

**Estimating and Inferring the Maximum Degree of Stimulus-Locked
Time-Varying Brain Connectivity Networks**

Kean Ming Tan

Department of Statistics, University of Michigan, Ann Arbor, Michigan, U.S.A.

email: keanming@umich.edu

and

Junwei Lu

Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, U.S.A.

email: junweilu@hsph.harvard.edu

and

Tong Zhang

Department of Computer Science and Engineering, The Hong Kong University of Science and Technology,
Clear Water Bay, Kowloon, Hong Kong.

email: tongzhang@tongzhang-ml.org

and

Han Liu

Department of Electrical Engineering and Computer Science, Northwestern University,
Evanston, Illinois, U.S.A.

email: hanliu@northwestern.edu

December 2008

SUMMARY: Neuroscientists have enjoyed much success in understanding brain functions by constructing brain connectivity networks using data collected under highly controlled experimental settings. However, these experimental settings bear little resemblance to our real-life experience in day-to-day interactions with the surroundings. To address this issue, neuroscientists have been measuring brain activity under natural viewing experiments in which the subjects are given continuous stimuli, such as watching a movie or listening to a story. The main challenge with this approach is that the measured signal consists of both the stimulus-induced signal, as well as intrinsic-neural and non-neuronal signals. By exploiting the experimental design, we propose to estimate stimulus-locked brain network by treating non-stimulus-induced signals as nuisance parameters. In many neuroscience applications, it is often important to identify brain regions that are connected to many other brain regions during cognitive process. We propose an inferential method to test whether the maximum degree of the estimated network is larger than a pre-specific number. We prove that the type I error can be controlled and that the power increases to one asymptotically. Simulation studies are conducted to assess the performance of our method. Finally, we analyze a functional magnetic resonance imaging dataset obtained under the Sherlock Holmes movie stimuli.

KEY WORDS: Gaussian multiplier bootstrap; Hypothesis testing; Inter-subject; Latent variables; Maximum degree; Subject specific effects.

1. Introduction

In the past few decades, much effort has been put into understanding task-based brain connectivity networks. For instance, in a typical visual mapping experiment, subjects are presented with a simple static visual stimulus and are asked to maintain fixation at the visual stimulus, while their brain activities are measured. Under such highly controlled experimental settings, numerous studies have shown that there are substantial similarities across brain connectivity networks constructed for different subjects (Hasson et al., 2003). However, such experimental settings bear little resemblance to our real-life experience in several aspects: natural viewing consists of a continuous stream of perceptual stimuli; subjects can freely move their eyes; there are interactions among viewing, context, and emotion (Hasson et al., 2004). To address this issue, neuroscientists have started measuring brain activity under continuous natural stimuli, such as watching a movie or listening to a story (Hasson et al., 2004; Simony et al., 2016; Chen et al., 2017). The main scientific question is to understand the dynamics of the brain connectivity network that are specific to the natural stimuli.

Graphical models have been used in modeling brain connectivity networks: graphical models encode conditional dependence relationships between each pair of brain regions, given the others. A graph consists of d nodes, each representing a random variable, and a set of edges joining pairs of nodes corresponding to conditionally dependent variables. We refer the reader to Drton and Maathuis (2017) for a review on learning the structure of undirected graphical models. Under natural continuous stimuli, it is often of interest to estimate a dynamic brain connectivity network, i.e., a graph that changes over time. A natural candidate for this purpose is the time-varying Gaussian graphical model (Zhou et al., 2010; Kolar et al., 2010). The time-varying Gaussian graphical model assumes

$$\mathbf{X}(z) \mid Z = z \sim N_d\{\mathbf{0}, \boldsymbol{\Sigma}_X(z)\}, \quad (1)$$

where $\boldsymbol{\Sigma}_X(z)$ is the covariance matrix of $\mathbf{X}(z)$ given $Z = z$, and $Z \in [0, 1]$ has a continuous

density. The inverse covariance matrix $\{\Sigma_X(z)\}^{-1}$ encodes conditional dependence relationships between pairs of random variables at time $Z = z$: $\{\Sigma_X(z)\}_{jk}^{-1} = 0$ if and only if the j th and k th variables are conditionally independent given the other variables at time $Z = z$.

In natural viewing experiments, the main goal is to construct a brain connectivity network that is locked to the processing of external stimuli, referred to as a *stimulus-locked network* (Simony et al., 2016; Chen et al., 2017; Regev et al., 2018; Musch et al., 2020). Constructing a stimulus-locked network can better characterize the dynamic changes of brain patterns across the continuous stimulus (Simony et al., 2016). The main challenge in constructing stimulus-locked network is the lack of highly controlled experiments that remove spontaneous and individual variations. The measured blood-oxygen-level dependent (BOLD) signal consists of not only signal that is specific to the stimulus, but also intrinsic neural signal (random fluctuations) and non-neuronal signal (physiological noise) that are specific to each subject. The intrinsic neural signal and non-neuronal signal can be interpreted as measurement error or latent variables that confound the stimuli-specific signal. We refer to non-stimulus-induced signals as subject specific effects throughout the manuscript. Thus, directly fitting (1) using the measured data will yield a time-varying graph that primarily reflects intrinsic BOLD fluctuations within each brain rather than BOLD fluctuations due to the natural continuous stimulus.

We exploit the experimental design aspect of natural viewing experiments and propose to estimate a dynamic stimulus-locked brain connectivity network by treating the intrinsic and non-neuronal signals as nuisance parameters. Our proposal exploits the fact that the same stimulus will be given to multiple independent subjects, and that the intrinsic neural and non-neuronal signals for different subjects are independent. This motivates us to estimate a brain connectivity network across two brains rather than within each brain. In fact, Simony

et al. (2016) considered the aforementioned idea where they estimated brain connectivity networks by calculating pairwise covariance for brain regions between two brains.

After estimating the stimulus-locked brain connectivity network, the next important question is to infer whether there are any regions of interest that are connected to many other regions during cognitive process (Hagmann et al., 2008). These highly connected brain regions are referred to as *hub nodes*, and the number of connections for each brain region is referred to as *degree*. Identifying hub brain regions that are specific to the given natural continuous stimulus will lead to a better understanding of the cognitive processes in the brain, and may shed light on various cognitive disorders. Several authors have proposed methods to estimate networks with hubs (see, for instance, Tan et al., 2014). In this paper, we instead focus on developing a novel inferential framework to test the hypothesis whether there exists at least one time point such that the maximum degree of the graph is greater than k .

The proposed inferential framework is motivated by two major components: (1) the Gaussian multiplier bootstrap for approximating the distribution of supreme of empirical processes (Chernozhukov et al., 2013, 2014), and (2) the step-down method for multiple hypothesis testing problems (Romano and Wolf, 2005). Neykov et al. (2019) proposed a framework for testing general graph structures on a static graph. In Appendix A, we show that our proposed method can be extended to testing graph structures similar to that of Neykov et al. (2019).

2. Stimulus-Locked Time-Varying Brain Connectivity Networks

2.1 A Statistical Model

Let $\mathbf{X}(z)$, $\mathbf{S}(z)$, $\mathbf{E}(z)$ be the observed data, stimulus-induced signal, and subject specific effects at time $Z = z$, respectively. Assume that Z is a continuous random variable with a continuous density. For a given $Z = z$, we model the observed data as the summation of

stimulus-induced signal and the subject specific effects:

$$\mathbf{X}(z) = \mathbf{S}(z) + \mathbf{E}(z), \quad \mathbf{S}(z) | Z = z \sim N_d\{\mathbf{0}, \boldsymbol{\Sigma}(z)\}, \quad \mathbf{E}(z) | Z = z \sim N_d\{\mathbf{0}, \mathbf{L}_X(z)\}, \quad (2)$$

where $\boldsymbol{\Sigma}(z)$ is the covariance matrix of the stimulus-induced signal, and $\mathbf{L}_X(z)$ is the covariance matrix of the subject specific effects. We assume that $\mathbf{S}(z)$ and $\mathbf{E}(z)$ are independent for all z . Thus, estimating the stimulus-locked brain connectivity network amounts to estimating $\{\boldsymbol{\Sigma}(z)\}^{-1}$. Fitting the model in (1) using $\mathbf{X}(z)$ will yield an estimate of $\{\boldsymbol{\Sigma}(z) + \mathbf{L}_X(z)\}^{-1}$, and thus, (1) fails to estimate the stimulus-locked brain connectivity network $\{\boldsymbol{\Sigma}(z)\}^{-1}$.

To address this issue, we exploit the experimental design aspect of natural viewing experiments. In many studies, neuroscientists often measure brain activity for multiple subjects under the same continuous natural stimulus (Chen et al., 2017; Simony et al., 2016). Let $\mathbf{X}(z)$ and $\mathbf{Y}(z)$ be measured data for two subjects at time point $Z = z$. Since the same natural stimulus is given to both subjects, this motivates the following statistical model:

$$\begin{aligned} \mathbf{X}(z) = \mathbf{S}(z) + \mathbf{E}_X(z), \quad \mathbf{Y}(z) = \mathbf{S}(z) + \mathbf{E}_Y(z), \quad \mathbf{S}(z) | Z = z \sim N_d\{\mathbf{0}, \boldsymbol{\Sigma}(z)\}, \\ \mathbf{E}_X(z) | Z = z \sim N_d\{\mathbf{0}, \mathbf{L}_X(z)\}, \quad \mathbf{E}_Y(z) | Z = z \sim N_d\{\mathbf{0}, \mathbf{L}_Y(z)\}, \end{aligned} \quad (3)$$

where $\mathbf{S}(z)$ is the stimulus-induced signal, and $\mathbf{E}_X(z)$ and $\mathbf{E}_Y(z)$ are the subject specific effects at $Z = z$. Model (3) motivates the calculation of *inter-subject covariance* between two subjects rather than the within-subject covariance. For a given time point $Z = z$, we have

$$\mathbb{E}[\mathbf{X}(z)\{\mathbf{Y}(z)\}^T | Z = z] = \mathbb{E}[\mathbf{S}(z)\{\mathbf{S}(z)\}^T | Z = z] + \mathbb{E}[\mathbf{E}_X(z)\{\mathbf{E}_Y(z)\}^T | Z = z] = \boldsymbol{\Sigma}(z).$$

That is, we estimate $\boldsymbol{\Sigma}(z)$ via the inter-subject covariance by treating $\mathbf{L}_X(z)$ and $\mathbf{L}_Y(z)$ as nuisance parameters. In the neuroscience literature, several authors have proposed to calculate an inter-subject covariance matrix to estimate marginal dependencies among brain regions that are stimulus-locked, and have found that such an approach better captures the stimulus-locked marginal relationships among pairs of brain regions (Simony et al., 2016).

For simplicity, throughout the paper, we focus on two subjects. When there are multiple subjects, we can split the subjects into two groups, and obtain an average of each group to

estimate the stimulus-locked brain network. We also discuss a U -statistic type estimator for the case when there are multiple subjects in Appendix B.

2.2 Inter-Subject Time-Varying Gaussian Graphical Models

We now propose inter-subject time-varying Gaussian graphical models for estimating stimulus-locked time-varying networks. Let $(Z_1, \mathbf{X}_1, \mathbf{Y}_1), \dots, (Z_n, \mathbf{X}_n, \mathbf{Y}_n)$ be n independent realizations of the triplets $(Z, \mathbf{X}, \mathbf{Y})$. Both subjects share the same Z_1, \dots, Z_n since they are given the same continuous stimulus. Let $K : \mathbb{R} \rightarrow \mathbb{R}$ be a symmetric kernel function. To obtain an estimate for $\Sigma(z)$, we propose the inter-subject kernel smoothed covariance estimator

$$\widehat{\Sigma}(z) = \frac{\sum_{i \in [n]} K_h(Z_i - z) \mathbf{X}_i \mathbf{Y}_i^T}{\sum_{i \in [n]} K_h(Z_i - z)}, \quad (4)$$

where $K_h(Z_i - z) = K\{(Z_i - z)/h\}/h$, $h > 0$ is the bandwidth parameter, and $[n] = \{1, \dots, n\}$. For simplicity, we use the Epanechnikov kernel $K(u) = 0.75 \cdot (1 - u^2) \cdot \mathbb{1}_{\{|u| \leq 1\}}$, where $\mathbb{1}_{\{|u| \leq 1\}}$ is an indicator function that takes value one if $|u| \leq 1$ and zero otherwise. The choice of kernel is not essential as long as it satisfies the regularity conditions in Section 5.1.

Let $\Theta(z) = \{\Sigma(z)\}^{-1}$. Given the kernel smoothed inter-subject covariance estimator in (4), there are multiple approaches to obtain an estimate of the inverse covariance matrix $\Theta(z)$. We consider the CLIME estimator proposed by Cai et al. (2011). Let \mathbf{e}_j be the j th canonical basis in \mathbb{R}^d . For a vector $\mathbf{v} \in \mathbb{R}^d$, let $\|\mathbf{v}\|_1 = \sum_{j=1}^d |v_j|$ and let $\|\mathbf{v}\|_\infty = \max_j |v_j|$. For each $j \in [d]$, the CLIME estimator takes the form

$$\widehat{\Theta}_j(z) = \underset{\boldsymbol{\theta} \in \mathbb{R}^d}{\operatorname{argmin}} \|\boldsymbol{\theta}\|_1 \quad \text{subject to} \quad \left\| \widehat{\Sigma}(z) \cdot \boldsymbol{\theta} - \mathbf{e}_j \right\|_\infty \leq \lambda, \quad (5)$$

where $\lambda > 0$ is a tuning parameter that controls the sparsity of $\widehat{\Theta}_j(z)$. We construct an estimator for the stimulus-locked brain network as $\widehat{\Theta}(z) = [\{\widehat{\Theta}_1(z)\}^T, \dots, \{\widehat{\Theta}_d(z)\}^T]^T$.

There are two tuning parameters in our proposed method: a bandwidth parameter h that controls the smoothness of the estimated covariance matrix, and a tuning parameter λ that controls the sparsity of the estimated network. The bandwidth parameter h can be selected according to the scientific context. For instance, in many neuroscience applications that

involve continuous natural stimuli, we select h such that there are always at least 30% of the time points that have non-zero kernel weights. In the following, we propose a L -fold cross-validation type procedure to select λ . We first partition the n time points into L folds. Let C_ℓ be an index set containing time points for the ℓ th fold. Let $\Theta(z)^{(-\ell)}$ be the estimated inverse covariance matrix using data excluding the ℓ th fold, and let $\Sigma(z)^{(\ell)}$ be the estimated kernel smoothed covariance estimated using data only from the ℓ th fold. We calculate the following quantity for various values of λ :

$$cv_\lambda = \frac{1}{L} \sum_{\ell=1}^L \sum_{i \in C_\ell} \|\widehat{\Sigma}(z_i, \lambda)^{(\ell)} \widehat{\Theta}(z_i, \lambda)^{(-\ell)} - \mathbf{I}_d\|_{\max}, \quad (6)$$

where $\|\cdot\|_{\max}$ is the element-wise max norm for matrix. Let λ_{\min} be the λ value that yields the minimum cross-validation error $cv_{\lambda_{\min}}$ across a range of values of λ . From performing extensive numerical studies, we find that picking λ_{\min} tends to lead to more false positives in terms of identifying the edges. We instead propose to pick the largest λ that yields a cv_λ that is less than $cv_{\lambda_{\min}}$ plus its corresponding two standard errors across the L folds.

2.3 Inference on Maximum Degree

We consider testing the hypothesis:

$$H_0 : \text{for all } z \in [0, 1], \text{ the maximum degree of the graph is not greater than } k, \quad (7)$$

$$H_1 : \text{there exists a } z_0 \in [0, 1] \text{ such that the maximum degree of the graph is greater than } k.$$

In the existing literature, many authors have proposed to test whether there is an edge between two nodes in a graph (see, Neykov et al., 2018, and the references therein). Due to the ℓ_1 penalty used to encourage a sparse graph, classical test statistics are no longer asymptotically normal. We employ the de-biased test statistic

$$\widehat{\Theta}_{jk}^{\text{de}}(z) = \widehat{\Theta}_{jk}(z) - \frac{\left\{ \widehat{\Theta}_j(z) \right\}^T \left\{ \widehat{\Sigma}(z) \widehat{\Theta}_k(z) - \mathbf{e}_k \right\}}{\left\{ \widehat{\Theta}_j(z) \right\}^T \widehat{\Sigma}_j(z)}, \quad (8)$$

where $\widehat{\Theta}_j(z)$ is the j th column of $\widehat{\Theta}(z)$. The subtrahend in (8) is the bias introduced by imposing an ℓ_1 penalty during the estimation procedure.

We use (8) to construct a test statistic for testing the maximum degree of a time-varying

graph. Let $G(z) = \{V, E(z)\}$ be an undirected graph, where $V = \{1, \dots, d\}$ is a set of d nodes and $E(z) \subseteq V \times V$ is a set of edges connecting pairs of nodes. Let

$$T_E = \sup_{z \in [0,1]} \max_{(j,k) \in E(z)} \sqrt{nh} \cdot \left| \widehat{\Theta}_{jk}^{\text{de}}(z) - \Theta_{jk}(z) \right| \cdot \left\{ \frac{1}{n} \sum_{i \in [n]} K_h(Z_i - z) \right\}. \quad (9)$$

The edge set $E(z)$ is defined based on the hypothesis testing problem. In the context of testing maximum degree of a time-varying graph as in (7), $E(z) = V \times V$, and therefore the maximum is taken over all possible edges between pairs of nodes. We will use the notation $E(z)$ to indicate some predefined known edge set. Note that the edge set will be different for testing different graph structures, and we refer the reader to Appendix A for details.

Since the test statistic (9) involves taking the supreme over z and the maximum over all edges in $E(z)$, it is challenging to evaluate its asymptotic distribution. To this end, we generalize the Gaussian multiplier bootstrap proposed in Chernozhukov et al. (2013) and Chernozhukov et al. (2014) to approximate the distribution of the test statistic T_E . Let $\xi_1, \dots, \xi_n \stackrel{\text{i.i.d.}}{\sim} N(0, 1)$. We construct the bootstrap statistic as

$$T_E^B = \sup_{z \in [0,1]} \max_{(j,k) \in E(z)} \sqrt{nh} \cdot \left| \frac{\sum_{i \in [n]} \left\{ \widehat{\Theta}_j(z) \right\}^T K_h(Z_i - z) \left\{ \mathbf{X}_i \mathbf{Y}_i^T \widehat{\Theta}_k(z) - \mathbf{e}_k \right\} \xi_i / n}{\left\{ \widehat{\Theta}_j(z) \right\}^T \widehat{\Sigma}_j(z)} \right|. \quad (10)$$

We denote the conditional $(1 - \alpha)$ -quantile of T_E^B given $\{(Z_i, \mathbf{X}_i, \mathbf{Y}_i)\}_{i \in [n]}$ as

$$c(1 - \alpha, E) = \inf \left(t \in \mathbb{R} \mid P \left[T_E^B \leq t \mid \{(Z_i, \mathbf{X}_i, \mathbf{Y}_i)\}_{i \in [n]} \right] \geq 1 - \alpha \right). \quad (11)$$

The quantity $c(1 - \alpha, E)$ can be calculated numerically using Monte-Carlo. In Section 5.2, we show that the quantile of T_E in (9) can be estimated accurately by the conditional $(1 - \alpha)$ -quantile of the bootstrap statistic.

We now propose an inference framework for testing the hypothesis problem of the form (7). Our proposed method is motivated by the step-down method in Romano and Wolf (2005) for multiple hypothesis tests. The details are summarized in Algorithm 1. Algorithm 1 involves evaluating all values of $z \in [0, 1]$. In practice, we implement the proposed method by discretizing values of $z \in [0, 1]$ into a large number of time points. We note that there

will be approximation error by taking the maximum over the discretized time points instead of the supremum of the continuous trajectory. The approximation error could be reduced to arbitrarily small if we increase the density of the discretization. The proposed method can be generalized to testing a wide variety of structures that satisfy the *monotone graph property*. Such a generalization will be presented in Appendix A.

Algorithm 1 Testing Maximum Degree of a Time-Varying Graph.

Input: type I error α ; pre-specified degree k ; de-biased estimator $\widehat{\Theta}^{\text{de}}(z)$ for $z \in [0, 1]$.

- (1) Compute the conditional quantile

$$c(1 - \alpha, E) = \inf [t \in \mathbb{R} \mid P(T_E^B \leq t \mid \{(Z_i, \mathbf{X}_i, \mathbf{Y}_i)\}_{i \in [n]}] \geq 1 - \alpha],$$

where T_E^B is the bootstrap statistic defined in (10).

- (2) Construct the rejected edge set

$$\mathcal{R}(z) = \left\{ e \in E(z) \mid \sqrt{nh} \cdot |\widehat{\Theta}_e^{\text{de}}(z)| \cdot \sum_{i \in [n]} K_h(Z_i - z)/n > c(1 - \alpha, E) \right\}.$$

- (3) Compute d_{rej} as the maximum degree of the dynamic graph based on the rejected edge set.

Output: Reject the null hypothesis if $d_{\text{rej}} > k$.

3. Simulation Studies

We perform numerical studies to evaluate the performance of our proposal using the inter-subject covariance relative to the typical time-varying Gaussian graphical model using within-subject covariance. To this end, we define the true positive rate as the proportion of correctly identified non-zeros in the true inverse covariance matrix, and the false positive rate as the proportion of zeros that are incorrectly identified to be non-zeros. To evaluate our testing procedure, we calculate the type I error rate and power as the proportion of falsely rejected H_0 and correctly rejected H_0 , respectively, over a large number of data sets.

To generate the data, we first construct the inverse covariance matrix $\Theta(z)$ for $z =$

$\{0, 0.2, 0.5\}$. At $z = 0$, we set $(d - 2)/4$ off-diagonal elements of $\Theta(0)$ to equal 0.3 randomly with equal probability. At $z = 0.2$, we set an additional $(d - 2)/4$ off-diagonal elements of $\Theta(0)$ to equal 0.3. At $z = 0.5$, we randomly select two columns of $\Theta(0.2)$ and add $k + 1$ edges to each of the two columns. This guarantees that the maximum degree of the graph is greater than k . To ensure that the inverse covariance matrix is smooth, for $z \in [0, 0.2]$, we construct $\Theta(z)$ by taking linear interpolations between the elements of $\Theta(0)$ and $\Theta(0.2)$. For $z \in [0.2, 0.5]$, we construct $\Theta(z)$ in a similar fashion based on $\Theta(0.2)$ and $\Theta(0.5)$. The construction is illustrated in Figure 1.

[Figure 1 about here.]

To ensure that the inverse covariance matrix is positive definite, we set $\Theta_{jj}(z) = |\Lambda_{\min}\{\Theta(z)\}| + 0.1$, where $\Lambda_{\min}\{\Theta(z)\}$ is the minimum eigenvalue of $\Theta(z)$. We then rescale the matrix such that the diagonal elements of $\Theta(z)$ equal one. The covariance $\Sigma(z)$ can be obtained by taking the inverse of $\Theta(z)$ for each value of z . Model (3) involves the subject specific covariance matrix $\mathbf{L}_X(z)$ and $\mathbf{L}_Y(z)$. For simplicity, we assume that these covariance matrices stay constant over time. We generate \mathbf{L}_X by setting the diagonal elements to be one and the off-diagonal elements to be 0.3. Then, we add random perturbations $\epsilon_m \epsilon_m^T$ to \mathbf{L}_X for $m = 1, \dots, 10$, where $\epsilon_m \sim N_d(\mathbf{0}, \mathbf{I}_d)$. The matrix \mathbf{L}_Y is generated similarly.

To generate the data according to (3), we first generate $Z_i \sim \text{Unif}(0, 1)$. Given Z_1, \dots, Z_n , we generate $\mathbf{S}(Z_i) \mid Z = Z_i \sim N_d\{\mathbf{0}, \Sigma(Z_i)\}$. We then simulate $\mathbf{E}_X(Z_i) \mid Z = Z_i \sim N_d(\mathbf{0}, \mathbf{L}_X)$ and $\mathbf{E}_Y(Z_i) \mid Z = Z_i \sim N_d(\mathbf{0}, \mathbf{L}_Y)$. Finally, for each value of Z , we generate

$$\mathbf{X}(Z_i) = \mathbf{S}(Z_i) + \mathbf{E}_X(Z_i) \quad \text{and} \quad \mathbf{Y}(Z_i) = \mathbf{S}(Z_i) + \mathbf{E}_Y(Z_i).$$

Note that both $\mathbf{X}(Z_i)$ and $\mathbf{Y}(Z_i)$ share the same generated $\mathbf{S}(Z_i)$ since both subjects are given the same natural continuous stimulus. In the following sections, we will assess the performance of our proposal relative to that of a typical approach for time-varying Gaussian

graphical models using the within-subject covariance matrix as input. We then evaluate the proposed inferential procedure in Section 2.3 by calculating its type I error and power.

3.1 Estimation

To mimic the data application we consider, we generate the data with $n = 945$, $d = 172$, and $k = 10$. Given the data $(Z_1, \mathbf{X}_1, \mathbf{Y}_1), \dots, (Z_n, \mathbf{X}_n, \mathbf{Y}_n)$, we estimate the covariance matrix at $Z = z$ using the inter-subject kernel smoothed covariance estimator as defined in (4). To obtain estimates of the inverse covariance matrices $\hat{\Theta}(Z_1), \dots, \hat{\Theta}(Z_n)$, we use the CLIME estimator as described in (5), implemented using the R package `clime`. There are two tuning parameters h and λ : we set $h = 1.2 \cdot n^{-1/5}$ and vary the tuning parameter λ to obtain the receiver operating characteristic (ROC) curve in Figure 2. The smoothing parameter h is selected such that there are always at least 30% of the time points that have non-zero kernel weights. We compare our proposal to time-varying Gaussian graphical models with the kernel smoothed within-subject covariance matrix. The true and false positive rates, averaged over 100 data sets, are in Figure 2.

[Figure 2 about here.]

From Figure 2, we see that our proposed method outperforms the typical approach for time-varying Gaussian graphical models by calculating the within-subject covariance matrix. This is because the typical approach is not estimating the parameter of interest, as discussed in Section 2.2. Our proposed method treats the subject specific effects as nuisance parameters and is able to estimate the stimulus-locked graph accurately.

3.2 Testing the Maximum Degree of a Time-Varying Graph

We evaluate Algorithm 1 by calculating its type I error and power. In all of our simulation studies, we consider $d = 50$ and $B = 500$ bootstrap samples, across a range of samples n . Similarly, we select the smoothing parameter to be $h = 1.2 \cdot n^{-1/5}$. The tuning parameter λ

is then selected using the cross-validation criterion defined in (6). The tuning parameter $\lambda = 0.9 \cdot [h^2 + \sqrt{\{\log(d/h)\}/(nh)}]$ is selected for one of the simulated data sets. For computational purposes, we use this value of tuning parameter across all replications.

We construct the test statistic T_E and the Gaussian multiplier bootstrap statistic T_E^B as defined in (9) and (10), respectively. Both the statistics T_E and T_E^B involve evaluating the supreme over $z \in [0, 1]$. In our simulation studies, we approximate the supreme by taking the maximum of the statistics over 50 evenly spaced grid $z \in [z_{\min}, z_{\max}]$, where $z_{\min} = \min \{Z_i\}_{i \in [n]}$ and $z_{\max} = \max \{Z_i\}_{i \in [n]}$. Our testing procedure tests the hypothesis

H_0 : for all $z \in [z_{\min}, z_{\max}]$, the maximum degree of the graph is no greater than k ,

H_1 : there exists a $z_0 \in [z_{\min}, z_{\max}]$ such that the maximum degree of the graph is greater than k .

For power analysis, we construct $\Theta(z)$ according to Figure 1 by randomly selecting two columns of $\Theta(0.2)$ and adding $k + 1$ edges to each of the two columns. This ensure that the maximum degree of the graph is greater than k . To evaluate the type I error under H_0 , instead of adding $k + 1$ edges to the two columns, we instead add sufficient edges such that the maximum degree of the graph is no greater than k . For the purpose of illustrating the type I error and power in the finite sample setting, we increase the signal-to-noise ratio of the data by reducing the effect of the nuisance parameters in the data generating mechanism described in Section 3. The type I error and power for $k = \{5, 6\}$, averaged over 500 data sets, are reported in Table 1. We see that the type I error is controlled and that the power increases to one as we increase the number of time points n . Note that the hypothesis problem (8) is a composite hypothesis. In general, a size α test is not achievable unless the true underlying parameter is at the boundary between the null and alternative hypotheses. In our numerical studies, the true underlying parameter is not generated such that it is at the boundary and therefore the size of the test is smaller than the specified level.

[Table 1 about here.]

4. Sherlock Holmes Data

We analyze a brain imaging data set studied in Chen et al. (2017). This data set consists of fMRI measurements of 17 subjects while watching audio-visual movie stimuli in an fMRI scanner. The subjects were asked to watch a 23-minute segment of BBC television series *Sherlock*, taken from the first episode of the series. The fMRI measurements were taken every 1.5 seconds, yielding $n = 945$ brain images for each subject. To understand the dynamics of the brain connectivity network under natural continuous stimuli, we partition the movie into 26 scenes (Chen et al., 2017). The data were pre-processed for slice time correction, motion correction, linear detrending, high-pass filtering, and coregistration to a template brain (Chen et al., 2017). Furthermore, for each subject, we attempt to mitigate issues caused by non-neuronal signal sources by regressing out the average white matter signal.

There are measurements for 271,633 voxels in this data set. For interpretation purposes, we reduce the dimension from 271,633 voxels to $d = 172$ regions of interest (ROIs) as described in Baldassano et al. (2015). We map the $n = 945$ brain images taken across the 23 minutes into the interval $[0, 1]$ chronologically. We then standardize each of the 172 ROIs to have mean zero and standard deviation one. Note that the statistical model is assumed on the standardized data.

We first estimate the stimulus-locked time-varying brain connectivity network. To this end, we construct the inter-subject kernel smoothed covariance matrix $\widehat{\Sigma}(z)$ as defined in (4). Since there are 17 subjects, we randomly split the 17 subjects into two groups, and use the averaged data to construct (4). Note that we could also construct a brain connectivity network for each pair of subjects separately. We then obtain estimates of the inverse covariance matrix using the CLIME estimator as in (5). We set the smoothing parameter $h = 1.2 \cdot n^{-1/5}$ so that at least 30% of the kernel weights are non-zero across all time points Z . For the sparsity tuning parameter, our theoretical results suggest picking

$\lambda = C \cdot \{h^2 + \sqrt{\log(d/h)/nh}\}$ to guarantee a consistent estimator. We select the constant C by considering a sequence of numbers using a 5-fold cross-validation procedure described in (6), and this yields $\lambda = 1.4 \cdot \{h^2 + \sqrt{\log(d/h)/(nh)}\}$. Heatmaps of the estimated stimulus-locked brain connectivity networks for three different scenes in Sherlock are in Figure 3.

[Figure 3 about here.]

From Figure 3, we see that there are quite a number of connections between brain regions that remain the same across different scenes in the movie. It is also evident that the graph structure changes across different scenes. We see that most brain regions are very sparsely connected, with the exception of a few ROIs. This raises the question of identifying whether there are hub ROIs that are connected to many other ROIs under audio-visual stimuli.

To answer this question, we perform a hypothesis test to test whether there are hub nodes that are connected to many other nodes in the graph across the 26 scenes. If there are such hub nodes, which ROIs do they correspond to? More formally, we test the hypothesis

H_0 : for all $z \in [0, 1]$, the maximum degree of the graph is no greater than 15,

H_1 : there exists a $z_0 \in [0, 1]$ such that the maximum degree of the graph is greater than 15.

The number 15 is chosen since we are interested in testing whether there is any brain region that is connected to more than 10% of the total number of brain regions. We apply Algorithm 1 with 26 values of z corresponding to the middle of the 26 scenes. Figure 4 shows the ROIs that have more than 12 rejected edges across the 26 scenes based on Algorithm 1. Since the maximum degree of the rejected nodes in some scenes are larger than 15, we reject the null hypothesis that the maximum degree of the graph is no greater than 15. In Figure 5, we plot the sagittal snapshot of the brain connectivity network, visualizing the rejected edges from Algorithm 1 and the identified hubs ROIs.

[Figure 4 about here.]

[Figure 5 about here.]

From Figure 4, we see that the rejected hub nodes (nodes that have more than 15 rejected edges) correspond to the frontal pole (7), temporal fusiform cortex (16, 100), lingual gyrus (17), and precuneus (102) regions of the brain. Many studies have suggested that the frontal pole plays significant roles in higher order cognitive operations such as decision making and moral reasoning (Okuda et al., 2003). The fusiform cortex is linked to face and body recognition (Iaria et al., 2008). In addition, the lingual gyrus is known for its involvement in processing of visual information about parts of human faces (McCarthy et al., 1999). Thus, it is not surprising that both of these ROIs have more than 15 rejected edges since the brain images are collected while the subjects are exposed to an audio-visual movie stimulus.

Compared to the lingual gyrus, temporal fusiform cortex, and the frontal pole, the precuneus is the least well-understood brain literature in the current literature. We see from Figure 4 that the precuneus is the most connected ROI across many scenes. This is supported by the observation in Hagmann et al. (2008) where the precuneus serves as a hub region that is connected to many other parts of the brain. In recent years, Lerner et al. (2011) and Ames et al. (2015) conducted experiments where subjects were asked to listen to a story under an fMRI scanner. Their results suggest that the precuneus represents high-level concepts in the story, integrating feature information arriving from many different ROIs of the brain. Interestingly, we find that the precuneus has the highest number of rejected edges during the first half of the movie and that the number of rejected edges decreases significantly during the second half of the movie. Our results correspond well to the findings of Lerner et al. (2011) and Ames et al. (2015) in which the precuneus is active when the subjects comprehend the story. However, it also raises an interesting scientific question for future study: is the precuneus active only when the subjects are trying to comprehend the story, that is, once the story is understood, the precuneus is less active.

5. Theoretical Results

We establish uniform rates of convergence for the proposed estimators, and show that the testing procedure in Algorithm 1 is a uniformly valid test. We study the asymptotic regime in which n , d , and s are allowed to increase. In the context of the Sherlock Holmes data set, n is the total number of brain images obtained under the continuous stimulus, d is the number of brain regions, and s is the maximum number of connections for each brain region in the true stimulus-locked network. The theoretical results assume that Z is a random variable with continuous density. Our results can be relaxed to the case when $\{Z_i\}_{i \in [n]}$ are fixed.

5.1 Theoretical Results on Parameter Estimation

Our proposed estimator involves a kernel function $K(\cdot)$: we require $K(\cdot)$ to be symmetric, bounded, unimodal, and compactly supported. More formally, for $l = 1, 2, 3, 4$,

$$\int K(u)du = 1, \quad \int uK(u)du = 0, \quad \int u^l K(u)du < \infty, \quad \int K^l(u)du < \infty. \quad (12)$$

In addition, we require the total variation of $K(\cdot)$ to be bounded, i.e., $\|K\|_{\text{TV}} < \infty$, where $\|K\|_{\text{TV}} = \int |\dot{K}|$. In other words, we require the kernel function to be a smooth function. A unimodal kernel function is extremely plausible in our setting: for instance, to estimate brain network in the “police press conference scene”, we expect the brain images within that scene to play a larger role than brain images that are far away from the scene. One practical limitation of the conditions on the kernel function is the symmetric kernel condition. When we are estimating a stimulus-locked brain network for a particular time point, the ideal case is to weight the previous images more heavily than the future brain images. The scientific reasoning is that there may be some time lag for information processing. In order to capture this effect, a carefully designed kernel function is needed and is out of the scope of this paper.

Next, we impose regularity conditions on the marginal density $f_Z(\cdot)$.

ASSUMPTION 1: There exists a constant \underline{f}_Z such that $\inf_{z \in [0,1]} f_Z(z) \geq \underline{f}_Z > 0$. Further-

more, f_Z is twice continuously differentiable and that there exists a constant $\bar{f}_Z < \infty$ such that $\max \{\|f_Z\|_\infty, \|\dot{f}_Z\|_\infty, \|\ddot{f}_Z\|_\infty\} \leq \bar{f}_Z$.

Next, we impose smoothness assumptions on the inter-subject covariance matrix $\Sigma(\cdot)$. Our theoretical results hold for any positive definite subject specific covariance matrices $\mathbf{L}_X(z)$ and $\mathbf{L}_Y(z)$, since these matrices are treated as nuisance parameters.

ASSUMPTION 2: There exists a constant M_σ such that

$$\sup_{z \in [0,1]} \max_{j,k \in [d]} \max \left\{ |\Sigma_{jk}(z)|, |\dot{\Sigma}_{jk}(z)|, |\ddot{\Sigma}_{jk}(z)| \right\} \leq M_\sigma.$$

In other words, we assume that the inter-subject covariance matrices are smooth and do not change too rapidly in neighboring time points. This assumption clearly holds in a dynamic brain network where we expect the brain network to change smoothly over time. Assumptions 1 and 2 on $f(z)$ and $\Sigma(z)$ are standard assumptions in the nonparametric statistics literature (see, for instance, Chapter 2 of Pagan and Ullah, 1999).

The following theorem establishes the uniform rates of convergence for $\widehat{\Sigma}(z)$.

THEOREM 1: Assume that $h = o(1)$ and that $\log^2(d/h)/(nh) = o(1)$. Under Assumptions 1–2, we have

$$\sup_{z \in [0,1]} \left\| \widehat{\Sigma}(z) - \Sigma(z) \right\|_{\max} = \mathcal{O}_P \left\{ h^2 + \sqrt{\frac{\log(d/h)}{nh}} \right\}.$$

Theorem 1 guarantees that our estimator always converges to the population parameter under the max norm, if the smoothing parameter h goes to zero asymptotically. For instance, this will satisfy if $h = C \cdot n^{-1/5}$ for some constant $C > 0$. The quantity $\sup_{z \in [0,1]} \|\widehat{\Sigma}(z) - \Sigma(z)\|_{\max}$ can be upper bounded by the summation of two terms: $\sup_{z \in [0,1]} \|\mathbb{E}[\widehat{\Sigma}(z)] - \Sigma(z)\|_{\max}$ and $\sup_{z \in [0,1]} \|\widehat{\Sigma}(z) - \mathbb{E}[\widehat{\Sigma}(z)]\|_{\max}$, which are known as the bias and variance terms, respectively, in the kernel smoothing literature (Pagan and Ullah, 1999). The terms h^2 and $\sqrt{\log(d/h)/(nh)}$ on the upper bound correspond to the bias and variance terms, respectively.

Next, we establish theoretical results for $\widehat{\Theta}(z)$. Recall that the stimulus-locked brain

connectivity network is encoded by the support of the inverse covariance matrix $\Theta(z)$: $\Theta_{jk}(z) = 0$ if and only if the j th and k th brain regions are conditionally independent given all of the other brain regions. We consider the class of inverse covariance matrices:

$$\mathcal{U}_{s,M} = \left\{ \Theta \in \mathbb{R}^{d \times d} \mid \Theta \succ 0, \|\Theta\|_2 \leq \rho, \max_{j \in [d]} \|\Theta_j\|_0 \leq s, \max_{j \in [d]} \|\Theta_j\|_1 \leq M \right\}. \quad (13)$$

Here, $\|\Theta\|_2$ is the largest singular value of Θ and $\|\Theta_j\|_0$ is the number of non-zeros in Θ_j .

Brain connectivity networks are usually densely connected due to the intrinsic-neural and non-neuronal signals. Our method allows the intrinsic brain network unrelated to the stimulus to be dense, and assume that the stimulus-locked brain network $\Theta(z)$ is sparse. The sparsity assumption on the stimulus-locked network is plausible since it characterizes brain activities that are specific to the stimulus. For instance, we may believe that only certain brain regions are active under cognitive process. The other conditions are satisfied since $\Theta(z)$ is the inverse of a positive definite covariance matrix. Given Theorem 1, the following corollary establishes the uniform rates of convergence for $\widehat{\Theta}(z)$ using the CLIME estimator as defined in (5). It follows directly from the proof of Theorem 6 in Cai et al. (2011).

COROLLARY 1: *Assume that $\Theta(z) \in \mathcal{U}_{s,M}$ for all $z \in [0, 1]$. Let $\lambda \geq C \cdot \{h^2 + \sqrt{\log(d/h)/(nh)}\}$ for $C > 0$. Under the same conditions in Theorem 1,*

$$\sup_{z \in [0,1]} \left\| \widehat{\Theta}(z) - \Theta(z) \right\|_{\max} = \mathcal{O}_P \left\{ h^2 + \sqrt{\frac{\log(d/h)}{nh}} \right\}; \quad (14)$$

$$\sup_{z \in [0,1]} \max_{j \in [d]} \left\| \widehat{\Theta}_j(z) - \Theta_j(z) \right\|_1 = \mathcal{O}_P \left[s \cdot \left\{ h^2 + \sqrt{\frac{\log(d/h)}{nh}} \right\} \right]; \quad (15)$$

$$\sup_{z \in [0,1]} \max_{j \in [d]} \left\| \left\{ \widehat{\Theta}_j(z) \right\}^T \widehat{\Sigma}(z) - \mathbf{e}_j \right\|_{\infty} = \mathcal{O}_P \left\{ h^2 + \sqrt{\frac{\log(d/h)}{nh}} \right\}. \quad (16)$$

Corollary 1 is helpful in terms of selecting the sparsity tuning parameter λ : it motivates a sparsity tuning parameter of the form $\lambda \geq C \cdot \{h^2 + \sqrt{\log(d/h)/(nh)}\}$ to guarantee statistically consistent estimated stimulus-locked brain networks. We consider a sequence of numbers and select the appropriate C using the cross-validation procedure in (6).

5.2 Theoretical Results on Topological Inference

In this section, we first show that the distribution of the test statistic T_E can be approximated by the conditional $(1 - \alpha)$ -quantile of the bootstrap statistic T_E^B . Next, we show that the proposed testing method in Algorithm 1 is valid in the sense that the type I error can be controlled at a pre-specified level α .

Recall from (11) the definition of $c(1 - \alpha, E)$. The following theorem shows that the Gaussian multiplier bootstrap is valid for approximating the quantile of the test statistic T_E in (9). Our results are based on the series of work on Gaussian approximation on multiplier bootstrap in high dimensions (see, e.g., Chernozhukov et al., 2013, 2014). We see from (9) that T_E involves taking the supremum over $z \in [0, 1]$ and a dynamic edge set $E(z)$. Due to the dynamic edge set $E(z)$, existing theoretical results for the Gaussian multiplier bootstrap methods cannot be directly applied. We construct a novel Gaussian approximation result for the supreme of empirical processes of T_E by carefully characterizing the capacity of the dynamic edge set $E(z)$.

THEOREM 2: *Assume that $\sqrt{nh^5} + s \cdot \sqrt{nh^9} = o(1)$. In addition, assume that $s\sqrt{\log^4(d/h)/(nh^2)} + \log^{22}(s) \cdot \log^8(d/h)/(nh) = o(1)$. Under the same conditions in Corollary 1, we have*

$$\lim_{n \rightarrow \infty} \sup_{\Theta(\cdot) \in \mathcal{U}_{s,M}} P_{\Theta(\cdot)} \{T_E \geq c(1 - \alpha, E)\} \leq \alpha.$$

Some of the scaling conditions are standard conditions in nonparametric estimation (Tsybakov, 2009). The most notable scaling conditions are $s\sqrt{\log^4(d/h)/(nh^2)} = o(1)$ and $\log^{22}(s) \cdot \log^8(d/h)/(nh) = o(1)$: these conditions arise from Gaussian approximation on multiplier bootstrap (Chernozhukov et al., 2013). These scaling conditions will hold asymptotically as long as the number of brain images n is much larger than the maximum degree in the graph s . This corresponds well with the data analysis where we expect only certain ROIs are active during information processing on the stimulus-locked network.

Recall the hypothesis testing problem in (7). We now show that the type I error of the

proposed inferential method for testing the maximum degree of a time-varying graph can be controlled at a pre-specified level α .

THEOREM 3: *Assume that the same conditions in Theorem 2 hold. Under the null hypothesis in (7), we have*

$$\lim_{n \rightarrow \infty} P_{\text{null}}(\text{Algorithm 1 rejects the null hypothesis}) \leq \alpha.$$

To study the power analysis of the proposed method, we define the signal strength of a precision matrix Θ as

$$\text{Sig}_{\text{deg}}(\Theta) := \max_{E' \subseteq E(\Theta), \text{Deg}(E') > k} \min_{e \in E'} |\Theta_e|, \quad (17)$$

where $\text{Deg}(E)$ is the maximum degree of graph $G = (V, E)$. Under the alternative hypothesis in (7), there exists a $z_0 \in [0, 1]$ such that the maximum degree of the graph is greater than k . We define the parameter space under the alternative:

$$\mathcal{G}_1(\theta) = \left[\Theta(\cdot) \in \mathcal{U}_{s,M} \mid \text{Sig}_{\text{deg}}\{\Theta(z_0)\} \geq \theta \text{ for some } z_0 \in [0, 1] \right]. \quad (18)$$

The following theorem presents the power analysis of Algorithm 1.

THEOREM 4: *Assume that the same conditions in Theorem 2 hold and select the smoothing parameter such that $h = o(n^{-1/5})$. Under the alternative hypothesis in (7) and the assumption that $\theta \geq C \sqrt{\log(d/h)/nh}$, where C is a fixed large constant, we have*

$$\lim_{n \rightarrow \infty} \inf_{\Theta \in \mathcal{G}_1(\theta)} \mathbb{P}_{\Theta}(\text{Algorithm 1 rejects the null hypothesis}) = 1, \quad (19)$$

for any fixed $\alpha \in (0, 1)$.

The signal strength condition defined in (17) is weaker than the typical minimal signal strength condition required on testing a single edge on a conditional independent graph, $\min_{e \in E(\Theta)} |\Theta_e|$. The condition in (17) requires only that there exists a subgraph whose maximum degree is larger than k and the minimal signal strength on that subgraph is above certain level. In our real data analysis, this requires only the edges for brain regions

that are highly connected to many other brain regions to be strong, which is plausible since these regions should have high brain activity.

6. Discussion

We consider estimating stimulus-locked brain connectivity networks from data obtained under natural continuous stimuli. Due to lack of highly controlled experiments that remove all spontaneous and individual variations, the measured brain signal consists of not only stimulus-induced signal, but also intrinsic neural signal and non-neuronal signal that are subject specific. Typical approach for estimating a time-varying Gaussian graphical model will fail to estimate the stimulus-locked brain connectivity network accurately due to the presence of subject specific effects. By exploiting the experimental design aspect of the problem, we propose a simple approach to estimating a stimulus-locked brain connectivity network. In particular, rather than calculating a within-subject smoothed covariance matrix as in the typical approach for modeling time-varying Gaussian graphical models, we propose to construct the inter-subject smoothed covariance matrix instead, treating the subject specific effects as nuisance parameters.

To answer the scientific question on whether there are any hub brain regions during the given stimulus, we propose an inferential method for testing the maximum degree of a stimulus-locked time-varying graph. In our analysis, we found that several interesting brain regions such as the fusiform cortex, lingual gyrus, and precuneus are highly connected. From the neuroscience literature, these brain regions are mainly responsible for high order cognitive operations, face and body recognition, and serve as control region that integrates information from other brain regions. We have also extended the proposed inferential framework to testing various topological graph structures in Appendix A.

The practical limitation of our proposed method is the Gaussian assumption on the data. While we focus on the time-varying Gaussian graphical model in this paper, our framework

can be extended to other types of time-varying graphical models such as the time-varying discrete graphical model or the time-varying nonparanormal graphical model (Kolar et al., 2010; Lu et al., 2018). Another limitation is the independence assumption on the data across time points. All of our theoretical results can be generalized to the case when the data across time points are correlated by imposing an α -mixing condition on Z , and we leave such a generalization for future work.

Acknowledgement

We thank the editor, associate editor, and two reviewers for their comments that improve earlier version of this paper. We thank Janice Chen for conversations on preprocessing the fMRI data. Kean Ming Tan is partially supported by NSF DMS-1949730 and NIH RF1-MH122833. Junwei Lu is partially supported by NSF DMS-1916211.

References

- Ames, D. L., Honey, C. J., Chow, M. A., Todorov, A., and Hasson, U. (2015). Contextual alignment of cognitive and neural dynamics. *Journal of Cognitive Neuroscience* **27**, 655–664.
- Baldassano, C., Beck, D. M., and Fei-Fei, L. (2015). Parcellating connectivity in spatial maps. *PeerJ* **3**, e784.
- Cai, T., Liu, W., and Luo, X. (2011). A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association* **106**, 594–607.
- Chen, J., Leong, Y. C., Honey, C. J., Yong, C. H., Norman, K. A., and Hasson, U. (2017). Shared memories reveal shared structure in neural activity across individuals. *Nature Neuroscience* **20**, 115–125.
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2013). Gaussian approximations and

- multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics* **41**, 2786–2819.
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2014). Gaussian approximation of suprema of empirical processes. *The Annals of Statistics* **42**, 1564–1597.
- Drton, M. and Maathuis, M. H. (2017). Structure learning in graphical modeling. *Annual Review of Statistics and Its Application* **4**, 365–393.
- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., and et al. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biol* **6**, e159.
- Hasson, U., Harel, M., Levy, I., and Malach, R. (2003). Large-scale mirror-symmetry organization of human occipito-temporal object areas. *Neuron* **37**, 1027–1041.
- Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., and Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science* **303**, 1634–1640.
- Iaria, G., Fox, C. J., Waite, C. T., Aharon, I., and Barton, J. J. (2008). The contribution of the fusiform gyrus and superior temporal sulcus in processing facial attractiveness: neuropsychological and neuroimaging evidence. *Neuroscience* **155**, 409–422.
- Kolar, M., Song, L., Ahmed, A., and Xing, E. P. (2010). Estimating time-varying networks. *The Annals of Applied Statistics* **4**, 94–123.
- Lerner, Y., Honey, C. J., Silbert, L. J., and Hasson, U. (2011). Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *The Journal of Neuroscience* **31**, 2906–2915.
- Lu, J., Kolar, M., and Liu, H. (2018). Post-regularization inference for time-varying nonparanormal graphical models. *Journal of Machine Learning Research* **18**, 1–78.
- McCarthy, G., Puce, A., Belger, A., and Allison, T. (1999). Electrophysiological studies of human face perception. ii: Response properties of face-specific potentials generated in occipitotemporal cortex. *Cerebral Cortex* **9**, 431–444.

- Musch, K., Himberger, K., Tan, K. M., Valiante, T. A., and Honey, C. J. (2020). Transformation of speech sequences in human sensorimotor circuits. *Proceedings of the National Academy of Sciences* **117**, 3203–3213.
- Neykov, M., Lu, J., and Liu, H. (2019). Combinatorial inference for graphical models. *The Annals of Statistics* **47**, 795–827.
- Neykov, M., Ning, Y., Liu, J. S., and Liu, H. (2018). A unified theory of confidence regions and testing for high-dimensional estimating equations. *Statistical Science* **33**, 427–443.
- Okuda, J., Fujii, T., Ohtake, H., Tsukiura, T., Tanji, K., Suzuki, K., and et al. (2003). Thinking of the future and past: The roles of the frontal pole and the medial temporal lobes. *Neuroimage* **19**, 1369–1380.
- Pagan, A. and Ullah, A. (1999). *Nonparametric Econometrics*. Cambridge university press.
- Regev, M., Simony, E., Lee, K., Tan, K. M., Chen, J., and Hasson, U. (2018). Propagation of information along the cortical hierarchy as a function of attention while reading and listening to stories. *Cerebral Cortex* **29**, 4017–4034.
- Romano, J. P. and Wolf, M. (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association* **100**, 94–108.
- Simony, E., Honey, C. J., Chen, J., Lositsky, O., Yeshurun, Y., Wiesel, A., and et al. (2016). Dynamic reconfiguration of the default mode network during narrative comprehension. *Nature Communications* **7**, 12141.
- Tan, K. M., London, P., Mohan, K., Lee, S.-I., Fazel, M., and Witten, D. (2014). Learning graphical models with hubs. *The Journal of Machine Learning Research* **15**, 3297–3331.
- Tsybakov, A. B. (2009). *Introduction to Nonparametric Estimation*. Springer.
- Zhou, S., Lafferty, J., and Wasserman, L. (2010). Time varying undirected graphs. *Machine Learning* **80**, 295–319.

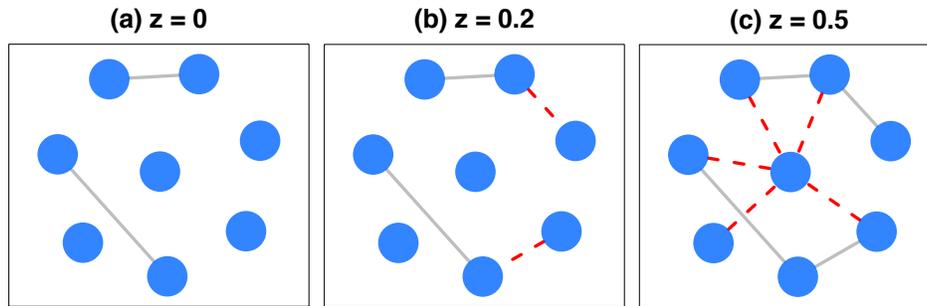


Figure 1. (a): A graph corresponding to $\Theta(0)$ with maximum degree no greater than four. (b): A graph corresponding to $\Theta(0.2)$ with maximum degree less than or equal to four. The red dash edges are additional edges that are added to $\Theta(0)$. (c): A graph corresponding to $\Theta(0.5)$ with maximum degree larger than four. The red dash edges are additional edges that are added to $\Theta(0.2)$ such that the maximum degree of the graph is larger than four.

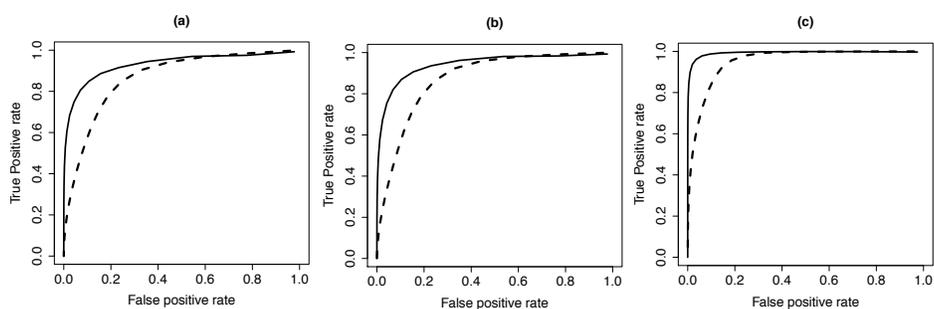


Figure 2. The true and false positive rates for the numerical study with $n = 952$, $d = 172$, and $k = 10$. Panels (a), (b), and (c) correspond to $Z = \{0.25, 0.50, 0.75\}$, respectively. The two curves represent our proposal (black solid line) and within-subject time-varying Gaussian graphical model (black dash), respectively.

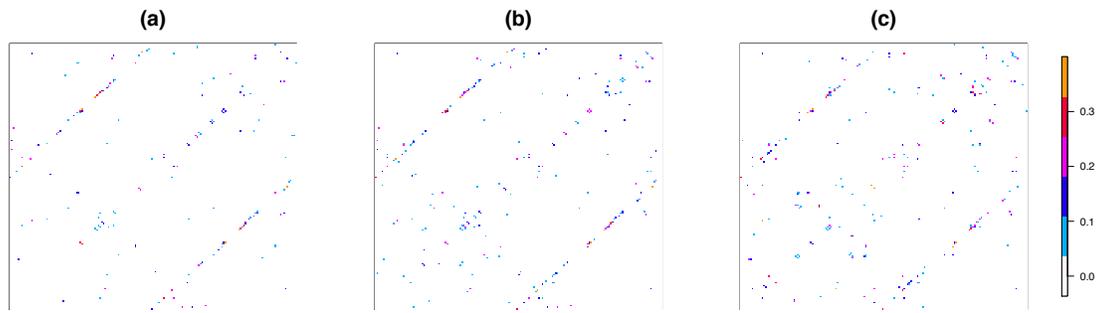


Figure 3. Heatmaps of the estimated stimulus-locked brain connectivity network for three different scenes in Sherlock. (a) Watson psychiatrist scene; (b) Park run in scene; and (c) Watson joins in scene. Colored elements in the heatmaps correspond to edges in the estimated brain network.

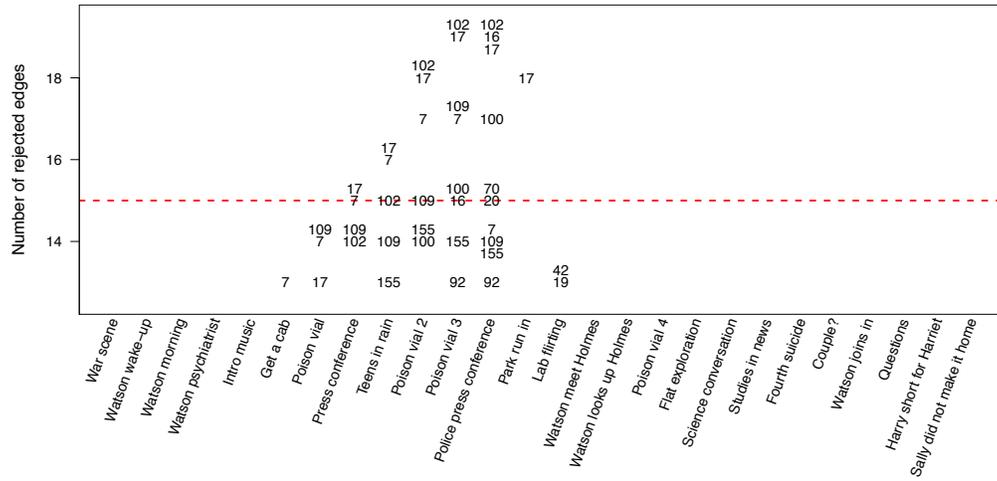


Figure 4. The x -axis displays the 26 scenes in the movie and the y -axis displays the number of rejected edges from Algorithm 1. The numbers correspond to the regions of interest (ROIs) in the brain. The ROIs correspond to frontal pole (7, 155), temporal fusiform cortex (16, 100), lingual gyrus (17), cingulate gyrus (19), cingulate gyrus (20), temporal pole (42), paracingulate gyrus (70), precuneus cortex (102), and postcentral gyrus (109).

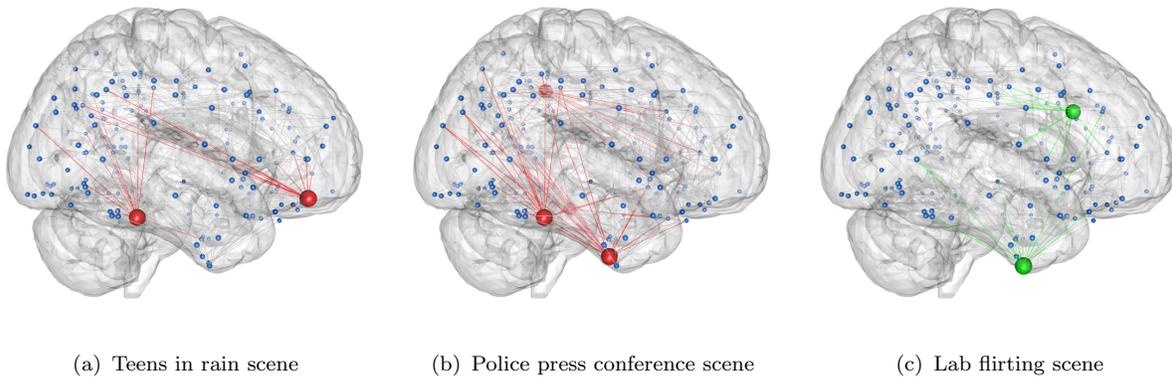


Figure 5. Sagittal snapshots of the rejected edges based on Algorithm 1. Panels (a)-(c) contain the snapshots for the “teens in rain”, “police press conference”, and “lab flirting” scenes, respectively. The red nodes and red edges are regions of interest that have more than 15 rejected edges. The grey edges are rejected edges from nodes that have no greater than 15 rejected edges. For (c), the green nodes and edges are regions of interest that have more than 12 rejected edges.

Table 1

The type I error and power for testing the maximum degree of the graph at the 0.05 significance level are calculated as the proportion of falsely rejected and correctly rejected null hypotheses, respectively, over 500 data sets.

Simulation results with $d = 50$ and $k = \{5, 6\}$, over a range of n are shown.

		$n = 400$	$n = 600$	$n = 800$	$n = 1000$	$n = 1500$
$k=5$	Type I error	0.014	0.024	0.030	0.034	0.028
	Power	0.068	0.182	0.690	0.976	1
$k=6$	Type I error	0.032	0.040	0.034	0.028	0.018
	Power	0.050	0.142	0.446	0.898	1