Published in final edited form as:

Med Image Comput Comput Assist Interv. 2014 ; 17(0 3): 193–200.

Discriminative Sparse Connectivity Patterns for Classification of fMRI Data

Harini Eavani¹, Theodore D. Satterthwaite², Raquel E. Gur², Ruben C. Gur², and Christos Davatzikos¹

¹Center for Biomedical Image Computing and Analytics, University of Pennsylvania

²Brain Behavior Laboratory, University of Pennsylvania

Abstract

Functional connectivity using resting-state fMRI has emerged as an important research tool for understanding normal brain function as well as changes occurring during brain development and in various brain disorders. Most prior work has examined changes in pair-wise functional connectivity values using a multi-variate classification approach, such as Support Vector Machines (SVM). While it is powerful, SVMs produce a dense set of high-dimensional weight vectors as output, which are difficult to interpret, and require additional post-processing to relate to known functional networks. In this paper, we propose a joint framework that combines network identification and classification, resulting in a set of networks, or Sparse Connectivity Patterns (SCPs) which are functionally interpretable as well as highly discriminative of the two groups. Applied to a study of normal development classifying children vs. adults, the proposed method provided accuracy of 76% (AUC= 0.85), comparable to SVM (79%, AUC=0.87), but with dramatically fewer number of features (50 features vs. 34716 for the SVM). More importantly, this leads to a tremendous improvement in neuro-scientific interpretability, which is specially advantageous in such a study where the group differences are wide-spread throughout the brain. Highest-ranked discriminative SCPs reflect increases in long-range connectivity in adults between the frontal areas and posterior cingulate regions. In contrast, connectivity between the bilateral parahippocampal gyri was decreased in adults compared to children.

1 Introduction

Functional connectivity, defined as the amount of correlation between observed BOLD time-series, has been widely applied to study the large-scale functional architecture of the human brain in both health, disease, and development. Many studies that examine changes in pair-wise connectivity use multi-variate methods, such as SVMs, with the vectorized correlation matrices as features [1]. Applied to classification, the l_2 -regularized SVM results in a list of connections or edges that are discriminative of the two groups. However, high-dimensional patterns are very difficult to parse and interpret, requiring additional processing and analysis based on prior knowledge of known functional network labels [2]. To address this issue, feature selection methods [3] or l_1 -regularized SVM can be used, which produces

a sparse discriminative pattern. However, it is known to ignore features that are highly correlated (and therefore redundant for classification), which could be neurobiologically relevant.

Alternately, a complementary set of analysis that can be performed begins with network identification using independent component analysis (ICA). ICA commonly identifies motor, visual, sub-cortical and default mode networks. These networks can then be analyzed for changes in spatial extent or average connectivity between the two groups. However, in such a method, network identification is performed independently of classification. This can be problematic since not all networks are necessarily different in the two groups - requiring an exhaustive search in order to find group differences. Such a method seeks to represent functional activity in general, and does not aim specifically to find networks that relate to the classification task.

In order to address the above limitations, in this paper, we propose a joint framework that combines network identification with classification, resulting in functionally meaningful networks that are highly discriminative of the two groups. Our method is akin to previous work in structural MRI [4], presenting a joint generative-discriminative formulation. We propose the use of sparse decompositions for network identification, resulting in a small set of *Sparse Connectivity Patterns* (SCPs) that are highly discriminative of the two groups, as well as functionally interpretable. Thus, this method reduces the high dimensionality of the connectivity data to a small set of strongly correlated regions, that are relevant to the classification task. In addition, as opposed to vectorizing the correlation matrices [1] the proposed method exploits the *positive-semi-definite* (PSD) property of the correlation matrix [5], in order to find stable SCPs.

To test its performance, we use the proposed method to investigate functional connectivity differences between children and young adults. Section 2 provides details of the proposed method. Section 3 describes the results and discusses the merits and drawbacks of the proposed method compared to other methods. Conclusions and future work are provided in section 4.

2 Identification of Discriminative Sparse Connectivity Patterns

A schematic diagram illustrating our method is shown in Figure 1. Given P regions, the input to the method is size $P \times P$ correlation matrices $\Sigma_n \succeq 0$, and the associated binary group membership y_n , for each subject n , $n = 1, 2, \dots, N$. We would like to find smaller networks, or SCPs common to all the subjects, such that the total connectivity within each network contributes to the two-group classification. Our formulation jointly optimizes two objectives: (1) Identification of SCPs (2) Learning discriminative SCPs. In the following sub-sections we provide the details of each, followed by the joint optimization strategy.

2.1 Identification of SCPs

We would like to find SCPs common to all the subjects, such that a non-negative combination of SCPs generates the correlation matrix Σ_n , for each subject n . We represent each SCP by a vector of region-weights $\mathbf{b}_k$, where $-1 \preceq \mathbf{b}_k \preceq 1$, $\mathbf{b}_k \in \mathbf{R}^P$, reflecting the

membership of the regions to SCP k . If two regions in $\mathbf{b}_k$ have the same sign, then they are positively correlated and opposing sign reflects anticorrelation. Thus, the rank-one matrix $\mathbf{b}_k \mathbf{b}_k^T$ reflects the correlation behavior of SCP k . To retain the P.S.D. nature of the correlation matrix, we would like to approximate each matrix $\{\Sigma_n\}_{n=1}^N$ by a non-negative combination of these sub-networks $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_K]$. Thus, we want

$\Sigma_n \approx \sum_{k=1}^K c_n(k) \mathbf{b}_k \mathbf{b}_k^T = \mathbf{B} \text{diag}(\mathbf{c}_n) \mathbf{B}^T$, where $\text{diag}(\mathbf{c}_n)$ denotes a diagonal matrix with values $\mathbf{c}_n \in \mathbf{R}_+^K$ along the diagonal.

We quantify the approximation above using the frobenius norm. Then the loss function $\mathcal{G}$ takes the form:

$$\mathcal{G}(\mathbf{B}, \mathbf{C}) = \sum_{n=1}^N \|\Sigma_n - \mathbf{B} \text{diag}(\mathbf{c}_n) \mathbf{B}^T\|_F^2 \quad (1)$$

In general, a *blind-decomposition* problem such as the above is ill-posed, i.e., multiple optima exist, and the solutions are not stable with respect to the noise in the data. To make the results stable, additional constraints need to be imposed on the matrix factors. Since known functional networks such as the visual or motor networks have small spatial extent compared to the whole brain, we use spatial sparsity as a constraint, similar to our prior work [6]. Hence we impose spatial sparsity on the SCPs by restricting the l_1 -norm of $\mathbf{b}_k$ to less than a constant value λ .

2.2 Learning Discriminative SCPs

Our objective is to find SCPs that can reconstruct that data as well as act as a discriminative basis that can classify the two groups. In other words, each SCP consists of regions which have the following properties (1) they are strongly positively or negatively correlated (2) the total absolute correlation between all the regions within an SCP contributes towards the group difference. The first property is modeled in the previous section; in this section we describe the discriminative term that models the second property.

Given an SCP $\mathbf{b}_k$, the scalar value $\mathbf{b}_k^T \Sigma_n \mathbf{b}_k$ measures the total absolute correlation between all the regions within the SCP for a given subject n . Computed for all SCPs, the K -dimensional vector $\text{diag}(\mathbf{B}^T \Sigma_n \mathbf{B})$ serves as the subject-specific measure that can be used in a multi-variate SVM framework. We use the squared hinge loss and l_2 regularization for the K -dimensional SVM hyperplane $\mathbf{w}$. The cost function for the discriminative term is:

$$\mathcal{D}(\mathbf{B}, \mathbf{w}) = \sum_{n=1}^N (1 - y_n \mathbf{w}^T \text{diag}(\mathbf{B}^T \Sigma_n \mathbf{B}))_+^2 \quad (2)$$

where y_n is the binary group label for subject n , and the subscript $+$ denotes the positive part of the argument.

2.3 Joint Optimization Framework

Bringing the terms in Eqns. 1 and 2 together, along with the l_2 regularizer for $\mathbf{w}$, we have the optimization problem

$$\begin{aligned} & \underset{\mathbf{B}, \mathbf{C}, \mathbf{w}}{\text{minimize}} (1 - \mu) \mathcal{G}(\mathbf{B}, \mathbf{C}) + \mu \mathcal{D}(\mathbf{B}, \mathbf{w}) + \|\mathbf{w}\|_2^2 \\ & \text{subject to } -1 \leq \mathbf{b}_k \leq 1, \|\mathbf{b}_k\|_1 \leq \lambda, k=1, \dots, K \quad (3) \\ & \mathbf{c}_n \geq 0, n=1, \dots, N \end{aligned}$$

where μ is the relative weighting fraction between the two terms. A value of $\mu = 0$ produces purely reconstructive SCPs. We use alternating minimization to iteratively solve for $\mathbf{B}$, $\mathbf{C}$ and $\mathbf{w}$. We use a projected gradient method [4] for $\mathbf{B}$ and $\mathbf{C}$ and the *libSVM* solver for $\mathbf{w}$ [7]. The parameter μ , which controls the trade-off between the two terms, will be linearly increased from a value of 0 to 1 during the iterative process. This ensures that the SCPs generated during the first few iterations are mainly reconstructive, which tend to be more stable.

Model Parameters—The free parameters of the proposed method are the number of SCPs K , and the sparsity level λ . Using grid search, for every pair of values in $K \in \{10, 20, \dots\}$ and $\lambda \in \{0.01, 0.02, \dots, 0.1\} * P$, we will use repeated five-fold cross-validation to find the optimal set of parameters.

3 Application to Study of Development

3.1 Data and Pre-processing

Data used here was drawn from the PNC database [8]. Our data consists of 91 children (age = 10.38 ± 1.01 yrs.) and 84 young adults (age = 20.21 ± 0.84 yrs.). As head motion is a known confound that correlates with age effects, we matched the two groups on motion, measured using the mean relative displacement [1,9] ($p = 0.79$). As described elsewhere in detail [8], Blood Oxygen Level Dependent (BOLD) fMRI was acquired using a whole-brain, echoplanar (EPI) sequence with the following parameters: 124 volumes, TR 3000 ms, TE 32 ms, effective voxel resolution $3.0 \times 3.0 \times 3.0$ mm.

Subject-level BOLD images fMRI images were affinely registered to the T1 image, followed by non-linear registration to the MNI 152 template. We used the 264 nodes defined in [10] for our experiments. Time-series data was preprocessed using a validated confound regression procedure [11]. Averaged time-courses corresponding to these 264 nodes were used to compute a symmetric Pearson correlation matrix $\Sigma_n \in \mathcal{S}_+^P$ for each subject.

3.2 Results Using Proposed Method

The results of the cross-validation provided an operating point of $K = 50$, $\lambda = 0.03P$ (roughly 10 nodes per SCP) at which the results are generalizable within sub-samples of the dataset. The cross-validation accuracy saturates at higher values of K . For these values, the proposed method gave an average classification accuracy of $76.3 \pm 7.08\%$ between children vs. young adults. The most discriminative SCPs are shown in Figures 2 and 3, rendered using *BrainNet Viewer* [12].

Children demonstrated stronger connectivity within the two SCPs, displayed in Figure 2. Positive correlations between the bilateral temporal poles and the parahippocampal gyri are significantly stronger in children, than in adults. This pattern is consistent with an increased degree of lateralization of temporal connectivity with development [13]. In contrast, adults demonstrated stronger connectivity within two SCPs including the motor network and the default mode network (Figure 3). These patterns are consistent with patterns of network segregation, whereby connectivity within certain major networks including the motor network and the default mode network increase with development [2,1]. Of the four SCPs, strongest p-value difference is exhibited by SCP 4, which shows increased anterior-posterior connectivity between two default-mode regions - medial Pre-Frontal Cortex (mPFC) and the Posterior Cingulate Cortex (PCC). This finding has been consistently found in other studies [2], further adding to the validity of our method.

3.3 Comparison with Other Methods

We compared our method with four alternate approaches: (1) l_2 -regularized l_2 -loss linear SVM [7] with pair-wise correlation values (2) l_1 -regularized l_2 -loss linear SVM (3) Principal Component Analysis (PCA), followed by classification (4) Independent Component Analysis (ICA), followed by classification. The first and second methods are purely discriminative, as they do not perform network identification. The third and fourth methods use unsupervised network identification methods, followed by classification using the total absolute connectivity values as features. The number of components K (in PCA and ICA) and the cost parameter for the SVM was chosen using cross-validation.

The classification performance for all the methods is reported in Table 1. The un-supervised PCA and ICA methods perform poorly. Of the three methods, l_2 -SVM provides a slightly better performance compared to the proposed method, although the difference in accuracies between the two methods is insignificant ($p = 0.0625$). The marginally higher accuracy provided by the SVM is due to the 1000-fold increase in the number of features used, leading to a complete loss of interpretability. This is illustrated in Figure 4, which displays the l_2 -SVM weight vector. The weight vector for l_1 -SVM is also shown in the same figure. While the l_1 penalty does dramatically reduce the number of features used, it does not necessarily alleviate the issue of non-interpretability. As explained earlier, strongly correlated features (connections) that are redundant to the classification are dropped. In contrast, the generative term $\mathcal{G}$ within the proposed method tends to retain these features by allocating them to the same SCP. Thus, a whole-brain discriminative pattern is split into multiple SCPs based on the dependencies between the connection strengths, allowing results to be interpreted within the context of known functional brain networks.

4 Conclusion

In this paper we presented a framework that performs supervised dimensionality reduction for functional connectivity data such that the discriminative pattern between two groups is preserved. Results demonstrate improved neurobiological interpretability compared to purely discriminative approaches without a significant loss of accuracy. SCPs that discriminate children from young adults are consistent with previously-described reports of increased temporal lateralization and functional network segregation. Future work will extend this method to continuous labels within a regression framework.

References

1. Satterthwaite TD, Wolf DH, Ruparel K, Erus G, Elliott MA, Eickhoff SB, Gennatas ED, Jackson C, Prabhakaran K, Smith A, et al. Heterogeneous impact of motion on fundamental patterns of developmental changes in functional connectivity during youth. *NeuroImage*. 2013; 83:45–57. [PubMed: 23792981]
2. Fair DA, Cohen AL, Power JD, Dosenbach NU, Church JA, Miezin FM, Schlaggar BL, Petersen SE. Functional brain networks develop from a local to distributed organization. *PLoS Computational Biology*. 2009; 5(5):e1000381. [PubMed: 19412534]
3. Venkataraman, A.; Kubicki, M.; Westin, CF.; Golland, P. Robust feature selection in resting-state fmri connectivity based on population studies. 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW); IEEE; 2010. p. 63-70.
4. Batmanghelich N, Taskar B, Davatzikos C. Generative-discriminative basis learning for medical imaging. *IEEE Transactions on Medical Imaging*. 2012; 31(1):51–69. [PubMed: 21791408]
5. Varoquaux, G.; Baronnet, F.; Kleinschmidt, A.; Fillard, P.; Thirion, B. Detection of brain functional-connectivity difference in post-stroke patients using group-level covariance modeling. In: Jiang, T.; Navab, N.; Pluim, JPW.; Viergever, MA., editors. MICCAI 2010, Part I. LNCS. Vol. 6361. Heidelberg: Springer; 2010. p. 200-208.
6. Eavani, H.; Satterthwaite, TD.; Gur, RE.; Gur, RC.; Davatzikos, C. Unsupervised learning of functional network dynamics in resting state fmri. In: Gee, JC.; Joshi, S.; Pohl, KM.; Wells, WM.; Zöllei, L., editors. IPMI 2013. LNCS. Vol. 7917. Heidelberg: Springer; 2013. p. 426-437.
7. Chang CC, Lin CJ. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*. 2011; 2:27:1–27:27.
8. Satterthwaite TD, Elliott MA, Ruparel K, Loughhead J, Prabhakaran K, Calkins ME, Hopson R, Jackson C, Keefe J, Riley M, et al. Neuroimaging of the philadelphia neurodevelopmental cohort. *NeuroImage*. 2014; 86:544–553. [PubMed: 23921101]
9. Van Dijk KR, Sabuncu MR, Buckner RL. The influence of head motion on intrinsic functional connectivity mri. *Neuroimage*. 2012; 59(1):431–438. [PubMed: 21810475]
10. Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL, Petersen SE. Functional network organization of the human brain. *Neuron*. 2011; 72(4):665–678. [PubMed: 22099467]
11. Satterthwaite TD, Elliott MA, Gerraty RT, Ruparel K, Loughhead J, Calkins ME, Eickhoff SB, Hakonarson H, Gur RC, Gur RE, et al. An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *Neuroimage*. 2013; 64:240–256. [PubMed: 22926292]
12. Xia M, Wang J, He Y. Brainnet viewer: A network visualization tool for human brain connectomics. *PloS One*. 2013; 8(7):e68910. [PubMed: 23861951]
13. Zuo XN, Kelly C, Di Martino A, Mennes M, Margulies DS, Bangaru S, Grzadzinski R, Evans AC, Zang YF, Castellanos FX, et al. Growing together and growing apart: regional and sex differences in the lifespan developmental trajectories of functional homotopy. *The Journal of Neuroscience*. 2010; 30(45):15034–15043. [PubMed: 21068309]

Fig. 1.

Schematic illustrating the joint framework. Panel to the left describes the SCP identification term, which factorizes connectivity matrices Σ_n of each subject n into a set of common SCPs $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_K]$ and its associated coefficients. Panel to the right illustrates a linear SVM, which uses the total absolute connectivity values $\text{diag}(\mathbf{B}^T \Sigma_n \mathbf{B})$ as input features to classify two groups, resulting in the hyperplane $\mathbf{w}$.

Fig. 2.

Top two SCPs whose total connectivity is stronger in children. 3D brain rendering (top) displays location of regions belonging to the SCP. Opposing colors (red/blue) indicate anti-correlation between the regions. Size of nodes reflect SCP values $\mathbf{b}_k$. Corresponding graphs (bottom) plot total connectivity within SCP for each subject vs. subject age. Uni-variate p-value scores comparing total connectivity between two groups are also shown.

Fig. 3.

Top two discriminative SCPs whose total connectivity is stronger in adults

Fig. 4.
SVM weight vector for l_2 -SVM (left) and l_1 -SVM (right)

Table 1

Accuracy and AUC values for all methods

Method	Acc (%)	AUC	# features
Proposed	76.3 ± 7.08	0.85 ± 0.06	50
l_2 -SVM	79 ± 1.45	0.87 ± 0.01	34716
l_1 -SVM	74 ± 2.25	0.81 ± 0.01	139
PCA+SVM	65.2 ± 2.30	0.67 ± 0.01	180
ICA+SVM	64.8 ± 1.42	0.71 ± 0.03	210