
Gender and age group classification using functional Magnetic Resonance Imaging and Gaussian Markov Random Fields

Roberto Vega
Department of Computing Science
University of Alberta
rvega@ualberta.ca

Kriti Khare
Department of Computing Science
University of Alberta
kriti@ualberta.ca

Sayem Mohammad Siam
Department of Computing Science
University of Alberta
siam@ualberta.ca

Abstract

Functional Magnetic Resonance Imaging (fMRI) is a non-invasive technique used for visualization of the brain activity. One particular application of fMRI is the study of the resting-state functional connectivity, whose focus is to study the spontaneous brain activity in resting condition. Information extracted from the resting-state functional connectivity might be used for further understanding of the brain, and to help in the diagnosis of mental diseases such as Schizophrenia, or Attention Deficit and Hyperactivity Disorder. In this work, we evaluate the application of Gaussian Markov Random Fields to build functional connectivity models that can be used as classifiers using resting-state fMRI data. These models present at least three advantages: 1) There is a natural mapping from brain data to the graphical model. 2) There are efficient algorithms to train the model. 3) It is possible to directly evaluate model fit of the data. We test our methodology in two binary problems: gender classification (male vs. female) and age group classification (people younger than 10 vs people of at least 17 years old). For the gender classification problem we achieved an accuracy of 59.46%, while for the age classification problem the accuracy was 79.60%. In both problems we presented a statistically significant improvement over a class-majority classifier.

1 Introduction

Functional Magnetic Resonance Imaging (fMRI) is a non-invasive technique used for visualization of the brain activity. The rationale behind fMRI is that active neurons consume oxygen at a higher rate than inactive ones, phenomenon known as blood oxygen level dependent (BOLD) signal: an increase in neuronal activity raises the oxygen consumption due to the higher metabolic requirements; this, in turn, changes the concentrations of oxyhemoglobin and deoxyhemoglobin in the blood. Because these two substances have different magnetic properties, they can be identified using a Magnetic Resonance scan. The changes in the concentrations of oxyhemoglobin and deoxyhemoglobin along time are reflected as changes in intensity of the images [1].

Resting state functional connectivity is an fMRI technique that studies the spontaneous brain activity during resting conditions. This approach enables the analysis of the brain regions that present synchrony in their neural activity even if they are not physically connected. By studying the resting state networks, it is possible to broaden the understanding of how multiple regions interact with

each other. This is important for at least two reasons: 1) The spontaneous activity consume more than 80% of the brain energy. 2) Resting-state studies do not rely on active participation of patients, which allow them to be used with uncooperative populations or in scenarios when it is difficult to maintain the attention of the patient [2].

There are two main approaches for analyzing resting-fMRI data from a group of people: model-driven approaches, and data-driven approaches [3]. In the first one, a group of regions of interest is selected a priori, and the analysis is limited to those regions. Data-driven approaches, like Independent Component Analysis (ICA), are exploratory and hypothesis-free, which allow them to generate results at a network level, and not only at inter-regional level [2].

Among the model-driven approaches, the use of probabilistic graphical models (PGM) is a natural way of finding the connectivity among the different regions of interest. PGM encode a complex distribution over high-dimensional spaces in the form of a graph [4]. These models are easily visualized and provide an inference tool that captures statistical independences among the different random variables [5]. Every region of interest in the brain might be represented as a node in the PGM, while the functional connectivity might be represented by the edges among the nodes. By using PGM, it is possible to create connectivity models for different populations and then use these model for classification purposes, or for finding statistically significant differences among them [6, 7].

In this paper we focus on the application of PGM, specifically Gaussian Markov Random Fields, for classification in two binary problems: gender (male vs. female), and age-group (younger than 9 vs. older than 17) using information extracted from resting-state fMRI scans of the publicly available dataset ADHD-200 [8].

The rest of the paper is organized as follows: Section 2 gives a summary of the work that has been developed in the analysis of resting-state fMRI and in the use of classifiers based on fMRI data. Section 3 describes the theory behind our approach, the dataset used for the experiments, and the methodology used for the experimentation. Section 4 shows the results achieved with the experiments. In Section 5 we analyze the results given by the experiments. Finally, Section 6 states our conclusions and future work.

2 Previous work

The objective of brain functional connectivity is to find statistical dependencies between neural activity in different parts of the brain [9, 10]. To find them, model driven approaches are commonly used. These methods require the definition of regions of interest to be analyzed. Then, the functional connectivity can be established using correlation, or partial correlation methods [3, 6].

Although correlation based analysis of fMRI data is a common technique, it lacks the ability of distinguishing between direct and indirect dependencies among the nodes. Since it works only with two nodes at the time, it cannot identify when two or more nodes are acting together [11]. Despite this limitation, it has successfully been used in the past for distinguishing between healthy controls from people with schizophrenia (93% accuracy) [12], or classifying males and females (71% accuracy) [13]; however, it requires the use of an independent classifier, like support vector machines, which diminishes the interpretability of the models.

These problems can be solved by Gaussian Markov Random Fields, which use partial correlation to find the connectivity among the nodes. In other words, they find the conditional independences between two random variables given the rest of the variables [4], solving the problem of several nodes acting together. At the same time, unlike correlation-based methods, we can easily estimate model fit by computing the likelihood of the data [6], avoiding the need of an independent classifier.

On the other side, they make strong assumptions about the distribution of the data: 1) That the samples are independent and identically distributed, 2) That the data follows a Gaussian distribution. Smith et al. [14] made an empirical study using simulated data where they found that, despite these strong assumptions, this method reliably identifies direct relationships between brain regions. Another drawback is that finding an optimal sparse model is an NP-hard problem [6]. Many research efforts have been focused on estimating the partial correlation and the corresponding graphical model, and very efficient solutions have been proposed [15, 16, 17]. These random fields were successfully

used by Narayan et al. [7] for finding statically significant differences in the brain connectivity in healthy controls and people with synesthesia; however, they did not perform any classification task.

A different problem in the construction of functional brain connectivity models is the definition of the regions of interest, since the use of inaccurate regions extremely damages the network estimation [14]. Power et al. [18] identified 264 regions of interest that represent functional "units of brain organization". In their work, they argue that other common approaches to define nodes for functional connectivity analysis, like parcels from anatomically based brain atlases, or the definition of nodes as voxels do not correspond to functional units of the brain, and their use might lead to inaccurate connectivity models.

Satterthwaite et al. [13] used the regions defined by Power et al. [18] and made an empirical study where they found that gender differences in cognitive profiles are related to multivariate patterns of resting-state functional connectivity MRI (rsfc-MRI). Specifically, they found that males display more between-module connectivity, while females demonstrated more within-module connectivity. They constructed a graph using a correlation based method and used SVM to classify between both genders. They achieved an accuracy of (71%; $P < 0.001$) using data from a sample of 674 persons obtained in a single site.

One question that might arise is how well results obtained in data obtained from a single site generalize to multisite data. Nielsen et al. [19], for example, obtained 60% accuracy for identifying healthy controls and people with autism when using multisite data, while the reported accuracy in data from a single site was 80%. In our work, we will perform a similar study to the one made by Satterthwaite et al. [13], but using data from 8 different sites, extending the problem to age group classification, and implementing Gaussian Markov Random Fields as a classifier.

3 Methods and experiments

3.1 Gaussian Markov Random Fields

A Gaussian Markov Random Field (GMRF) is similar to Markov Random Fields, with the difference that now the random variables take continuous values [4]. A multivariate Gaussian distribution over the random variables X_1, X_2, \dots, X_n is parametrized by an n dimensional mean vector μ and a symmetric $n \times n$ covariance matrix Σ . The density function is then defined as:

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} \exp \left[-\frac{1}{2} (\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu) \right]$$

Any Gaussian distribution can be represented as a pairwise Markov Network with quadratic node and edge potentials. These kind of networks are known as Gaussian Markov Random Fields [4]. One advantage of these models is that two Gaussian variables, X_i and X_j are conditionally independent given the rest of the variables if and only if their corresponding entries $\Sigma_{i,j}^{-1} = \Sigma_{j,i}^{-1} = 0$ [20]. Therefore, learning the structure of a GMRF reduces to the problem of finding zero entries on the inverse of the covariance matrix Σ^{-1} , which is also known as the precision matrix $\Omega = \Sigma^{-1}$. Figure 1 exemplifies this behavior.




Figure 1: The zeros in the precision matrix $\Omega = \Sigma^{-1}$ represent conditional independences on the graphical model.

A common approach when finding probabilistic graphical models is to chose the simplest one that adequately explains the data [5]. One metric that can be used for evaluating the model fit is the

likelihood function, which measures the probability of the data given the model. Since a fully connected model will output the highest likelihood function, a regularization term for penalizing complex models is required.

For a dataset D with n independent and identically distributed samples $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, where each variable in \mathbf{x}_i follows a Gaussian distribution with zero mean, its log-likelihood function can be expressed as:

$$L(D) = \frac{n}{2} [\log(|\Omega|) - \text{tr}(A\Omega)] + \text{const} \quad (1)$$

$$A = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^T$$

Since the constant term is independent of the mean and covariance, it can be ignored in the optimization function. The objective is then to find the model that maximizes the penalized likelihood function with a restriction in the number of parameters, which is known as the *spase inverse covariance selection problem*:

$$\max_{\Omega} \log(|\Omega|) - \text{tr}(X^T X \Omega) - \lambda \|\Omega\|_1 \quad (2)$$

where X is the a matrix of $n \times m$, n is the number of samples in the dataset, m is the number of random variables, λ is the regularization term, and $\|\Omega\|_1$ is the l_1 norm of the precision matrix.

An efficient algorithm for solving this optimization problem is graphical lasso, which uses a coordinate descent procedure for the lasso. One of the advantages of graphical lasso is its computational speed: it is 30 to 4,000 times faster than competing methods [16] We used the freely available graphical lasso package for our experiments [21].

3.2 Dataset

ADHD-200 [8] is a dataset collected with the objective of studying Attention Deficit and Hyperactivity Disorder (ADHD). It contains 973 resting-state fMRI scans collected across 8 independent imaging sites. The age range of the entire sample is 7-21 years. The ADHD-200 dataset was pre-processed for Brown et al. [22], and we used these preprocessed data for our experiments.

Out of the 973 scans, we used only the ones corresponding to healthy subjects (301 males, 259 females) for the gender classification problem. For the age group problem, we used the scans of the 45 males and 59 females younger than 10 years old, and 24 males and 34 females of at least 17 years old.

Every fMRI scan is represented as a 4-dimensional matrix of 79x95x68x91. The first three dimensions represent the 3D representation of the brain, while the 4th dimension is the time steps. In other words, we represent every brain with 510,340 voxels and there is a time series of 91 points for every voxel.

3.3 Experiment design

Our methodology involves the 7 steps depicted in Figure 2. In general terms, it consists in creating a graphical model for every class that we are interested in classifying. Then, for every subject to be classified, we will compute the likelihood of the data given a particular model. The subject will be classified with the class whose model has the highest likelihood.

1. *Extract regions of interest.* Instead of working with the more than 500,000 voxels, which might become an intractable problem from the computational perspective, we used the 264 regions of interest defined by Power et al. [18]. Each region of interest is defined as a sphere of 5 mm radius, so we computed the average of the time signals of the voxels inside this sphere. The result of this process is 264 time series per subject. Since every time series consists in 91 time points, every subject is represented by a matrix of 91 x 264.




Figure 2: The 7 steps of our methodology for classifying the resting-state fMRI scans.

2. *Normalize data.* We normalized the data of every subject and region of interest independently. Every time series was normalized to mean $\mu = 0$ and standard deviation $\sigma = 1$.
3. *Separate data by class.* Separate the dataset in groups of interest. For this particular case the classes are male or female for the gender classification task, and younger than 10 years old or at least 17 years old for the age group classification task.
4. *Create training and test set.* Randomly select 70% of the data available in each class for training; the remaining 30% will be used for testing purposes and will not be used for taking any decision in the selection of the model.
5. *Concatenate time series.* Steps 5 and 6 are implemented using 5-fold cross validation and the training set from the last step. Using the 4 subsets of each round, concatenate the time series corresponding to the same region of interest. This will result in a matrix X of $N * 91 \times 264$ per class, where N is the number of subjects present in the 4 subsets.
6. *Construct the model.* Using Eq. 2, construct a model for each class. Then, test the performance of the model on the remaining subset of data (of the cross validation process). Performance was measured as the percentage of samples correctly classified on the subset not used for training. To perform classification we simply compute the likelihood of the data given each model using Eq. 1, and assign the subject the class whose model gave the highest likelihood. Since the constant term is independent of the model, it can be ignored for the classification step. We experimented with the following values for the regularization term: $\lambda = [0.001, 0.003, 0.01, 0.03, 0.1, 0.3]$, and selected the value of λ with the highest cross-validation accuracy.
7. *Classify new data.* After selecting the best model for each class (steps 5 and 6), test the performance of the model on the test set created on step 4. Report the accuracy on this test set.

4 Results

Since the step 4 in our methodology involves the creation of random training and test sets, different runs of the experiment will output a different accuracy. In order to obtain a better estimate of the performance of the methodology, we repeated the experiment 30 times. Table 1 shows the statistics for the gender classification problem, while Table 2 shows the results for the age group problem. Each table presents the value of the average accuracy over the 30 runs of the training set (the one created on step 5 of the methodology using 4 subsets), the cross-validation set (the remaining subset of step 5, used to select the best model), and the test step (the one created on step 4 of the methodology). Additionally, we report the accuracy of a class-majority model that would classify all the subjects with the class with the highest number of samples.

In both experiments, the difference between the mean accuracy in the test set obtained with our methodology and the class-majority model is statistically significant ($p < .001$) using a one-sample t-test. The 95% confidence interval for the mean accuracy in the test set of the gender classification problem is: 57.59% - 61.33%, giving an improvement of at least 4% when compared with the class-majority model. For the case of the age group classification problem, the 95% confidence interval for the mean on the test set is: 77.41% - 81.77%, giving an improvement of at least 13.21%.

We also made a third experiment in which we solved 4 binary problems. We separated the data in 4 different groups: females younger than 10, males younger than 10, females of at least 17 and males of at least 17 years old. We then compared the classification accuracy among the possible combinations of these 4 groups. The results are shown on Table 3.

Table 1: Classification results for the male vs. female problem: mean, standard deviation (std), minimum value, and maximum value of the accuracy over the 30 runs, $\lambda = 0.1$

| | Mean | Std | Min | Max |
|-----------------------|--------|-------|--------|--------|
| Train | 89.61% | 4.38% | 77.34% | 96.26% |
| CV | 60.41% | 2.32% | 56.48% | 66.73% |
| Test | 59.46% | 4.93% | 49.11% | 70.54% |
| Class-majority | 53.57% | - | - | - |

Table 2: Classification results for age group problem: mean, standard deviation (std), minimum value, and maximum value of the accuracy over the 30 runs, $\lambda = 0.01$

| | Mean | Std | Min | Max |
|-----------------------|--------|-------|--------|--------|
| Train | 97.37% | 0.96% | 95.16% | 99.42% |
| CV | 81.28% | 1.86% | 78.29% | 85.90% |
| Test | 79.60% | 5.75% | 69.7% | 96.97% |
| Class-majority | 64.20% | - | - | - |

Table 3: Average classification accuracy in the age group problem over 30 runs. In parenthesis we report the class-majority model accuracy.

| | Male < 9 | Male > 17 | Female < 9 | Female > 17 |
|-----------------------|-----------------|-----------------|-----------------|-----------------|
| Male < 9 | - | 77.86% (65.22%) | 52.54% (56.73%) | 73.81% (56.96%) |
| Male > 17 | 77.86% (65.22%) | - | 81.76% (71.08%) | 58.89% (58.62%) |
| Female < 9 | 52.54% (56.73%) | 81.76% (71.08%) | - | 83.33% (63.44%) |
| Female > 17 | 73.81% (56.96%) | 58.89% (58.62%) | 83.33% (63.44%) | - |

The precision matrix, Ω , of every model represents the conditional independences of the graphical model. In every model, each of the 264 regions of interest is represented by node, while the functional connectivity is represented by the edges between the nodes. For the gender classification problem, the models with highest prediction accuracy generated 22,968 connections for the female brain and 22,738 connections for the male brain. Analogously, the best models for the age classification problem found 29,862 connections for the brain of people younger than 10 years old, and 25,942 connections for the brain of people of at least 17 years old. Figure 3 shows the model found in the gender classification problem, while Figure 4 shows the analogous model for the age group classification problem.

5 Discussion

In both problems, gender classification and age group classification, the Gaussian Markov Random Field classifier performed statistically significant better than a simple model that would label all the subjects with the mode of the classes. Despite this difference, it is not difficult to see that the improvement on the gender classification problem is practically imperceptible: slightly less than 6%. One important aspect to note is that even when the test accuracy is low, the training accuracy is high: 89.61%, so the model is clearly overfitting. We tried to increase the weight of the regularization term; however, the training accuracy dropped without an increment in the test accuracy.

One natural way to try to improve the performance is to subdivide the total sample in age groups, since it might be possible that a model for each gender cannot appropriately capture the differences present among different age groups; however, the results shown in Table 3 showed that it was not the case. Using the current approach we were unable of correctly classifying males and females within the same age group. The analysis of the male and female models created might bring some intuition




Figure 3: Left: Thresholded connectivity model for females. Right: Thresholded connectivity model for males.




Figure 4: Left: Thresholded connectivity model for people younger than 10 years old. Right: Thresholded connectivity model of at least 17 years old.

about why that happened: both models had basically the same number of connections (22,968 for females and 22,738 for males). At the same time, by analyzing the Figure 3 we can appreciate that many of the most important connections are present in both models, including the strongest one!

On the other side, Figure 4 shows important differences between the brain model for people younger than 10 years old, and the one for people of at least 17 years old. The number of connections above the selected threshold is significantly higher in the model of older people. At the same time, the difference between the total number of connections in both models was of almost 4,000 (compared to only 230 in the gender problem). Here is important to note that even when the brain model for younger people had more number of connections, the model for older people had more number of "strong" connections. This difference in the models is reflected in the classification accuracy, where we could correctly classify 79.60% of the subjects, a difference of 15.20% when compared to a class-majority model. This improvement was kept when we subdivided the age groups by gender, as can be appreciated on Table 3. It would be interesting to analyze, from the perspective of neuroscience, if these differences in the number of connections might have a physiological basis or are only relevant features found by the graphical model and cannot be directly interpreted.

Another important thing to remark is the computational efficiency of graphical LASSO [21], algorithm that was used for finding the different brain models. This algorithm could find a model for two classes in less than 400 seconds when run in a processor Intel Core i3 at 2.13 Ghz x 4. This efficiency allowed the execution of the 30 different runs that we used to obtain statistics about the performance and compute the statistical significance analysis.

6 Conclusions and future work

In this work we implemented Gaussian Markov Random Fields to solve two bi-class problems: gender (male vs. female) and age group (more than 10 years old vs. at least 17 years old). We used graphical LASSO to train the random field. This approach had at least three advantages: 1) It was computationally efficient. 2) It enforced sparsity, which diminished the complexity of the model and facilitated visualization. 3) It was possible to assess model fit in a natural and simple way.

The proposed methodology had a good performance in the age group problem; however, it could not find a good classifier for the male vs. female problem. The current approach makes some strong assumptions that could be relaxed in an effort to improve the performance of the classifiers. In first

place, we are assuming that the time signals of the regions of interest, across all the subjects, follow a gaussian distribution. A more profound analysis of these signals is needed in order to evaluate if a mixture of gaussians could be a better design approach. On the other side, we are assuming that every time point in the signal is independent of the other ones. A different approach could be to extract a different set of features from every signal using signal processing technique and evaluate if the results improve.

Finally, for the age group problem we used only two different groups. As a future work we will also identify how many different age groups are present in this particular dataset, and then implement the methodology in different dataset to evaluate the generalization capabilities.

References

- [1] David J Heeger and David Ress. What does fmri tell us about neuronal activity?. *Nature Reviews Neuroscience*, 3(2):142 – 151, 2002.
- [2] Cristina Rosazza and Ludovico Minati. Resting-state brain networks: Literature review and clinical applications. *Neurological Sciences*, 32(5):773–785, 2011.
- [3] David M Cole, Stephen M Smith, and Christian F Beckmann. Advances and pitfalls in the analysis and interpretation of resting-state fmri data. *Frontiers in Systems Neuroscience*, 4(8), 2010.
- [4] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009.
- [5] Irina Rish and Genady Ya Grabarnik. *Sparse modeling : theory, algorithms, and applications*. Chapman & Hall/CRC machine learning & pattern recognition series. Boca Raton, FL : CRC Press : Taylor & Francis Group, 2015., 2015.
- [6] Varoquaux G, A Gramfort, J B Poline, and B Thirion. Markov models for fmri correlation structure: Is brain functional connectivity small world, or decomposable into networks? *Journal of Physiology-Paris*, 106(56):212 – 221, 2012. New trends in neurogeometrical approaches to the brain and mind problem.
- [7] Manjari Narayan, Genevera I Allen, and Steffie Tomson. Two sample inference for populations of graphical models with applications to functional connectivity. 2015.
- [8] The adhd-200 sample http://fcon_1000.projects.nitrc.org/indi/adhd200/. [Online; accessed 17-March-2015].
- [9] Stephen M Smith, Diego Vidaurre, Christian F Beckmann, Matthew F Glasser, Mark Jenkinson, Karla L Miller, Thomas E Nichols, Emma C Robinson, Gholamreza Salimi-Khorshidi, Mark W Woolrich, et al. Functional connectomics from resting-state fmri. *Trends in cognitive sciences*, 17(12):666–682, 2013.
- [10] Jeanette A Mumford and Joseph D Ramsey. Bayesian networks for fmri: a primer. *Neuroimage*, 86:573–582, 2014.
- [11] R Cameron Craddock, Saad Jbabdi, Chao-Gan Yan, Joshua T Vogelstein, F Xavier Castellanos, Adriana Di Martino, Clare Kelly, Keith Heberlein, Stan Colcombe, and Michael P Milham. Imaging human connectomes at the macroscale. *Nature methods*, 10(6):524–539, 2013.
- [12] Schizophrenia as a network disease: Disruption of emergent brain function in patients with auditory hallucinations. *PLoS ONE*, 8(1):1 – 15, 2013.
- [13] Theodore D Satterthwaite, Daniel H Wolf, David R Roalf, Kosha Ruparel, Guray Erus, Simon Vandekar, Efsthios D Gennatas, Mark A Elliott, Alex Smith, Hakon Hakonarson, Ragini Verma, Christos Davatzikos, Raquel E Gur, and Ruben C Gur. Linked sex differences in cognition and functional connectivity in youth. *Cerebral Cortex*, 2014.
- [14] Stephen M Smith, Karla L Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for fmri. *NeuroImage*, 54(2):875 – 891, 2011.
- [15] Onureena Banerjee, Laurent El Ghaoui, and Alexandre d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *The Journal of Machine Learning Research*, 9:485–516, 2008.

- [16] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- [17] Cho-Jui Hsieh, Inderjit S Dhillon, Pradeep K Ravikumar, and Mátyás A Sustik. Sparse inverse covariance matrix estimation using quadratic approximation. In *Advances in Neural Information Processing Systems*, pages 2330–2338, 2011.
- [18] Jonathan D Power, Alexander L Cohen, Steve M Nelson, Gagan S Wig, Kelly A Barnes, Jessica A Church, Alecia C Vogel, Timothy O Laumann, Fran M Miezin, Bradley L Schlaggar, and Steven E Petersen. Functional network organization of the human brain. *Neuron*, 72(4):665 – 678, 2011.
- [19] Jared A Nielsen, Brandon A Zielinski, P Thomas Fletcher, Andrew L Alexander, Nicholas Lange, Erin D Bigler, Janet E Lainhart, and Jeffrey S Anderson. Multisite functional connectivity mri classification of autism: Abide results. *Frontiers in human neuroscience*, 7, 2013.
- [20] Steffen L. Lauritzen. *Graphical Models*. Oxford University Press, 1996.
- [21] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Graphical lasso in r and matlab <http://statweb.stanford.edu/~tibbs/glasso/>. [Online; accessed 17-March-2015].
- [22] Matthew R G Brown, Gagan S Sidhu, Russell Greiner, Nasimeh Asgarian, Meysam Bastani, Peter H Silverstone, Andrew J Greenshaw, and Serdar M Dursun. Adhd-200 global competition: Diagnosing adhd using personal characteristic data can outperform resting state fmri measurements. *Frontiers in Systems Neuroscience*, 6(69), 2012.

Table 4: Responsibilities of the team members

| Student | Responsibilities |
|----------------------------|---|
| Roberto Vega | Experiment design and implementation |
| Kriti Khare | Review and use of Graphical LASSO |
| Sayem Mohammad Siam | Literature review and implement data whitening |
| Shared | Prepare and edit manuscript Prepare and edit presentations |