

Modularity Reinforcement for Improving Brain Subnetwork Extraction

Chendi Wang¹(✉), Bernard Ng², and Rafeef Abugharbieh¹

¹ Biomedical Signal and Image Computing Lab, UBC, Vancouver, Canada
`{chendiw,rafeef}@ece.ubc.ca`

² Department of Statistics, UBC, Vancouver, Canada
`bernardng@gmail.com`

Abstract. Functional subnetwork extraction is commonly employed to study the brain’s modular structure. However, reliable extraction from functional magnetic resonance imaging (fMRI) data remains challenging. As representations of brain networks, brain graph estimates are typically noisy due to the pronounced noise in fMRI data. Also, confounds, such as region size bias, motion artifacts, and signal dropout, introduce region-specific bias in connectivity, e.g. a node in a signal dropout area tends to display lower connectivity. The traditional approach of global thresholding might thus remove relevant edges that have low connectivity due to confounds, resulting in erroneous subnetwork extraction. In this paper, we present a modularity reinforcement strategy that deals with the above two challenges. Specifically, we propose a local thresholding scheme that accounts for region-specific connectivity bias when pruning noisy edges. From the resulting thresholded graph, we derive a node similarity measure by comparing the adjacency structure of each node, i.e. its connection fingerprint, with that of other nodes. Drawing on the intuition that nodes belonging to the same subnetwork should have similar connection fingerprints, we refine the brain graph with this similarity measure to reinforce its modularity structure. On synthetic data, our strategy achieves higher accuracy in subnetwork extraction compared to using standard brain graph estimates. On real data, subnetworks extracted with our strategy attain higher overlaps with well-established brain systems and higher subnetwork reproducibility across a range of graph densities. Our results thus demonstrate that modularity reinforcement with our strategy provides a clear gain in subnetwork extraction.

Keywords: Brain graph estimation · Connection fingerprint · fMRI · Local thresholding · Subnetwork extraction

1 Introduction

The human brain naturally befits a graphical representation, where brain regions and their pair-wise interactions constitute graph nodes and weighted edges, respectively. An important attribute of the brain is its modular structure, in which specific subnetworks of brain regions work in tandem to execute various functions. Functional magnetic resonance imaging (fMRI) is widely used

for studying this modular structure of the brain. However, reliable subnetwork extraction from fMRI data remains challenging. First, the brain network topology may be obscured by noisy connectivity estimates [1]. Second, confounds, such as region size bias [2], effects of motion artifacts [3], and signal dropouts due to susceptibility artifacts (especially in regions like the orbitofrontal cortex and the inferior temporal lobe) [4], introduce region-specific biases to the connectivity estimates.

The conventional way for dealing with noisy connectivity matrices is to apply global thresholding (GT) by either keeping only connections with values above a certain threshold or keeping a certain graph density [1]. Due to region-specific connectivity biases, e.g. brain regions in signal dropout locations tend to display lower connectivity, certain regions that do belong to a subnetwork might not appear as such based on the fMRI measurements, especially after GT, which prunes weak edges. To mitigate this overlooked problem, a local thresholding (LT) method based on the minimal spanning tree and k-nearest neighbors (MST-kNN) has been proposed [5]. The idea in [5] was to build a single connected graph using the MST and expand the tree by adding edges from each node to its nearest neighbors until a desired graph density is reached. However, both key steps of enforcing a single connected graph and adding edges to all nodes when expanding the tree lack neuro-scientific justifications. A few studies have explored spectral graph wavelet transform for graph de-noising [6], but this approach does not explicitly handle region-specific connectivity biases. In fact, most existing connectivity estimation and subnetwork extraction techniques [1, 7] do not account for these biases.

In this paper, we propose a modularity reinforcement strategy for improving brain subnetwork extraction. To deal with noisy edges and region-specific connectivity biases, we propose a local thresholding scheme that normalizes the connectivity distribution of each node prior to thresholding (Sect. 2.1). Also, since node pairs belonging to the same subnetwork presumably connect to a similar set of brain regions, i.e. have similar connection fingerprints, we derive a node similarity measure from the thresholded graph by comparing the adjacency structure of each node pair, and refine the graph with this similarity measure to reinforce its modularity structure (Sect. 2.2). More reliable subnetwork extraction is consequently facilitated on the refined graph (Sect. 2.3). To set the number of subnetworks, we adopt an automated technique based on graph Laplacian [8], and compare that against the conventional modularity-maximization approach [9]. We validate our modularity reinforcement strategy on both synthetic data and real data from the Human Connectome Project (HCP).

2 Methods

2.1 Local Thresholding

Due to region-specific connectivity biases, conventional GT might prune relevant connections with weak edge strength. To account for these biases, we present here a LT scheme. The idea is to first normalize the connectivity distribution of each

node into a uniform interval to rectify the biases. Subsequent global thresholding on this normalized graph would have the effect of applying local thresholding on each node. Specifically, let \mathbf{C} be an $n \times n$ connectivity matrix, where n is the number of nodes in the brain graph. We normalize the connectivity distribution by mapping each row of \mathbf{C} from $[\min(\mathbf{C}_{i,:}), \max(\mathbf{C}_{i,:})]$ to $[0, 1]$, where $\mathbf{C}_{i,:}$ denotes row i of \mathbf{C} corresponding to the connectivity between brain region i and all other regions in the brain. A threshold is then applied to generate a binary adjacency matrix, \mathbf{G} , which we then symmetrize by taking the union of \mathbf{G} and \mathbf{G}^T : $\mathbf{A} = \mathbf{G}_{i,j} \cup \mathbf{G}_{j,i}$. This binary adjacency matrix \mathbf{A} is used to mask out the noisy edges from \mathbf{C} : $\hat{\mathbf{C}}_{i,j} = \mathbf{A}_{i,j} \mathbf{C}_{i,j}$, which is equivalent to applying a local threshold to $\mathbf{C}_{i,:}$ for all i . We note that in the event that noisy nodes are accidentally included, some of the connections to these noise nodes (that might not be kept by GT) would be kept by LT due to the normalization step.

2.2 Modularity Reinforcement

Since nodes working in tandem are expected to have similar connection fingerprints, given \mathbf{A} , where $\mathbf{A}_{i,:}$ is the connection fingerprint of node i , we define the similarity between a pair of nodes (i, j) as the number of common adjacent nodes they share, normalized by the minimum node degree of the node pair:

$$\mathbf{S}_{i,j} = \frac{\sum_{k=1}^n \mathbf{A}_{i,k} \mathbf{A}_{j,k}}{\min(d_i, d_j)} \quad (1)$$

where $d_i = \sum_{k=1}^n \mathbf{A}_{i,k}$. We use the minimal degree for normalization, instead of e.g. the average degree, so that connections associated with hub nodes (nodes with more edges) will not be overly down-weighted. Since nodes within a subnetwork are expected to share more adjacent neighbors than nodes belonging to different subnetworks, \mathbf{S} boosts the within-subnetwork edges while suppresses the between-subnetwork edges, which highlights the modular pattern inherent in $\hat{\mathbf{C}}$: Hence, we use \mathbf{S} to refine $\hat{\mathbf{C}}$ to reinforce its modular structure: $\hat{\mathbf{C}}_{i,j}^{\mathbf{S}} = \mathbf{S}_{i,j} \hat{\mathbf{C}}_{i,j}$.

2.3 Subnetwork Extraction

For subnetwork extraction, we employ normalized cuts (Ncuts), chosen due to its wide use by the fMRI community. To set the number of subnetworks, m , we adopt an automated technique based on the spectral properties of the graph Laplacian: $\mathbf{L} = \mathbf{D} - \mathbf{W}$, where \mathbf{W} is a connectivity matrix, $\mathbf{D}_{ii} = \sum_{k=1}^n \mathbf{W}_{i,k}$. Specifically, an eigenvalue of 1 has been shown to correspond to the transition where single isolated nodes would no longer be declared as a subnetwork [8]. We thus set m to the number of eigenvalues of \mathbf{L} with values less than 1.

3 Materials

3.1 Synthetic Data

To illustrate our strategy, we synthesized a small-scale network consisting of $n = 13$ nodes in Fig. 1. We also generated synthetic data that cover 100 random

network configurations with n set to 100 nodes. For each network configuration, the number of subnetworks, N , was randomly selected from [10, 20]. The number of regions within each subnetwork was set to round $(n/N) + r$, where r was randomly selected from $[-2, 2]$. With the resulting configuration, we created the corresponding adjacency matrix, Σ , and drew time courses with 4,800 samples (analogous to real data) from $N(0, \Sigma)$. We then added Gaussian noise to the time courses with signal-to-noise ratio randomly set between $[-6\text{dB}, -3\text{dB}]$. Sample covariance was then estimated from these time courses with correlation values associated with $q\%$ of the nodes reduced by $z\%$, where q was randomly selected from $[20\%, 30\%]$ and z was randomly selected from $[30\%, 40\%]$ to simulate region-specific connectivity biases for smaller brain regions [2].

3.2 Real Data

We used the resting state fMRI scans of 77 healthy subjects (36 males and 41 females, ages ranging from 22 to 35) from the HCP Q3 dataset [10]. The data comprised two sessions, each having a 30 min acquisition with a TR of 0.72 s and an isotropic voxel size of 2 mm. Preprocessing already applied to the data by HCP [11] included gradient distortion correction, motion correction, spatial normalization to MNI space, and intensity normalization. Additionally, we regressed out motion artifacts, mean white matter and cerebrospinal fluid signals, and principal components of high variance voxels [12], followed by band-pass filtering with cutoff frequencies of 0.01 and 0.1 Hz. We used the Will90fROI atlas [13] and the Harvard-Oxford (HO) atlas [14] to define regions of interest (ROIs). The Will90fROI and HO atlas have 90 and 112 ROIs, respectively. Voxel time courses within ROIs were averaged to generate region time courses. The region time courses were demeaned, normalized by the standard deviation, and concatenated across subjects for extracting group subnetworks. The Pearson’s correlation values between the region time courses were taken as estimates of connectivity. Negative elements in the connectivity matrix were set to zero due to the currently unclear interpretation of negative connectivity [15].

4 Results and Discussion

We compared our strategy (LT with modularity reinforcement - LTMR) against GT, LT, GT with modularity reinforcement (GTMR) and MST-kNN in [5]. LT was implemented using our proposed scheme (Sect. 2.1). GTMR was implemented by deriving adjacency matrices with global thresholding, and subsequently executing our proposed modularity reinforcement strategy (Sect. 2.2). Instead of using a specific threshold, we examine a range of graph densities to test the robustness of our proposed strategy. For synthetic data, evaluation was based on the accuracy of subnetwork extraction. To estimate accuracy, we matched the extracted subnetworks to the ground truth subnetworks using Hungarian clustering [16] with the Dice coefficient: $DC = 2|X \cap Y| / (|X| + |Y|)$, where X is the set of regions of an extracted subnetwork and Y is the set of regions of

a ground truth subnetwork. The average DC over matched subnetworks was taken as accuracy. For real data, we assessed the overlap between the extracted subnetworks and fourteen well-established brain systems [13] and subnetwork reproducibility for a range of graph densities [14] using DC.

4.1 Synthetic Data

An example of the various steps of our strategy is shown in Fig. 1c–f to demonstrate how our strategy highlights the modular structure of the graph. With GT (Fig. 1c), node 2 was isolated from subnetwork 1. In contrast, our LT scheme (Fig. 1d) was able to preserve node 2. Also, with our LT (Fig. 1d), one of between-subnetwork edges (i.e. edges between nodes 6 and 7 & nodes 6 and 9) was pruned, which would help prevent the two subnetworks from being declared as one, whereas none of between-network edges was pruned using GT (Fig. 1c). Further, refining the graph (Fig. 1c, d) with our similarity helped to highlight the modular pattern (Fig. 1e, f), e.g. the between-network edges which were similar to or higher than some within-network edges (especially those edges between nodes 12 and 13, node 2 and 1 & nodes 2 and 5) in Fig. 1c, d were suppressed by our similarity to be the lowest values in Fig. 1e, f.





Fig. 1. Schematic illustrating our method using small scale example having two subnetworks with each subnetwork having a provincial hub (blue) and linked by a connector hub (orange). In (b), warmer color indicates higher connectivity and black dots indicate the ground truth adjacency matrix. We denote \bar{C} as global thresholded, and \hat{C} as local thresholded connectivity matrix. At a graph density of 0.25, the GT generated isolated node 2 in (c), while our LT preserved two edges linked to node 2 in (d). Refining the graph (c) and (d) suppressed the between-network edges (edges between nodes 6 and 7 & nodes 6 and 9) to be the lowest connectivity in (e) and (f).

On the 100 synthetic dataset with 100 nodes over a density range of [0.005, 0.5] at an interval of 0.01, LTMR achieved significantly higher accuracy (average DC = 0.6735) than GT (average DC = 0.6216, $p = 7.56e-10$), LT (average DC = 0.6537, $p = 2.89e-7$), and MST-kNN (average DC = 0.6327, $p = 7.38e-8$) based on Wilcoxon signed rank test. LTMR also achieved higher DC than GTMR (average DC = 0.6610, $p = 0.34$), though did not reach significance.

4.2 Real Data

We first evaluated our strategy by examining the overlap between our extracted subnetworks and 14 well-established brain systems presented in [13], which we



(a) Overlap with established subnetworks (b) Reproducibility over density range

Fig. 2. Subnetwork extraction on real data at graph densities from 0.05 to 0.5 at interval of 0.05. Blue = GT, green = LT, black = GTMR, cyan = MST-kNN, and red = our proposed LTMR strategy. Dash lines indicate average value. In (b), the DC of the reference density of 0.2 was left blank, since inclusion of DC = 1 might mislead the reader. In both (a) and (b), local thresholding outperforms the global thresholding, and modularity reinforcement further increases DC compared to using connectivity alone. Our proposed strategy attained the highest DC overall.

used as ground truth, Fig. 3a. For this assessment, we only considered connectivity matrices based on the Will90fROI atlas [13]. Our proposed LTMR achieved an average DC of 0.6222, which was significantly higher than GT (average DC = 0.5384, $p = 0.002$), MST-kNN (average DC = 0.4567, $p = 0.002$), GTMR (average DC = 0.5422, $p = 0.006$), and higher than LT (average DC = 0.5936, $p = 0.063$), as shown in Fig. 2a. At a graph density of 0.5041, corresponding to no thresholding except negative correlation removal, a DC of 0.5667 was attained, suggesting that some thresholding to remove noisy edges is beneficial. We note that although some node-wise variations in connectivity distribution might have a neuronal basis, we postulate that these variations would be overwhelmed by the various confound-induced connectivity biases, as supported by how local thresholding outperforms global thresholding. We further note that an average m of 11 was estimated with the Laplace approach, whereas an average m of 4 was estimated with modularity maximization. This result shows the resolution limits of modularity maximization [9], i.e. it tends to underestimate the number of subnetworks in favoring network partitions with groups of modules combined into larger communities. This suggests the need to explore alternative techniques for estimating the number of subnetworks.

We next evaluated the subnetwork reproducibility over a range of graph densities. We used connectivity matrices based on the HO atlas, which has larger brain coverage than the Will90fROI atlas but does not have subnetwork labels assigned to the regions. We set subnetworks corresponding to an edge density of 0.2 as the reference. Based on the Laplace approach, the optimal number of subnetworks was found to be 11 ± 5 over the range of graph density examined. Our proposed strategy achieved an average DC of 0.7302, which is significantly higher than that of GT (DC = 0.6121, $p = 0.004$), LT (DC = 0.6677, $p = 0.027$), MST-kNN (DC = 0.5737, $p = 0.003$), and higher than GTMR (DC = 0.7004, $p = 0.262$), Fig. 2b. The results hold with other densities used as reference.




Fig. 3. Subnetwork visualization. 11 subnetworks were extracted from graphs with a density of 0.2. (a) Well-established brain systems [13] (b) Two subnetwork formed by isolated nodes and false inclusion of premotor-related regions into auditory system was observed using global thresholding. (c) Local thresholding failed to detect one region of known visual systems and falsely detected four unrelated regions into dorsal default mode system. (d) Our strategy correctly detect most of the subnetworks found in [13].

Qualitatively, with GT (Fig. 3b), we observed two subnetworks comprising only isolated nodes in the left and right Pallidum (yellow and light grey node in the blue circle). We also observed that a region in the right premotor area was falsely grouped into the auditory subsystem (the light green region with a red arrow). With GTMR, two subnetworks comprising single nodes were found. As for LT (Fig. 3c), we observed the left and right insular cortex as well as the right Frontal Operculum Cortices (orange nodes with red arrows) were falsely grouped with Dorsal Default Mode regions and the left paracingulate gyrus was excluded. In contrast, our proposed strategy correctly identified known Dorsal Default Mode regions, such as paracingulate gyrus, anterior division of cingulate gyrus, and Accumbens, as a single subnetwork. Further, LT excluded the left Cuneal Cortex in the visual system (blue arrow in Fig. 3c). Other found subnetworks with our strategy, such as left and right executive control subnetworks (red and yellow), Fig. 3d, also conform well to known brain systems as was quantitatively demonstrated in Fig. 3a.

5 Conclusions

We proposed a modularity reinforcement strategy for improving brain subnetwork extraction. By applying local thresholding in combination with modularity reinforcement based on connection fingerprint similarity, we attained higher accuracy in subnetwork extraction compared to conventional global thresholding and local thresholding. Higher overlap with established brain systems and higher subnetwork reproducibility were also shown on the real data. Our results thus demonstrate clear benefits of refining conventional connectivity estimates with our strategy for subnetwork extraction. In fact, our strategy can be extended to applications beyond subnetwork extraction by deriving features based on the extracted subnetworks, e.g. within-subnetwork connectivity computed from the original connectivity estimates, and using those features for group analysis and behavioural association studies.

References

1. Fornito, A., Zalesky, A., Breakspear, M.: Graph analysis of the human connectome: promise, progress, and pitfalls. *Neuroimage* **80**, 426–444 (2013)
2. Achard, S., Coeurjolly, J.F., Marcillaud, R., Richiardi, J.: fMRI functional connectivity estimators robust to region size bias. In: *Statistical Signal Processing Workshop*, pp. 813–816. IEEE (2011)
3. Spisák, T., Jakab, A., Kis, S.A., Opposits, G., Aranyi, C., Berényi, E., Emri, M.: Voxel-wise motion artifacts in population-level whole-brain connectivity analysis of resting-state fMRI. *PLoS ONE* **9**(9), e104947 (2014)
4. Weiskopf, N., Hutton, C., Josephs, O., Turner, R., Deichmann, R.: Optimized EPI for fMRI studies of the orbitofrontal cortex: compensation of susceptibility-induced gradients in the readout direction. *Magn. Reson. Mater. Phys. Biol. Med.* **20**(1), 39–49 (2007)
5. Alexander-Bloch, A.F., Gogtay, N., Meunier, D., Birn, R., Clasen, L., Lalonde, F., Lenroot, R., Giedd, J., Bullmore, E.T.: Disrupted modularity and local connectivity of brain functional networks in childhood-onset schizophrenia. *Front. Syst. Neurosci.* **4**, 147 (2010)
6. Hammond, D.K., Vandergheynst, P., Gribonval, R.: Wavelets on graphs via spectral graph theory. *Appl. Comput. Harmonic Anal.* **30**(2), 129–150 (2011)
7. Fortunato, S.: Community detection in graphs. *Phys. Rep.* **486**(3), 75–174 (2010)
8. Niu, J., Fan, J., Stojmenovic, I.: JLMC: a clustering method based on Jordan-Form of Laplacian-Matrix. In: *Performance Computing and Communications Conference*, pp. 1–8. IEEE (2014)
9. Fortunato, S., Barthelemy, M.: Resolution limit in community detection. *Proc. Nat. Acad. Sci.* **104**(1), 36–41 (2007)
10. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al.: The WU-Minn human connectome project: an overview. *Neuroimage* **80**, 62–79 (2013)
11. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., et al.: The minimal preprocessing pipelines for the human connectome project. *Neuroimage* **80**, 105–124 (2013)
12. Behzadi, Y., Restom, K., Liao, J., Liu, T.T.: A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage* **37**(1), 90–101 (2007)
13. Shirer, W., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.: Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cereb. Cortex* **22**(1), 158–165 (2012)
14. Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, 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)
15. Skudlarski, P., Jagannathan, K., Calhoun, V.D., Hampson, M., Skudlarska, B.A., Pearlson, G.: Measuring brain connectivity: diffusion tensor imaging validates resting state temporal correlations. *Neuroimage* **43**(3), 554–561 (2008)
16. Munkres, J.: Algorithms for the assignment and transportation problems. *J. Soc. Ind. Appl. Math.* **5**(1), 32–38 (1957)