Discovering Sparse Functional Brain Networks Using Group Replicator Dynamics (GRD)

Bernard Ng[1], Rafeef Abugharbieh[1], and Martin J. McKeown[2]

[1]Biomedical Signal and Image Computing Lab, Department of Electrical Engineering,
[2]Department of Medicine (Neurology), Pacific Parkinson's Research Center
The University of British Columbia, Vancouver, BC, Canada
bernardn@ece.ubc.ca, rafeef@ece.ubc.ca, mmckeown@interchange.ubc.ca

Abstract. Functional magnetic resonance imaging (fMRI) has become increasingly used for studying functional integration of the brain. However, the large inter-subject variability in functional connectivity renders detection of representative group networks very difficult. In this paper, we propose a new iterative method that we refer to as "group replicator dynamics," for detecting sparse functional networks that are common across subjects within a group. The proposed method uses replicator dynamics, which we show to be equivalent to non-negative sparse PCA, and incorporates group information for identifying common networks across subjects with subject-specific weightings of the identified brain regions reflecting individual differences. Finding a separate network for each subject, as opposed to employing traditional averaging approaches, permits statistical testing of group significance. We validated our method on synthetic data, and applying it to real fMRI data detected task-specific group networks that conform well with prior neuroscience knowledge.

Keywords: fMRI, functional connectivity, group analysis, replicator dynamics

1 Introduction

In recent years, an increasing interest in the functional integration of the brain has emerged for studying higher order cognitive functions [1-3] and neurological diseases [3-7]. Functional integration refers to the interactions between different brain regions and is typically characterized in terms of functional connectivity [8] and effective connectivity [9]. Functional connectivity is defined as the correlations between brain regions, whereas effective connectivity is the influence of a brain region over another brain region. The focus of this paper will be on inferring functional connectivity from functional magnetic resonance imaging (fMRI) data.

The classical approach of analyzing functional connectivity is to pick a seed voxel or region and examine the linear correlations between the seed and all other voxels in the brain [10]. This approach has been widely used for detecting resting-state networks with promising results [10, 11], but suffers from the drawback of neglecting the joint interactions between multiple voxels [12]. To account for such interactions, multivariate techniques such as clustering [13-15], principal component analysis

(PCA) [8], and independent component analysis (ICA) [12, 16] have previously been explored. To explicitly exploit the sparse nature of human brain networks [17], we have adopted replicator dynamics [18-21] in this paper, which we show to be a fast solution to the non-negative sparse PCA problem [22, 23] under certain constraints. Enforcing sparsity eases the diffused weighting problem seen in classical PCA [24, 25], where non-zero loadings (i.e. spatial component maps) are often assigned to the majority of voxels within brain, which complicates network identification.

The multivariate methods described above provide effective means of analyzing functional connectivity within a single subject, but the optimal way for performing group analysis is not trivial. The traditional way is to either concatenate [26, 27] or average [27, 28] the voxel time courses across subjects in detecting representative group networks. Alternatively, one can perform connectivity analysis separately on each subject and average the results. For instance, Lohmann and Bohn performed group inference by first estimating the pairwise correlations between voxels for each subject and subsequently applying replicator dynamics on the average correlation matrix [18]. Similarly, van de Heuvel et al. in their study of resting-state networks applied normalized cut to cluster voxels within each subject and used the consistency of the clustermaps to infer group clusters [15]. Provided the functional organization is similar across subjects, the average networks would be reasonable representations of the group. However, if the connection strengths between brain regions vary across subjects, the average group networks may not reflect the features in the individual networks of each subject [29]. Moreover, collapsing the individual networks into a single group network does not permit statistical testing for group significance.

Another challenge associated with group analysis is the need for establishing a correspondence across subjects. Connectivity analysis performed at the voxel level, as often done for most of the aforementioned methods, typically entails spatial normalization, which is prone to mis-registration [30-32] and spatial distortions [33]. An alternative approach to create a subject correspondence without spatial normalization is to define regions of interest (ROIs) in each subject's native space and examine the connectivity between these brain regions. For instance, Marrelec et al. used partial correlation to measure the connectivity between anatomically defined ROIs [34]. Similarly, Ng et al. applied replicator dynamics on anatomically defined ROIs to investigate the effects of task frequency on functional connectivity [21]. To avoid the drawbacks associated with spatial normalization, this ROI-based approach is thus employed in our analysis.

In this paper, we propose a new iterative method for detecting sparse functional networks consistently recruited among a group of subjects. To encourage sparsity, replicator dynamics is used to iteratively update a weight vector that represents the degree to which an ROI belongs to the most coherent network. To facilitate statistical testing of group significance, we propose to learn the dominant network of each subject *individually* (as opposed to applying replicator dynamics to the average correlation matrix of the group [18, 21]) but with group information incorporated within each iteration, so that subject variability can be modeled while ensuring *group* networks are being detected. This iterative scheme, as we will demonstrate with synthetic and real fMRI data, results in functional networks consisting of the same ROIs across subjects yet with different ROI weightings for each subject reflecting the individual differences, thus enabling group significance to be assessed.

2 Materials

After obtaining informed consent, fMRI data were collected from 10 healthy subjects (3 men, 7 women, mean age 57.4 ± 14 years) and 10 Parkinson's disease (PD) patients on and off medication (4 men, 6 women, mean age 66 ± 8 years). Each subject was required to first use their right-hand and then their left hand to squeeze a bulb with sufficient pressure such that a horizontal bar shown on a screen was kept within an undulating pathway. The pathway remained straight during baseline periods, which required a constant pressure to be applied. During the time of stimulus, the pathway became sinusoidal at a frequency of 0.25 Hz (slow), 0.5 Hz (medium) or 0.75 Hz (fast) presented in a pseudo-random order. Each run lasted 260s, alternating between baseline and stimulus of 20 s duration with an extra 20 s baseline period at the beginning and end of each run.

Functional MRI was performed on a Philips Gyroscan Intera 3.0 T scanner (Philips, Best, Netherlands) equipped with a head-coil. T2*-weighted images with blood oxygen level dependent (BOLD) contrast were acquired using an echo-planar (EPI) sequence with an echo time of 3.7 ms, a repetition time of 1985 ms, a flip angle of 90°, an in plane resolution of 128×128 pixels, and a pixel size of 1.9×1.9 mm. Each volume consisted of 36 axial slices of 3 mm thickness with a 1 mm gap. A 3D T1-weighted image consisting of 170 axial slices was further acquired to facilitate anatomical localization of activation.

Each subject's fMRI data was pre-processed using Brain Voyager's (Brain Innovation B.V.) trilinear interpolation for 3D motion correction and sinc interpolation for slice time correction. Further motion correction was performed using motion corrected independent component analysis (MCICA) [35]. The voxel time courses were then high-pass filtered to account for temporal drifts. No spatial warping or smoothing was performed.

Eighteen motor-related ROIs were manually drawn by an expert on each subject's structural scan in their native space based upon anatomical landmarks and guided by a neurological atlas [36] using AMIRA (Mercury Computer Systems, San Diego, USA). ROIs included the bilateral putamen, caudate, globus pallidus, thalamus (THA), cerebellum (CER), primary motor cortex (M1), supplementary motor area (SMA), prefrontal cortex (PFC), and anterior cingulated cortex (ACC). The segmented ROIs were resliced at the fMRI resolution and used to extract the preprocessed voxel time courses within each ROI for subsequent analysis.

3 Methods

We present here a novel iterative method, which we refer to as group replicator dynamics (GRD), for detecting sparse functional networks that are consistent across subjects. Given an initial weight vector for each subject, where the magnitude of the elements represent the degree to which the corresponding ROIs belong to the most coherent network, replicator dynamics is iteratively applied to update these weight vectors, as described in Section 3.1. In each iteration, after all subjects' weight vectors

are updated, these updated weight vectors are adjusted based on the group information as described in Section 3.2. These two steps are repeated until convergence.

3.1 Replicator Dynamics

Replicator dynamics is a well-known concept in theoretical biology for modeling the evolution of interacting species [37]. In this model, each species is assigned a fitness value, which is used to determine the fittest species that survives over time. In the context of functional connectivity analysis, each brain region is considered a species, with its fitness measured by the correlation between brain regions. If we let $w^i(0)$ be the initial weight vector of the i^{th} subject as defined above, the most coherent network can be determined by iteratively applying the following replicator equation [37]:

$$w^i(k+1) = \frac{w^i(k).*C^i w^i(k)}{w^i(k)^T C^i w^i(k)} \quad , \tag{1}$$

where .* represents element-wise multiplication, k is the iteration counter, and C^i is the correlation matrix of the i^{th} subject with each element corresponding to the pair-wise correlation between average ROI intensity time courses. To avoid self connections, the main diagonal of C^i is set to zero. Furthermore, to avoid bias during initialization, $w^i(0)$ is set to $1/N_r$, where N_r is the number of brain regions being examined. Using (1), w^i is guaranteed to converge to a local maximum based on the fundamental theorem of natural selection [37], provided C^i is real-value, non-negative, and symmetric. In this paper, positive and negative correlations are assumed to be of equal importance, thus the absolute value of the correlations is used to ensure C^i is non-negative. Nevertheless, networks with positively correlated regions can be separately analyzed by nulling out the negative elements in C^i and vice versa. Upon convergence, only elements corresponding to the most correlated brain regions will be non-zero due to the sparse nature of replicator dynamics as we discuss next.

The objective function that (1) maximizes is in fact the same as PCA, i.e. $w^{iT}C^i w^i$. Thus, one might argue that using PCA will provide similar results as replicator dynamics. However, the results obtained from replicator dynamics can in fact be very different from PCA for the following reason. The first principle component of PCA, w^i, is the solution to the optimization problem:

$$\arg\max_{w^i} w^{iT} C^i w^i, \quad subject\ \ to \left\| w^i \right\|_2 = 1 \quad , \tag{2}$$

which does not encourage sparsity [24, 25], hence all elements of w^i will likely be non-zero. This problem of diffused weighting complicates network identification, since no clear threshold will necessarily be present for separating out the dominant network from other less correlated brain regions. To impose sparsity, the constraint in (2) needs to be modified to $\|w^i\|_1 \le K$, where K is a user-defined parameter [38]. The resulting optimization problem is often referred to as sparse PCA:

$$\arg\max_{w^i} w^{iT} C^i w^i, \quad subject\ \ to \left\| w^i \right\|_1 \le K \quad . \tag{3}$$

Various solutions for (3) have been proposed. For instance, d' Aspremont et al. solved (3) using semidefinite programming [25] and Zou et al. solved (3) as a regression problem with lasso (elastic net) constraints [24]. If we further constrain w^i to be non-negative, the resulting problem will be of the form:

$$\arg\max_{w^i} {w^i}^T C^i w^i, \quad subject\ \ to \sum_j w_j^i \le K, w^i \ge 0 \ , \tag{4}$$

which is known as non-negative sparse PCA [22, 23]. Examining (1), since all elements of C^i are restricted to be non-negative, w^i will also be non-negative by construction. Moreover, (1) enforces the elements of w^i to sum to one. Thus, replicator dynamics is in fact a solution of the non-negative sparse PCA problem! We note that previous studies [18] have observed an unexplained tendency of the detected networks to comprise of mutually correlated voxels. This observation can actually be explained by our finding of the equivalence between replicator dynamics and non-negative sparse PCA, since imposing sparsity given limited weights (i.e. $\sum {w_j}^i = 1$) will encourage the weights to be assigned to mutually correlated brain regions.

3.2 Group Replicator Dynamics

If we simply apply (1) separately for each subject until convergence, the detected networks will likely be different across subjects [29]. However, if there exists a core functional network that is common across subjects, we can exploit group information to identify this core network. The underlying assumption is that the majority of subjects use the same core functional network to perform the same task, but in addition, each subject may recruit extra brain regions, which constitutes inter-subject variability. Also, the common core network may not necessarily be the most coherent network for every subject, potentially due to different processing strategies or simply due to different levels of noise [29]. Given $w^i(k)$, to encourage networks consisting of the same brain regions across subjects, one could adjust $w^i(k)$ such that the group entropy is minimized [39, 40]. For analytic simplicity, if we assume $w^i(k)$ are instances of a normal random variable, w, then the entropy of w is given by:

$$\ln\sqrt{(2\pi e)^{N_r}|\Sigma|} \propto \ln|\Sigma| \ , \tag{5}$$

where Σ is the covariance of w. To minimize the entropy, a gradient descent approach is used with $w^i(k)$ updated as follows [39, 40]:

$$W(k) = W(k) - \lambda(W_c(k)^T (W_c(k) W_c(k)^T + \alpha I)^{-1}) \ , \tag{6}$$

where $W(k)$ is a $N_r \times N_s$ matrix with $w^i(k)$ of each subject along the columns and $W_c(k)$ is the same as $W(k)$ but with the subject mean removed from each row. λ governs the learning rate and must be set below α to ensure numerical stability [39]. Since the number of subjects, N_s, will typically be much less than the number of brain regions, N_r, $W_c W_c^T$ will be rank deficient. Hence, αI (with $\alpha = 0.1$) is added to regularize the inversion of $W_c W_c^T$. At each iteration k, after updating all subjects' $w^i(k)$ using (1), (6) is applied to adjust $w^i(k)$ so that elements of $w^i(k)$ that are consistently large across

subjects are reinforced, while other elements are suppressed. (1) is then applied to the adjusted $w^i(k)$ at the next iteration, k+1, and the same process is repeated until convergence. Since (1) searches for the most coherent network within a subject, while (6) encourages networks to be similar across subjects, the combined effect of (1) and (6) results in highly coherent networks that comprise of the same brain regions across subjects, but with subject-specific ROI weightings reflecting individual differences.

To test if the detected functional network is significantly different from that obtained from random noise, we first approximate the network coherence of each subject, c^i, as $w^{iT}C^iw^i$. A t-test is then applied to the Fisher's Z-transform of the c^i's, which we denote as z^i, to estimate the probability that the z^i's are drawn from a null distribution generated from random noise. To estimate the mean of the null distribution, we temporally permute all the average ROI intensity time courses 10,000 times, calculate the mean z averaged across subjects for each permutation, and take the largest mean z as the null mean. We note that approximating c^i as $w^{iT}C^iw^i$ is legitimate only if the same ROIs are detected as the core functional network across subjects. Otherwise, the estimated c^i's would represent different networks and declaring the presence of a significant common group network under this scenario would be invalid. Also, although applying replicator dynamics to the average correlation matrix [18, 21] may detect a similar group network, collapsing all subjects' networks into a single group network does not permit testing of group significance. We note that subsequent networks can be detected by removing those ROIs detected in the dominant network and reapplying the procedures above [18, 21].

3.3 Empirical Evaluation

To test our proposed connectivity analysis framework, we generated 1,000 synthetic datasets, each consisting of 10 subjects and 20 ROIs for each subject, as summarized in Fig. 1(m). ROIs 1 to 4 were designed to be mutually correlated and form the core functional network. These ROIs were also set to be connected to ROI 10 for the subject 1, ROI 11 for subject 2, etc. to introduce inter-subject variability. Furthermore, we have set ROIs 5 to 9 to be mutually correlated but with lower correlation except for subjects 1 and 2 to test the case where a secondary network is present and the core functional network is not the most coherent network for some of the subjects. The synthetic ROI time courses were generated by convolving a box-car with the hemodynamic response and adding Gaussian noise (Fig. 1(k)) at a level similar to the real data (Fig. 1(l)). The timing of the box-car for ROIs 1 to 4 and the extra subject-dependent ROI is set to be exactly the same in our experiments. For ROIs 5 to 9, the stimulus period is shrunk by half to distinguish this secondary network from the core. Also, except for subjects 1 and 2, the noise level is set to be slightly higher than the core to ensure this secondary network is less correlated. For the remaining ROIs, random Gaussian noise is used as the ROI time courses.

The results obtained by applying the proposed group replicator dynamics as well as other related schemes are summarized in Fig. 1(a) to (j), where the average w^i of each subject over 1,000 synthetic datasets are plotted. Applying replicator dynamics separately to each subject (curve with dots) detected ROIs 1 to 4 to be part of the most

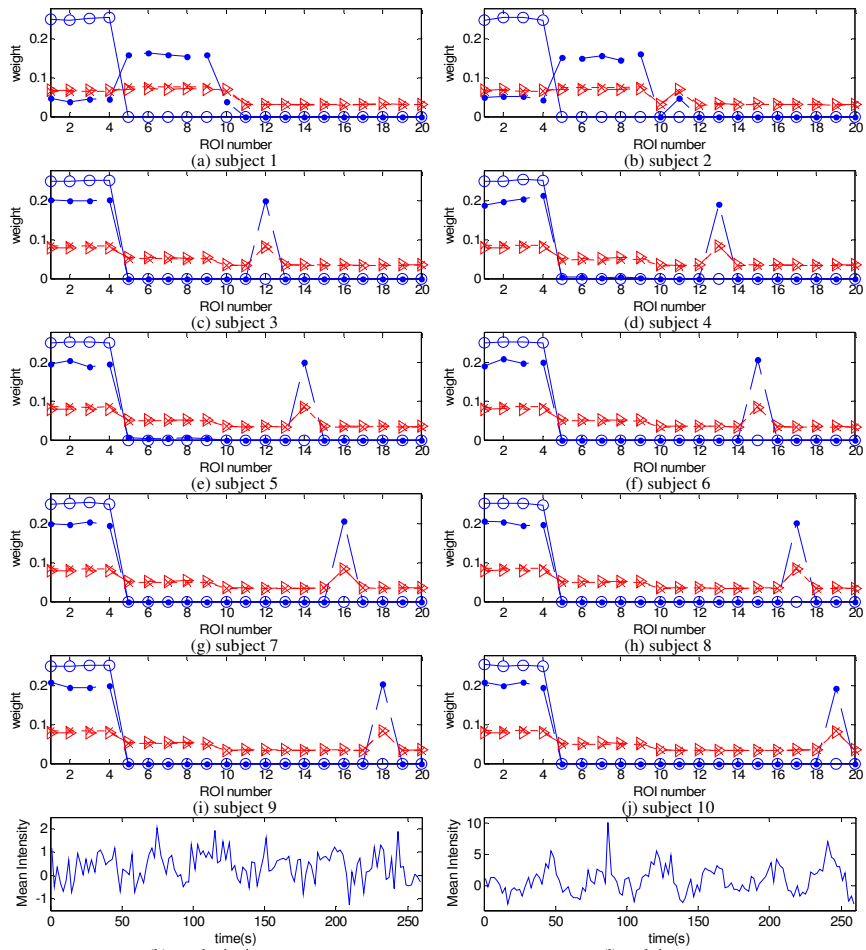

Sub/ROI	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	S	S	S	S	D	D	D	D	D	S	x	x	x	x	x	x	x	x	x	x
2	S	S	S	S	D	D	D	D	D	x	S	x	x	x	x	x	x	x	x	x
3	D	D	D	D	S	S	S	S	S	x	x	D	x	x	x	x	x	x	x	x
4	D	D	D	D	S	S	S	S	S	x	x	x	D	x	x	x	x	x	x	x
5	D	D	D	D	S	S	S	S	S	x	x	x	x	D	x	x	x	x	x	x
6	D	D	D	D	S	S	S	S	S	x	x	x	x	x	D	x	x	x	x	x
7	D	D	D	D	S	S	S	S	S	x	x	x	x	x	x	D	x	x	x	x
8	D	D	D	D	S	S	S	S	S	x	x	x	x	x	x	x	D	x	x	x
9	D	D	D	D	S	S	S	S	S	x	x	x	x	x	x	x	x	D	x	x
10	D	D	D	D	S	S	S	S	S	x	x	x	x	x	x	x	x	x	D	x

(m)

Fig. 1. Synthetic data results. (a) to (j) Average weight vectors of subjects 1 to 10 over 1,000 datasets detected using group replicator dynamics (circle), replicator dynamics (dot), group PCA (triangle), and PCA (cross). (k) Sample synthetic ROI time course with similar noise level as in real data (l). (m) Ground truth. "D" and "S" indicate if an ROI belongs to the dominant or secondary network. "x" implies random noise. Applying group replicator dynamics detected group networks that perfectly matched with the ground truth, but not so for the other methods.

coherent network for subjects 3 to 10, but not for subjects 1 and 2. This result is expected since ROIs 1 to 4 were designed to have lower mutual correlations than ROIs 5 to 9 for these two subjects. The extra subject-dependent ROI was also assigned roughly equal weights as ROIs 1 to 4. Hence, replicator dynamics alone provides a powerful means of detecting sparse subject-specific networks, but it is not as effective in identifying group networks, where certain heuristics will be required to prune out ROIs 1 to 4 as being the core functional network. We note that computationally estimating the weight vector using (1) took less than a millisecond for each subject.

Applying group replicator dynamics (curve with circle) detected ROIs 1 to 4 as the core functional network for all subjects even for subjects 1 and 2. Also, no weight was falsely assigned to the extra subject-dependent ROI for all 1,000 datasets. Furthermore, the p-values obtained for all datasets were well below 0.05. We note that applying replicator dynamics on the average correlation matrix also detected ROIs 1 to 4 as being the core network, but statistical significance could not be assessed since all subjects' information was collapsed in this approach.

For comparison, we also applied PCA to each subject separately (curve with crosses). Weights were assigned to all ROIs even for random noise, although slightly higher weights were on average given to ROIs 1 to 4 and the subject-dependent ROI for subjects 3 to 10, and ROIs 5 to 9 for subjects 1 and 2. However, if we instead examine the weights of each dataset separately, the weights given to ROIs 1 to 9 and the subject-dependent ROI were in fact very similar with weights assigned to ROIs 5 to 9 occasionally being higher than that of ROIs 1 to 4 for subjects 3 to 10, and vice versa for subjects 1 and 2. Thus, defining a subjective threshold to identify the most coherent network would be very difficult and ROIs not belonging to the most coherent network might be falsely detected.

To incorporate group information into PCA, we replaced (1) with $w^i(k+1) = C^i w^i(k)$, $\|w^i(k)\|_2 = 1$ (i.e. estimating the first principal component using the Power method) in our proposed iterative scheme. We refer to this algorithm as group PCA. Using this algorithm resulted in similar networks as obtained by applying PCA to each subject separately. We suspect the reason for this result was that PCA assigned similar weights to every ROI, thus the group entropy of the w^i's would be minute, leading to negligibly small adjustments applied to w^i by (6).

4 Results and Discussion

Fig. 2 summarizes the core functional networks detected by applying group replicator dynamics to our real experimental fMRI data for each combination of subject group and task (e.g. healthy and slow). The correlation matrices for each task were estimated using the corresponding segments of the average ROI intensity time courses. Only the left-handed results are presented due to space limitation.

For healthy subjects, the bilateral cerebellar hemispheres, bilateral SMA, and right M1 were detected as the core network during slow frequency task. At the medium frequency, the bilateral thalami were further detected. Since the thalamus is known to serve as a relay between the cerebellum and the cortex, detecting the

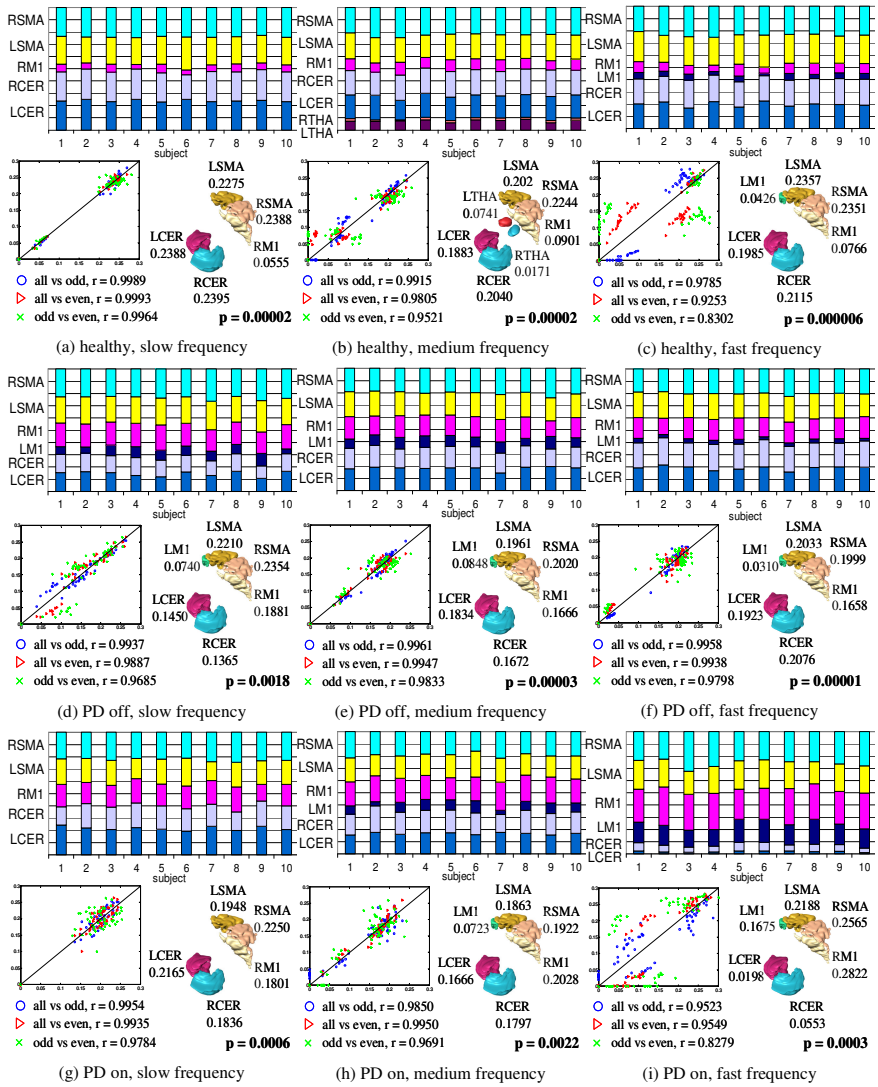

Fig. 2. Real fMRI data results. At the top of each sub-figure is a bar graph of the estimated weight vector of each subject using group replicator dynamics. The ROIs given non-zeros weights are labeled on the side of the bar graph. The bottom left of each sub-figure is a consistency plot. The bottom right of each sub-figure is a 3D rending of the detected ROIs and the average weights of these ROIs over the subjects. All networks showed high consistency and were detected to be significant at a p-value of 0.05 with Bonferroni correction (i.e. 0.05/18 = 0.0028). Networks detected in PD off medication were similar to those detected in healthy subjects during the fast frequency, demonstrating the effects of brain deficits in PD patients. Upon medication, the network appeared to have normalized back to that of the healthy subjects during slow frequency, but not at medium and fast.

thalamus may be a result of increased communication. At the fastest frequency, the left M1 instead of the bilateral thalami were found to be part of the core network. Prior studies suggested that the ipsilateral M1 becomes more involved as task difficulty increases [41], thus detecting the left M1 during a left-handed task at the highest frequency conforms to prior neuroscience findings. For PD patients off medication, the core network detected at the slow frequency was exactly the same as that of the healthy subjects' during the fast frequency. This result suggests that the brain deficits in PD patients made the slow frequency task appeared as difficult as the fast condition, hence the need for the left M1. The same network was detected during medium and fast frequencies. For PD patients on medication, the core network appeared to have normalized to that of healthy subjects at the slow frequency, but with more involvement of the right M1, which also matches findings from other prior studies [42]. At the medium frequency, the left M1 again became involved, and at the fast frequency, substantially more involvement of the bilateral M1 and bilateral SMA were found. Increased activity in the SMA upon medication was also observed in prior studies [42]. We note that all detected core networks were significant at a p-value of 0.05 with Bonferroni correction for the number of comparisons (i.e. 0.05/18 = 0.0028). Also, to test for consistency, we divided the time courses into odd and even time points and re-applied our proposed method. The weights obtained by using all time points and only the odd and even time points were very similar with correlations well above 0.9 for most cases. Regarding the constraint of non-negative correlation matrix required for applying replicator dynamics, the correlations of the detected ROIs were all positive. Hence, this constraint was not a major issue for the current experiment. For comparisons, we also applied group PCA to the data, but the results were analogous to the synthetic case, where similar weights were assigned to all ROIs. Thus, no distinct group networks could easily be identified.

5 Conclusions

In this paper, we proposed a new iterative method for detecting common networks within a group of subjects. By exploiting replicator dynamics, which we showed to be equivalent to non-negative sparse PCA, and incorporating group information, our method was able to detect sparse group networks that perfectly matched with the ground truth when applied to synthetic data. In contrast, classical PCA resulted in diffused weightings, which greatly hindered network identification. Applying group replicator dynamics to real fMRI data detected significant task-specific group networks that matched well with prior neuroscience knowledge. Our results thus demonstrated that the combination of sparsity and group information provides an effective means to combat against inter-subject variability in detecting group functional networks.

References

1. Schlosser, R.G., Wagner, G., Sauer, H.: Assessing the Working Memory Network: Studies with Functional Magnetic Resonance Imaging and Structural Equation Modeling. Neurosci.; 139, 91--103 (2006)
2. Au Duong, M.V., et al.: Modulation of Effective Connectivity inside the Working Memory Network in Patients at the Earliest Stage of Multiple Sclerosis. NeuroImage. 24, 533--538 (2005)
3. Addis, D.R., Moscovitch, M., McAndrews, M.P.: Consequences of Hippocampal Damage across the Autobiographical Memory Network in Left Temporal Lobe Epilepsy. Brain. 130(9), 2327--2342 (2007)
4. He, B.J., et al.: The Role of Impaired Neuronal Communication in Neurological Disorders. Curr. Opin. Neurol. 20, 655--660 (2007)
5. Waites, A.B., et al.: Functional Connectivity Networks are Disrupted in Left Temporal Lobe Epilepsy. Ann. Neurol. 59, 335--343 (2006)
6. Au Duong, M.V., et al.: Altered Functional Connectivity Related to White Matter Changes inside the Working Memory Network at the very Early Stage of MS. J. Cereb. Blood Flow Metab. 25, 1245--1253 (2005)
7. Cader, S., et al.: Reduced Brain Functional Reserve and Altered Functional Connectivity in Patients with Multiple Sclerosis. Brain. 129(2), 527--537 (2006)
8. Friston, K.J., et al.: Functional Connectivity: The Principal Component Analysis of Large (PET) Data Sets. J. Cereb. Blood Flow Metab. 13, 5--14 (1993)
9. Friston, K.J., Frith, C.D., Frackowiak, R.S.J.: Time-dependent Changes in Effective Connectivity Measured with PET. Hum. Brain Mapp. 1, 69--80 (1993)
10. Biswal, B., et al.: Functional Connectivity in the Motor Cortex of Resting Human Brain using Echo-planar MRI. Magn. Reson. Med. 34, 537--541 (1995)
11. Greicius MD, Krasnow B, Reiss AL, Menon V. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proc Natl Acad Sci U S A 2003; 100:253–258.
12. van de Ven, V.G., et al.: Functional Connectivity as Revealed by Spatial Independent Component Analysis of fMRI Measurements during Rest. Hum. Brain. Mapp. 22, 165--178 (2004)
13. Goutte, C., et al.: On Clustering fMRI Time Series. NeuroImage. 9(3), 298--310 (2002)
14. Cordes, D., et al.: Hierarchical Clustering to Measure Connectivity in fMRI Resting State Data. Magn. Reson. Imaging. 20, 305--317 (2002)
15. van den Heuvel, M., Mandl, R., Hulshoff Pol, H.: Normalized Cut Group Clustering of Resting-state fMRI Data. PLoS ONE. 3(4), e2001 (2008)
16. McKeown, M.J., et al.: Analysis of fMRI Data by Blind Separation into Independent Spatial Components. Hum. Brain. Mapp. 6, 160 --188 (1998)
17. Bassett, B.S.: Small-World Brain Networks. The Neuroscientist. 12(6), 512--523 (2006)
18. Lohmann, G., Bohn, S.: Using Replicator Dynamics for Analyzing fMRI Data of the Human Brain. Trans. Med. Imaging. 21, 485--492 (2002)
19. Neumann, J., et al.: Meta-analysis of Functional Imaging Data Using Replicator Dynamics. Hum. Brain Mapp. 25, 165--173 (2005)
20. Neumann, J., et al.: The Parcellation of Cortical Areas Using Replicator Dynamics in fMRI. NeuroImage. 32, 208--219 (2006)
21. Ng, B., Abugharbieh, R., McKeown, M.J.: Inferring Functional Connectivity using Spatial Modulation Measures of fMRI Signals within Brain Regions of Interest. In: 5th IEEE International Symposium on Biomedical Imaging: From Nano to Marco. 572--275 (2008)
22. Zass, R., Shashua, A.: Nonnegative Sparse PCA. In: Advances in Neural Information Processing Systems. 1561--1568 (2006)

23. Sigg, C.D., Buhmann, J.M.: Expectation Maximization for Sparse and Non-Negative PCA. In: 25th International Conference on Machine Learning, Helsinki, Finland, 2008.
24. Zou, H., Hastie, T., Tibshirani, R.: Sparse Principal Component Analysis. J. Computational and Graphical Statistics. 15(2), 265--286 (2004)
25. d'Aspremont, A., Bach, F., El Ghaoui, L.: Full Regularization Path for Sparse Principal Component Analysis. In: 24th International Conference on Machine Learning. 227, 177--184. ACM, New York (2007)
26. Calhoun, V.D., et al.: A Method for Making Group Inferences from Functional MRI Data using Independent Component Analysis. Hum. Brain. Mapp. 14, 140--151 (2001)
27. Schmithorst, V.J., Holland, S.K.: Comparison of Three Methods for Generating Group Statistical Inferences from Independent Component Analysis of Functional Magnetic Resonance Imaging Data. J. Magn. Reson. Imaging. 19(3), 365--368 (2004)
28. Bokde, A.L., et al.: Functional Interactions of the Inferior Frontal Cortex during the Processing of Words and Word-like Stimuli. Neuron. 30, 609--617 (2001)
29. Goncalves, M. S., et al: Can Meaningful Effective Connectivities be Obtained between Auditory Cortical Regions. NeuroImage. 14, 1353--1360 (2001)
30. Uylings, H.B.M., et al.: Consequences of Large Interindividual Variability for Human Brain Atlases: Converging Macroscopical Imaging and Microscopical Neuroanatomy. Anat. Embryol. 210, 423--431 (2005)
31. Samanez-Larkin, G.R., D'Esposito, M.: Group Comparisons: Imaging the Aging Brain. Soc. Cogn. Affect. Neurosci. 3(3), 290--297 (2008)
32. Wilke, M., Schmithorst, V.J., Holland, S.K.: Assessment of Spatial Normalization of Whole-Brain Magnetic Resonance Images in Children. Hum. Brain Mapp. 17(1), 48--60 (2002)
33. Ng, B., Abugharbieh, R., McKeown, M.J.: Adverse Effects of Template-based Warping on Spatial fMRI Analysis. SPIE Conference on Medical Imaging, Orlando, Florida (2009)
34. Marrelece, G., et al.: Partial Correlation for Functional Brain Interactivity Investigation in Functional MRI. NeuroImage. 32, 228--237 (2006)
35. Liao, R., Krolik, J.L., McKeown, M.J.: An Information-theoretic Criterion for Intrasubject Alignment of fMRI Time Series: Motion Corrected Independent Component Analysis. Trans. Med. Imaging. 24(1), 29--44 (2005)
36. Talairach, J., Tournoux, P.: Co-Planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System - an Approach to Cerebral Imaging. Thieme Medical Publishers, New York. (1988)
37. Schuster, P., Sigmund, K.: Replicator dynamics. J. Theor. Biol. 100, 533--538 (1983)
38. Tibshirani R: Regression Shrinkage and Selection via the LASSO. J. Royal Stat. Soc. series B. 58, 267--288 (1996)
39. Cates, J., et al.: Shape Modeling and Analysis with Entropy-Based Particle Systems. In: Karssenemeijer, N. Lelievedldt, B. (eds.) IPMI 2007, LNCS 4584, 333--345 Springer-Verlag Berlin Heidelberg (2007)
40. Oguz, I., et al.: Cortical Correspondence Using Entropy-based Particle Systems and Local Features. In: 5th IEEE International Symposium on Biomedical Imaging: From Nano to Marco. 1637--1640 (2008)
41. Verstynen, T., et al.: Ipsilateral Motor Cortex Activity during Unimanual Hand Movements Relates to Task Complexity. J. Neurophysiol. 93, 1209-1222 (2005)
42. Haslinger, B., et al.: Event-related Functional Magnetic Resonance Imaging in Parkinson's Disease before and after Levodopa. Brain. 124(3), 558-570 (2001)