

Detecting Brain Activation in fMRI Using Group Random Walker

Bernard Ng¹, Ghassan Hamarneh², and Rafeef Abugharbieh¹

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

² Medical Image Analysis Lab, Simon Fraser University

{Bernardn, rafeef}@ece.ubc.ca, hamarneh@cs.sfu.ca

Abstract. Due to the complex noise structure of functional magnetic resonance imaging (fMRI) data, methods that rely on information within a single subject often results in unsatisfactory functional segmentation. We thus propose a new graph-theoretic method, “Group Random Walker” (GRW), that integrates group information in detecting single-subject activation. Specifically, we extend each subject’s neighborhood system in such a way that enables the states of both intra- and inter-subject neighbors to be regularized without having to establish a one-to-one voxel correspondence as required in standard fMRI group analysis. Also, the GRW formulation provides an exact, unique closed-form solution for jointly estimating the probabilistic activation maps of all subjects with global optimality guaranteed. Validation is performed on synthetic and real data to demonstrate GRW’s superior detection power over standard analysis methods.

Keywords: fMRI, graphical models, group analysis, random walker.

1 Introduction

Functional magnetic resonance imaging (fMRI) has become one of the most widely-used modality for studying human brain activity. The standard approach for analyzing fMRI data involves separately comparing each voxel’s intensity time course against an expected response to generate statistics that reflect the likelihood of activation [1]. The drawback to this univariate approach is that voxel interactions are ignored despite that each voxel is unlikely to function in isolation. To remedy this limitation, methods based on Markov random fields (MRF) [2] and Bayesian statistics [3] have been proposed to incorporate voxel interactions in the form of neighborhood information. These methods help suppress false declaration of isolated voxels as being active. However, the inherently low signal-to-noise (SNR) of fMRI data limits the reliability of the neighbors, which reduces the effectiveness of the currently-used regularization methods. The core of the problem is that there may just be insufficient information within a single subject’s data to obtain satisfactory functional segmentation. Additional information is thus needed to disambiguate the state of noisy voxels.

Most fMRI studies focus on identifying common patterns across subjects, and thus exploiting the *group* dimension presents a direct, intuitive means of enhancing single-subject segmentations [4]. In standard fMRI group analysis, brain images of all

subjects are first warped onto a common template to create a voxel correspondence [1]. Activation statistics are then compared across subjects to generate a group map. The underlying assumption is that a perfect one-to-one voxel correspondence is established after whole-brain warping. However, the vast anatomical variability renders this assumption questionable. In fact, even if perfect anatomical alignment is achieved, whether a one-to-one functional correspondence exists between voxels is debatable. Past studies have shown considerable functional inter-subject variability [5], which suggests that such one-to-one voxel correspondence is rather unlikely. However, active voxels are typically observed within the same anatomical regions across subjects [5]. Thus, integrating inter-subject neighborhood evidence is likely to help regularize single-subject segmentations and better distinguish signal from noise.

In this paper, we propose a new graph-theoretic method, “Group Random Walker” (GRW) that extends our previous work [6] on RW for estimating single-subject probabilistic activation maps. Treating each voxel as a graph vertex, we extend edges to inter-subject in addition to intra-subject neighboring voxels to jointly exploit group information and voxel interactions. Integrating group information into each subject’s activation map, as opposed to estimating a group map, also facilitates inter-subject commonalities as well as differences to be modeled. Moreover, GRW draws upon a RW formulation [7] that provides an exact, unique closed-form solution for computing probabilistic activation maps with global optimality guaranteed.

2 Proposed Method

We propose extending the single-subject neighborhood system to other subjects within a group to disambiguate the state of noisy voxels. The intuition behind this approach is that true brain activation should appear in similar proximal locations across subjects [5], whereas false positives are more randomly scattered across the brain. Hence, regularizing inter-subject neighbors reinforces brain areas that are consistently recruited across subjects, while suppressing the false positives. Since only voxels that are spatially-proximal to each other are encouraged to be in similar state in our framework, the stringent one-to-one voxel correspondence requirement in standard fMRI group analysis is mitigated. Establishing an inter-subject neighborhood system requires first aligning the brain structures of all subjects. However, the vast anatomical variability renders accurate whole-brain warping difficult, especially for diseased subjects. Therefore, we instead employ a region-based approach, where we extract anatomical regions of interest (ROIs) and perform alignment at the regional level. This approach ensures that no brain structures will be mistakenly taken as part of another structure which has shown to improve activation localization [8].

2.1 Group Random Walker

In the original RW framework, each voxel is represented as a graph vertex with weighted edges added between spatial neighbors to bias the paths for which a random walker may transverse. Voxels are labeled (e.g. active or non-active) based on the probability that a random walker starting at each voxel location will first reach a pre-labeled seed. This framework, however, not only requires specifying seed voxels but

also does not model unary voxel information, such as activation effects in the context of fMRI. Therefore, we adopt an augmented RW formulation [7] that facilitates incorporation of unary information as label priors. This formulation is equivalent to adding an artificial seed vertex for each label and connecting these seeds to every vertex in the original graph with label priors being the edge weights [7]. The corresponding energy functional is as follows:

$$E(x^s) = x^{sT} Lx^s + \sum_{k=1, k \neq s}^K x^{kT} \Lambda^k x^k + (x^s - 1)^T \Lambda^s (x^s - 1), \tag{1}$$

where x^s are the unknown posterior probabilities of the voxels belonging to label class s , K is the number of labels, Λ^s is a diagonal matrix containing prior probabilities of the voxels belonging to label class s (Section 2.2), and L is a weighted graph Laplacian matrix (Section 2.3). This construction is analogous to graph cuts, where the first term in (1) models voxel interactions, while the second term models unary voxel information. The main difference is that RW minimizes (1) over real-valued probabilities instead of binary numbers, which has an exact, unique closed-form solution with global optimality guaranteed for an arbitrary number of labels [7]. Specifically, x^s can be estimated by solving [7]:

$$\left(L + \sum_{k=1}^K \Lambda^k \right) x^s = \lambda^s, \tag{2}$$

where λ^s are the diagonal elements of Λ^s . To incorporate group information, we extend the augmented RW graph structure by inserting edges between each subject’s voxels and their inter-subject neighbors. By treating all subjects’ voxels as a single set and adding edges in the manner described later in Section 2.3, (2) can be directly applied to jointly estimate probabilistic activation maps of all subjects. GRW hence inherits all desired properties of the RW formulation. Globally optimal labeling can be obtained by assigning voxels to the labels associated with the highest probability.

2.2 Label Priors

To compute label priors λ^s , we first estimate the ROI activation statistics t_j of each subject using the standard general linear model (GLM) [1]:

$$y_j = X\beta_j + \omega_j, \quad t_j = \beta_j / se(\beta_j), \tag{3}$$

where y_j is the time course of voxel j , ω_j is assumed to be white Gaussian noise after preprocessing, β_j are the estimated activation effects, and $se(\beta_j)$ is the standard error of β_j . X is a design matrix with boxcar functions (time-locked to stimulus) convolved with the hemodynamic response (HDR) as regressors [1]. We model t_j using a constrained Gaussian mixture model (CGMM) [9]. Specifically, t_j is assumed to be generated from a mixture of K Gaussian distributions with mixing coefficients π_k , means μ_k , and variance σ_k^2 . Conjugate priors are used to constrain these parameters:

$$t_j \sim \sum_{k=1}^K \pi_k N(\mu_k, \sigma_k), \quad \pi \sim Dir(\alpha), \quad \mu_k \sim N(\eta, \tau^2), \quad \sigma_k^2 \sim IG(a, b), \tag{4}$$

where $IG(a,b)$ and $Dir(\alpha)$ denote inverse Gamma and Dirichlet distributions. Adding priors enables us to integrate our knowledge into the model. In particular, we know that t_j of non-active voxels should theoretically be 0 and t -threshold for active voxels is typically set between 3 and 4 based on Gaussian Random Field (GRF) theory [1]. We thus encode this prior knowledge on t -values of active and non-active voxels through η with τ^2 set to 1 to model uncertainty in η . We set K to 2 to classify voxels as active or non-active [9]. As for σ_k^2 , we use an uninformative prior by setting a and b to 0.5 [9], since little is known about σ_k^2 . α is set to $1/K$ assuming equal prior class probabilities. Gibbs sampling is employed to estimate the probability of voxel j belonging to each of the K labels [9], which we use as label priors λ^s .

2.3 Weighted Graph Laplacian

Treating voxels of all subjects as a single set, we define the weighted graph Laplacian L based on functional connectivity $f_{ij} = correlation(Y_i, Y_j)$ and spatial distance d_{ij} :

$$L_{ij} = \begin{cases} -w_{ij}, & e_{ij} \in E_{\text{intra}} \cup E_{\text{inter}} \\ \sum_i w_{ij}, & i = j \\ 0, & \text{otherwise} \end{cases}, \quad (5)$$

where $w_{ij} = f_{ij} \exp(-d_{ij})$, Y_j is the magnitude spectrum of the Fourier transform of voxel j 's time course y_j , e_{ij} denote a graph edge between voxels i and j , and E_{intra} and E_{inter} denote the sets of intra- and inter-subject edges. This choice of w_{ij} is motivated by the well-known bilateral filtering technique, which enables closer spatial neighbors with higher functional connectivity to exert greater influence on each voxel, thus de-weighting contributions from outliers. Edges are added between every given voxel of subject p and its 6-connected intra-subject spatial neighbors and c closest inter-subject spatial neighbors for every subject pairs (p,q) , $p \neq q$. c is empirically set to 3. Note that voxel interactions are modeled using $correlation(Y_i, Y_j)$, instead of temporal correlations, since the temporal profile of HDR is known to vary across subjects [9]. In contrast, magnitude spectrums, Y_i , of active voxels would likely display higher similarity across subjects since all subjects are guided by the same stimulus.

2.4 Empirical Evaluation

500 synthetic datasets were generated to validate our proposed method. Each dataset consisted of 10 subjects with artificial activation injected within real, manually-segmented anatomical ROIs (Section 3). Voxels within a radius of 8 mm from the anatomical centroid were defined as active (circled in red in Fig. 1(b)-(f)). Synthetic time courses of the active voxels were generated by convolving a box-car function, having the same stimulus timing as our experiment (Section 3), with a canonical HDR [1] and adding low frequency drifts and Gaussian noise. To simulate functional inter-subject variability, signal intensity of the active voxels was set to decrease exponentially as a function of distance from the activation centroid, whose location was randomly varied across subjects (Fig. 1(a)). This emulates the situation where true active regions highly overlap across subjects, but the apparent overlap appears




Fig. 1. Synthetic data results. (a) t -maps. (b) iGLM, (c) iCGMM, (d) iRW, (e) gGLM, and (f) GRW results shown. Blue dots = detected active voxels. Red circles = ground truth. For all SNR, (g) GRW achieved the highest DSC. Note iGLM resulted in a DSC of 0 at a SNR of 0.25. (h) GRW's DSC increased with number of subjects.

much less due to variation in location at which the fMRI signal concentrates. Maximal SNR was set as 0.5 in Fig. 1(a). For comparison, we also tested (i) GLM with spatial smoothing using a 8 mm FWHM Gaussian kernel and a threshold based on GRF theory for a p -value of 0.05 [1], (ii) CGMM, (iii) RW, and (iv) second level GLM, which involved taking the union of all subjects' ROI point sets to generate an ROI template, interpolating spatially smoothed β_j onto the template, applying GLM on the resulting β_j , thresholding based on GRF theory [1], and interpolating the thresholded group map back onto the subjects' native ROI space for comparison purposes. We refer to methods (i), (ii), (iii), and (iv) as individual GLM (iGLM), individual CGMM (iCGMM), individual RW (iRW), and group GLM (gGLM).

Qualitative results for the various contrasted methods are shown in Fig. 1(b)-(f). Only half of the subjects for one of the synthetic datasets are displayed due to space limitation. iGLM detected only a few active voxels, whereas iCGMM detected majority of the active voxels but also declared many false positives. Imposing intra-subject regularization using iRW reduced the number of false positives, but an ample amount remained due to lack of reliable intra-subject neighbors. Using gGLM detected all the active voxels, but also falsely declared many nearby voxels as active. Using GRW detected almost all the active voxels, while exerting much stricter control on false positives than gGLM. To quantify the performance, we computed the average Dice similarity coefficient (DSC) over the 500 synthetic datasets for a range of SNR.

$$DSC = \frac{2TP}{2TP + FP + FN}, \tag{6}$$

where TP, FP, and FN denote the number of true positives, false positives, and false negatives, respectively. As evident from Fig. 1(g), the poor DSC for iGLM, iCGMM, and iRW again illustrates that solely relying on single-subject information may be inadequate to obtain satisfactory segmentation at low SNR. However, naively incorporating group information can also be problematic as apparent from the gGLM results, where increasing SNR reduced DSC. This counter-intuitive result arose from the increase in signals leaking into the non-active voxels as a consequence of spatial smoothing, as required in gGLM for employing GRF theory and increasing functional overlap across subjects. In contrast, increasing SNR resulted in higher DSC for GRW, since GRW does not blindly smooth the data. Instead, intra- and inter-subject neighborhood information is adaptively controlled based on functional connectivity with contributions from dissimilar voxels de-weighted. As a result, GRW achieved the highest DSC for all SNR compared to the other examined methods. Also, adding group information improved performance over using iGLM, iCGMM, and iRW even at higher SNR, where reliable intra-subject information is available. Furthermore, increasing the number of subjects increased DSC as shown in Fig. 1(h).

3 Materials

After obtaining informed consent, fMRI data were collected from 10 Parkinson's disease (PD) patients off and on medication (4 men, 6 women, mean age 66 ± 8 years) and 10 healthy controls (3 men, 7 women, mean age 57.4 ± 14 years). Each subject used their right hand to squeeze a bulb with sufficient pressure to maintain a bar shown on a screen within an undulating pathway. The pathway remained straight during baseline periods and became sinusoidal at a frequency of 0.25 Hz (slow), 0.5 Hz (medium) or 0.75 Hz (fast) during time of stimulus. Each session lasted 260 s, alternating between baseline and stimulus of 20 s duration. 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 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 T1-weighted image consisting of 170 axial slices was also acquired. For each subject's data, slice timing and motion correction were performed using Brain Voyager's (Brain Innovation B.V.). Further motion correction was then applied using motion corrected independent component analysis (MCICA) [10]. The voxel time courses were high-pass filtered to account for temporal drifts and temporally whitened using an autoregressive AR(1) model. No whole-brain warping or spatial smoothing was performed. For testing our proposed method, we selected the left primary motor cortex (LM1), which is known to activate during right-hand movements. Delineation of LM1 was performed by an expert based on anatomical landmarks and guided by a neurological atlas. The segmented ROIs were resliced at fMRI resolution for extracting preprocessed voxel time courses within each ROI and non-rigidly aligned using "Coherent Point Drift", which has shown greater robustness to noise and outliers than conventional techniques such as iterative closest point [11].




Fig. 2. Real data results for 5 PD subjects pre- and post-medication and 5 controls. Blue in (b), (d) & (e) indicates detected active voxels. (a) t -maps. (b) iGLM detected only a few active voxels in the (c) hand region, whereas (d) gGLM detected the hand region but falsely included the hip and leg areas for PDpre. (e) GRW correctly identified the hand region in all subjects.

4 Results and Discussion

Results obtained with iGLM, gGLM, and GRW on real data are shown in Fig. 2. iCGMM and iRW results were similar to the synthetic case with many isolated false positives detected, and were thus excluded. Also, only results for 5 controls and 5 PD subjects during the fast condition are displayed due to space limitation, but consistent results were observed across all subjects. iGLM detected few active voxels in the hand region, whereas gGLM detected the hand region, but mistakenly included the hip and leg areas for PD pre-medication. In contrast, GRW correctly identified the hand region in all subjects without falsely declaring the hip and leg areas as active. In addition, the GRW results suggest a very interesting trend across subject groups. Specifically, PD pre-medication seemed to require recruiting a wider area of LM1, which normalized back to an extent similar to the controls upon medication. Such spatial focusing effect of levo-dopa medication has been observed in past studies [12], thus further confirming the validity of our results. This trend is also noticeable in the gGLM results, but whether levo-dopa truly over-normalized the extent of activation in PD is unclear, since the wider active region in control subjects could have simply arisen from gGLM’s weak control over false positives (Fig. 1(e) and Fig. 2(d) PD pre-medication). In contrast, the stronger control GRW has on false positives provides us more confidence with our findings, which greatly eases result interpretation.

5 Conclusion

We proposed a novel graph-theoretic method for enhancing single-subject fMRI activation detection. GRW expands the single-subject graph structure to include inter-subject neighbours, which enables group information to propagate into each subject's activation map without having to establish a one-to-one voxel correspondence. Also, the proposed GRW formulation permits joint estimation of all subjects' probabilistic activation maps with global optimality guaranteed. Superior detection power over standard techniques was shown on synthetic data for a range of SNR. When applied to real data, GRW consistently detected activation in regions implicated with the experimental task employed, whereas methods based on single-subject information failed. Our results thus demonstrate the effectiveness of incorporating group information for dealing with noisy fMRI data in single-subject analysis.

References

1. Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.B., Frith, C.D., Frackowiak, R.S.J.: Statistical Parametric Maps in Functional Imaging: A General Linear Approach. *Hum. Brain Mapp.* 2, 189–210 (1995)
2. Descombes, X., Kruggel, F., von Cramon, D.Y.: Spatio-Temporal fMRI Analysis Using Markov Random Fields. *IEEE Trans. Med. Imaging.* 17, 1028–1039 (1998)
3. Penny, W.D., Trujillo-Barreto, N.J., Friston, K.J.: Bayesian fMRI Time Series Analysis with Spatial Priors. *NeuroImage* 24, 350–362 (2005)
4. Bathula, D.R., Staib, L.H., Tagare, H.D., Papademetris, X., Schultz, R.T., Duncan, J.S.: Multi-Group Functional MRI Analysis Using Statistical Activation Priors. In: *MICCAI fMRI Data Analysis Workshop* (2009)
5. Vandenbroucke, M.W.G., Goekoop, R., Duschek, E.J.J., Netelenbos, J.C., Kuijter, J.P.A., Barkhof, F., Scheltens, P., Rombouts, S.A.R.B.: Interindividual Differences of Medial Temporal Lobe Activation during Encoding in an Elderly Population Studied by fMRI. *NeuroImage* 21, 173–180 (2004)
6. Ng, B., Abugharbieh, R., Hamarneh, G., McKeown, M.J.: Random Walker Based Estimation and Spatial Analysis of Probabilistic fMRI Activation Maps. In: *MICCAI fMRI Data Analysis Workshop* (2009)
7. Grady, L.: Multilabel Random Walker Image Segmentation Using Prior Models. In: *IEEE Comp. Soc. Conf. Comp. Vision Pattern Recog.*, vol. 1, pp. 763–770 (2005)
8. Stark, C.E., Okado, Y.: Making Memories without Trying: Medial Temporal Lobe Activity Associated with Incidental Memory Formation during Recognition. *J. Neurosci.* 23, 6748–6753 (2003)
9. Makni, S., Ciuciu, P., Idier, J., Poline, J.B.: Joint Detection-Estimation of Brain Activity in Functional MRI: A Multichannel Deconvolution Solution. *IEEE Trans. Sig. Processing.* 53, 3488–3502 (2005)
10. Liao, R., Krolik, J.L., McKeown, M.J.: An Information-theoretic Criterion for Intrasubject Alignment of fMRI Time Series: Motion Corrected Independent Component Analysis. *IEEE Trans. Med. Imaging.* 24(1), 29–44 (2005)
11. Myronenko, A., Song, X., Carreira-Perpinan, M.: Non-rigid Point Set Registration: Coherent Point Drift. In: *20th NIPS*, pp. 1009–1016. MIT Press, Cambridge (2006)
12. Monchi, O., Petrides, M., Doyon, J., Postuma, R.B., Worsley, K., Dagher, A.: The Neural Bases of Set-shifting Deficits in Parkinsons Disease. *J. Neurosci.* 24, 702–710 (2004)