

Functional Segmentation of fMRI Data Using Adaptive Non-negative Sparse PCA (ANSPCA)

Bernard Ng¹, Rafeef Abugharbieh¹, and Martin J. McKeown²

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. We propose a novel method for functional segmentation of fMRI data that incorporates multiple functional attributes such as activation effects and functional connectivity, under a single framework. Similar to PCA, our method exploits the structure of the correlation matrix but with neighborhood information adaptively integrated to encourage detection of spatially contiguous clusters yet without falsely pooling non-active voxels near the functional boundaries. In addition, our method adaptively combines PCA and replicator dynamics, which we show to be equivalent to non-negative sparse PCA, based on the sparsity of the activation pattern. We validate our method quantitatively on synthetic data and demonstrate that it outperforms methods including replicator dynamics, PCA, Gaussian mixture models, and general linear models. Furthermore, when applied to real fMRI data, our method successfully segmented the Brodmann area 6 into its known functional sub-regions, whereas other conventional methods that we examined failed to attain such delineation.

1 Introduction

Segmentation of functional magnetic resonance imaging (fMRI) data has by far been dominated by univariate analysis approaches. These methods examine each voxel in isolation, thus voxel interactions are ignored. To account for spatial correlations, Descombes et al. proposed modeling fMRI data using Markov random fields (MRF) [1], whereas Woolrich et al. proposed using a spatio-temporal autoregressive model [2]. Due to computational complexity, only local spatial correlations are typically modeled. Another approach for functionally segmenting the brain relies on identifying voxels with temporal responses similar to a pre-selected seed region [3], which directly models the correlations between spatially disconnected voxels. However, pre-specifying a seed region can be difficult. To automatically identify seed regions, Golland et al. proposed using Gaussian mixture models (GMM) under a hierarchical framework [4], which alleviates the need to pre-define the number of clusters. Instead, expert knowledge is exploited to determine the necessary level of decomposition. The limitations to these seed-based approaches are that the detected clusters may not be spatially contiguous and the detected voxels may not necessarily pertain to task-related responses. To detect spatially contiguous clusters, Woolrich et al. proposed

applying spatial mixture models on the activation statistics [5]. Similarly, Thirion et al. proposed using spectral clustering with activation statistics and physical distances between voxels as similarity metrics [6]. Using activation statistics encourages detection of task-related clusters, but unexpected responses will be neglected. Thus, jointly optimizing both functional connectivity and activation effects is desired.

To identify functional clusters without seeding, Friston et al. proposed using principal component analysis (PCA), which exploits the structure of the covariance matrix [7]. Thus, correlations between spatially disconnected voxels are also modeled. However, PCA often results in diffused weightings (i.e. spatial component maps with non-zero weights assigned to the majority of the voxels), which complicates cluster identification [8]. Another covariance-based method employing replicator dynamics was proposed by Lohmann et al. [9]. The authors noted that replicator dynamics has the interesting property of detecting clusters with *mutually* correlated voxels. This property, as we have shown previously [10], is in fact a result of the equivalence between replicator dynamics and non-negative sparse PCA. Thus, replicator dynamics can be used to handle the problem of diffused weightings in classical PCA. However, connections within a brain region tend to be dense [11], hence direct application of replicator dynamics may not be suitable for functional segmentation. A balance between diffused and sparse weightings is thus needed.

In this paper, we propose a new iterative method for functional segmentation of fMRI data that integrates the above desired characteristics, namely incorporation of activation effects, functional connectivity, neighborhood information, spatial continuity, and a balance between sparse and diffused weightings, under a single framework. Similar to PCA, the proposed method exploits the structure of the full correlation matrix, where correlations between spatially disconnected voxels are modeled. However, as opposed to computing voxel correlations in a pair-wise manner, our method incorporates neighborhood information into the correlation estimates. Employing a similar approach, Neumann et al. showed that incorporating neighborhood information encourages detection of spatially contiguous clusters, but may pool voxels near the functional boundaries into the clusters [12]. Therefore, we instead devise our method to *adaptively* incorporate neighborhood information based on activation dissimilarity, which we demonstrate in Section 3.4 to be an effective means of mitigating non-active voxels from being falsely pooled. To account for activation effects, we replace the diagonal of the correlation matrix (which is simply a set of ones) with the correlation between each voxel and the expected response. Voxels within the detected clusters will thus be highly correlated as well as activated. To draw a balance between diffused and sparse weightings, we propose combining the weight estimates from PCA and replicator dynamics in an iterative manner with the relative contributions adaptively adjusting to the sparsity of the activation pattern. We thus refer to this method as adaptive non-negative sparse PCA (ANSPCA).

2 Materials

After obtaining informed consent, fMRI data were collected from 10 Parkinson's disease (PD) patients on and off medication (4 men, 6 women, mean age 66 ± 8 years). Each subject used their right-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, and became sinusoidal at a frequency of 0.25 Hz, 0.5 Hz or 0.75 Hz during time of stimulus. Each run lasted 260s, alternating between baseline and stimulus of 20 s duration.

fMRI 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 timing correction. Further motion correction was performed using motion corrected independent component analysis (MCICA) [13]. The voxel time courses were high-pass filtered to account for temporal drifts and temporally whitened using an autoregressive AR1 model. No spatial warping or smoothing was performed. For testing our proposed method, we have selected Brodmann Area 6 (BA6), which is known to consist of multiple functional subdivisions, as the region of interest (ROI). Anatomical delineation of this ROI was performed by an expert based on anatomical landmarks and guided by a neurological atlas. 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

This section presents our new iterative method for functional segmentation of an ROI. A modified correlation matrix incorporating activation effects and neighborhood information is first estimated. ANSPCA is then applied to detect the most correlated *and* activated cluster. Subsequent clusters are detected by removing the previously identified voxels from the modified correlation matrix, and repeating the procedure.

3.1 Modified Correlation Matrix

Let N_i and N_j be the neighborhood of voxels i and j (including voxels i and j). We compute a modified correlation estimate between voxels i and j , $C(i, j)$, as follows:

$$C(i, j) = \begin{cases} \text{corr}(I_{N_i}(t), I_{N_j}(t)), & i \neq j \\ \text{corr}(I_{N_i}(t), \text{ref}(t)), & i = j \end{cases}, \quad (1)$$

$$I_{N_m}(t) = \sum_p w_{mp} I_p(t), \quad w_{mp} = \exp(-|\Delta_{mp}| d_{mp}^2 / 2), \quad \text{st} \quad \sum_p w_{mp} = 1, \quad p \in N_m \quad (2)$$

where the matrix $I_p(t)$ consists of voxel time courses belonging to N_m ($m \in \{i, j\}$) along the rows. d_{mp} is the Euclidean distance between voxels m and p . Δ_{mp} is the difference in activation statistics (t-values) between voxels m and p . The t-values are estimated by applying a general linear model (GLM) to each voxel with a column of ones and a box-car convolved with the hemodynamic response, $\text{ref}(t)$, as regressors.

We note that naively incorporating neighborhood information may pool voxels near functional boundaries into the clusters. Therefore, we have specifically designed w_{kp} to adaptively reduce the influence from neighboring voxels with dissimilar activation level, which we demonstrate in Section 3.4 to be an effective way of moderating non-active voxels from being mistakenly declared as part of a functional cluster.

3.2 Replicator Dynamics

Replicator dynamics is a well known concept that originated from theoretical biology for modeling the evolution of different species. In our context, each voxel corresponds to a species with its fitness measured by its correlations to other voxels. Let w_{RD} be a weight vector with the i^{th} element representing the degree of which the i^{th} voxel belongs to the most correlated cluster. w_{RD} can be estimated by [14]:

$$w_{RD}(k+1) = \frac{w_{RD}(k) \cdot C w_{RD}(k)}{w_{RD}^T(k) C w_{RD}(k)}, \quad (3)$$

where \cdot represents element-wise product and k is the iteration number. Based on the fundamental theorem of natural selection [14], w_{RD} is guaranteed to converge provided C is real-value, non-negative, and symmetric. Since voxels belonging to the same cluster will presumably display positive correlations, we null out the negative elements to ensure C is non-negative [9]. Restricting C to be non-negative enforces w_{RD} to be non-negative. Also, (4) constrains the elements of w_{RD} to sum to one. Moreover, (4) maximizes the same objective function as PCA, i.e. $w_{RD}^T C w_{RD}$. Thus, replicator dynamics is in fact a solution to the non-negative sparse PCA problem [8]:

$$\arg \max_{w_{RD}} w_{RD}^T C w_{RD}, \quad st \quad \sum_i w_{RD}^i \leq K, w_{RD}^i \geq 0. \quad (4)$$

Prior studies have noted that replicator dynamics has the desirable property of detecting clusters with mutually correlated voxels [9], [12]. This property can actually be explained by the fact that imposing sparsity given limited weights (i.e. $\sum w_{RD}^i = 1$) encourages weights to be assigned to mutually correlated voxels [10].

3.3 Adaptive Non-negative Sparse PCA

Let C be the modified correlation matrix as described in Section 3.1 and let w be a weight vector with the i^{th} element corresponding to the degree of which the i^{th} voxel belongs to the most correlated and activated cluster. To adaptively adjust the sparsity of w , we propose to iteratively estimate w as follows:

$$w(k) = \gamma \cdot w_{PCA}(k) + (1 - \gamma) \cdot w_{RD}(k), \quad (5)$$

$$w_{PCA}(k) = C w(k-1), \quad st \quad \|w_{PCA}(k)\|_1 = 1 \quad (6)$$

$$w_{RD}(k) = w(k-1) \cdot C w(k-1) / (w^T(k-1) C w(k-1)) \quad (7)$$

where $w_{PCA}(k)$ is the PCA estimate of $w(k)$ (i.e. using the Power method), but with $\|w_{PCA}(k)\|_1=1$ to ensure that $\|w(k)\|_1=1$ as required for computing $w_{RD}(k)$. γ is the percentage of activated voxels estimated as the number of voxels with t-values above a

user-specified threshold over the total number of voxels in the ROI. The typical t-threshold of 1.96 is used. We note that the t-threshold only serves to estimate γ , and not to remove voxels with t-values below the threshold. In fact, voxels with t-values falling below the t-threshold but functionally connected to their neighbors can still be declared as part of a functional cluster. Using the percentage of activated voxels as γ is particularly suitable for drawing a balance between sparse weighting, $w_{RD}(k)$, and diffused weighting, $w_{PCA}(k)$, since a sparse activation pattern will result in a low percentage of activated voxels, which place greater influence from $w_{RD}(k)$ in estimating w and vice versa. To avoid bias, we initialize $w(0)$ to $1/N_r$, where N_r is the number of voxels within the ROI. Upon convergence, elements of w corresponding to the most correlated and activated cluster will rise above $1/N_r$, but the detected voxels may not form a spatially connected patch. Therefore, we perform connected component analysis to first group the detected voxels into spatially connected clusters. We then find the cluster with the highest $w^T C w$, remove the other clusters from C , and reapply the above procedure until no spatially disconnected clusters are detected. The resulting cluster upon convergence will thus consist of voxels that are highly activated, functionally connected, and spatially connected. To identify subsequent clusters, we remove only voxels in the previously detected functional clusters from C and repeat the procedure above.

3.4 Empirical Evaluations

To test our proposed method on data with ground truth, we generated 1,000 synthetic datasets with simulated activation patterns consisting of two clusters that were one voxel apart (Fig.1). Also, the signal intensity of the activated voxels was set to decrease as a function of their distances from cluster centroids. The time courses of the activated voxels in the larger cluster were generated by convolving a box-car function having the same stimulus timing as in our experiment with the hemodynamic response and adding Gaussian noise. The smaller cluster was generated in a similar manner but with the box-car delayed by 2 seconds.

For comparisons with the state-of-the-art, we also tested the following methods: replicator dynamics [9] with Pearson's correlation, PCA [7], GMM [4] assuming two clusters and background, and GLM with Gaussian spatial smoothing and a threshold estimated from Gaussian random field (GRF) theory for an uncorrected p-value of 0.05. For PCA, we renormalized the PCs such that $\|w_{PCA}\|_1=1$ and used the same threshold as ANSPCA (i.e. $1/N_r$). Also, only the first PC was used, since the second PC mainly detected non-active voxels (squares in Fig. 1b). Fig. 1 contains the synthetic data results with the average true positive rate (TP) and false positive rate (FP) indicated. We note that TP and FP are computed only based on whether the voxels are correctly labeled as active, and not based on the cluster labels.

Replicator dynamics (Fig. 1a) did not falsely declare any non-activated voxels as active (FP = 0.00), but many activated voxels were missed (TP = 0.15). In contrast, PCA (Fig. 1b) detected most of the activated voxels with its first PC (TP = 0.83), but also detected many non-active voxels (FP = 0.04). Neither of these methods was able to separate the two clusters. GMM behaved similarly to PCA with one of its mixtures encompassing both clusters, and the other mixture including mainly non-active voxels (Fig. 1c). Hence, a high FP of 0.13 with only a TP of 0.77 was obtained.

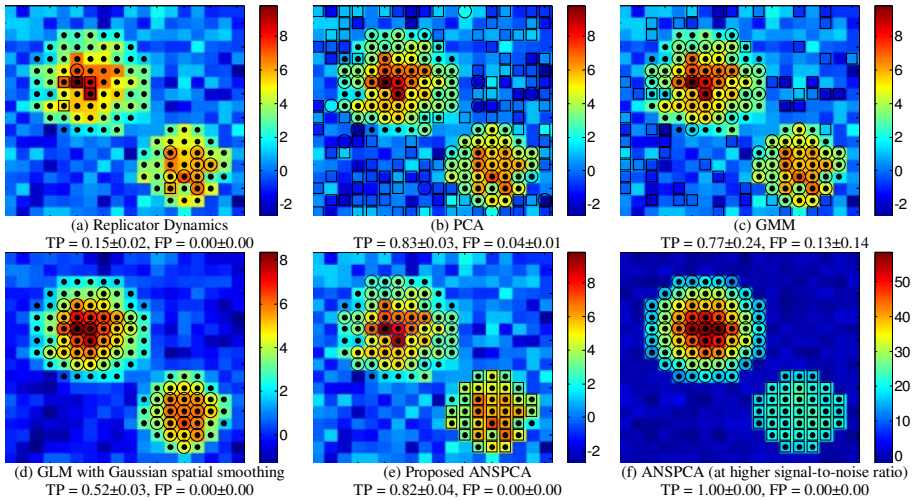


Fig. 1. Synthetic data results. t-map estimated using GLM with “dots” indicating the ground truth activated voxels. Voxels with a circle (square) correspond to the first (second) cluster detected. Note how ANSPCA was able to separate the two clusters and achieve a FP of 0.00.




Fig. 2. Real fMRI data results. (a) & (d) Left BA6 t-map of a PD subject before and after medication. (b) & (e) Proposed ANSPCA segments the BA6 into its known sub-regions, whereas (c) & (f) GLM appears to over-divide the BA6.

GLM resulted in a FP of 0.00 with a TP of 0.52. Examining Fig. 1d, post-processing the voxels detected by GLM with connected component analysis could have separated the two clusters. Using ANSPCA, we were able to separate the two clusters as shown in Fig. 1e. Compared to PCA, with a mere decrease of 0.01 in TP, ANSPCA achieved a FP of 0.00 without any non-active voxels near the functional boundaries being falsely

declared. Being able to incorporate neighborhood information without falsely pooling the non-active boundary voxels is especially important for functionally segmenting ROIs with unclear sub-region boundaries such as the BA6 [15]. We thus further tested our method by varying the signal-to-noise ratio (SNR). For the case shown in Fig. 1f, we increased the SNR, which presumably would increase the correlation between the active and non-active boundary voxels if their information is mistakenly pooled in the correlation estimates. This increase in correlation would result in a higher chance of falsely declaring the non-active boundary voxels as part of the functional clusters [12]. Yet, as shown in Fig. 1f, ANSPCA was able to attain a TP of 1 without including any of the non-active voxels. We note that by incorporating neighborhood information and functional connectivity in addition to activation effects, ANSPCA was able to detect many voxels with t-values below the t-threshold determined from GRF. Also, upon detecting the larger cluster and removing the corresponding voxels from C , ANSPCA was able to adapt to the sparser activation pattern without falsely declaring voxels near the functional boundaries as part of the second cluster.

4 Results and Discussion

Results obtained by applying ANSPCA on the BA6 are shown in Fig. 2. We only included the left BA6 results for an exemplar PD subject due to space limitation. Nevertheless, ANSPCA did correctly separate the SMA and PM in all subjects, although the SMA and PMd appeared partly joined in 3 subjects. For comparison, we applied replicator dynamics, PCA, GMM, and GLM to the data. However, except for GLM, the results obtained were similar to that in the synthetic data experiments with no spatially contiguous clusters identified. Thus, we only included the GLM results in favour of space. Also, the t-threshold estimated using GRF with Gaussian spatial smoothing of 8mm FWHM was found to be too stringent. Hence, we instead present the thresholded t-maps for an uncorrected p-value of 0.05 with FDR correction.

For PD predrug, despite the unclear functional boundaries (Fig. 2a), ANSPCA was able to delineate the left BA6 into its known functional sub-regions, namely pre-SMA, SMA proper, dorsal premotor cortex (PMd), and ventral premotor cortex (PMv) [15]. We argue that this delineation was attained by the additional functional connectivity information included in our proposed method, which was not modeled in the activation statistics. Also, incorporating neighborhood information enabled ANSPCA to detect PMd and PMv as two spatially contiguous clusters, whereas GLM over-divided PMd and PMv into multiple sub-regions. For PD postdrug, ANSPCA was again able to delineate the left BA6 into its constituent sub-regions. Interestingly, our results suggest that the extent of activation within the pre-SMA (Fig. 2b) reduced upon medication (Fig. 2e). This “focusing” effect conforms to prior findings in computational model studies, where increased dopamine level was found to be associated with more focused activation patterns. In contrast, GLM split the SMA proper into multiple pieces with islands of activation scattered across the BA6.

5 Conclusions

We proposed a novel method that integrates multiple functional attributes such as activation effects, functional connectivity, and neighborhood information under a

single framework for functionally segmenting fMRI data. By adaptively combining PCA and replicator dynamics, our method facilitates functional segmentation of activation patterns with various degree of sparsity. Applying our method to synthetic data outperformed all other examined methods including replicator dynamics, PCA, GMM, and GLM. When applied to real data, the integration of the various functional attributes enabled our method to segment the BA6 into its constituent functional sub-regions, whereas other examined methods failed to attain such delineation.

References

1. Descombes, X., Kruggel, F., von Cramon, D.Y.: Spatio-Temporal fMRI Analysis Using Markov Random Fields. *Trans. Med. Imaging* 17, 1028–1039 (1998)
2. Woolrich, M.W., Jenkinson, M., Brady, J.M., Smith, S.M.: Fully Bayesian Spatio-temporal Modeling of fMRI Data. *Trans. Med. Imaging* 28, 213–231 (2004)
3. Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S.: Functional Connectivity in the Motor Cortex of Resting Human Brain Using Echo-planar MRI. *Magn. Reson. Med.* 34, 537–541 (1995)
4. Golland, P., Golland, Y., Malach, R.: Detection of Spatial Activation Patterns as Unsupervised Segmentation of fMRI Data. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) *MICCAI 2007, Part I. LNCS*, vol. 4791, pp. 110–118. Springer, Heidelberg (2007)
5. Woolrich, M.W., Behrens, T.E., Beckmann, C.F., Smith, S.M.: Mixture Models with Adaptive Spatial Regularization for Segmentation with an Application to fMRI Data. *Trans. Med. Imaging* 24, 1–11 (2005)
6. Thirion, B., Flandin, G., Pinel, P., Roche, A., Ciuciu, P., Poline, J.B.: Dealing with the Shortcomings of Spatial Normalization: Multi-subject Parcellation of fMRI Datasets. *Hum. Brain Mapp.* 27, 678–693 (2006)
7. Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S.: Functional Connectivity: The Principal Component Analysis of Large (PET) Data Sets. *J. Cereb. Blood Flow Metab.* 13, 5–14 (1993)
8. Zass, R., Shashua, A.: Nonnegative Sparse PCA. In: *Advances in Neural Information Processing Systems*, pp. 1561–1568 (2006)
9. Lohmann, G., Bohn, S.: Using Replicator Dynamics for Analyzing fMRI Data of the Human Brain. *Trans. Med. Imaging* 21, 485–492 (2002)
10. Ng, B., Abugharbieh, R., McKeown, M.J.: Discovering Sparse Functional Brain Networks Using Group Replicator Dynamics (GRD). In: 21st Int. Conf. IPMI (accepted 2009)
11. Bassett, D.S., Bullmore, E.: Small-World Brain Networks. *Neuroscientist* 12, 512–523 (2006)
12. Neumann, J., von Cramon, D.Y., Forstmann, B.U., Zysset, S., Lohmann, G.: The Parcellation of Cortical Areas Using Replicator Dynamics in fMRI. *Neuroimage* 32, 208–219 (2006)
13. 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)
14. Schuster, P., Sigmund, K.: Replicator Dynamics. *J. Theor. Biol.* 100, 533–538 (1983)
15. Mayka, M.A., Corcos, D.M., Leurgans, S.E., Vaillancourt, D.E.: Three-dimensional Locations and Boundaries of Motor and Premotor Cortices as Defined by Functional Brain Imaging: A Meta-analysis. *Neuroimage* 31, 1452–1474 (2006)