

Classifying Phenotypes Based on the Community Structure of Human Brain Networks

Anvar Kurmukov¹(✉), Marina Ananyeva², Yulia Dodonova¹, Boris Gutman³,
Joshua Faskowitz³, Neda Jahanshad³, Paul Thompson³, and Leonid Zhukov²

¹ Kharkevich Institute for Information Transmission Problems, Moscow, Russia
kurmukovai@gmail.com

² National Research University Higher School of Economics, Moscow, Russia

³ Imaging Genetics Center, Stevens Neuroimaging and Informatics Institute,
University of Southern California, Marina del Rey, USA

Abstract. Human anatomical brain networks derived from the analysis of neuroimaging data are known to demonstrate modular organization. Modules, or communities, of cortical brain regions capture information about the structure of connections in the entire network. Hence, anatomical changes in network connectivity (e.g., caused by a certain disease) should translate into changes in the community structure of brain regions. This means that essential structural differences between phenotypes (e.g., healthy and diseased) should be reflected in how brain networks cluster into communities. To test this hypothesis, we propose a pipeline to classify brain networks based on their underlying community structure. We consider network partitionings into both non-overlapping and overlapping communities and introduce a distance between connectomes based on whether or not they cluster into modules similarly. We next construct a classifier that uses partitioning-based kernels to predict a phenotype from brain networks. We demonstrate the performance of the proposed approach in a task of classifying structural connectomes of healthy subjects and those with mild cognitive impairment and Alzheimer’s disease.

1 Introduction

Understanding disease-related changes in human brains has always been a challenge for neuroscience. A growing field of network science provides a powerful framework to study these changes [5]. This is because any shifts in brain anatomy or functioning are rarely confined to a single locus but rather affect the entire network system.

Human brain networks have been extensively studied in a recent decade. These networks, called connectomes, are constructed from neuroimaging data and represent either anatomical or functional connectivity between cortical brain regions. Several aspects of typical brain network organization have been described, including their modular structure. Modular structure of a network

means that its nodes tend to group into modules, or communities, with close within-group connections and sparse between-group connectivity. Meunier et al. [10] discuss why it is reasonable for human brains to be modular, and also review studies on the community structure of human connectomes. Alexander-Bloch et al. [1] demonstrate that brain network community structure differs between phenotypes (healthy subjects and those with childhood-onset schizophrenia).

This suggests that brain network community structure captures enough information about network topology to classify phenotypes associated with certain diseases. To test this hypothesis, one needs a framework to classify networks based on similarity in their partitions into communities. Recently, Kurmukov et al. [8] proposed such an algorithm. Its basic idea was to detect non-overlapping brain network communities, measure pairwise distances between the obtained network partitions and use these distances in a kernel classification framework. However, [8] only considered non-overlapping brain network communities and demonstrated the performance of the proposed method on a small dataset.

Although non-overlapping communities are more commonly studied in network neuroscience, a model of community structure that allows for overlapping offers a more realistic model of brain-network organization [13]. Some cortical areas are known to be heteromodal and to have a role in multiple networks; consistently with this, current theories on brain organization suggest that cognitive functions are organized into widespread, segregated, and overlapping networks. Thus, clarifying the overlapping structure of brain network communities remains a challenging and relatively unexplored research area.

In this study, we generalize the classification approach [8] by considering both non-overlapping and overlapping communities of cortical brain regions. We show how both types of partitions may be used to estimate distances between brain networks and run a kernel classifier on these distances. Based on a large Alzheimer’s Disease Neuroimaging Initiative dataset, we question whether similarity in brain modular structure can help to differentiate subjects with different diagnoses and tackle this question with the proposed approach.

2 Similarity of Brain Network Community Structures

Clustering networks into communities has attracted much attention in graph theory. Here, we only briefly describe the algorithms that we used for partitioning brain networks into communities (both non-overlapping and overlapping), and discuss how community structures of different brain networks may be quantitatively compared.

2.1 Detecting Communities in Structural Brain Networks

We use two approaches to detect brain network community structure. Both approaches aim to identify communities, or groups of tightly anatomically connected cortical regions. The major difference is that the first approach separates brain network regions into unique, non-overlapping modules, while the

second algorithm allows for nodes belonging to more than one community. Algorithms of the former type are much more common in graph theory, and hence much more widely used in applications including brain network analysis [10]. However, as discussed above, overlapping community structures offer more powerful description of human brain organization, although they are much rarer evaluated [13].

In this study, we use the Louvain method [2] to produce non-overlapping partitions of structural connectomes. Given a graph $G(E, V)$ with a set of edges E , a set of nodes V , and the adjacency matrix A , the algorithm divides nodes V into groups $\{V_1, V_2, \dots, V_k\}$ so that $V_1 \cup V_2 \cup \dots \cup V_k = V$. Similarly to many other graph partitioning methods, it optimizes the so-called modularity by maximizing the number of intra-community connections and minimizing the number of inter-community links. The Louvain algorithm is a two-step iterative procedure. It starts with each node assigned to a separate cluster. In the first step, it moves each node i to a cluster of one of its neighbors j so that the gain in modularity is maximal. Once there is no such move that improves modularity, the algorithm proceeds to the second step, builds a new graph wherein nodes are clusters from the previous step, and reapplies the first step. Importantly, the Louvain method does not require any a-priori defined number of communities to be detected.

Second, we aim to estimate overlapping communities of structural brain networks. Two types of algorithms can accommodate this, differing in whether they use crisp or fuzzy assignment of nodes into communities. The former means that each node either belongs to each of the possible clusters or not, while the latter allows for a strength of belonging to a community. We detect fuzzy communities based on non-negative matrix factorization (NMF) [7]. Given a non-negative graph adjacency matrix A of size $n \times n$ (n being the number of nodes in brain network), we find its low-rank approximation

$$A \simeq WH, \quad (1)$$

where W is of size $n \times k$ and H is $k \times n$. A parameter k is usually selected to be much smaller than n and stands for a number of communities to be detected. Elements h_{ij} of a normalized matrix H denote probability of a node i being in a community j . Unlike the first method, the NMF algorithm requires specifying the number of communities. In our computational experiments, we show results obtained for different values of k .

2.2 Measuring Distance Between Community Structures

We aim to evaluate similarity in community structure of brain networks stemming from different subjects, possibly with different diagnoses. Hence, we need to introduce a measure of distance between two partitions obtained from different brain networks. This becomes possible because nodes in connectomes (i.e., cortical regions) are uniquely labeled, and the set of labels is the same across connectomes obtained with the same parcellation atlas.

To estimate pairwise similarity of partitions of different brain networks we use two modifications of mutual information (MI) score. Let $U = \{U_1, U_2, \dots, U_l\}$

and $V = \{V_1, V_2, \dots, V_k\}$ be partitions of two networks G_U and G_V with the same sets of node labels, l and k be the number of clusters in the partitions U and V , respectively. MI between the partitions U and V is defined by:

$$MI(U, V) = \sum_{i=1}^l \sum_{j=1}^m P(i, j) \log \frac{P(i, j)}{P(i)P'(j)}, \quad (2)$$

For brain network partitions into non-overlapping communities, we use adjusted mutual information, AMI [12]. We measure similarity between partitions into overlapping communities based on normalized mutual information (NMI, [9]). A property of the latter measure is that it only accepts partitions into overlapping modules with crisp node assignment. To accommodate this, we binarize the community membership matrix H (1) using a threshold parameter; we demonstrate how the results of our computational experiments change depending on this parameter.

Both measures take values in $[0, 1]$, with the value of 1 indicating exactly the same partitions. We thus define a distance $\omega(G_U, G_V)$ between the community structures of networks G_U and G_V by:

$$\omega(G_U, G_V) = 1 - I(U, V), \quad (3)$$

where $I(U, V)$ is the index of similarity (AMI or NMI). Networks with the same community structure now have zero distance, and the maximum distance is close to 1.

3 Classifying Connectomes Based on their Community Structure

Since we obtained an optimal partition of each brain network into communities and introduced a measure of difference between community structures, we can proceed to the question of whether community structure of cortical brain regions provides enough information for differentiating between phenotypic classes. This question can be addressed in a machine learning framework.

Given a set of brain networks G_i (each with known community structure), class labels y_i , a training set of pairs (G_i, y_i) and the test set of input objects G_j , the task is to make a best possible prediction of the unknown class label y_j . Provided that we already defined a matrix of pairwise distances $\omega(G_U, G_V)$ (3), the most straightforward approach to classification is to convert the obtained distance matrix into a kernel and feed it to a kernel classifier. We accommodate this by exponentiating the obtained distances:

$$K(G_U, G_V) = e^{-\alpha\omega(G_U, G_V)}, \quad (4)$$

and run the support vector machines (SVM) classifier with the obtained kernel.

4 Experiments: Network-Based Alzheimer’s Disease Classification

We argue that if the community structure of anatomical brain networks is affected by a disease in a certain manner, it should be possible to differentiate between healthy and diseased brain networks solely based on similarity in their community structures. In other words, brain networks stemming from the same class (e.g., obtained for healthy participants) should be more similar in their community structure than brain networks from different phenotypic classes (e.g., normal and diseased brains). Using the approach described in the previous sections, we test this hypothesis in a task of classifying Alzheimer’s disease (AD), late- and early-stage mild cognitive impairment (LMCI and EMCI), and healthy participants (normal controls, NC).

4.1 Data and Network Construction

We use the Alzheimer’s Disease Neuroimaging Initiative (ADNI2) database which comprises a total of 228 individuals (756 scans), with a mean age at baseline visit 72.9 ± 7.4 years, 96 females. Each individual has at least 1 brain scan and at most 6 scans. The data include 47 people with AD (136 AD scans), 40 individuals with LMCI (147 LMCI scans), 80 individuals with EMCI (283 EMCI scans), and 61 healthy participants (190 scans).

Corrected T1-weighted images were processed with Freesurfer’s [4] recon-all pipeline to obtain a triangle mesh of the grey-white matter boundary registered to a shared spherical space, as well as corresponding vertex labels per subject. We used cortical parcellation based on the Desikan-Killiany (DK) atlas [3] which includes 68 cortical brain regions. T1w images were aligned (6-dof) to the 2 mm isotropic MNI 152 template. These were used as the template to register the average b_0 of the DWI images, in order to account for EPI related susceptibility artifacts. DWI images were also corrected for eddy current and motion related distortions. Rotation of b-vectors was performed accordingly. Tractography for ADNI data was then conducted using the distortion corrected DWI in 2-mm isotropic MNI 152 space. Probabilistic streamline tractography was performed using the Dipy [6] LocalTracking module and implementation of constrained spherical deconvolution (CSD) [11] with a spherical harmonics order of 6. Streamlines longer than 5 mm with both ends intersecting the cortical surface were retained.

Edge weights in the original cortical connectivity matrices were proportional to the number of streamlines detected by the algorithm. We binarize these weights by:

$$a_{ij}^{\text{binarized}} = \begin{cases} 1 & \text{if } a_{ij} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

Thus, we only work with non-weighted graphs throughout the paper.

4.2 Experimental Setup

We obtain the best partition of each network into non-overlapping communities using the Louvain algorithm and compute a matrix of pairwise distances between partitions with the AMI metric. In parallel, we cluster each network into overlapping communities based on NMF and produce a matrix of pairwise NMI distances between these clusterings. This second algorithm requires two parameters (the number of communities and the cluster membership threshold), we report how the results of the overall pipeline change depending on their particular values. For purposes of comparison, we also compute pairwise distances between connectomes using the L_2 (Frobenius) norm.

For each of the three distance matrices, we compute a kernel by (4) and run an SVM classifier with this kernel. We vary the values of α in (4) from 0.01 to 10 and the penalty parameter of the classifier from 0.1 to 50; we only report the results obtained for the optimal values of these technical parameters.

We consider four binary classification tasks: AD versus NC, AD versus LMCI, LMCI versus EMCI, EMCI versus NC. We find optimal values for all parameters in the simplest task of classifying AD versus NC and keep them fixed in the remaining tasks. We use 10-fold cross-validation to train SVM on a subsample and make predictions for an unseen part of a sample. As the data include several networks for each subject, we use subjects rather than networks to split data into train and test and put all networks of the same subject into a respective category (thus avoiding data leakage).

We train the models on networks and next make a subject-based prediction as an average of predictions obtained for individual networks; this method of evaluation (subject-based rather than network-based) does not affect the reported results in any systematic way. We repeat the procedure 50 times with different data splits and report ROC AUC as a quality metric. All scripts are available at <https://github.com/kurmukovai/GRAIL2017-communities>.

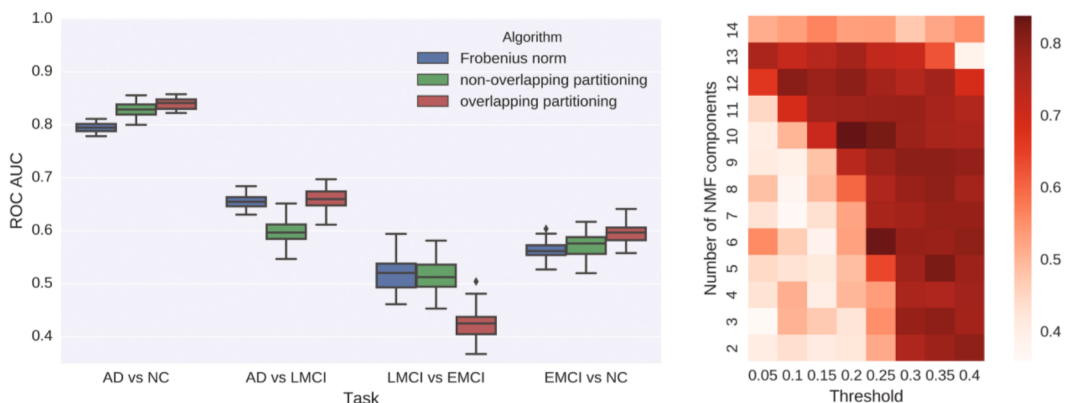


Fig. 1. Left: Classification results. Right: Results of classifying AD versus NC based on the overlapping community detection algorithm, depending on the number of components and the membership threshold; colorbar shows average ROC AUC values.

4.3 Results and Discussion

Figure 1 (left) shows the results of classifying AD, LMCI, EMCI and healthy controls based on L_2 -distance between the structural connectivity matrices of brain networks and on the distances representing similarity in brain community structures.

As expected, classifying AD versus NC was the simplest task, while for EMCI versus LMCI all algorithms only performed at chance level. For the tasks with reasonable overall classification quality, an algorithm based on overlapping community structures slightly outperformed the other algorithms. For AD versus NC, the model with overlapping communities provides an ROC AUC of 0.840 ± 0.010 ; the one based on non-overlapping communities gives an ROC AUC 0.828 ± 0.013 . For this task, Fig. 1 (right) shows how the outcomes of the

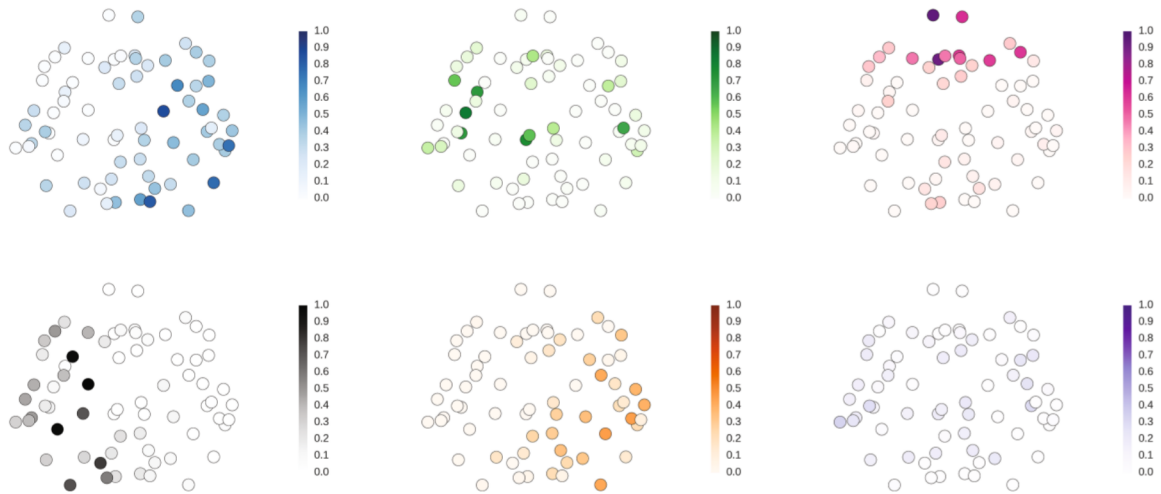


Fig. 2. Six overlapping communities: an example of a single network (healthy subject) with the nodes shown in their original 3D coordinates (axial view); color intensity is proportional to the strength of belonging to the respective community



Fig. 3. Comparison of the non-overlapping (left) and overlapping (right) community structures obtained for the same example graph as in Fig. 3; node size is proportional to its degree (the number of edges coming from the respective node). Right plot is produced by selecting a single community for each node based on the maximal membership probability.

best-performing algorithm depend on the predefined number of clusters and the threshold of cluster membership used in computing the NMI distance. The best classification results are obtained with the community structure of six overlapping components, with membership probability thresholded at 0.25.

Figure 2 illustrates the obtained community structure based on a single example graph. Figure 3 compares the non-overlapping and the simplified overlapping community structures obtained for the same graph. The two algorithms seem to identify similar communities, but the outcome of the overlapping community detection algorithm retains more information on the underlying brain network structure.

5 Conclusions

Human brain networks show modular structure which arises based on the entire system of connections between cortical brain regions. Systematic shifts in connectivity patterns, for example those caused by a brain disease, may be expected to induce changes in the community structure of the macroscale brain networks. If true, that would produce similar modular structure in brain networks of individuals with the same phenotype (e.g., Alzheimer’s disease) and different community structures in brain networks from different phenotypes (e.g., patients versus healthy controls).

In this study, we explored whether the community structure of anatomical human brain networks provides enough information to differentiate phenotypes of the respective individuals. We proposed a framework to compare both overlapping and non-overlapping community structures of brain networks within the machine learning settings. We demonstrated the performance of the proposed pipeline in a task of classifying Alzheimer’s disease, mild cognitive impairment, and healthy participants. Algorithms based on the distances between partitions of brain networks slightly outperformed the baseline. Models that made full use of overlapping community structures performed slightly better than those based on non-overlapping community structures.

To sum up, the modular structure of anatomical brain networks seems to capture important information about the underlying network structure and can be useful in classifying phenotypes. Further studies are needed to study this idea on other phenotypic categories, and to specifically explore overlapping community structure of cortical regions in human anatomical brain networks.

Acknowledgments. The data used in preparing this paper were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database. A complete listing of ADNI investigators and imaging protocols may be found at www.adni.loni.usc.edu.

The results of Sects. 2–5 are based on the scientific research conducted at IITP RAS and supported by the Russian Science Foundation under grant 17-11-01390.

References

1. Alexander-Bloch, A.F., Gogtay, N., Meunier, D., Birn, R., et al.: Disrupted modularity and local connectivity of brain functional networks in childhood-onset schizophrenia. *Front. Syst. Neurosci.* **4** (2010)
2. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *J. Stat. Mech. Theory Exp.* (2008)
3. Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., et al.: An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *Neuroimage* **31**(3), 968–980 (2006)
4. Fischl, B.: Freesurfer. *Neuroimage* **62**(2), 774–781 (2012)
5. Fornito, A., Zalesky, A., Breakspear, M.: The connectomics of brain disorders. *Nature Reviews. Neurosci.* **16**, 159–172 (2015)
6. Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., et al.: Dipy, a library for the analysis of diffusion mri data. *Front. Neuroinformatics* **8**, 8 (2014)
7. Kuang, D., Ding, C., Park, H.: Symmetric nonnegative matrix factorization for graph clustering. In: *The 12th SIAM International Conference on Data Mining*, pp. 106–117 (2012)
8. Kurmukov, A., Dodonova, Y., Zhukov, L.E.: Classification of normal and pathological brain networks based on similarity in graph partitions. In: *2016 IEEE 16th International Conference Data Mining Workshops (ICDMW)*, pp. 107–112 (2016)
9. McDaid, A.F., Greene, D., Hurley, N.: Normalized mutual information to evaluate overlapping community finding algorithms (2011)
10. Meunier, D., Lambiotte, R., Bullmore, E.T.: Modular and hierarchically modular organization of brain networks. *Frontiers of Neuroinformatics* **4** (2010)
11. Tax, C.M., Jeurissen, B., Vos, S.B., Viergever, M.A., Leemans, A.: Recursive calibration of the fiber response function for spherical deconvolution of diffusion mri data. *Neuroimage* **86**, 67–80 (2014)
12. Vinh, N.X., Epps, J., Bailey, J.: Information theoretic measures for clusterings comparison: variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.*, 2837–2854 (2010)
13. Wu, K., Taki, Y., Sato, K., et al.: The overlapping community structure of structural brain network in young healthy individuals. *PLoS One* **6** (2011)