Stable Overlapping Replicator Dynamics for Brain Community Detection

Burak Yoldemir, Bernard Ng, and Rafeef Abugharbieh, *Senior Member, IEEE*

***Abstract*—A fundamental means for understanding the brain's organizational structure is to group its spatially disparate regions into functional subnetworks based on their interactions. Most community detection techniques are designed for generating partitions, but certain brain regions are known to interact with multiple subnetworks. Thus, the brain's underlying subnetworks necessarily overlap. In this paper, we propose a technique for identifying overlapping subnetworks from weighted graphs with statistical control over false node inclusion. Our technique improves upon the replicator dynamics formulation by incorporating a graph augmentation strategy to enable subnetwork overlaps, and a graph incrementation scheme for merging subnetworks that might be falsely split by replicator dynamics due to its stringent mutual similarity criterion in defining subnetworks. To statistically control for inclusion of false nodes into the detected subnetworks, we further present a procedure for integrating stability selection into our subnetwork identification technique. We refer to the resulting technique as stable overlapping replicator dynamics (SORD). Our experiments on synthetic data show significantly higher accuracy in subnetwork identification with SORD than several state-of-the-art techniques. We also demonstrate higher test-retest reliability in multiple network measures on the Human Connectome Project data. Further, we illustrate that SORD enables identification of neuroanatomically-meaningful subnetworks and network hubs.**

***Index Terms*—Brain, connectivity, fMRI, overlapping community detection, replicator dynamics, stability selection.**

I. Introduction

The brain is known to consist of spatially disparate regions that are functionally connected as subnetworks [1]. Unraveling this modular structure of the brain is a challenging endeavor. The prevalent neuroimaging modality for tackling this challenge is functional magnetic resonance imaging (fMRI). To identify subnetworks from fMRI data, seed-based correlation [2] and independent component analysis (ICA) [3] are primarily used. More recently, graph theoretical approaches are gaining popularity as they facilitate compact brain representation and enable the reduction of complex brain interactions into intuitive summary network measures. Under this formalism, the brain is modeled as a graph, where brain regions and their connections correspond to nodes and edges, respectively. Particularly relevant for investigating the modular structure of the brain are the numerous techniques for detecting subnetworks (also commonly referred to as communities or modules) in graphs [4], such as the Louvain method [5], the fast greedy modularity maximization algorithm (ModMax) [6], and normalized cuts (NC) [7]. A major drawback of these graph partitioning techniques in our context is that each node can only be assigned to one subnetwork, which contradicts with how a brain region might interact with different parts of the brain at different times [8]. Also, these techniques force all nodes to be assigned to a subnetwork, which prohibits the possibility that some nodes simply do not belong to any subnetwork. For instance, a region containing mostly white matter due to segmentation errors should not be assigned to a subnetwork of grey matter regions. Thus, to veraciously characterize the modular structure of the brain, developing techniques that permit subnetwork overlaps without forced assignments is essential.

Recently, a number of techniques have been put forth for identifying overlapping brain subnetworks. In [9], the technique involves first performing maximal clique detection and treating each maximal clique as a node of a new graph. Subnetworks are then extracted by applying modularity optimization [5] on this new graph. Since the maximal cliques might share nodes, overlaps between subnetworks are enabled. The main shortcoming of this technique is that it requires graph edges to be binarized since it cannot handle weighted graphs, but brain connectivity lies on a continuum. Edge binarization could thus result in important information being discarded. A technique adopted from social network analysis called connected iterative scan (CIS) has also been explored for brain subnetwork identification [10]. Taking each node as a candidate subnetwork, CIS determines if other nodes belong to this candidate subnetwork based on their influence on a graph density metric. A limitation of CIS is that it is sensitive to a weighting factor that controls subnetwork size. Another technique for finding overlapping brain subnetworks is temporal ICA (tICA), which is more suited for finding overlaps than the more widely-used spatial ICA (sICA), since maximizing spatial independence tends to result in little overlaps [8]. However, the proposed implementation uses an

Manuscript received December 30, 2014; revised July 23, 2015 and September 13, 2015; accepted September 19, 2015.

B. Yoldemir is with the Biomedical Signal and Image Computing Lab, University of British Columbia, Vancouver, Canada (e-mail: buraky@ece.ubc.ca).

B. Ng is with the Functional Imaging in Neuropsychiatric Disorders Lab, Stanford University, USA and the Parietal Team, INRIA, France (e-mail: bernardyng@gmail.com).

R. Abugharbieh is with the Biomedical Signal and Image Computing Lab, University of British Columbia, Vancouver, Canada (e-mail: rafeef@ece.ubc.ca).

ad hoc threshold for extracting interacting nodes from the spatial maps, which limits the statistical interpretability of the results. Grouping the edges between nodes instead of the nodes themselves has also been proposed for detecting subnetwork overlaps [11]. In that technique, edges are grouped using e.g. hierarchical clustering with the set of nodes associated with any edge of a given edge cluster declared as a subnetwork [12]. The main limitation of this technique is scalability since in weighted graphs where all nodes are connected by an edge, the number of pairwise relationships between edges grows quartically with the number of nodes. Multiple relational embedding has also been adopted for brain subnetwork identification [13]. This method embeds the adjacency matrices from multiple task paradigms into a single space to unravel the joint connectivity structure. As such, this method is not applicable for datasets from single task studies, such as resting state fMRI (RS-fMRI) experiments. A related strategy is to find brain areas associated with different clinical indices and look for overlaps [14]. The limitation is that the type of subnetworks that can be identified is restricted by the clinical indices available. Several other overlapping community detection techniques have been proposed in the general context, a review of which can be found in [15].

A serious limitation in most community detection techniques is their sole focus on optimizing a criterion, such as modularity, without considering the variability of the subnetwork assignment under data perturbation or what subnetworks one would get from random graphs to assess statistical significance. In the area of brain activation detection, rigorous statistics has long been employed to control for false positives [16]. It is known that even random networks can have high modularity [17], thus statistically controlling for false node inclusion in subnetwork identification is crucial. A notable exception is the order statistics local optimization method (OSLOM) [18]. OSLOM uses a statistical criterion based on a null model derived from random graphs for assessing the significance of the identified subnetworks, and is shown to outperform many state-of-the-art community detection techniques [15].

In this paper, we propose a subnetwork identification technique that: (i) can operate on weighted graphs, (ii) allows for subnetwork overlaps, (iii) has an intrinsic criterion for setting the number of subnetworks, (iv) does not force all nodes to be assigned to a subnetwork, and (v) statistically controls for false node inclusion in the detected subnetworks.[1] Our technique extends the replicator dynamics (RD) formulation [19] (Section II-A) by incorporating a graph augmentation strategy [20] (Section II-B) to enable overlaps between subnetworks and a graph incrementation scheme [21] (Section II-C) to facilitate merging of subnetwork components that are falsely split by RD due to its strict mutual similarity criterion [22]. The graph incrementation scheme involves setting a parameter that governs the sparsity level of the identified subnetworks. If set incorrectly, over-merging might occur, which would result in inclusion of false nodes in addition to those arising from noise. To jointly deal with over-merging and noise, we present a procedure for integrating stability selection [23] into RD (Section II-C). The underlying idea is to declare a node as part of a subnetwork only if its probability of being selected over bootstraps for various sparsity levels is higher than a theoretical threshold that provably controls for the expected number of false nodes. The basis of stability selection is that consistent selection of the same false nodes over a large number of bootstraps is statistically unlikely. Importantly, stability selection is insensitive to the choice of sparsity levels. Thus, incorporating stability selection enables removal of false nodes arising from noise and over-merging in a data-driven fashion. Integrating stability selection into subnetwork identification is generally nontrivial due to the need for establishing subnetwork correspondence across bootstraps. This problem is especially challenging when the subnetworks are allowed to overlap and are estimated from noisy graphs. We resolve this problem by exploiting an intrinsic property of RD, namely that the initialization of the RD process governs the subnetwork to which RD converges. We refer to our proposed technique as stable overlapping replicator dynamics (SORD). Flowchart of the overall algorithm is presented in Fig. 1. To evaluate SORD, we compare it against a number of state-of-the-art techniques using synthetic data comprising a wide array of network configurations. We also assess its test-retest reliability on real data from the Human Connectome Project (HCP) [24].

[1] Preliminary versions of this work appeared in [20, 21]. We have substantially extended our previous work by: 1) expanding the literature survey, 2) thoroughly discussing parameter selection, 3) comparing SORD against eight existing methods as opposed to four, 4) comparing the computational complexity of the contrasted methods, 5) discussing bootstrapping procedures for subject and group-level, 6) presenting a proof of the graph augmentation strategy, and 7) performing test-retest reliability analysis at subject and group-level.

II. Methods

A. Replicator Dynamics: An Overview

RD is a concept that originated from theoretical biology for modeling the evolution of interacting and self-replicating entities [19]. We adopt RD in this work since it groups graph nodes based on their mutual similarity as explained below. This property makes RD a suitable model for brain subnetwork identification, where the aim is to identify groups of brain regions that closely interact with each other.

Let $\mathbf{w}(k) = (\mathbf{w}_1(k), \ldots, \mathbf{w}_n(k))^T$ be a $n\times 1$ vector, where $\mathbf{w}_j(k)$ is the proportion of allele that are of type j in the gene pool during generation k, and n is the number of alleles. Further, let $\mathbf{C}$ be a $n\times n$ matrix with each element, $\mathbf{C}_{ij}$, reflecting the fitness of a genotype, i.e. a pair of alleles i and j. Under the assumption of natural selection, $\mathbf{w}(k+1)$ is given by the replicator equation [19]:

$$\mathbf{w}(k+1) = \frac{\mathbf{w}(k) \circ \mathbf{C}\mathbf{w}(k)}{\mathbf{w}(k)^T \mathbf{C}\mathbf{w}(k)}, \tag{1}$$

where $\circ$ denotes element-wise multiplication. Eq. (1) is shown to solve the following non-convex optimization problem [25]:

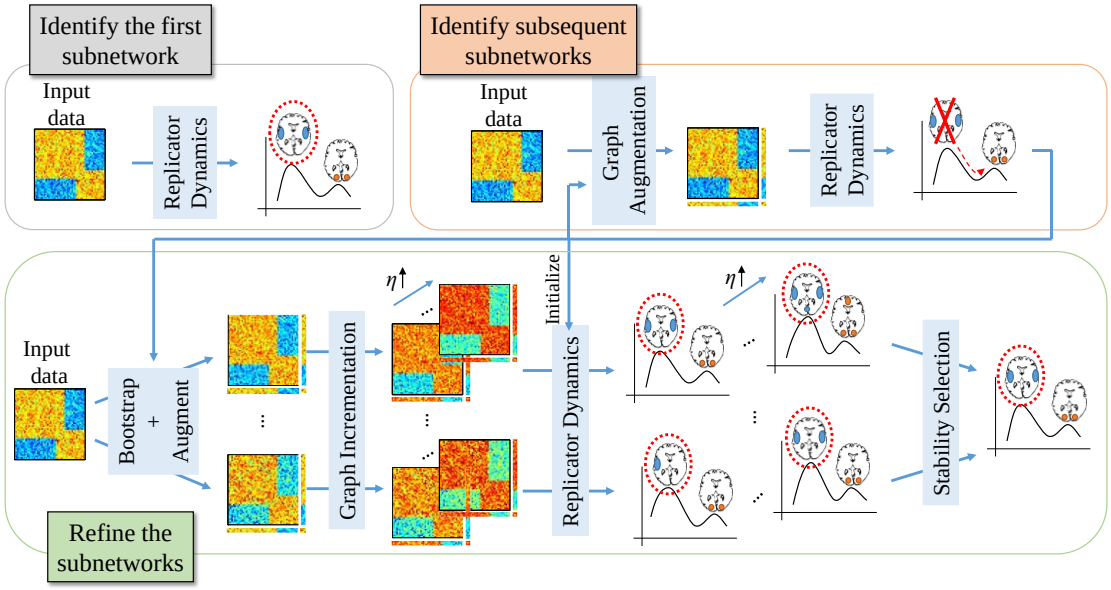

Fig. 1. Flowchart of the proposed algorithm. The first subnetwork is identified using RD (Section II-A). Subsequent subnetworks are found by augmenting the graph, which destabilizes previously found solutions, and reapplying RD (Section II-B). Due to RD's stringent mutual similarity criterion, resulting subnetworks are often split into multiple components. To refine these initial subnetworks, graph incrementation and stability selection are performed (Section II-C) to adjust the sparsity of the subnetworks and statistically control for false node inclusion in the subnetworks, respectively.

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

The fundamental theorem of natural selection states that if $\mathbf{C}$ is real-valued, symmetric, and non-negative, $\mathbf{w}(k)^T\mathbf{C}\mathbf{w}(k)$ strictly increases with k along any nonstationary trajectory $\mathbf{w}(k)$ until it converges to a strict local maximum as (1) is iterated [26]. In the more general case of any real-valued $\mathbf{C}$, local maxima of $\mathbf{w}(k)^T\mathbf{C}\mathbf{w}(k)$ are guaranteed to attract nearby trajectories $\mathbf{w}(k)$ asymptotically as (1) is iterated [29]. $\mathbf{w}_j(k+1)$ will equal $\mathbf{w}_j(k)$ upon convergence thus $(\mathbf{C}\mathbf{w}(k))_j = \mathbf{w}(k)^T\mathbf{C}\mathbf{w}(k)$ for all j, where $(\cdot)_j$ denotes the j^{th} element of the corresponding vector. In graph-theoretic terms with $\mathbf{C}$ being a similarity matrix, this implies that all selected nodes (i.e. nodes with non-zero $\mathbf{w}(k)$) will have the same weighted average correlation with one another. Hence, the grouping of the nodes is based on mutual similarity among the nodes. With $\mathbf{w}$ constrained to lie on the standard simplex, i.e. $\|\mathbf{w}\|_1 = 1$, $\mathbf{w} \geq 0$, sparse $\mathbf{w}$ is encouraged since restricting the l_1 norm induces sparsity [27]. $\mathbf{w}_j(k) > 0$ hence indicates that allele j persisted in generation k. The local maximum to which RD converges highly depends on $\mathbf{w}(0)$. Since the vicinity at which the local maxima reside is unknown a priori, an unbiased way of setting $\mathbf{w}(0)$ is to assign all of its elements to $1/n$, which tends to find the community with the highest mutual similarity among nodes [25]. Other local maxima can be found using an extension of RD given in Section II-B.

In the context of brain subnetwork identification, alleles correspond to brain regions of interest (ROIs), and $\mathbf{C}_{ij}$ corresponds to the functional connectivity between ROIs i and j. Nonzero elements in $\mathbf{w}$ indicate the corresponding ROIs form a subnetwork. We estimate functional connectivity using the conventional Pearson's correlation between ROI time courses. To ensure $\mathbf{C}$ is non-negative, as required for the properties of RD to hold, one can either take the absolute value of $\mathbf{C}$ or ignore negative $\mathbf{C}_{ij}$ [28]. Since the interpretation of negative functional connectivity is currently unclear, we set negative $\mathbf{C}_{ij}$ to zero. Also, we set $\mathbf{C}_{ii}$ to zero since Pearson's correlation does not provide estimates of self-connections.

B. Overlapping Replicator Dynamics

Since the replicator equation is deterministic, reapplying it with the same initialization will converge to the same local maximum. To find the other local maxima, one might modify $\mathbf{C}$ by removing the nodes [28] or the edges [25] in the identified subnetworks and reapply RD, but these "peeling strategies" would alter the graph structure, hence affecting the location of the remaining local maxima. Instead, we adopt a more principled strategy based on graph augmentation that permits spatial overlaps among subnetworks and preserves the local maxima of the original graph [20]. Reapplication of RD on this augmented graph would converge to one of the remaining local maxima. We refer to this extension of RD as overlapping replicator dynamics (ORD). The required graph augmentation is as follows [29]:

$$\mathbf{C}_{ij}^{aug} = \begin{cases} \mathbf{C}_{ij} & \text{if } i, j \leq n \\ \alpha & \text{if } j > n \text{ and } i \notin S_{j-n} \\ \beta & \text{if } i, j > n \text{ and } i = j \\ \gamma_{ij} & \text{if } i > n \text{ and } j \in S_{i-n} \\ 0 & \text{otherwise} \end{cases}, \tag{3}$$

$$\gamma_{ij} = \sigma_{ij} \frac{1}{|S_{i-n}|} \sum_{m \in S_{i-n}} \mathbf{C}_{mj}, \quad \sigma_{ij} > 1,$$

where $\alpha > \beta$, $\beta = \max_{i \neq j} \mathbf{C}_{ij}$, S_l is the set of nodes of the l^{th} identified subnetwork, and $|\cdot|$ denotes cardinality. In effect, this is equivalent to adding artificial nodes to the graph by appending $\mathbf{C}$ with new rows and new columns, and extending

weighted edges to the original nodes such that the previously-found local maxima provably become unstable. α and γ are mathematically engineered to ensure that: (i) selecting the artificial nodes is more favorable than the previously extracted subnetworks based on the optimization criterion of the replicator equation, (ii) the artificial nodes are not stable solutions, and (iii) there exist other solutions that are more favorable than the artificial nodes (see Appendix for proof). Given S_1, finding S_2 proceeds by adding a new row and a new column to $\mathbf{C}$ based on (3) and reapplying RD. S_3 is then found by adding a new row and a new column to the $\mathbf{C}^{aug}$ that is used for identifying S_2 and reapplying RD, and the remaining S_l can be similarly extracted. Note that (1) is guaranteed to converge to a local maximum even for an asymmetric $\mathbf{C}$ [29]. We note that an intrinsic criterion for terminating further subnetwork extraction is to stop if $\mathbf{w}(k)^{\mathrm{T}}\mathbf{C}\mathbf{w}(k) \leq \mathbf{w}(0)^{\mathrm{T}}\mathbf{C}\mathbf{w}(0)$ after convergence, since this suggests no further solutions of (2) are present. In practice, we empirically found on synthetic data that $\mathbf{w}(k)^{\mathrm{T}}\mathbf{C}\mathbf{w}(k) \leq a\cdot\mathbf{w}(0)^{\mathrm{T}}\mathbf{C}\mathbf{w}(0)$, a>1, is more robust to noisy $\mathbf{C}$. A rigorous way of choosing a is described in Section II-D.

C. Stable Overlapping Replicator Dynamics

Graph incrementation scheme. With mutual similarity as the criterion for grouping nodes, applying RD on a noisy $\mathbf{C}$ could easily result in a subnetwork being falsely split into components [22], since even small perturbations to $\mathbf{C}$ would render certain nodes non-mutually connected. To reduce this effect, one strategy is to add a constant, η, to the off-diagonal elements of $\mathbf{C}$. Using an extreme example for intuition, if we add 10,000 to a $\mathbf{C}$ with values ranging from 0 to 1, the original differences between elements of $\mathbf{C}$ would be negligible. Applying RD to such a $\mathbf{C}$ would result in all nodes being assigned to a single subnetwork, i.e. non-zero values for all elements of $\mathbf{w}$. If we instead use a η of magnitude similar to values in $\mathbf{C}$, only the smaller differences in $\mathbf{C}$ would be negligible. Thereby, we can adjust the sparsity level of $\mathbf{w}$ by changing η, as illustrated in Fig. 2 for an exemplar synthetic dataset (Section III-A), where the percentage of nodes included in the first subnetwork identified by ORD is plotted with respect to η. The dashed red line indicates the percentage of nodes in the ground truth subnetwork. The original ORD solution (η = 0) has much fewer nodes than the ground truth subnetwork, illustrating the false subnetwork splitting problem of RD. As η is increased, the number of nodes in the identified subnetwork gradually increases with all nodes declared as a single subnetwork at extreme values of η. Choosing the optimal η is non-trivial. Furthermore, noise in $\mathbf{C}$ could result in false nodes being included in the subnetworks. To jointly deal with these two problems, we propose a procedure for integrating stability selection into ORD.

Stability Selection. The idea behind stability selection is that if we bootstrap the data many times and perform subnetwork identification on each bootstrap sample, the nodes that truly belong to the same subnetwork are likely to be jointly selected over a large fraction of bootstrap samples, whereas false nodes are unlikely to be persistently selected [23]. To incorporate stability selection in refining each subnetwork identified by ORD, we first generate a set of matrices $\{\mathbf{C}_\eta\}$ by adding η to the off-diagonal elements of $\mathbf{C}$ for a range of η from 0 to $n\cdot\max_{ij}|\mathbf{C}_{ij}|$ at a step size of κ, where κ is empirically chosen small enough to ensure that no more than one node is added to a subnetwork at each increment of η. We note that the exact value of κ used has no effect on the results as long as it is substantially smaller than the original values in $\mathbf{C}$. Specifically, we have experimented with values ranging from $0.01\cdot c_{\max}$ to $0.05\cdot c_{\max}$, where $c_{\max}$ is the maximum value in $\mathbf{C}$, and observed no difference in results. We then apply RD to each $\mathbf{C}_\eta$ with $\mathbf{w}(0)$ set to the RD solution of the previous η increment (with $1/n$ added to all elements and renormalized to sum to 1), which induces RD to converge to a local maximum in the vicinity of the previous η increment for retaining subnetwork correspondence. Note that adding $1/n$ is critical since if $\mathbf{w}_j(0) = 0$, it cannot become nonzero as (1) is iterated, which prohibits new nodes from being added. We discard subnetworks that contain >10% of the nodes assuming a subnetwork would not span >10% of the brain.

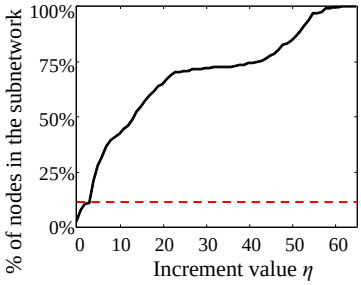

Fig. 2. Graph incrementation: Percentage of nodes included in the first subnetwork identified by ORD on an exemplar synthetic dataset. Artificially increasing the similarity of the nodes enables the sparsity level of the subnetwork to be adjusted.

Next, let $\mathbf{X}$ be a $t\times n$ fMRI time series matrix, where t is the number of time samples. We generate 100 bootstrap samples of $\mathbf{X}$, each denoted as $\mathbf{X}^b$, by sampling with replacement. The spatiotemporal correlations in fMRI time series [30] warrant special care during bootstrapping. For extracting group subnetworks, we randomly select S subjects from the pool of S subjects with replacement. Time series of all nodes are then temporally concatenated across the randomly drawn subjects so that the spatiotemporal structure in each subject's time series matrix is retained. For extracting subject-specific subnetworks, we employ parametric bootstrapping [31]. Given a subject's time series matrix $\mathbf{X}$, we randomly draw bootstrap samples with the same size as $\mathbf{X}$ from a multivariate normal distribution with zero mean and covariance estimated from $\mathbf{X}$. Given $\mathbf{X}^b$, we compute $\{\mathbf{C}^b{}_\eta\}$ and apply RD to each $\mathbf{C}^b{}_\eta$ to estimate $\mathbf{w}^b{}_\eta$. For each η, we estimate the selection probability of each node as the proportion of the bootstrap samples in which the given node attains a nonzero weight [23]. Finally, we threshold the selection probabilities to identify nodes that belong to a given subnetwork that we are refining. A threshold τ that bounds the expected number of nodes falsely included in a subnetwork, E, is given by [23]:

$$\tau \leq \left(\frac{q^2}{En}+1\right)\Big/2, \tag{4}$$

where q is the average number of nodes per subnetwork, which can be estimated as the sum of selection probabilities of

all nodes averaged over η. We set E to 1 so that the average number of false nodes included in a subnetwork is statistically controlled to be less than or equal to one. We declare the set of nodes with selection probabilities higher than the resulting τ for any η to be a subnetwork [23]. Note that each subnetwork identified by ORD is refined separately using this procedure.

D. Parameter Selection

Except for the upper bound on the percentage of nodes that a subnetwork is allowed to include and the scaling factor, a, which is used to define the threshold to terminate subnetwork identification, SORD has no other free parameters in theory. However, the choice of α and σ in (3) can affect the results in practice. Theoretical lower bounds of α and σ are known ($\alpha > \beta$ and $\sigma > 1$), and any value above these lower bounds should result in identical subnetworks. However, we empirically observed slight differences in the identified subnetworks for different α and σ. To jointly select α and σ, we generate 500 synthetic datasets over a wide assortment of network configurations (Section III-A) and use the mean Omega index [15] to quantify the agreement between the estimated and ground truth subnetworks over the 500 synthetic datasets for finding the optimal α-σ pair. Our strategy is thus to try to exhaust the possible networks that are typically encountered and use the α-σ pair that provides the best performance overall. Let N_{est} and N_{gnd} be the number of estimated and ground truth subnetworks, respectively. Omega index is defined as $\omega = (\omega_u - \omega_e)/(1 - \omega_e)$, where [15]:

$$\omega_u = \frac{1}{M} \sum_{s=0}^{\max(N_{est}, N_{gnd})} |\, t_s^{est} \bigcap t_s^{gnd} \,|, \tag{5}$$

$$\omega_e = \frac{1}{M^2} \sum_{s=0}^{\max(N_{est}, N_{gnd})} |\, t_s^{est} \,| \cdot |\, t_s^{gnd} \,|, \tag{6}$$

where M equals to $n(n\text{-}1)/2$ represents the number of node pairs, and t_s^{est} is the set of node pairs that appear in exactly s of the estimated subnetworks. t_s^{gnd} is defined similarly for the ground truth subnetworks. ω_e is the expected Omega index under the null, and ω_u is the unadjusted Omega index. The adjustment on ω_u accounts for the agreements in node pairs resulting from chance alone. Larger ω values indicate better matching between the two sets of subnetworks, and ω attains a value of one when the two sets perfectly match. A plot of the Omega index averaged over 500 synthetic datasets for different α-σ pairs is shown in Fig. 3. The values of α and σ that provide the best overall performance are found to be 1000 and 50, respectively, but any α and σ proximal to these values provide similar Omega index. We emphasize that the synthetic datasets used for optimizing the choice of α and σ are separate from the synthetic datasets used for comparing different subnetwork identification techniques (Section IV-A). Using separate datasets for parameter selection and technique assessment circumvents introducing bias to our results.

An open problem in subnetwork identification is the choice of the number of subnetworks. Although $\mathbf{w}(k)^T\mathbf{C}\mathbf{w}(k) \leq \mathbf{w}(0)^T\mathbf{C}\mathbf{w}(0)$ is a theoretically justified criterion, noisy $\mathbf{C}$ could lead to spurious subnetworks with $\mathbf{w}(k)^T\mathbf{C}\mathbf{w}(k)$ slightly above $\mathbf{w}(0)^T\mathbf{C}\mathbf{w}(0)$, hence the need for a stricter threshold, $a \cdot \mathbf{w}(0)^T\mathbf{C}\mathbf{w}(0)$, $a \geq 1$. Since fewer time samples (e.g. subject-level versus group-level analysis) would result in noisier $\mathbf{C}$, the optimal a would vary with the number of time samples. To account for this effect, we generate synthetic datasets using different numbers of time points, N_t, and identify subnetworks over a range of a from 1 to 10. 500 synthetic datasets are generated for each N_t with the value of a that minimizes $|N_{est} - N_{gnd}|$ declared as the optimal for each N_t (Fig. 4). Again, the synthetic datasets used for the selection of a are separate from those used for evaluating the performance of SORD.

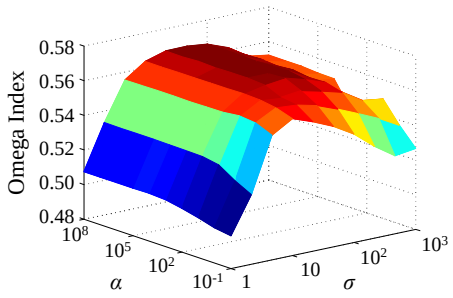

Fig. 3. Omega index averaged over 500 synthetic datasets. Values proximal to the optimal provide similar Omega index. Thus, performance is robust to a broad range of α and σ.

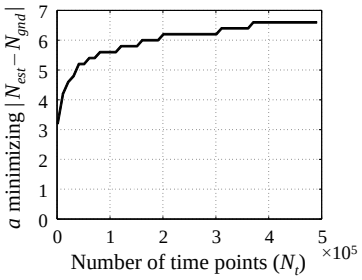

Fig. 4. Scaling factor, a, used for terminating further subnetwork extraction.

The upper bound on the percentage of nodes that a subnetwork is allowed to include is analogous to the lower bound on the sparsity parameter that needs to be selected in l_1 regularization techniques [27]. We tested SORD's sensitivity to this parameter by identifying group-level subnetworks from real data using an upper limit of 20% instead of 10%, which led to less than 2% of the nodes being additionally included in the resulting subnetworks. Hence, SORD is only marginally sensitive to the choice of this parameter.

III. Materials

A. Synthetic Data

We simulated 500 synthetic datasets spanning a wide variety of network configurations to assess the performance of SORD. A separate 500 datasets were created to jointly set α and σ, and another 500 datasets for each N_t in Fig. 4 to set a. Each dataset consisted of n = 200 ROIs, with an associated time series of 4800 time points per ROI as in the real data of each subject. For each dataset, the number of subnetworks, N, was randomly selected between 10 and 20. The number of

ROIs in each subnetwork was set to $\lceil n/N \rceil + c$ with c being a random number between 1 and 5, which ensures that the total number of ROIs across subnetworks is greater than n, hence guarantees the presence of subnetwork overlaps. ROIs within each subnetwork were randomly drawn from the pool of n ROIs and were allowed to repeat across subnetworks. We opted to used random network configurations since the ground truth number of subnetworks, number of nodes within each subnetwork, and amount of subnetwork overlap are unknown. Our strategy is thus to try to exhaust the possibilities. For each network configuration, we created a corresponding $n \times n$ binary adjacency matrix $\mathbf{\Sigma}$. Using this binary matrix, $\mathbf{\Sigma}$, we create a positive definite approximation, $\mathbf{\Sigma}^{pd}$, by setting the non-positive eigenvalues of $\mathbf{\Sigma}$ to 10^{-16}. Simulated time series are then generated by drawing random samples from a multivariate normal distribution with zero mean and covariance matrix $\mathbf{\Sigma}^{pd}$. Gaussian noise was then added to the time series with SNR randomly set between -6 to -3 dB, corresponding to the typical levels seen in task-based fMRI data, i.e. between 0.2 and 0.5 [32, 33] (SNR of RS-fMRI data is hard to determine).

B. Real Data

Minimally preprocessed RS-fMRI data from 40 unrelated healthy subjects were obtained from the Human Connectome Project (HCP) database [24]. The acquisition involved two sessions per subject, each having two scans of 1200 time frames acquired over a period of 14 min and 33 s with the following parameters: TR=720 ms, TE=33.1 ms, FA=52°, 2×2×2 mm voxels, FOV=208×180 mm. The preprocessing steps already applied on the data included gradient distortion correction, motion correction, spatial normalization, intensity normalization, and bias field correction [34]. We additionally regressed out motion artifacts, white matter and cerebrospinal fluid confounds, and principal components of high variance voxels found using CompCor [35]. Band-pass filtering (0.01 Hz – 0.1 Hz) was further applied to reduce low-frequency drifts and high-frequency respiratory and cardiac noise. To define the ROIs, we functionally divided the brain into n = 200 parcels by temporally concatenating voxel time series across the 21 subjects (two scans each) of the multimodal Kirby dataset [36], and applying spatially-constrained, hierarchical Ward clustering [37]. We opted to use an independent dataset for functional parcellation to avoid introducing any potential bias by "seeing" our data prior to analysis. The value of n used is the minimum recommended in [38], which provides ROIs with higher functional homogeneity than typical anatomical atlases, while maintaining reasonable number of ROIs to time samples ratio for reliable connectivity estimation. ROI time series were then generated by averaging voxel time series within each ROI, with the mean removed and normalized by the standard deviation. For group-level analysis, ROI time series were concatenated across subjects prior to estimating $\mathbf{C}$. For subject-level analysis, $\mathbf{C}$ was directly estimated from each subject's ROI time series.

IV. Results and Discussion

A. Synthetic Data Experiments

We compared SORD against ORD [20], Girvan's technique based on edge betweenness (EB) [39], temporal ICA (tICA) [8], ModMax [6], NC [7], Louvain method [5], OSLOM [18], and InfoMap [40]. Among these techniques, SORD, ORD, tICA, OSLOM, and InfoMap allow subnetwork overlaps, while the others can only find disjoint subnetworks. When an intrinsic criterion to determine the number of subnetworks is not available, as in NC and tICA, we used maximum modularity as the termination criterion [41]. tICA was chosen over spatial ICA since it enables more subnetwork overlap [8].

We generated 500 synthetic datasets and assessed the performance of each contrasted technique by computing the Omega index between the estimated and ground truth subnetworks. InfoMap collapsed all nodes into a single subnetwork for every dataset, thus was not subject to further analysis and excluded from the results. We observed that InfoMap starts finding multiple subnetworks only when the SNR is above 8 dB, which is significantly higher than the typical SNR levels observed in fMRI studies. Overall, SORD outperformed all other techniques tested (Fig. 5). The pairwise difference between the Omega indices of SORD and every other technique tested was found to be statistically significant at $p<10^{-10}$ based on Wilcoxon signed rank test.

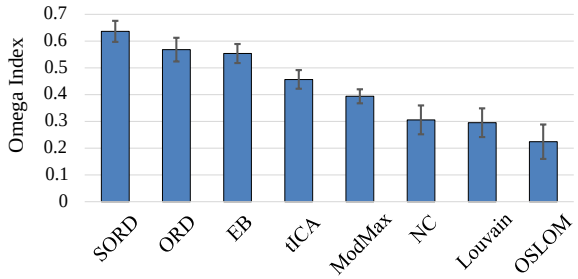

Fig. 5. Subnetwork identification accuracy assessed by average Omega Index over 500 synthetic datasets.

To more directly demonstrate that SORD alleviates the false subnetwork splitting problem in ORD, we also computed average Dice coefficient (DC) of each technique over the 500 synthetic datasets. DC is based on the number of true positive nodes within each subnetwork, hence better reflects whether nodes belonging to the same subnetwork are jointly selected. Specifically, DC is defined as: $DC=2|S_{est} \cap S_{gnd}|/(S_{est}+S_{gnd})$, where S_{est} is the set of nodes in estimated subnetwork, and S_{gnd} is the set of nodes in the ground truth subnetwork matched to S_{est} using Hungarian clustering [42]. DC of unmatched subnetworks was set to zero. As shown in Fig. 6, SORD achieved the best performance, with a significantly higher DC over each contrasted technique at $p<10^{-10}$. Our results in Figs. 5 and 6 thus show that SORD is robust to various network configurations and SNR levels.

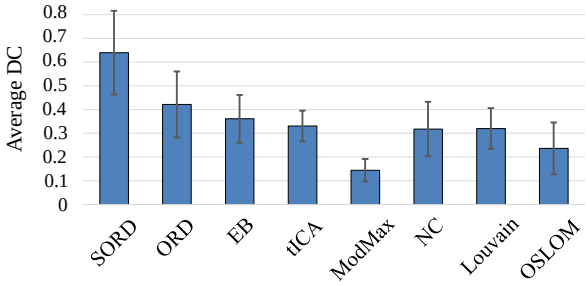

Fig. 6. Subnetwork identification accuracy assessed by average DC over 500 synthetic datasets.

To evaluate the accuracy of the detected overlapping nodes, we used F-score [15]:

$$F = \frac{2 \cdot precision \cdot recall}{precision + recall}, \tag{7}$$

where *precision* is the number of correctly detected overlapping nodes divided by the total number of detected overlapping nodes, and *recall* is the number of correctly detected overlapping nodes divided by the true number of overlapping nodes. SORD achieved the highest F-score among the techniques that allow subnetwork overlaps (Fig. 7). The pairwise difference between the F-score of SORD and each contrasted technique was statistically significant at $p<10^{-10}$.

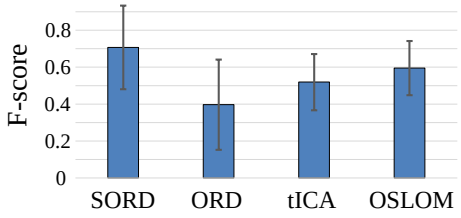

Fig. 7. Overlapping node accuracy assessed by F-score over 500 synthetic datasets.

A caveat with F-score is that it does not reflect whether the subnetwork memberships of the overlapping nodes are accurate. For example, having a node assigned to two subnetworks or all detected subnetworks does not change the F-score. Therefore, we additionally computed the DC between the estimated and the true set of subnetworks in which each overlapping node resides, denoted as DC_{ol}. Averaged over the 500 synthetic datasets, DC_{ol} of SORD was significantly higher than all other tested techniques at $p<10^{-10}$ (Fig. 8).

Across all subnetworks identified by SORD, the average and median number of false nodes were found to be 1.19 and 0, respectively, which approximately matches the bound we set for the expected number of false nodes, i.e. *E* was set to 1. In terms of estimating the true number of subnetworks, SORD was on par with EB and tICA and outperformed all the other contrasted methods as demonstrated in Table SI of the Supplementary Materials.

B. Real Data Experiments

Since ground truth subnetworks are unknown for real data, we quantitatively evaluated the contrasted techniques based on test-retest reliability as estimated using the intra-class

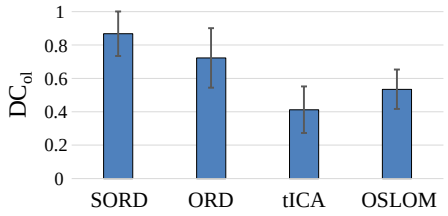

Fig. 8. Accuracy of the subnetwork memberships of the overlapping nodes assessed by average DC_{ol} over 500 synthetic datasets.

correlation (ICC). For graph characterization, we focused on clustering coefficient, transitivity, and global efficiency, which are arguably some of the most widely-used metrics [43]. Since these metrics do not depend on the learned subnetwork structure, we estimated these metrics for each identified subnetwork instead of on the whole graph, so that the learned subnetwork structure was taken into account. Specifically, we first computed these metrics for each subnetwork from data of each scan session of a given subject. We then averaged the values of each metric across the identified subnetworks, and computed the ICC value for each graph metric. The number of identified subnetworks might be different across scans and subjects. Hence, we normally would need to draw a one-to-one correspondence between the subnetworks to compare them, a problem which we bypassed by taking the subnetwork average for each metric. The implicit assumption here is that by capturing local graph properties through estimating multiple metrics for each identified subnetwork and combining the values of each metric over subnetworks, the resulting summary metrics in aggregate would reflect the learned subnetwork structure. There exist different definitions of ICC. Here, we used the two-way ANOVA model, with random subject effects and fixed session effects, which is defined as ICC(3,1) = (*BMS* - *EMS*)/(*BMS* + *EMS*) [44], where *BMS* is between-subject mean squares and *EMS* is error mean squares. As shown in Fig. 9, SORD exhibited superior test-retest reliability overall. Notably, EB resulted in negative ICC values, which indicates that within-subject variation in summary metrics of the identified subnetworks is greater than between-subject variation. We note that the trend remains similar when the other ICC definitions [44] are used.

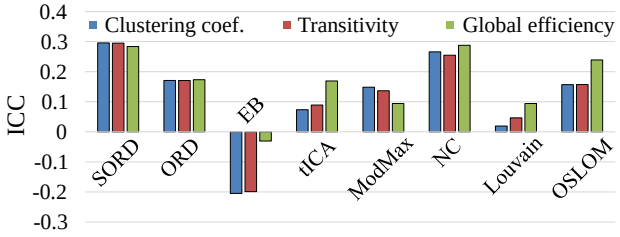

Fig. 9. Test-retest reliability on real data as assessed by ICC. Three graph measures were used to summarize the identified subnetworks.

To assess group-wise test-retest reliability, we compared the group-level subnetworks extracted separately from the two scan sessions of RS-fMRI data using the Omega index without extracting summary statistics from the subnetworks. The subnetworks identified from the first session's data are taken as the "ground truth", against which the subnetworks from the

second session are compared. SORD is found to be on par with ModMax in terms of group-wise test-retest reliability, and it outperforms all other methods tested (Fig. 10). The number of subnetworks identified using each contrasted method at the subject and group-level are given in Tables SII and SIII of the Supplementary Materials. Note that we did not observe a trend between the number of subnetworks identified and test-retest reliability.

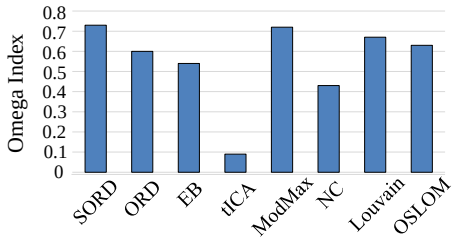

Fig. 10. Group-level test-retest reliability on real data as assessed by the Omega index.

Qualitatively, SORD found all the widely-observed subnetworks reported in the literature [45]. The extracted auditory subnetwork, posterior default mode subnetwork, and right executive control subnetwork are shown in Fig. 11 as exemplar results. To illustrate the concept of overlap, two overlapping visual subnetworks identified by SORD are shown in Fig. 12. These two subnetworks correspond to the primary visual and extra-striate visual subnetworks, which have been widely observed in previous studies [45]. The overlap between the two subnetworks (yellow) likely reflects the information propagation path from the primary visual areas to the higher visual areas. The subnetworks detected in this study exhibit substantial similarities with the subnetworks reported in Yeo et al. [46]. Specifically, primary visual, extra-striate visual and auditory subnetworks SORD detected correspond to subnetworks #1, #5 and #4 in [46], respectively. In contrast to [46], the default mode subnetwork identified by SORD did not include the medial prefrontal cortex as was the case in [47]. SORD also identified two lateralized executive control subnetworks, which were reported as a single subnetwork in [46]. To demonstrate that there exist fundamental differences in the topology of subnetworks identified using different methods, the primary visual and auditory subnetworks detected by all contrasted methods are shown in Figs. S1 and S2 of the Supplementary Materials. We note that the identified group subnetworks well reflect the commonalities across subjects, but there are definitely some individual differences that are not captured by the group subnetwork (Figs. S3 and S4 of the Supplementary Materials).

The dorsal attention subnetwork (DAS) was also among the subnetworks identified by SORD. DAS is involved in the executive control of attention, and there has been conflicting reports in the literature as to whether DAS comprises extrastriate visual areas [48], [49]. Interestingly, SORD identified two subnetworks that contain key regions in DAS, where the extrastriate visual areas were included in one but not the other (Fig. 13). Our results thus support the notion that DAS can be manifested in different forms depending on the subjects' attentional states [50]. Using a subnetwork identification technique that does not permit subnetwork overlaps would only find one of these two subnetworks, which might explain the diverging results in the literature.

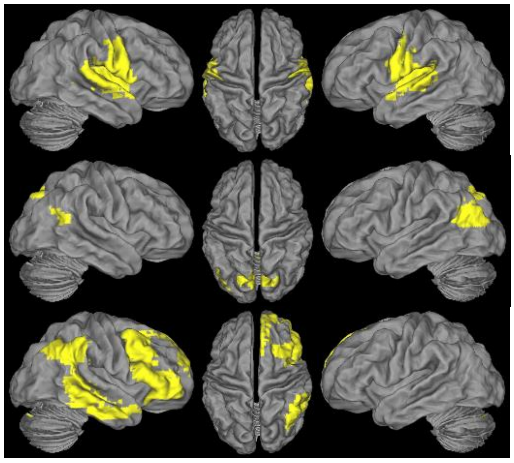

Fig. 11. Exemplar group-level subnetworks identified by SORD. Top to bottom: auditory subnetwork, posterior default mode subnetwork, right executive control subnetwork. Left to right: right lateral view, top view, left lateral view. Visualization was carried out using Caret software [52].

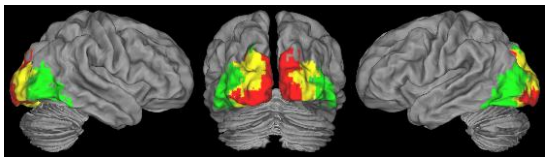

Fig. 12. Primary (red) and extra-striate visual (green) subnetworks identified by SORD. Yellow represents subnetwork overlaps. Left to right: right lateral view, posterior view, left lateral view.

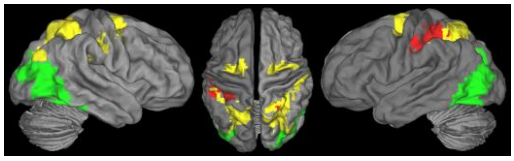

Fig. 13. Two subnetworks identified by SORD that correspond to DAS. The subnetworks are shown in red and green. Yellow represents subnetwork overlaps. Left to right: right lateral view, top view, left lateral view.

Subnetwork overlaps were largely found within precentral gyrus, medial superior frontal cortex, superior parietal lobule, precuneous, posterior cingulate, superior/middle temporal gyrus, lateral occipital cortex, occipital pole and frontal orbital cortex. These regions match well with functional hubs previously identified by graph-theoretical analysis based on the degree of the voxels [51]. We also computed the betweenness centrality of each node to further evaluate the detected overlapping regions. The betweenness centrality of these regions was found to be significantly higher than the rest of the nodes at $p<10^{-4}$ based on Wilcoxon rank sum test. This finding confirms that the overlapping regions detected by SORD indeed participate in a large number of shortest paths, facilitating functional integration of the brain.

Given that HCP dataset is of much higher quality compared to typical datasets, it is not clear whether the evaluations presented in this paper will hold when applied to typical data. To this end, we further assessed the performance of the contrasted methods on lower-quality RS-fMRI data. Specifically, we temporally subsampled the first 6 min of each

scan in the HCP dataset by a factor of 3, resulting in a TR of 2.16 s. Additionally, we created 20 synthetic datasets of 180 time points each, simulating 6 min scans with a TR of 2 s. Subnetwork identification accuracy on synthetic data was assessed using Dice coefficient and Omega index, and test-retest reliability on real data was evaluated using ICC. As expected, performance of all tested methods degraded with reduced number of time points. However, the trend remained similar with SORD outperforming the contrasted methods (Figs. S3-S5 in the Supplementary Materials).

C. Computational Complexity Comparison

Worst-case complexity of all community detection methods tested is presented in Table I. Eq. (1) only involves vector-matrix multiplications, hence the complexity of RD is $O(n^2)$ for each iteration, where n is the number of nodes. Due to the augmentation of one artificial node for each of the N subnetworks extracted, the complexity of ORD is $O((n+N)^2)$. With the addition of n_η increments of η in the graph incrementation scheme and n_b bootstrap samples in stability selection, the complexity of SORD is $O(n_b n_\eta (n+N)^2)$. In Table I, m is the number of edges, and d is the depth of the dendrogram describing the community structure. We note that $m=O(n^2)$ and $d=O(\log n)$ in most of the weighted graphs [6,15].

TABLE I
WORST-CASE COMPUTATIONAL COMPLEXITY OF THE METHODS TESTED

SORD	$O(n_b n_\eta (n+N)^2)$
ORD [20]	$O((n+N)^2)$
EB [39]	$O(m^2 n)$
tICA [8]	$O(n^3)$
ModMax [6]	$O(md \log n)$
NC [7]	$O(n^{3/2})$
Louvain [5]	$O(m)$
OSLOM [18]	$O(n^2)$

V. CONCLUSION

We proposed SORD for the identification of overlapping subnetworks from weighted graphs with statistical control over false node inclusion. On synthetic datasets spanning a wide variety of network configurations, we demonstrated that SORD outperforms multiple state-of-the-art techniques including ORD, Girvan's technique based on edge betweenness, tICA, ModMax, NC, Louvain method, OSLOM, and InfoMap. When applied to real fMRI data, SORD exhibited superior test-retest reliability over the contrasted techniques. Moreover, SORD detected resting state subnetworks and functional hubs that match well with prior neuroscience knowledge, and provided new insights for the diverging definitions of DAS in the literature.

APPENDIX: PROOF OF (3)

We leverage the mathematical foundation of evolutionary game theory (EGT) to prove (3). In EGT, $\mathbf{C}$ is called a *payoff matrix*, and each allele is viewed as a *pure strategy* available to individuals playing a two-person game. A combination of pure strategies, $\mathbf{w} = (\mathbf{w}_1,\ldots,\mathbf{w}_n)^T$, is called a *mixed strategy*. The *support* of a mixed strategy $\mathbf{w}$, denoted by $\Omega_\mathbf{w}$, is the set of pure strategies with nonzero probability. The *best replies* against a mixed strategy $\mathbf{w}$, represented by $\varphi(\mathbf{w})$, is the set of strategies that maximize the expected payoff against $\mathbf{w}$ [53]:

$$\varphi(\mathbf{w}) = \arg\max_{\mathbf{z}} \mathbf{z}^T \mathbf{C}\mathbf{w}. \tag{A1}$$

A mixed strategy $\mathbf{w}$ is a *Nash equilibrium* if $\mathbf{w}$ is a best reply to itself [53]. If we set $\mathbf{w}(0)$ inside the standard simplex, (1) converges on a Nash equilibrium [54]. To prevent RD from converging on a previously identified solution of (1), $\mathbf{w}_s$, we augment $\mathbf{C}$ to introduce a new strategy i such that $\mathbf{w}_i^T \mathbf{C}^{aug} \mathbf{w}_s > \mathbf{w}_s^T \mathbf{C}^{aug} \mathbf{w}_s$, where $\mathbf{w}_i$ denotes playing the pure strategy i. $\mathbf{w}_s$ is thus no longer a Nash equilibrium since it is not the best reply to itself. To see this, consider the expected payoff of pure strategy i against $\mathbf{w}_s$:

$$\mathbf{w}_i^T \mathbf{C}^{aug} \mathbf{w}_s = (\mathbf{C}^{aug} \mathbf{w}_s)_i > \frac{1}{|\Omega_{\mathbf{w}_s}|} \sum_{k\in\Omega_{\mathbf{w}_s}} \sum_{j\in\Omega_{\mathbf{w}_s}} \mathbf{C}_{kj}^{aug} (\mathbf{w}_s)_j, \tag{A2}$$

for $\mathbf{C}^{aug}$ given in (3), where $(\cdot)_i$ denotes the i^{th} element of a given vector. For all $k \in \Omega_\mathbf{w}$:

$$\sum_{j\in\Omega_{\mathbf{w}_s}} \mathbf{C}_{kj}^{aug} (\mathbf{w}_s)_j = (\mathbf{C}^{aug} \mathbf{w}_s)_j = \mathbf{w}_s^T C^{aug} \mathbf{w}_s, \tag{A3}$$

where the last equality follows from the fact that all pure strategies in the support of a Nash equilibrium yield the same payoff [53]. Combining (A2) and (A3) gives:

$$(\mathbf{C}^{aug} \mathbf{w}_s)_i > \frac{1}{|\Omega_{\mathbf{w}_s}|} \sum_{k\in\Omega_{\mathbf{w}_s}} (\mathbf{C}^{aug} \mathbf{w}_s)_k = \mathbf{w}_s^T \mathbf{C}^{aug} \mathbf{w}_s, \tag{A4}$$

which implies the expected payoff of pure strategy i against $\mathbf{w}_s$ is greater than that of $\mathbf{w}_s$ against itself. $\mathbf{w}_s$ is hence no longer a Nash equilibrium, which ensures that it will not be identified by reapplication of RD. RD cannot converge on $\mathbf{w}_i$ either, since it is not a Nash equilibrium: $\mathbf{w}_j^T \mathbf{C}^{aug} \mathbf{w}_i = \alpha$ is greater than $\mathbf{w}_i^T \mathbf{C}^{aug} \mathbf{w}_i = \beta$ for any pure strategy $j \notin \Omega_\mathbf{w}$.

REFERENCES

[1] O. Sporns, *Networks of the Brain*. Cambridge, MA: MIT Press, 2010.

[2] B. Biswal, F.Z. Yetkin, V.M. Haughton, and J.S. Hyde, "Functional connectivity in the motor cortex of resting human brain using echo-planar MRI," *Magnet. Reson. Med.*, vol. 34, no. 4, pp. 537–541, 1995.

[3] M.J. McKeown, S. Makeig, G.G. Brown, T.-P. Jung, S.S. Kindermann, A.J. Bell, and T.J. Sejnowski, "Analysis of fMRI data by blind separation into independent spatial components," *Hum. Brain Mapp.*, vol. 6, pp. 160–188, 1998.

[4] S. Fortunato, "Community detection in graphs," *Phys. Rep.*, vol. 486, no. 3, pp. 75–174, 2010.

[5] V.D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *J. Stat. Mech: Theory Exp.*, vol. 2008, no. 10, P10008, 2008.

[6] A. Clauset, M.E.J. Newman, and C. Moore, "Finding community structure in very large networks," *Phys. Rev. E*, vol. 70, no. 6, 066111, 2004.

[7] M. Van Den Heuvel, R. Mandl, and H.H. Pol, "Normalized cut group clustering of resting-state fMRI data," *PLoS ONE*, vol. 3, no. 4, e2001, 2008.

[8] S.M. Smith, K.L. Miller, S. Moeller, J. Xu, E.J. Auerbach, M.W. Woolrich, C.F. Beckmann, M. Jenkinson, J. Andersson, M.F. Glasser, D.C. Van Essen, D.A. Feinberg, E.S. Yacoub, and K. Ugurbil, "Temporally-independent functional modes of spontaneous brain

activity," *Proc. Nat. Acad. Sci. USA*, vol. 109, no. 8, pp. 3131–3136, 2012.

[9] K. Wu, Y. Taki, K. Sato, Y. Sassa, K. Inoue, R. Goto, K. Okada, R. Kawashima, Y. He, A.C. Evans, and H. Fukuda, "The overlapping community structure of structural brain network in young healthy individuals," *PLoS ONE*, vol. 6, no. 5, e19608, 2011.

[10] X. Yan, S. Kelley, M. Goldberg, and B.B. Biswal, "Detecting overlapped functional clusters in resting state fMRI with Connected Iterative Scan: A graph theory based clustering algorithm," *J. NeuroSci. Meth.*, vol. 199, no. 1, pp. 108–118, 2011.

[11] T. Madhyastha, Y. Cao, S. Sujitnapisatham, G. Presnyakov, W.A. Chaovalitwongse, and T. Grabowski, "Link clustering to explore brain dynamics using resting state functional MRI," *J. Radiol. Radiat. Ther.*, vol. 1, no. 2, pp. 1012:1–1012:3, 2013.

[12] Y.-Y. Ahn, J.P. Bagrow, and S. Lehmann, "Link communities reveal multiscale complexity in networks," *Nature*, vol. 466, no. 7307, pp. 761–764, 2010.

[13] K.H. Nenning, V. Schoepf, D. Prayer, and G. Langs. "Overlapping functional networks in multiple fMRI paradigms," in *NIPS Workshop on Machine Learning and Interpretation in Neuroimaging*, 2012.

[14] I.I. Galynker, Z.S. Yaseen, C. Katz, X. Zhang, G.Jennings-Donovan, S. Dashnaw, J. Hirsch, H. Mayberg, L.J. Cohen, and A. Winston, "Distinct but overlapping neural networks subserve depression and insecure attachment," *Soc. Cogn. Affect. Neurosci.*, vol. 7, no. 8, pp. 896-908, 2012.

[15] J. Xie, S. Kelley, and B.K. Szymanski, "Overlapping community detection in networks: The state-of-the-art and comparative study," *ACM Comput. Surv.*, vol. 45, no. 4, 43:1–43:37, 2013.

[16] T. Nichols, S. Hayasaka, "Controlling the familywise error rate in functional neuroimaging: A comparative review," *Stat. Methods Med. Research*, vol. 12, no. 5, pp. 419–446, 2003.

[17] R. Guimera, M. Sales-Pardo, and L.A.N. Amaral, "Modularity from fluctuations in random graphs and complex networks," *Phys. Rev. E*, vol. 70, no. 2, 025101, 2004.

[18] A. Lancichinetti, F. Radicchi, J.J. Ramasco, and S. Fortunato, "Finding statistically significant communities in networks," *PLoS ONE*, vol. 6, no. 4, e18961, 2011.

[19] P. Schuster and K. Sigmund, "Replicator dynamics," *J. Theoretical Biol.*, vol. 100, pp. 533–538, 1983.

[20] B. Yoldemir, B. Ng, and R. Abugharbieh, "Overlapping replicator dynamics for functional subnetwork identification," in *Proc. Med. Image Computing Computer Assist. Intervent.*, 2013, pp. 682–689.

[21] B. Ng, B. Yoldemir, and R. Abugharbieh, "Stable overlapping replicator dynamics for subnetwork identification." in *NIPS Workshop on Networks: From Graphs to Rich Data*, 2014.

[22] B. Thirion, P. Pinel, A. Tucholka, A. Roche, P. Ciuciu, J.-F. Mangin, and J.B. Poline, "Structural analysis of fMRI data revisited: Improving the sensitivity and reliability of fMRI group studies," *IEEE Trans. Med. Imaging*, vol. 26, no. 9, pp. 1256–1269, 2007.

[23] N. Meinshausen and P. Bühlmann, "Stability selection," *J. Roy. Statist. Soc. Ser. B*, vol. 72, no. 4, pp. 417–473, 2010.

[24] D.C. Van Essen, S.M. Smith, D.M. Barch, T.E.J. Behrens, E. Yacoub, and K. Ugurbil, "The WU-Minn Human Connectome Project: An overview," *NeuroImage*, vol. 80, pp. 62–79, 2013.

[25] B. Ng, M.J. McKeown, and R. Abugharbieh, R., "Group replicator dynamics: A novel group-wise evolutionary approach for sparse brain network detection," *IEEE Trans. Med. Imaging*, vol. 31, no. 3, pp. 576–585, 2012.

[26] J. Hofbauer and K. Sigmund, *The Theory of Evolution and Dynamical Systems*. Cambridge, U.K.: Cambridge Univ. Press, 1988.

[27] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Royal. Statist. Soc. B.*, vol. 58, no. 1, pp 267–288, 1996.

[28] G. Lohmann and S. Bohn, "Using replicator dynamics for analyzing fMRI data of human brain," *IEEE Trans. Med. Imaging*, vol. 21, no. 5, pp. 485–492, 2002.

[29] A. Torsello, S.R. Bulò, and M. Pelillo, "Beyond partitions: Allowing overlapping groups in pairwise clustering," in *Proc. Int. Conf. Pattern Recogn*, 2008, pp. 1–4.

[30] E. Bullmore, B. Horwitz, G. Honey, M. Brammer, S. Williams, and T. Sharma, "How good is good enough in path analysis of fMRI data?," *NeuroImage*, vol. 11, no. 4, pp. 289–301, 2000.

[31] B. Efron, *The jackknife, the bootstrap, and other resampling plans*. Society of Industrial and Applied Mathematics CBMS-NSF Monographs, 1982.

[32] G.E. Sarty, *Computing Brain Activity Maps from fMRI Time-series Images*. New York: Cambridge University Press, 2007.

[33] V. Calhoun. Preprocessing III & IV: Spatial Normalization and Physiological Artifacts. Available: http://ece-research.unm.edu/vcalhoun/courses/fMRI_Spring10/

[34] M.F. Glasser, S.N. Sotiropoulos, J.A. Wilson, T.S. Coalson, B. Fischl, J.L. Andersson, J. Xu, S. Jbabdi, M. Webster, J.R. Polimeni, D.C. Van Essen, and M. Jenkinson, "The minimal preprocessing pipelines for the Human Connectome Project," *NeuroImage*, vol. 80, pp. 105–124, 2013.

[35] Y. Behzadi, K. Restom, J. Liau, and T.T. Liu, "A component based noise correction method (CompCor) for BOLD and perfusion based fMRI," *NeuroImage*, vol. 37, no. 1, pp. 90–101, 2007.

[36] B.A. Landman, A.J. Huang, A. Gifford, D.S. Vikram, I.A.L. Lim, J.A.D. Farrell, J.A. Bogovic, J. Hua, M. Chen, S. Jarso, S.A. Smith, S. Joel, S. Mori, J.J. Pekar, P.B. Barker, J.L. Prince, and P.C.M. van Zijl. "Multi-parametric neuroimaging reproducibility: A 3T resource study," *NeuroImage*, vol. 54, no. 4, pp. 2854–2866, 2011.

[37] B. Thirion, G. Varoquaux, E. Dohmatob, and J.-B. Poline, "Which fMRI clustering gives good brain parcellations?," *Front. Neurosci.*, vol. 8, pp. 167:1–167:13, 2014.

[38] B. Thyreau, B. Thirion, G. Flandin, and J.-B. Poline, "Anatomo-functional description of the brain: A probabilistic approach," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Proc.*, 2006, pp. 14–19.

[39] M. Girvan and M. E. J. Newman, "Community structure in social and biological networks," *Proc. Nat. Acad. Sci. USA*, vol. 99, no. 12, pp. 7821–7826, 2002.

[40] A.V. Esquivel and M. Rosvall, "Compression of flow can reveal overlapping modular organization in networks," *Phys. Rev. X*, vol. 1, 021025, 2011.

[41] Q. Wang and E. Fleury, "Fuzziness and overlapping communities in large-scale networks," *J. Univers. Comput. Sci.*, vol. 18, no. 3, pp. 457–486, 2012.

[42] H.W. Kuhn, "The Hungarian method for the assignment problem," *Nav. Res. Logist. Q.*, vol. 2, pp. 83–97, 1955.

[43] M. Rubinov and O. Sporns, "Complex network measures of brain connectivity: Uses and interpretations," *NeuroImage*, vol. 52, no. 3, pp. 1059–1069, 2010.

[44] P. Shrout and J. Fleiss, "Intraclass correlations: uses in assessing rater reliability," *Psychol. Bull.*, vol. 86, no. 2, pp. 420–428, 1979.

[45] M. Van Den Heuvel, H.E. Hulshoff Pol, "Exploring the brain network: A review on resting-state fMRI functional connectivity," *Eur. Neuropsychopharm.*, vol. 20, no. 8, pp. 519–534, 2010.

[46] B.T. Yeo, F.M. Krienen, J. Sepulcre, M.R. Sabuncu, D. Lashkari, M. Hollinshead, J.L. Roffman, J.W. Smoller, L. Zöllei, J. R. Polimeni, B. Fischl, H. Liu, and R.L. Buckner, "The organization of the human cerebral cortex estimated by intrinsic functional connectivity," *J. Neurophysiol.*, vol. 106, no. 3, pp. 1125–1165, 2011.

[47] R.L. Buckner, J.R. Andrews-Hanna, D.L. Schacter, "The brain's default network: Anatomy, function, and relevance to disease," *Ann. N. Y. Acad. Sci.*, vol. 1124, pp. 1–38, 2008.

[48] N.B. Turk-Browne, J.D. Golomb, and M.M. Chun, "Complementary attentional components of successful memory encoding," *NeuroImage*, vol. 66, pp. 553–562, 2013.

[49] N.D. Woodward, B. Rogers, and S. Heckers, "Functional resting-state networks are differentially affected in schizophrenia," *Schizophr. Res.*, vol. 130, no. 1, pp. 86–93, 2011.

[50] T.J. Ozaki, "Frontal-to-parietal top-down causal streams along the dorsal attention network exclusively mediate voluntary orienting of attention," *PLoS ONE*, vol. 6, no. 5, e20079, 2011.

[51] R.L. Buckner, J. Sepulcre, T. Talukdar, F.M. Krienen, H. Liu, T. Hedden, J.R. Andrews-Hanna, R.A. Sperling, and K.A. Johnson, "Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to Alzheimer's disease," *J. Neurosci.*, vol. 29, no. 6, pp. 1860–1873, 2009.

[52] D.C. Van Essen, J. Dickson, J. Harwell, D. Hanlon, C.H. Anderson, and H.A. Drury, "An integrated software system for surface-based analyses of cerebral cortex," *J. Am. Med. Inform. Assoc.*, vol. 8, no. 5, pp. 443–459, 2001.

[53] L. Dugatkin, *Game Theory and Animal Behavior*. Oxford, U.K.: Oxford Univ. Press, 1998.

[54] S.R. Bulò, A. Torsello, and M. Pelillo, "A game-theoretic approach to partial clique enumeration," *Image Vision Comput.*, vol. 27, no. 7, pp. 911–922, 2009.