DETECTING STIMULUS DRIVEN CHANGES IN FUNCTIONAL BRAIN CONNECTIVITY

Pingmei Xu, Hao Xu, Peter J. Ramadge

Department of Electrical Engineering,
Princeton University, Princeton NJ
{pingmeix,haoxu,ramadge}@princeton.edu

ABSTRACT

We consider the problem of detecting stimulus driven changes in brain functional connectivity. Estimating functional connectivity from fMRI data sampled over a small time period is difficult - there is simply not enough data to permit reliable estimates. We investigate the use of a sparse Gaussian graphical model regularized by a graph learned from data sampled over a longer time period. We establish a framework to identify the changes in brain connectivity driven by short-term stimuli. Results of experiments on both synthetic and real fMRI data illustrate the attributes of our methods as well as the difficulty of the problem.

***Index Terms*—** Brain Connectivity, Graphical Lasso, fMRI, Time-Dependent Network

1. INTRODUCTION

An undirected graphical model is a mathematical structure that can be used to represent pair-wise relationships among a set of variables by a joint distribution. It has applications in many domains including information extraction, computer vision and gene networks. One of the most widely used undirected graphical models is the Gaussian graphical model. This assumes the variables to have a multivariate Gaussian distribution. This model has been recently used to analyse functional interactions between different brain regions in human brain - this is known as a functional brain connectivity network. In the graph, each node represents a voxel or set of voxels forming a region of interest (ROI). The graph edges encode dependencies between voxels or ROIs. The graphical model thus provides a functional connectivity model within the human brain. Previous studies based on anatomical brain databases suggest that the true brain connectivity network is sparse [1–3], and many recent efforts have shown the promise of using sparse Gaussian graphical models to analyse brain connectivity [4–7].

These studies are applied to steady state time series data, and assume that the brain connectivity network does not change during the experiment. However, when the experimental conditions (such as the stimuli) vary with time, it is of interest to learn the dynamical manner of interactions between different brain regions [8]. Some previous work has addressed this problem. The topological changes in the connectivity network is analyzed by statistical parametric network in [9]. But the timing and duration of different states, which are usually difficult to estimate, are specified as prior knowledge. The approach in [10] estimates connectivity networks for predefined non-overlapping windows with different temporal scales (days, hours and minutes), but it loses the ability to adaptively choose the location of the windows. A dynamic connectivity regression method in [8] automatically detects the change points of different states by estimating a sparse Gaussian graphical model, but it requires an exhaustive search so that both computational complexity and sampling complexity can be issues for fine partitions of the time series.

In this paper, we take a different approach. We propose a method to identify changes in the brain connectivity network driven by short-term stimuli. Unlike [10], we use a sliding window to extract overlapping subsets of the time series data. To mitigate the issue of insufficient data, a temporal average network is learnt over a larger time scale and is then used to design a regularization penalty for the estimation of the connectivity network over a smaller time scale. The assumption that graph structure is the same across subjects as in [8] is not required for our method so that different sets of voxels or ROIs can be selected for different subjects. We conducted experiments on both synthetic and real data to demonstrate the the attributes and deficiencies of the proposed method.

2. MAP-PENALIZED GRAPHICAL LASSO

The graph of an undirected graphical model is denoted by $G = (V, E)$ where V is the set of nodes and E the set of edges. Each edge in E models the relationship between the two variables in V that it connects. The problem of learning a graph for Gaussian graphical model is equivalent to estimating the precision matrix (inverse covariance matrix), since the non-zero off-diagonal elements of the precision matrix represent edges in the graph [11]. Specifically, given n independent samples $\{y_1, y_2, \ldots, y_n\}$ drawn from a p-variate Gaussian distribution such that $y_i \sim N(\mu, \Sigma)$, the task is to estimate its inverse covariance matrix $\Theta = \Sigma^{-1}$. A missing edge between two nodes in the graph indicates conditional independence between the corresponding variables. It gives rise to a zero element in the precision matrix. Therefore, learning a sparse Gaussian graphical model is equivalent to estimating a sparse precision matrix. A sparse estimate of Θ can be obtained by minimizing the penalized log-likelihood over positive definite matrices [11]:

$$\arg\min_{\Theta} - \log(\det(\Theta)) + \text{trace}(S\Theta) + \lambda\|\Theta\|_1 \quad (1)$$

where $S = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})(x_i - \bar{x})^T$ is the empirical covariance matrix, $\|\Theta\|_1$ is the ℓ_1-norm – the sum of the absolute values of the elements in Θ, and $\lambda > 0$ is the regularization parameter. This problem is also referred to as graphical lasso (GLasso) problem [12], and much recent research has focused on efficiently estimating Θ [11–18]. However, when the sample size is much less than the sparsity of the precision matrix, GLasso fails to recover the correct graph (see more details in §4.1). In this paper, we employ an extra map-penalized regularization to learn the precision matrix of a subset of the data. We formulate the problem as:

$$\arg\min_{\Theta} - \log(\det(\Theta)) + \text{trace}(S\Theta) + \lambda\|\Theta\|_1 + \alpha\|M \odot \Theta\|_1 \quad (2)$$

where $\alpha > 0$, the map M is a binary matrix determined by the whole data and $\odot$ denotes element-wise multiplication. The quality of M depends on the number of available data samples, but at least we are using all the data to estimate M. The idea is to use M to regularize Θ for the subset of the data. Problem (2) is equivalent to a GLasso with a generalized regularization term $\|L \odot \Theta\|_1$ [19], where $L = \lambda\mathbf{1} + \alpha M$ and $\mathbf{1}$ is a matrix of ones. We adapted the QUIC algorithm proposed in [17] to solve this optimization problem.

Given a matrix A, define an indicator matrix B where $b_{ij} = 1$ for nonzero elements a_{ij}, otherwise zero. Given A and a baseline matrix A_0, we get corresponding indicator matrices B and B_0 and denote $\neg B$ and $\neg B_0$ as their element-wise negation not matrices. Then we define:
(1) $sp(A) = \|B\|_0$, the number of non-zero elements in A.
(2) $nn(A) = \|B \odot B_0\|_0$. The number of non-zero elements in A_0 that are also non-zero in A.
(3) $nz(A) = \|\neg B \odot B_0\|_0$. The number of non-zero elements in A_0 that are zero in A.
(4) $nz(A) = \|B \odot \neg B_0\|_0$. The number of zero elements in A_0 that are non-zero in A.

3. DETECTING STIMULUS DRIVEN CHANGES

Let $D \in \mathbb{R}^{p \times n}$ be the fMRI time-series data of a subject, where p is the number of voxels and n is the number of TR (time to repetition = time needed to loop through all of the slices). We proposed the following approach to detect changes in the brain connectivity network driven by short-term stimuli.

Step One: obtain a *temporal average* graph by learning the precision matrix Y from D via graphical lasso (GLasso):

$$Y = \arg\min_{\Theta} - \log(\det(\Theta)) + \text{trace}(S\Theta) + \lambda\|\Theta\|_1$$

Here the λ is chosen so that $sp(Y)$ is equal to a fixed sparsity. Define a binary regularization map $M \in \mathbb{R}^{p \times p}$ with non-zero element $M_{ij} = 1$ if Y_{ij} is zero.
Step Two: obtain *moving average* graphs from subsets of the data set D. First select k overlapping subsets $\{D^{(i)}\}_{i=1\ldots k}$ by a sliding window of size w and step size d, where w is chosen to be in the same time scale of the duration of the interested stimulus. For each subset $D^{(i)}$, learn a precision matrix $Y^{(i)}$ by solving a map-penalized GLasso (MP-GLasso) (problem (2)) with M determined in step one and λ is chosen to achieve a fixed sparsity $sp(Y^{(i)})$. Compute $nn(Y^{(i)})$, $nz(Y^{(i)})$ and $zn(Y^{(i)})$ based on the *temporal average* precision matrix.
Step Three: obtain *random average* graphs from random subsets of the data set D for a statistical test in step four. Randomly select subsets from data $\{D^{(r_i)}\}_{i=1\ldots k_r}$ with size equal to the window size in step two. Repeat the computation in step two on random subsets to get precision matrices $\{Y^{(r_i)}\}$ and corresponding $\{nn(Y^{(r_i)})\}$, $\{nz(Y^{(r_i)})\}$ and $\{zn(Y^{(r_i)})\}$.
Step Four: detect the changes in the brain connectivity network. First fit the Gaussian distribution by $\{nn(Y^{(r_i)})\}$, $\{nz(Y^{(r_i)})\}$ and $\{zn(Y^{(r_i)})\}$. Second compute the p-value of $\{nn(Y^{(i)})\}$, $\{nz(Y^{(i)})\}$ and $\{zn(Y^{(i)})\}$ based on the corresponding distributions. If the p-value is less than a predefined significance level, then a change in the brain connectivity is detected.

4. EXPERIMENT RESULTS

4.1. Sampling Complexity

We conducted experiments on synthetic data to show that map-penalized graphical lasso (MP-GLasso) exhibits improved performance over graphical lasso (GLasso) for data with small sample size. We generated graphs with random sparsity structure as in [17, 20]. Specifically, to generate an inverse covariance matrix with random non-zero patterns, we first generated a sparse square matrix $U \in \mathbb{R}^{p \times p}$ by randomly setting its non-zero elements as ± 1. Then we set the true precision matrix Θ^* to be $U^T U$ and add an identity matrix to ensure the positive definiteness of Θ^*. We controlled the number of non-zeros in U so that the resulting Θ^* has approximately $10p$ non-zero elements.

We first tested GLasso. In Fig. 1(a) for a fixed regularization parameter λ, the structure learned from a small sample size such as 200 contains a large number of incorrect edges, while the structure estimated by data with a larger sample size recovers the true graph much better by reducing this number. In Fig. 1(b), we adaptively choose λ so that the resulting sparsity of the estimated precision matrix is around $\|\Theta^*\|_0$.

Though we use the ground truth to limit the total number of edges, data of sample size 200 still estimates more than 400 incorrect edges. The same trend also applies to the number of missing edges since $nz(Y) = \|\Theta^*\|_0 - nn(Y) = \|\Theta^*\|_0 - (sp(Y) - zn(Y)) \approx zn(Y)$, where Y is the learnt structure. We then apply the MP-GLasso to the same data set as shown in Fig. 1(c). Specifically, we first learn a *temporal average* structure by estimating the precision matrix Y_{5000} using a random data set with 1000 samples. Then we use Y_{5000} to regularize the estimation of precision matrices for data with fewer samples. Fig. 1(c) indicates that MP-GLasso increases the number of correct edges as well as decreases the number of wrong edges given the same λ. Intuitively, MP-GLasso imposes the estimated "mean connectivity" to improve learning from data with a small sample size.

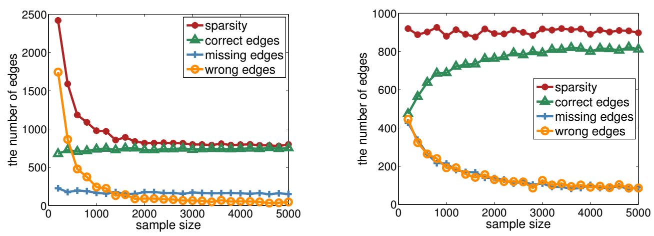

(a) GLasso: Fixed λ (b) GLasso: Fixed Sparsity

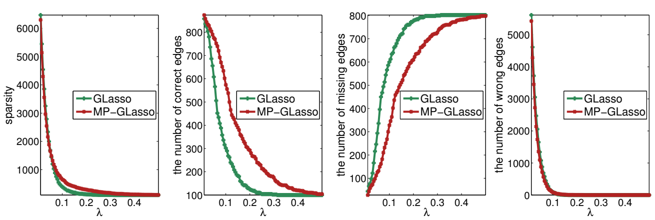

(c) GLasso and MP-GLasso

Fig. 1. GLasso on synthetic data set with sample size varying from 200 to 5000: (a) shows metrics of the resulting precision matrices by setting a fixed regularization parameter $\lambda = 0.04$; (b) shows metrics of the resulting precision matrices by adaptively choosing λ to obtain the correct number of non-zero elements $\|\Theta^*\|_0$. Comparison between graphical lasso (GLasso) and GLasso with map penalty (MP-GLaaso): (c) from left to right – sparsity, the number of correct edges, the number of missing edges and the number of wrong edges. $p = 100$, $\|\Theta^*\|_0 = 902$, $\alpha = \lambda$. The sample size in (c) is 200, and the penalty map is obtained from data consisting of 5000 random samples with $\lambda = 0.035$.

4.2. Synthetic Dataset

To demonstrate the effectiveness of our method for detecting connectivity changes, we generated a set of graphs similar to [6] in the following way: we set a reference precision matrix Θ^0 as $0.25 I_{p\times p}$, where $p = 100$. To add an edge (i, j) to the graph, we add σ to Θ^0_{ii} and Θ^0_{jj}, and subtract σ from Θ^0_{ij} and Θ^0_{ji} to keep Θ positive definite. To delete an edge, reverse the above steps with $\sigma = \Theta^0_{ij}$. We randomly assign 200 edges for Θ^0. Then generate a set of precision

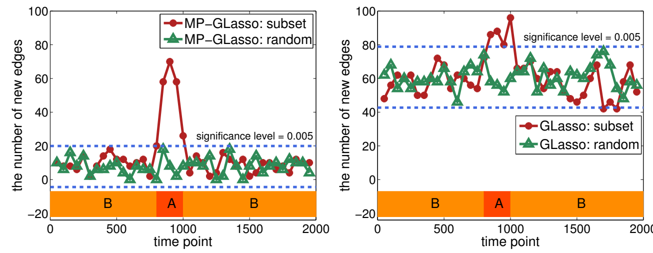

(a) MP-GLasso (b) GLasso

Fig. 2. Detecting stimulus-driven changes: comparison between (a) map-penalized graphical lasso (MP-GLasso) and (b) graphical lasso (GLasso) on synthetic data. $p = 100$, $\|\Theta^{(k)}\|_0 = 300$ for $k = 1, \dots, 10$, block size = 200. The block indicated in orange-red color is in state **A**, the rest of blocks in dark orange color are in state **B**. Window size = 100, step size = 20, $\alpha = \lambda$, significance level = 0.005. Iteratively choose the value of λ such that the sparsity for *temporal average* graph = 300 and that for *moving average* graph = 200.

matrices $\{\Theta^{(1)}, \dots, \Theta^{(K)}\}$ ($K = 10$), where each $\Theta^{(k)}$ is obtained based on Θ^0 by adding $n_a^{(k)}$ new edges and deleting $n_d^{(k)}$ existing edges. We set $n_a^{(k)} = n_d^{(k)}$ to 90 for $k = 5$ and $n_a^{(k)} = n_d^{(k)}$ as 5 for the rest. Then we generate K blocks of data. For each block, we draw 200 samples of data from Gaussian distribution with the corresponding precision matrix $\Theta^{(k)}$. Since the graph for $k = 5$ deviates a lot more from the reference precision matrix, we refer the 5th block as state **A** and the rest as state **B**. We expect to detect the occurrence of state **A** by checking the *moving average* precision matrix.

The dashed lines in Figure 2 indicate the number of new edges with p-value equal to the significance level 0.005. Graphs with p-value outside this region are identified as significantly different from the *temporal average* graph. Figure 2(a) shows that our method correctly detected when the *moving average* graph deviates significantly from the *temporal average* graph. Compared with the results of GLasso, by imposing a map penalty, the mean and variance of the number of new edges obtained by MP-GLasso are both reduced for *moving graphs* in state **B** and for *random graphs*. More importantly, since the *moving average* graphs and the random average graphs vary too much, the changing pattern of different states suggested by GLasso in Figure 2(b) is also not as clear as by MP-GLasso in Figure 2(a).

4.3. Hitchcock Movie Dataset

We tested our method on the fMRI dataset from [21]. 11 female subjects watched a 20-mins movie *Bang! You're Dead* directed by Alfred Hitchcock, while the whole brain volumes (58/40/46 slices on Anteroposterior/ Dorsoventral/ Mediolateral axes, respectively) were scanned every 2 seconds (TR = 2 sec). The total number of TR is 618 for each subject. The movie contains several tense scenes, each lasting from tens of seconds to several minutes. We preprocessed the data set as in [22] to extract 200*2 voxels from (left and right) ventral temporal (VT) for each scan image. Then we use a matrix $D \in \mathbb{R}^{200\times 618}$ to detect the connectivity changes in (left or right) the brain connectivity for each subject.

Figure 3 shows the results for left and right hemispheres of a single subject. Since the sets of voxels for two hemispheres are selected separately, the similar trends of two red curves suggests the existence of consistent across hemisphere temporal brain connectivity changes. Window samples with p-value less than the significance level 0.005 are also confirmed as deviation from the average state statistically. So for this subject, the *moving average* connectivity deviates from the *temporal average* connectivity both at the beginning and the last third of the movie.

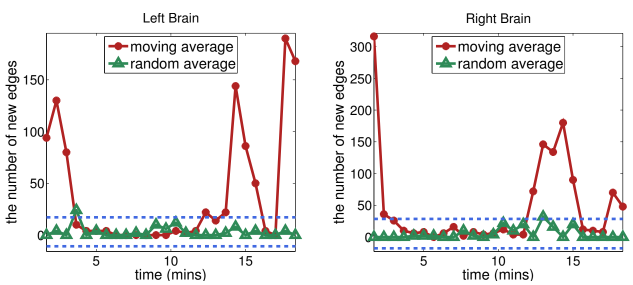

Fig. 3. Detecting stimulus driven changes in the left and right brain connectivity network: the number of new edges of *moving average* (red) or *random average* (green) brain connectivity network compared to the *temporal average* brain connectivity network. The (blue) dashed lines indicate the number of new edges with p-value equal to the significance level 0.005. The length of the movie = 20mins, TR=2s. $p = 200$, window size = 100, step size = 20, $\alpha = 0.4\lambda$. Iteratively choose the value of λ such that the sparsity for *temporal average* graph = 1200 and the sparsity for *moving average* graph = 1200.

The estimated patterns of connectivity changes differ across subjects. The counting of subjects with detected changes over 11 female subjects is shown in Figure 4. We see changes along the entire time line, with some consistency in the temporal pattern of connectivity changes in the two hemispheres as visualized by dashed lines in Figure 4 (a). And the counts in the right hemisphere varies more frequently than the left hemisphere, which may be due to the fact that all subjects in this experiments are right-handed. To anaylze how the connectivity changes relate to the stimulus – in this case the movie, we defined two types of events and annotated the movie by an event matrix $E \in \mathbb{R}^{11\times 60}$, where $e_{ij} = 1$ if event i happened in the 20 seconds non-overlapping movie segments j, and zero otherwise. The type *I* contains 3 events which mark transitions between two different scenes: 1) indoor and outdoor, 2) different indoor backgrounds and 3) people and objects. The type *II* contains the rest 8 semantic events listed in Figure 4 (b). For the type *II* events, we replaced each row vector in E with its first order difference to capture the transition of events. To simulate the averaging

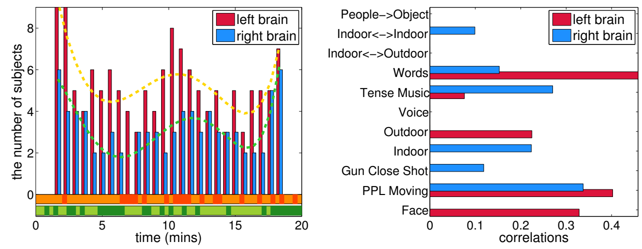

Fig. 4. (a) The counting of subjects with detected stimulus driven brain connectivity changes. Yellow and green dashed curves are polynomial fitting (7 degree) of the histograms for visualization. Two examples of event vectors: object close shots marked by orange red blocks; people moving shots marked by dark green.(b) The correlation coefficients with positive values between the counting histogram in (a) with the semantic vectors learned from the movie.

effect of window sampling, we applied a Gaussin filter with $\sigma = 2.5$ (such that $\pm 2\sigma$ corresponding to time interval 200 seconds = 100 TRs) to row vectors in E and calculate the correlation between the histograms in Figure 4 (a) and the resulting 11 smoothed vectors. Consistent results with relatively large correlation in the two hemispheres, such as the 'people moving', 'words' and 'tense music' events, may suggest their relevance to the cause of connectivity changes. So the subjects may pay more attention to the movie when actors started or stopped moving. Since VT mainly processes visual information, the reason that the 'tense music' event has large correlation can be possibly interpreted as: during the movie, tense music was always played with the most exciting scenes so that the music raises the visual attention of our subjects. For the rest of the semantic events, since no consistent results for two hemispheres are shown, it needs further investigation of their relationship to connectivity changes.

5. CONCLUSION

In this paper, we considered the problem of detecting changes in functional brain connectivity network driven by short-term stimuli. We presented a new scheme called map-penalized graphical lasso to address the issue of insufficient sample size and to improve the performance of detection. Experiment results on synthetic data indicated improved performance over graphical lasso. Experiment results on real fMRI data suggested similar trends of changes in the left and right hemispheres, and the detected changes in the right hemisphere varied more frequently than the left hemisphere for right-handed subjects. Our results suggest the difficulty of analyzing the causality between the detected changes and the stimulus, so a proper design for new experiments will consider stimulus with simpler semantic meaning. We would like to test more methods on the Hitchcock dataset in the future work. A future direction we are pursuing is to explicitly analyze how the connectivity network evolves with time as the stimulus changes.

References

[1] C. C. Hilgetag, R. Kotter, K. E. Stephan, and O. Sporns, "Computational methods for the analysis of brain connectivity," *Computational Neuroanatomy*, 2002.

[2] R. Kötter and K. E. Stephan, "Network participation indices: characterizing component roles for information processing in neural networks," *Neural Networks*, 2003.

[3] O. Sporns, D. R. Chialvo, M. Kaiser, and C. C. Hilgetag, "Organization, development and function of complex brain networks," *Trends in cognitive sciences*, 2004.

[4] S. Huang, J. Li, L. Sun, J. Ye, A. Fleisher, T. Wu, K. Chen, and E. Reiman, "Learning brain connectivity of Alzheimer's disease by sparse inverse covariance estimation," *NeuroImage*, 2010.

[5] G. Varoquaux, A. Gramfort, J. B. Poline, and B. Thirion, "Brain covariance selection: better individual functional connectivity models using population prior," *Advances in Neural Information Processing Systems (NIPS)*, 2010.

[6] S. Yang, Z. Pan, X. Shen, P. Wonka, and J. Ye, "Fused Multiple Graphical Lasso," *arXiv preprint arXiv:1209.2139*, 2012.

[7] J. Honorio, D. Samaras, I. Rish, and G. Cecchi, "Variable Selection for Gaussian Graphical Models," *AISTATS*, vol. XX, 2012.

[8] I. Cribben, R. Haraldsdottir, L. Y. Atlas, T. D. Wager, and M. A. Lindquist, "Dynamic connectivity regression: Determining state-related changes in brain connectivity," *NeuroImage*, 2012.

[9] C. E. Ginestet and A. Simmons, "Statistical parametric network analysis of functional connectivity dynamics during a working memory task.," *NeuroImage*, vol. 55, no. 2, pp. 688–704, Mar. 2011.

[10] D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton, "Dynamic reconfiguration of human brain networks during learning," *Proceedings of the National Academy of Sciences*, 2011.

[11] O. Banerjee, L. E. Ghaoui, and A. D'Aspremont, "Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data," *Journal of Machine Learning Research*, 2008.

[12] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, 2008.

[13] N. Meinshausen and P. Bühlmann, "High-dimensional graphs and variable selection with the lasso," *The Annals of Statistics*, 2006.

[14] M. Yuan and Y. Lin, "Model selection and estimation in the Gaussian graphical model," *Biometrika*, 2007.

[15] J. Duchi, S. Gould, and D. Koller, "Projected subgradient methods for learning sparse gaussians," *Conference on Uncertainty in Artificial Intelligence*, 2008.

[16] K. Scheinberg, S. Ma, and D. Goldfarb, "Sparse Inverse Covariance Selection via Alternating Linearization Methods," *Advances in Neural Information Processing Systems (NIPS)*, , no. 1, pp. 1–9, 2010.

[17] C. J. Hsieh and M. A. Sustik, "Sparse inverse covariance matrix estimation using quadratic approximation," *Advances in Neural Information Processing Systems (NIPS)*, 2011.

[18] R. Mazumder and T. Hastie, "Exact covariance thresholding into connected components for large-scale graphical lasso," *The Journal of Machine Learning Research*, pp. 1–19, 2012.

[19] J. Bien and R. J. Tibshirani, "Sparse estimation of a covariance matrix.," *Biometrika*, vol. 98, no. 4, pp. 807–820, Dec. 2011.

[20] L. Li and K. C. Toh, "An inexact interior point method for l 1-regularized sparse covariance selection," *Mathematical Programming Computation*, 2010.

[21] U. Hasson, O. Landesman, B. Knappmeyer, I. Vallines, N. Rubin, and D. J. Heeger, "Neurocinematics: The neuroscience of film," *Journal for Movies and Mind*, 2008.

[22] H. Xu, A. Lorbert, P. J. Ramadge, J. S. Guntupalli, and J. V. Haxby, "Regularized Hyperalignment of Multi-Set fMRI Data," in *SSP*, 2012, number 1.