

Overlapping Replicator Dynamics for Functional Subnetwork Identification

Burak Yoldemir¹, Bernard Ng², and Rafeef Abugharbieh¹

¹ Biomedical Signal and Image Computing Lab, The University of British Columbia, Canada

² Parietal Team, INRIA Saclay, France

buraky@ece.ubc.ca

Abstract. Functional magnetic resonance imaging (fMRI) has been widely used for inferring brain regions that tend to work in tandem and grouping them into subnetworks. Despite that certain brain regions are known to interact with multiple subnetworks, few existing techniques support identification of subnetworks with overlaps. To address this limitation, we propose a novel approach based on replicator dynamics that facilitates detection of sparse overlapping subnetworks. We refer to our approach as overlapping replicator dynamics (RDOL). On synthetic data, we show that RDOL achieves higher accuracy in subnetwork identification than state-of-the-art methods. On real data, we demonstrate that RDOL is able to identify major functional hubs that are known to serve as communication channels between brain regions, in addition to detecting commonly observed functional subnetworks. Moreover, we illustrate that knowing the subnetwork overlaps enables inference of functional pathways, e.g. from primary sensory areas to the integration hubs.

Keywords: connectivity, fMRI, overlapping network detection, resting state.

1 Introduction

The interactions between specialized regions in the brain are known to play a key role in mediating brain function [1]. In recent years, functional magnetic resonance imaging (fMRI) has become the dominant means for studying this integrative property of the brain. One of the main approaches for gaining an understanding of the brain's organization is to group brain regions into functional subnetworks (also commonly referred to as “modules” or “communities”) based on the temporal dependencies between their fMRI observations. The most widely-used techniques for this include the seed-based approach and independent component analysis (ICA) [2]. In the seed-based approach, functional subnetworks are estimated based on the temporal correlations between a seed region and all other regions in the brain [2]. Despite its simplicity and intuitiveness, this approach has the drawback that results are restricted by the choice of seeds [3]. In contrast, ICA provides a data-driven means for extracting subnetworks by decomposing fMRI observations into maximally independent spatial components and associated time courses [4]. At the intra-subject level, brain regions assigned weights larger than a threshold within the same spatial component

are grouped as a subnetwork. However, the threshold is generally chosen in an ad hoc manner, which precludes statistical interpretation [4]. In practice, ICA is typically performed at the group level, in which a statistically meaningful threshold can be set [5]. Graph-theoretic techniques, such as normalized cut, that partition the brain into disjoint functional subnetworks have also been explored [6].

A major drawback of the aforementioned methods is that they cannot model how certain brain regions interact with multiple subnetworks [7]. These methods thus only give a partial depiction of the brain's functional architecture [7]. One way for identifying overlapping subnetworks is to use temporal ICA as opposed to spatial ICA. Though this approach has shown promising results, temporal ICA is statistically less robust than spatial ICA since the number of voxels is typically far greater than the number of time points [7]. A few methods based on graph theory have recently been proposed [8, 9], in which overlapping subnetworks can be found by analyzing the correlation structure of the fMRI observations. These methods, however, tend to be sensitive to parameters, such as the number of subnetworks to extract [8], the minimum clique size [8], the threshold used for binarizing the correlation matrix [8], and the weighting factor for controlling the size of the extracted subnetworks [9], which are nontrivial to set a priori, and data-driven means for their selection, e.g. through cross validation, can be computationally expensive.

In this paper, we propose a novel approach based on replicator dynamics (RD) [10] for identifying overlapping subnetworks. RD is a concept that originated in theoretical biology for modeling the survival of different species based on a fitness measure of their genotypes [10]. This concept has been previously adopted for detecting disjoint subnetworks with estimates of functional connectivity taken as the fitness measure [11]. We extend the RD formulation for detecting overlapping subnetworks, which we refer to as RDOL in the remainder of this paper. RDOL has a number of properties that are particularly advantageous for overlapping subnetwork extraction. First, RDOL can handle positively-weighted graphs, which mitigates information loss from binarization. Second, RDOL does not require setting the number of subnetworks to extract, which is generally unknown. Instead, a data-driven criterion is inherent in the RDOL formulation. Third, sparsity is implicitly imposed on the weight vectors representing the subnetwork membership of the brain regions with the level of sparsity automatically determined based on mutual correlation. This property circumvents the need for weight thresholding in contrast to ICA. We validate RDOL on synthetic data and apply it to fMRI data acquired during resting state (RS).

2 Methods

We start this section by briefly reviewing the RD formulation and its properties, followed by a discussion on how we extend RD to RDOL.

2.1 Replicator Dynamics

Let \mathbf{C} be a $d \times d$ matrix with each element reflecting the fitness of a genotype (a pair of alleles) and let $\mathbf{w}(k)$ be a $d \times 1$ vector with the j^{th} element, $w_j(k)$, being the probability

that allele j remains in the gene pool in the k^{th} generation. Assuming the survival of the alleles is governed by natural selection, $\mathbf{w}(k)$ can be iteratively estimated by [10]:

$$\mathbf{w}(k+1) = \frac{\mathbf{w}(k) \cdot * \mathbf{C}\mathbf{w}(k)}{\mathbf{w}(k)^T \mathbf{C}\mathbf{w}(k)}, \quad (1)$$

where $\cdot *$ denotes element-wise multiplication. Based on the fundamental theorem of natural selection [10], if \mathbf{C} is real, symmetric, and non-negative, $\mathbf{w}(k)^T \mathbf{C}\mathbf{w}(k)$ will strictly increase with increasing k until it converges to a strict local maximum of $\mathbf{w}^T \mathbf{C}\mathbf{w}$. The multiplicative form of the $\mathbf{w}(k)$ update restricts $\mathbf{w}(k)$ to be non-negative given a non-negative \mathbf{C} . Also, $\sum_j \mathbf{w}_j(k) = 1$, since if we denote the j^{th} element of $\mathbf{C}\mathbf{w}(k)$ as $(\mathbf{C}\mathbf{w}(k))_j$, the sum of elements in the numerator of (1), i.e. $\sum_j \mathbf{w}_j(k) (\mathbf{C}\mathbf{w}(k))_j = \mathbf{w}(k)^T \mathbf{C}\mathbf{w}(k)$, equals exactly the denominator. Thus, as discussed in [11], RD solves the non-negative sparse principal component analysis (PCA) optimization problem:

$$\max_{\mathbf{w}} \mathbf{w}^T \mathbf{C}\mathbf{w} \quad \text{s.t.} \quad \|\mathbf{w}\|_1 = 1, \mathbf{w} \geq 0, \quad (2)$$

with \mathbf{w} constrained to lie on the standard simplex, i.e. $\|\mathbf{w}\|_1 = 1, \mathbf{w} \geq 0$, which induces sparsity since restricting the l_1 norm encourages sparse \mathbf{w} estimates [12]. In addition, upon convergence, $\mathbf{w}_j(k+1) = \mathbf{w}_j(k)$ thus $(\mathbf{C}\mathbf{w}(k))_j = \mathbf{w}(k)^T \mathbf{C}\mathbf{w}(k)$ for all j . In PCA terminology with \mathbf{C} being a correlation matrix, all selected features (i.e. features with non-zero $\mathbf{w}_j(k)$) will have the same weighted average correlation with one another. Hence, the level of sparsity in $\mathbf{w}(k)$ is implicitly determined based on mutual correlations among the features. In the context of functional subnetwork identification, alleles correspond to brain regions of interest (ROIs) and each element of \mathbf{C} , denoted as \mathbf{C}_{ij} , corresponds to the functional connectivity estimate between ROIs i and j . We use the conventional Pearson's correlation between ROI time courses as the connectivity estimate [1]. To ensure \mathbf{C} is non-negative, as required for the properties of RD to hold, we null the negative elements of \mathbf{C} under the assumption that ROIs in the same subnetwork would presumably display positive correlations [13]. Also, we set \mathbf{C}_{ii} to 0 to avoid self-connections [11]. For initialization, we set all elements of $\mathbf{w}(0)$ to $1/d$ to mitigate bias [11]. We define convergence of (1) as $\mathbf{w}(k)^T \mathbf{C}\mathbf{w}(k)$ changing by $< 10^{-15}$.

2.2 Overlapping Subnetwork Extraction

Let S_1 be the set of nodes of the subnetwork found by applying RD on \mathbf{C} with the corresponding weight vector denoted as \mathbf{w}_{S_1} . Let S_1' be the remaining nodes that are not in S_1 . To identify the other subnetworks in the graph encoded by \mathbf{C} , we propose adopting a strategy based on a simple augmentation of \mathbf{C} [14] that renders \mathbf{w}_{S_1} unstable as a solution of (2). The intuition is that (2) is non-convex with multiple local maxima, each of which corresponds to a subnetwork [14]. By destabilizing \mathbf{w}_{S_1} as a solution of (2), re-application of RD would identify the subnetwork, S_2 , that has the next highest mutual correlation. To destabilize \mathbf{w}_{S_1} as an RD solution, we first create

an artificial node, v , and extend (directed) edges from v to all nodes in S_1 with edge weights chosen such that v and all the nodes of S_1 are mutually correlated. This augmentation trick replaces S_1 by $\{v, S_1\}$ as a solution of (2). If we then extend (directed) edges from all nodes of S_1' to v with edge weights selected such that v is more correlated to the nodes of S_1' than of S_1 , it can be shown that $\{v, S_1\}$ would become unstable as a solution under RD [14]. With S_1 no longer a solution and $\{v, S_1\}$ unstable, S_2 would be detected next, which could potentially overlap with S_1 since no orthogonality constraints are imposed between subnetworks. In practice, creating an artificial node is equivalent to appending \mathbf{C} with a new row and a new column. We denote this augmented \mathbf{C} by \mathbf{C}^{aug} . Extending weighted edges from/to the artificial node is accomplished by setting the elements of the last row/column of \mathbf{C}^{aug} to the edge weights. Weights that can provably destabilize \mathbf{w}_{S_1} as a solution of (2) are given by [14]:

$$\mathbf{C}_{ij}^{aug} = \begin{cases} \mathbf{C}_{ij} & \text{if } i, j \leq d \\ \alpha & \text{if } j > d \text{ and } i \notin S_{j-d} \\ \beta & \text{if } i, j > d \text{ and } i = j \\ \frac{1}{|S_{i-d}|} \sum_{k \in S_{i-d}} \mathbf{C}_{kj} & \text{if } i > d \text{ and } j \in S_{i-d} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where $\alpha > \beta$, $\beta = \max_{i,j} \mathbf{C}_{ij}$ for $i \neq j$, and $| \cdot |$ denotes cardinality. S_2 can thus be found by applying (1) to \mathbf{C}^{aug} , and other subnetworks can be extracted by iterating the same matrix augmentation procedure for each newly-discovered subnetwork and applying (1). We terminate the subnetwork extraction process if $\mathbf{w}(k)^T \mathbf{C} \mathbf{w}(k)$ after convergence is $\leq \mathbf{w}(0)^T \mathbf{C} \mathbf{w}(0)$ as this suggests that no further solutions of (2) can be found. Note that (1) is guaranteed to converge to a local maximum even for an asymmetric \mathbf{C} [14].

3 Materials

Synthetic Data. We generated 500 synthetic datasets consisting of 85 ROIs with 75 ROIs divided among 3 subnetworks and the remaining 10 ROIs used for modeling noise. Each dataset comprised 42 scans of 210 time points as in the real data. ROIs 1 to 40 were defined as the primary network, ROIs 34 to 63 formed the secondary network, and ROIs 31 to 36 as well as ROIs 62 to 75 were set as the tertiary network. The three resulting subnetworks composed of 40, 30 and 20 ROIs, respectively, in which 3 ROIs were in all 3 subnetworks and 9 ROIs belonged to 2 subnetworks. A binary 85×85 correlation matrix, Σ , matching the described network configuration was created for generating ROI time courses by drawing samples from $N(0, \Sigma)$. Gaussian noise was added to the time courses for a range of signal-to-noise ratio (SNR) between 0.1 to 2 dB. The noisy time courses across scans were temporally concatenated with their sample correlation matrix, Σ_s , used as input to RDOL. To further investigate the performance of RDOL under more idealistic conditions, we simulated another 500 synthetic datasets with 200 scans each using the above procedure, in which substantially more accurate sample correlation estimates could be obtained.

Real Data. The multimodal Kirby 21 dataset was obtained from the NITRC database¹, which comprises 21 healthy subjects with two RS-fMRI scans of 7 min acquired for each subject. Acquisition details can be found in [15]. Preprocessing steps included motion correction, spatial normalization, bandpass filtering at 0.01 and 0.1 Hz, and removal of white matter and cerebrospinal fluid confounds from gray matter voxel time courses. We divided the brain into 500 parcels by temporally concatenating voxel time courses across scans and applying Ward clustering [16]. Each parcel was taken as an ROI. Voxel time courses within each ROI were averaged to generate ROI time courses. We concatenated the ROI time courses of both scans of all subjects and computed the sample correlation matrix, which was used for group subnetwork identification. Concatenating the scans ensures that there are enough samples to reasonably estimate a 500×500 correlation matrix.

4 Results and Discussion

Synthetic Data. We compared RDOL against NC [6] and RD [11] applied to Σ_s and temporal ICA [7] applied to time courses drawn from $N(0, \Sigma)$. We set the number of subnetworks to extract to the ground truth value of 3 for all methods. Performance was assessed using the average Dice coefficient (DC) over the 500 synthetic datasets:

$$DC = \frac{2|S_{GT} \cap S|}{|S_{GT}| + |S|}, \quad (4)$$

where S_{GT} is the ground truth subnetwork, S is the estimated subnetwork, and $|\cdot|$ denotes cardinality. As shown in Fig. 1, RDOL consistently outperformed RD, temporal ICA, and NC for all noise levels tested. Increasing the number of scans improved the performance of RDOL and RD considerably (Fig. 1(b)), but only slightly for temporal ICA. NC performed essentially the same for both conditions. It is worth noting that in the absence of noise, RDOL identified all subnetworks perfectly, whereas the average DC values for RD, temporal ICA, and NC were 0.87, 0.61, and 0.57, respectively.




Fig. 1. Synthetic data results. (a) 42 scans as in real data. (b) Idealistic case with 200 scans.

¹ The Kirby 21 dataset is available online at: <http://www.nitrc.org/projects/multimodal/>




Fig. 2. Exemplar RS subnetworks detected by RDOL. (a) Sensorimotor network. (b) Default-mode network (split into subnetworks similar to recent studies [17, 18]). (c) Auditory network.

Real Data. Exemplar subnetworks found by RDOL are shown in Fig. 2 (each color corresponds to a distinct subnetwork). Detected functional subnetworks include the sensorimotor network, the default-mode network, and the auditory network among other commonly observed subnetworks. Similar to recent findings [17, 18], the default-mode network was divided into subnetworks. The average voxel overlap over all overlapping subnetwork pairs was 16%.

To gain further insights from the detected subnetworks, we performed an analysis similar to that in [19] for inferring different functional pathways. The underlying idea is that for two subnetworks to communicate, there must be certain brain areas that serve as relays, i.e. spatial overlaps between subnetworks. Drawing on this idea, the analysis proceeds by first selecting a seed region and finding the set of subnetworks, N_1 , that overlaps with the seed. We then find the set of subnetworks, N_2 , that overlaps with N_1 and repeat the process until no overlap remains. With the well-studied visual system as an example (Fig. 3), we placed a seed in the primary visual cortex (V1), and observed that information from V1 would pass through V2, propagate to the higher visual areas (through both the ventral and dorsal stream), and then reach the precuneus and posterior cingulate cortex, which are regions known for sensory integration [20]. This trajectory matches exactly the well known pathways of the visual system. Progressing from these hubs, information would spread to the medial prefrontal cortex and other sensory areas, but in fact, prior neuroscience knowledge tells us that the direction of information flow is normally the reverse for sensory regions. Thus, it is important to acknowledge that the presented analysis can only infer the path for information transfer, but not the directionality. We performed the same analysis on the subnetworks found by temporal ICA for comparison, since temporal ICA also permits identification of overlapping subnetworks. The number of subnetworks to extract was set to the value implicitly determined by RDOL. The obtained functional pathway remained similar to that of RDOL, but regions such as precentral gyrus, thalamus, and brain stem were incorrectly included during the first two steps. These spurious connections resulted in artifactual streams of information during subsequent steps.

Furthermore, we identified the prominent functional hubs based on the overlaps between the subnetworks extracted by RDOL. Detected functional hubs include the posterior cingulate cortex, precuneus, lateral temporal cortex, lateral parietal cortex, and medial/lateral prefrontal cortex, which conform to previous findings based on

graph analysis on node degree [21]. In addition, parts of the lateral occipital cortex were also identified as functional hubs, which are known for facilitating communication between the primary and higher visual areas.




Fig. 3. Functional pathway inference with a seed placed in the primary visual cortex. (a) Blue = seed, red and green = steps 1 and 2. (b) Violet, yellow, cyan = steps 3 to 5. (c) Orange, purple, dark green = steps 6 to 8. (d) Pink, light blue, crimson = steps 9 to 11.

5 Conclusions

We proposed a novel approach for identifying sparse subnetworks from fMRI data. In contrast to most existing subnetwork identification methods, RDOL allows for the identified subnetworks to overlap, which provides a more realistic portrayal of the brain's functional architecture. On synthetic data, we demonstrated that RDOL outperforms all other examined methods, including temporal ICA, NC, and RD. Applying RDOL to real RS-fMRI data detected group subnetworks that matched well with prior neuroscience knowledge. Furthermore, the identified subnetworks overlapped in brain regions that are known to serve as functional hubs. RDOL thus enables the untangling of the overlapping structure of brain subnetworks, which may provide insights that are pivotal in understanding human brain organization.

References

1. Friston, K.: Functional and Effective Connectivity: A Review. *Brain Connectivity* 1(1), 13–36 (2011)
2. Li, K., Guo, L., Nie, J., Li, G., Liu, T.: Review of Methods for Functional Brain Connectivity Detection using fMRI. *Comput. Med. Imaging Graph.* 33, 131–139 (2009)
3. Ma, L., Wang, B., Chen, X., Xiong, J.: Detecting Functional Connectivity in the Resting Brain: A Comparison Between ICA and CCA. *Magn. Reson. Imaging* 25, 47–56 (2007)

4. McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J.: Analysis of fMRI Data by Blind Separation into Independent Spatial Components. *Hum. Brain Mapp.* 6, 160–188 (1998)
5. Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J.: A Method for Making Group Inferences from Functional MRI Data Using Independent Component Analysis. *Hum. Brain Map.* 14(3), 140–151 (2001)
6. Van Den Heuvel, M., Mandl, R., Hulshoff Pol, H.: Normalized Cut Group Clustering of Resting-State fMRI Data. *PLoS ONE* 3, e2001 (2008)
7. Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., Van Essen, D.C., Feinberg, D.A., Yacoub, E.S., Ugurbil, K.: Temporally-independent Functional Modes of Spontaneous Brain Activity. *PNAS* 109(8), 3131–3136 (2012)
8. Wu, K., Taki, Y., Sato, K., Sassa, Y., Inoue, K., Goto, R., Okada, K., Kawashima, R., He, Y., Evans, A.C., Fukuda, H.: The Overlapping Community Structure of Structural Brain Network in Young Healthy Individuals. *PLoS ONE* 6(5), e19608 (2011)
9. Yan, X., Kelley, S., Goldberg, M., Biswal, B.B.: Detecting Overlapped Functional Clusters in Resting State fMRI with Connected Iterative Scan: A Graph Theory Based Clustering Algorithm. *J. Neurosci. Methods* 199(1), 108–118 (2011)
10. Schuster, P., Sigmund, K.: Replicator Dynamics. *J. Theor. Biol.* 100, 533–538 (1983)
11. Ng, B., Abugharbieh, R., McKeown, M.J.: Group Replicator Dynamics: A Novel Group-wise Evolutionary Approach for Sparse Brain Network Detection. *IEEE Trans. Med. Imag.* 31(3), 576–585 (2012)
12. Tibshirani, R.: Regression Shrinkage and Selection via the LASSO. *J. Royal Stat. Soc. Series B* 58, 267–288 (1996)
13. Lohmann, G., Bohn, S.: Using Replicator Dynamics for Analyzing fMRI Data of the Human Brain. *IEEE Trans. Med. Imag.* 21(5), 485–492 (2002)
14. Torsello, A., Bulò, S.R., Pelillo, M.: Beyond Partitions: Allowing Overlapping Groups in Pairwise Clustering. In: *IEEE Int. Conf. Patt. Recog.*, pp. 1–4 (2008)
15. Landman, B.A., Huang, A.J., Gifford, A., Vikram, D.S., Lim, I.A.L., Farrell, J.A.D., Bogovic, J.A., Hua, J., Chen, M., Jarso, S., Smith, S.A., Joel, S., Mori, S., Pekar, J.J., Barker, P.B., Prince, J.L., van Zijl, P.C.M.: Multi-Parametric Neuroimaging Reproducibility: A 3T Resource Study. *NeuroImage* 54, 2854–2866 (2011)
16. Thyreau, B., Thirion, B., Flandin, G., Poline, J.-B.: Anatomic-Functional Description of the Brain: A Probabilistic Approach. In: *IEEE Int. Conf. Acoustics, Speech and Signal Proc.*, vol. 5, pp. 14–19 (2006)
17. Buckner, R.L., Andrews-Hanna, J.R., Schacter, D.L.: The Brain’s Default Network: Anatomy, Function, and Relevance to Disease. *Ann. N. Y. Acad. Sci.* 1124, 1–38 (2008)
18. Zuo, X.N., Kelly, C., Adelman, J.S., Klein, D.F., Castellanos, F.X., Milham, M.P.: Reliable Intrinsic Connectivity Networks: Test–retest Evaluation Using ICA and Dual Regression Approach. *NeuroImage* 49, 2163–2177 (2010a)
19. Sepulcre, J., Sabuncu, M.R., Yeo, B.T., Liu, H., Johnson, K.A.: Stepwise Connectivity of the Modal Cortex Reveals the Multimodal Organization of the Human Brain. *J. Neurosci.* 32(31), 10649–10661 (2012)
20. Cavanna, A.E., Trimble, M.R.: The Precuneus: A Review of its Functional Anatomy and Behavioural Correlates. *Brain* 129(3), 564–583 (2006)
21. Buckner, R.L., Sepulcre, J., Talukdar, T., Krienen, F.M., Liu, H., Hedden, T., Andrews-Hanna, J.R., Sperling, R.A., Johnson, K.A.: Cortical Hubs Revealed by Intrinsic Functional Connectivity: Mapping, Assessment of Stability, and Relation to Alzheimer’s Disease. *J. Neurosci.* 29(6), 1860–1873 (2009)