "Mozart's music in children with drug-refractory epileptic encephalopathies," *Epilepsy and Behaviour*, vol. 50, pp. 18-22, 2015.

[13] W. Verrusio, E. Ettorre, E. Vicenzini, "The Mozart effect: a quatitative EEG study," *Consciousness and Cognition*, vol. 35, pp. 150-155, 2015.

[14] S. Ogawa T.M. Lee A.R. Kay D.W. Tank, "Brain magnetic resonance imaging with contrast dependent on blood oxygenation," *Proc. Natl. Acad. Sci.*, vol. 87, pp. 9868-9872, 1990.

[15] A. Hyvärinen J. Karhunen E. Oja, "Independent Component Analysis," *Wiley*, New York, 2001.

[16] C.F. Beckmann C.E. Mackay, N. Filippini S.M. Smith, "Group comparison of resting-state fMRI data using multi-subject ICA and dual regression", Hum. Brain Mapp. Conf., p. 181, 2009.

[17] M. Tipping C. Bishop, "Mixtures of probabilistic principal component analyzers," *Neural Computation*, vol. 11, no. 2, pp. 443-482, 1999.

[18] S. Bressler A. Seth, "Wiener-Granger Causality: A well-established methodology," *Neuroimage 58*, pp. 323-329, 2011.

[19] A. Seth P Chorley L. Barnett, "Granger causality analysis of fMRI BOLD signals is invariant to hemodynamic convolution but not downsampling," *Neuroimage*, vol. 65, pp. 540-555, 2013.

[20] H. Akaike, "A new look at the statistical Model Identification," *IEEE Transactions on automatic control*, vol. ac-19, no. 6, 1974.

[21] R. Yates D. Goodman, "Probability and stochastic processes," *Wiley*, Second Edition, 2005.

[22] L. Barnett A. Seth, "The MVGC multivariate Granger causality toolbox: A new approach to Granger-causal inference," *Journal of Neuroscience Methods 223*, pp. 50-68, 2014.

[23] C.M. Bishop, *Pattern recognition and Machine learning*, Springer, 2009.

[24] D.M. Chickering, "Optimal structure identification with greedy search," *Journal of Machine Learning Research*, vol. 3, pp. 507-554, november, 2002.

[25] T. Nichols A. Holmes, "Nonparametric permutation tests for functional neuroimaging: a primer with examples," *Human Brain Mapping 15*, pp. 1-25, 2001.

[26] R. Buxton E. Wong L. Frank, "Dynamics of blood flow and oxygenation changes during brain activation: The balloon model," *MRM 39*, pp. 855-864, 1998.

[27] F.H. Rauscher, G.L. Shaw, K.N. Ky, "Music and spatial task performance," *Nature*, 1993;365:611.

[28] L. Ljung, *System identification: Theory for the User*, Prentice-Hall, 1987.

APPENDIX

A. Table of brain regions in the Mozart study

TABLE II. ALL BRAIN REGIONS OF WHICH THE CORRESPONDING ICA TIME SERIES WERE USED IN THE MOZART STUDY.

#	***Shorthand***	***Name***
1	DMN_ANT	Default mode network, anterior
2	MED_VISU	Medial visual cortex
3	OCC_LAT_VISU	Occipital and lateral visual cortex
4	DAN	Dorsal attention network
5	FPR	Fronto-parietal right
6	SM_LAT_AUDI	Sensori-motor lateral (with auditory parts)
7	FPL	Fronto-parietal left
8	VAN	Ventral attention network
9	LING_FUS	Lingual fusiform cortex
10	DMN_POS	Default mode network, posterior
11	Noise	n/a
12	SAL_AUDI	Salience-auditory network, basal ganglia
13	ANG	Angular gyrus
14	Noise	White matter
15	SUPTEMP	Superior temporal gyrus
16	CEN	Central executive network
17	Noise	n/a
18	SM_SUP	Sensori motor medial superior
19	CEREB	Cerebellum
20	Noise	Borders and Movements

B. Maps of brain regions

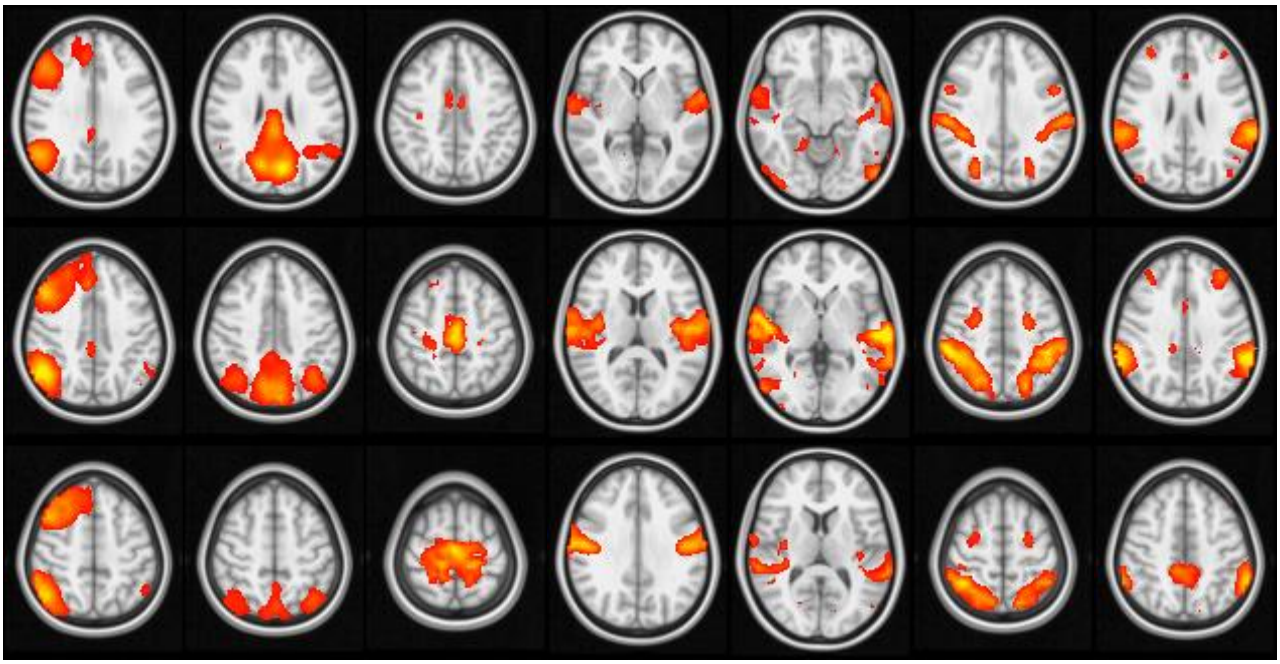

Figure 14: Brain regions of the connections shown in Figures 10 to 13. In each column from left to right: FPR, DMN_ANT, SM_SUP, SM_LAT_AUDI, SUPTEMP, DAN, ANG.