

# Spatial Patterns and Functional Profiles for Discovering Structure in fMRI Data

(Invited Paper)

Polina Golland, Danial Lashkari and Archana Venkataraman

Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA  
{polina,danial,pega85}@csail.mit.edu

**Abstract**—We explore unsupervised, hypothesis-free methods for fMRI analysis in two different types of experiments. First, we employ clustering to identify large-scale functionally homogeneous systems. We formulate a generative mixture model, derive the EM algorithm and apply it to delineate functional systems. We also investigate spectral clustering in application to this problem and demonstrate that both methods give rise to similar partitions of the brain based on resting state fMRI data. Second, we demonstrate how to extend this approach to include information about the experimental protocol. Specifically, we formulate a mixture model in the space of possible profiles of brain response to stimuli. In both applications, our methods confirm previously known results in brain mapping and point to new research directions for exploratory analysis of fMRI data.

## I. INTRODUCTION

With advancements in the field of functional MRI (fMRI), the functional imaging studies have moved from simple questions of localizing brain regions strongly modulated by the experimental protocol to more complex problems that call for novel analysis techniques. Here we present our work in developing methods for exploratory fMRI analysis. This paper provides an overview of several different projects; we refer the readers to previously published papers [16, 22, 31] for further details on the methods and experimental validation.

### A. Spatial Patterns of co-Activation

Functional connectivity analysis [4, 5, 9] is widely used in fMRI studies to detect and characterize large networks that co-activate with a user-selected ‘seed’ region of interest. The method typically uses the Pearson correlation coefficient as a measure of similarity between a time course of an individual voxel and the mean time course of the selected seed region. Since no alternative hypothesis for correlation values is formulated, the user must select a threshold, or significance level, for rejecting the null

hypothesis that assumes zero correlation with the seed time course. This approach is highly useful for analyzing co-activation patterns in fMRI data. However, high variability of inter-voxel correlation values across scans presents a serious challenge for integrating results across different runs or across subjects. Furthermore, in some studies it is unclear how to select the seed region. Instead, we would prefer to discover the interesting ‘seeds’ and the associated networks in an unsupervised way. In this work, we formulate the problem of identifying functionally homogeneous systems as clustering and explore mixture modeling and spectral clustering in application to functional connectivity analysis.

Our approach is based on a model that parcelates the brain into disjoint sub-regions. Principal Component Analysis (PCA) and Independent Component Analysis (ICA) [3] provide an alternative model of functional connectivity that treats the data as a linear combination of spatial maps with associated time courses. ICA, PCA and clustering have been extensively explored in the context of regression-based detection [1, 2, 8, 10, 12, 17, 24, 26, 29], as a way to identify and remove noise-induced components of the data. Application of clustering in fMRI analysis has traditionally focused on grouping voxels into small, functionally homogeneous regions in paradigm-based studies [10, 17, 29]. Recently, clustering was also demonstrated in application to full-brain scans in resting state fMRI experiments [6, 30], revealing anatomically meaningful regions of high functional coherency. In this work, we aim to construct top-down representations of global patterns of activation spanning the entire brain.

### B. Activation Profiles in Multi-Category Experiments

One natural extension of the proposed method is to incorporate the information about the experimental protocol into the analysis. Our motivation for this work comes from fMRI studies of category selectivity in the visual cortex where subjects are presented with images from several categories of visual stimuli. Using hypothesis-driven localization methods [14], investigators discovered regions with specific category selectivity that consistently

This work was supported in part by the National Alliance for Medical Image Analysis (NIH NIBIB NAMD U54-EB005149), the Neuroimaging Analysis Center (NIH NCRR NAC P41-RR13218), the NSF CAREER Grant 0642971 and the MIT McGovern Institute Neurotechnology Program. A. Venkataraman is supported by the National Defense Science and Engineering Graduate Fellowship (NDSEG).

appear in most subjects. For instance, the well-known fusiform face area (FFA) is associated with higher response to faces when compared to other visual stimuli. In addition, the parahippocampal place area (PPA), and extrastriate body area (EBA) exhibit high selectivity for places, and body parts, respectively [21].

With the increasing number of conditions or tasks, the number of potential hypotheses grows exponentially. To address this challenge, we demonstrate an exploratory clustering method that operates in the space of all possible activation profiles and identifies robust patterns of activation in experiments with a large number of stimulus categories.

Spatial consistency of the localization maps across subjects has traditionally served as evidence for the validity of the corresponding hypothesis. Rather than rely on spatial consistency, we employ functional consistency across subjects to evaluate the robustness, and therefore relevance, of the detected profiles. Thus, we obtain a fully functional characterization of the data.

This work bridges the gap between exploratory methods and hypothesis-driven localization approaches. Most exploratory techniques work on the raw fMRI time courses and use clustering or ICA to decompose the data into a set of distinct time courses of interest and their localization maps. However, this framework offers no clear mechanism to characterize the relationship between the multitude of experimental conditions and the noisy representative time courses identified through such analysis. Some exploratory methods use the information on the experimental protocol indirectly to define a measure of similarity between voxels, effectively projecting the original high dimensional time courses onto a low dimensional feature space. Clustering is then performed in the new space [17, 18, 29]. However, these methods mainly focus on identifying the “active” voxels in simple experiments. In contrast, we employ mixture modeling in the space of activation profiles that explicitly capture the effects of the experimental protocol on the observed fMRI time courses. This approach enables hypothesis-free search for robust activation profiles.

Our experimental results confirm many different findings previously established in fMRI studies. The results also suggest future directions of research in exploratory analysis of fMRI signals.

## II. METHODS

Given fMRI time courses  $\{\mathbf{y}_1, \dots, \mathbf{y}_N\}$  of  $N$  voxels, our goal is to find a partition of the data set into  $K$  functionally homogeneous sets. In this section, we review two variants of clustering algorithms that achieve such partitioning. We also present a modified algorithm that incorporates information about the experimental protocol into the clustering framework. Finally, we discuss our approach to group analysis.

### A. Mixture Modeling for Functional Connectivity

We model fMRI time courses as noisy instantiations of one of the  $K$  representative hypotheses, or time courses,  $\mathbf{m}_1, \dots, \mathbf{m}_K$ . This assumption gives rise to the generative mixture model [25]:

$$p(\mathbf{y}) = \sum_{k=1}^K \lambda_k p_k(\mathbf{y}) = \sum_{k=1}^K \lambda_k \mathcal{N}(\mathbf{y}; \mathbf{m}_k, \Sigma_k), \quad (1)$$

where  $p_k(\mathbf{y})$  is the class-conditional likelihood of the signal in system  $k$ , and  $\lambda_k$  is the prior probability that a voxel belongs to system  $k$ . Following the commonly used approach in fMRI analysis, we model the class-conditional densities as normal distributions centered around the system mean time course. The high dimensionality of the fMRI data makes modeling a full covariance matrix impractical. We take a simpler approach of restricting the covariance matrix to be diagonal, i.e., modeling variance components only, and note that the mixture model estimation can be straightforwardly extended to include more sophisticated signal dynamics.

We use the EM algorithm [7] to estimate the model parameters. We let  $\hat{p}_{nk}$  be the posterior probability that voxel  $n$  belongs to system  $k$  and arrive at the familiar update rules:

E-step:

$$\hat{p}_{nk}^{(\tau)} = \frac{\lambda_k^{(\tau)} \mathcal{N}(\mathbf{y}_n; \mathbf{m}_k^{(\tau)}, \Sigma_k^{(\tau)})}{\sum_{k'} \lambda_{k'}^{(\tau)} \mathcal{N}(\mathbf{y}_n; \mathbf{m}_{k'}^{(\tau)}, \Sigma_{k'}^{(\tau)})}, \quad (2)$$

M-step:

$$\lambda_k^{(\tau+1)} = \frac{1}{N} \sum_n \hat{p}_{nk}^{(\tau)}, \quad \mathbf{m}_k^{(\tau+1)} = \frac{\sum_n \hat{p}_{nk}^{(\tau)} \mathbf{y}_n}{\sum_n \hat{p}_{nk}^{(\tau)}}, \quad (3)$$

$$\Sigma_k^{(\tau+1)} = \frac{\sum_{n,t} \hat{p}_{nk}^{(\tau)} \left( \mathbf{y}_n(t) - \mathbf{m}_k^{(\tau+1)}(t) \right)^2}{T \sum_n \hat{p}_{nk}^{(\tau)}}, \quad (4)$$

where  $T$  is the length of the time courses and  $\{\hat{p}_{nk}^{(\tau)}, \lambda_k^{(\tau)}, \mathbf{m}_k^{(\tau)}, \Sigma_k^{(\tau)}\}$  are the estimates at step  $\tau$  of the algorithm. We initialize the algorithm by randomly selecting  $K$  time courses from the original data set as an initial guess for the cluster means. To ensure that the algorithm properly explores the non-convex space of the solutions, we perform multiple runs of the algorithm using different random initializations and select the solution that achieves the maximum likelihood of the data. When the algorithm converges, the estimates  $\{\hat{p}_{nk}\}_{n=1}^N$  represent probabilistic segmentation for system  $k$ .

The well known k-means clustering algorithm replaces the probabilistic assignments  $\hat{p}_{nk}$  with hard binary assignments in each step of the algorithm. This variant can be shown to minimize the sum of  $L^2$  distances between the time courses and their corresponding cluster means.

Since for functional connectivity analysis, the time courses are typically normalized to have zero mean and unit variance, minimizing  $L^2$  distance is equivalent to maximizing correlation. The k-means algorithm can therefore be viewed as a natural data-driven extension of the seed-based connectivity analysis.

### B. Spectral Clustering for Functional Connectivity

To investigate the sensitivity of the resulting partitions to the assumptions made by the generative model in Equation (1), we also applied spectral clustering [27] to the same data. For spectral clustering, we construct a pairwise affinity matrix  $W$ ,

$$W_{ij} = e^{-d^2(\mathbf{y}_i, \mathbf{y}_j)/2\sigma^2} \quad (5)$$

where  $d^2(\mathbf{y}_i, \mathbf{y}_j)$  is the distance between the time courses  $\mathbf{y}_i$  and  $\mathbf{y}_j$ , and  $\sigma^2$  is the kernel width parameter. If we use  $L^2$  distance, Equation (5) corresponds to the standard Gaussian kernel often used in spectral clustering.

Given the affinity matrix  $W$ , spectral clustering seeks a partitioning of the voxel set based on a spectral decomposition of  $W$ . In this work, we use the Normalized Cut variant of spectral clustering [28]. We construct a continuous relaxation of the original combinatorial optimization, which leads to the eigenvalue problem

$$D^{-1/2}WD^{-1/2}\mathbf{v} = \lambda\mathbf{v} \quad (6)$$

where  $D$  is a diagonal matrix such that  $D_{ii} = \sum_j W_{ij}$ . We define a vector of row sums  $\mathbf{d}$ , i.e.,  $d_i = D_{ii}$ . The left and right multiplications by  $D^{-1/2}$  in Equation (6) correspond to a symmetric normalization of  $W$  where each entry  $W_{ij}$  is divided by  $\sqrt{d_i d_j}$ . To find the partition of the original data, we cluster the rows of the matrix

$$V = [D^{-1/2}\mathbf{v}_1 \dots D^{-1/2}\mathbf{v}_{K+1}]. \quad (7)$$

A typical fMRI data set contains hundreds of thousands of voxels, leading to extremely large eigenvector problems in spectral clustering. To overcome this difficulty, we adopt the standard approximation of the leading eigenvalues and eigenvectors of the matrix  $D^{-1/2}WD^{-1/2}$  via the Nyström Method [11].

Given a set of  $N$  time courses, we first select  $M \ll N$  samples at random. The  $N \times N$  affinity matrix  $W$  can be represented as

$$W = \begin{bmatrix} A & B \\ B^T & C \end{bmatrix}, \quad (8)$$

where  $A$  is the  $M \times M$  matrix of affinities between the randomly selected samples, and  $B$  is the  $M \times (N - M)$  matrix of affinities between the random samples and the remaining data points.  $C$  is a large  $(N - M) \times (N - M)$  matrix of remaining affinities that we want to avoid computing.

We first normalize  $W$  by the matrix  $D^{-1/2}$ . As shown in [11], we can approximate the row sum vector  $\mathbf{d}$  via

$$\hat{\mathbf{d}} = \begin{bmatrix} A\mathbf{1}_M + B\mathbf{1}_{N-M} \\ B^T\mathbf{1}_M + B^T A^{-1}B\mathbf{1}_{N-M} \end{bmatrix} \quad (9)$$

where  $\mathbf{1}_M$  denotes an all-ones vector of length  $M$ . The normalized matrices  $\tilde{A}$  and  $\tilde{B}$  are given by

$$\tilde{A}_{ij} = \frac{A_{ij}}{\sqrt{d_i d_j}} \quad \tilde{B}_{ij} = \frac{B_{ij}}{\sqrt{d_i d_{j+M}}}$$

The Nyström Method approximates the eigenvectors of  $\tilde{W} = D^{-1/2}WD^{-1/2}$  using  $\tilde{A}$  and  $\tilde{B}$ . Let  $U\Lambda U^T$  denote the SVD of the  $M \times M$  symmetric matrix  $\tilde{A} + \tilde{A}^{-1/2}\tilde{B}\tilde{B}^T\tilde{A}^{-1/2}$ . The  $M$  leading eigenvectors of  $\tilde{W}$  are then computed as

$$\tilde{V} = \begin{bmatrix} \tilde{A} \\ \tilde{B}^T \end{bmatrix} \tilde{A}^{-1/2}U\Lambda^{-1/2} \quad (10)$$

The data set partitioning is obtained by clustering rows of the matrix  $\hat{D}^{-1/2}\tilde{V}$ .

### C. Group Analysis of Spatial Patterns of co-Activation

Any permutation of cluster indices leaves the clustering cost function unchanged. This symmetry leads to inherent ambiguity in cluster label assignments. However, a correspondence among labels assigned to each voxel across runs is required for group-wise analysis of the cluster patterns. Unfortunately, a naïve approach reduces to a combinatorial search over all possible labeling combinations. In this work, we employ a greedy algorithm to match cluster labels across subjects. Given two clusterings, we perform an iterative label matching procedure. In each iteration, we assign the next available cluster index to one cluster in each clustering in such a way as to maximize the number of consistently labeled voxels. We then remove the newly labeled voxels from further consideration. The algorithm stops when all voxels have been assigned a new cluster index. While this approach may not yield the globally optimal alignment across runs given arbitrary data, the cluster patterns in our application are often similar enough for this method to accurately match corresponding systems across subjects and across experiments.

### D. Mixture Modeling in the Space of Activation Profiles

We define an activation profile to be a vector whose components describe selectivity to different categories. Given a set of raw fMRI time courses, we apply a General Linear Model (GLM) analysis [14] at each voxel and form a vector containing the estimated regression coefficients of the experiment stimuli. The norm of these vectors is mainly a byproduct of irrelevant variables such as distance from major vessels or the overall magnitude of response to the type of stimuli used in the experiment. Moreover, it is widely accepted that only relative values

of responses are important in characterizing selectivity to different stimuli. To remove the effect of the magnitude of activation while preserving the relative strength of activation across categories, we normalize the activation profiles to be unit length vectors. With  $D$  categories of visual stimuli present in the experiment, our space of activation profiles is a unit sphere  $S^{D-1}$  in a  $D$ -dimensional space. When represented in the space of activation profiles, an fMRI data set becomes a population of vectors on a unit sphere. Naturally, the interesting patterns of selectivity in this population correspond to the directions with highest concentration of data points around them. It is easy to see that finding these patterns can be thought of as clustering the activation profiles and estimating the corresponding cluster means.

In this case, our data  $\{\mathbf{y}_n\}_{n=1}^N$  is a set of activation profiles from  $N$  voxels on a  $S^{D-1}$  sphere. We devise a mixture model based on the inner product as the natural measure of similarity for unit vectors. Here we choose von Mises-Fisher distribution [23] to model the class-conditional likelihood of the data:

$$p_k(\mathbf{y}) = C_D(\mu) e^{\mu \langle \mathbf{m}_k, \mathbf{y} \rangle} \quad (11)$$

where  $\langle \cdot, \cdot \rangle$  denotes the inner product of two vectors and the normalizing constant

$$C_D(\mu) = \frac{\mu^{D/2-1}}{(2\pi)^{D/2} I_{D/2-1}(\mu)} \quad (12)$$

is defined in terms of the  $\gamma$ -th order modified Bessel function of the first kind  $I_\gamma$ . The concentration parameter  $\mu$  controls the concentration of the distribution around the mean direction  $\mathbf{m}$ . In general, mixture components can have distinct concentration parameters but in this work, we use the same parameter for all the clusters to ensure a more robust estimation.

With a bit of algebra, we derive the EM algorithm for this likelihood model:

E-step:

$$\hat{p}_{nk}^{(\tau)} = \frac{\lambda_k^{(\tau)} e^{\mu^{(\tau)} \langle \mathbf{m}_k^{(\tau)}, \mathbf{y}_n \rangle}}{\sum_{k'} \lambda_{k'}^{(\tau)} e^{\mu^{(\tau)} \langle \mathbf{m}_{k'}^{(\tau)}, \mathbf{y}_n \rangle}} \quad (13)$$

M-step:

$$\lambda_k^{(\tau+1)} = \frac{1}{N} \sum_n \hat{p}_{nk}^{(\tau)}, \quad \mathbf{m}_k^{(\tau+1)} \propto \sum_n \hat{p}_{nk}^{(\tau)} \mathbf{y}_n, \quad (14)$$

$$\mu^{(\tau+1)} : \frac{I_{D/2}(\mu^{(\tau+1)})}{I_{D/2-1}(\mu^{(\tau+1)})} = \frac{1}{N} \sum_{n,k} \hat{p}_{nk}^{(\tau)} \langle \mathbf{m}_k^{(\tau)}, \mathbf{y}_n \rangle. \quad (15)$$

We normalize vectors  $\mathbf{m}_k^{(\tau+1)}$  in each step to unit length. The nonlinear equation (15) for the estimation of  $\mu^{(\tau+1)}$  can be solved with a simple zero-finding algorithm.

### E. Group Analysis in the Space of Activation Profiles

Since we aim to discover activation profiles that robustly appear in brain response to stimuli, it is reasonable to assume that the space of activation profiles is shared across subjects. For the group-wise analysis based on  $S$  subjects, we let  $\mathbf{y}_n^s$  be the time course of voxel  $n$  in the data from subject  $s$ ,  $s \in \{1, \dots, S\}$ . We model each voxel  $\mathbf{y}_n^s$  as an independent sample from the same mixture distribution and combine the data from several subjects to perform our analysis across subjects. Applying clustering to the group data, the resulting posterior probabilities  $\{\hat{p}_{nk}^s\}_{n=1}^{N_s}$  define the localization map of activation profile  $k$  in subject  $s$ .

To assess the stability of the detected clusters across subjects, we estimate the group-wise model for the pooled data and the subject-specific model for each subject separately. For each profile in the group-wise model, we find the closest profile in each subject-specific model, using correlation as a measure of similarity. We then use the average correlation of the subject-specific profiles with the group-wise profile as a consistency score for that profile. The consistency score provides an indication of how persistent the detected activation profile is in the population, representing an initial step towards a fully functional model of consistency across subjects that does not assume perfect spatial alignment.

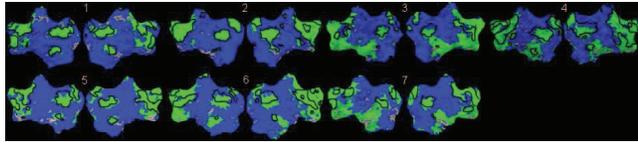
## III. EXPERIMENTAL RESULTS

We demonstrate our approach in three different fMRI studies. Each study focused on a different aspect of functional organization of the brain and allowed us to explore the potential of the algorithms described in the previous section.

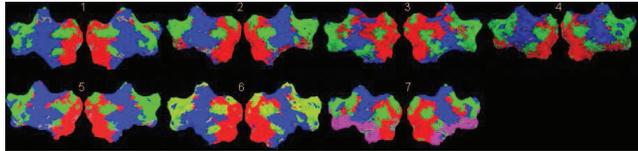
### A. Spatial co-Activation Patterns in Diverse fMRI Data

This study of functional connectivity included 7 subjects. We used previously collected fMRI scans in a large set of visual experiments, from simple localizer tasks to viewing continuous stimuli (movies), as well as a resting state scan. The total amount of fMRI data per subject was close to one hour. In the movie viewing experiments, the traditional seed-based functional connectivity analysis revealed two systems. The first system contained sensory-motor cortices and was strongly correlated with the seed region in the visual cortex. The second, ‘default’ system showed little correlation with the visual seed, but exhibited high intra-system correlation [15].

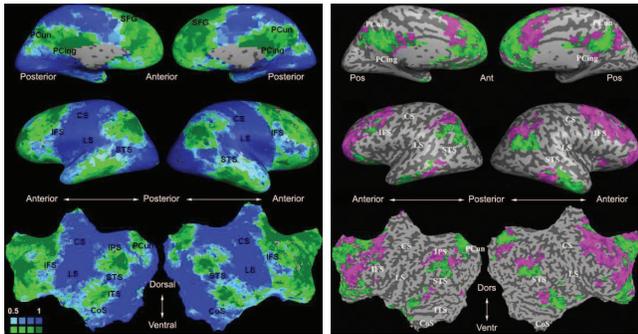
Figure 1a shows the 2-system partition determined through clustering in each subject independently of all others, based on all data available for the subjects. It also displays the boundaries of the default system determined through the traditional seed selection, showing good agreement between the two partitions. In contrast to the



(a) 2-system segmentation, subject-specific maps



(b) 3-system segmentation, subject-specific maps



(c) group average of 2-system subject-specific maps in (a) (d) group average of 2-system partition of the default system

Fig. 1. Functional segmentation examples. (a,b) Subject-specific segmentation results for two and three systems respectively (flattened view). Green: default system, blue: stimulus-driven cortex, red: visual cortex. Solid lines show the boundaries of the default system determined through seed selection. (c) Group average of the subject-specific 2-system maps. Color shading shows the proportion of subjects whose clustering agreed with the majority label. (d) Group average of the subject-specific segmentation of the default system into two sub-systems. Only voxels consistently labeled across subjects are shown.

difficulties associated with the subject-specific threshold selection in group analysis within the standard functional connectivity framework, the clustering-based decomposition produces highly repeatable maps that do not involve subject-specific adjustments. Figure 1c shows a group-level label map that summarizes the maps from Figure 1a, further illustrating the stability of the decomposition. Subsequent subdivision of the cortical gray matter into three systems produced the results in Figure 1b. With the exception of one subject, the 3-system segmentation reveals visual cortex. In subject 7, the visual cortex separated in segmentation into 4 systems (shown in the figure). We also performed a subdivision of just the default system in each subject. This subdivision produced a stable partition across subjects; the corresponding group-level map is shown in Figure 1d. The overlap of the smaller sub-systems is weaker than that of the default system itself, but it clearly represents a coherent division of the default system.

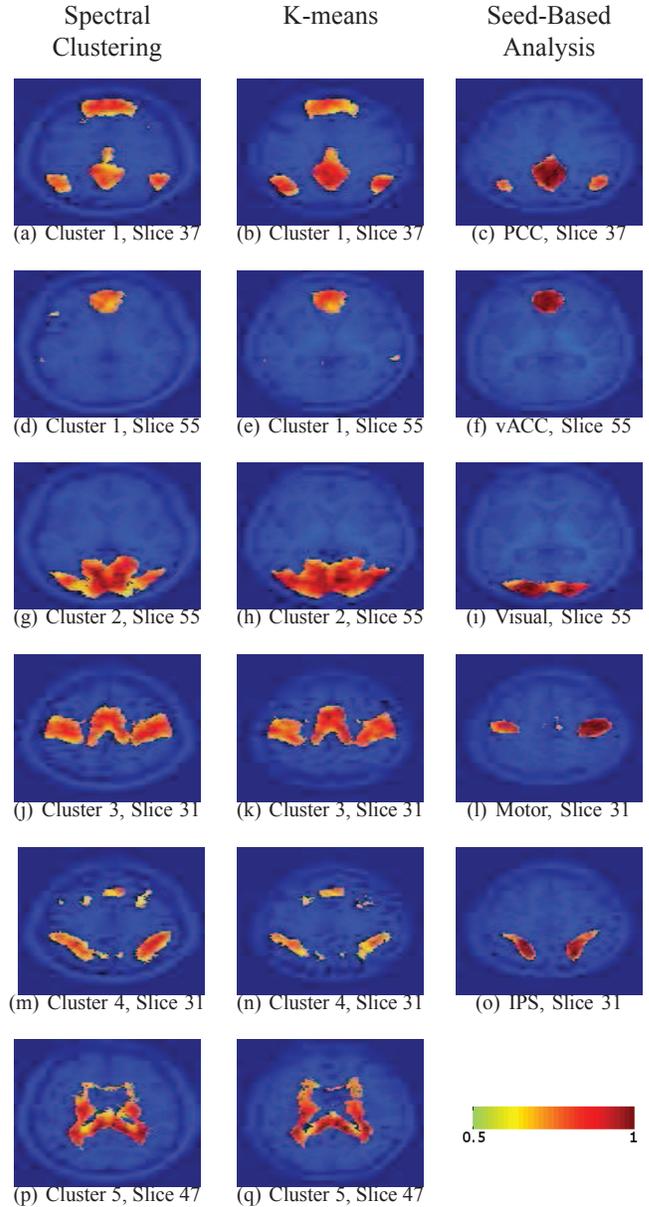


Fig. 2. Clustering results across subjects. The brain is partitioned into 5 clusters using spectral clustering (left) and k-means (middle); various seed are selected for seed-based analysis (right). Color indicates the proportion of subjects for whom the voxel was included in the detected system.

### B. Spatial Activation Patterns in Rest fMRI

In the second experiment, we further validated our approach using resting state fMRI data obtained from 45 healthy young adults [20]. The main goal of this experiment was to compare the performance of spectral clustering and mixture-model clustering to that of seed-based connectivity analysis. For seed-based analysis, we selected five seeds, corresponding to the motor and visual cortices, the ventral anterior cingulate cortex (vACC), the posterior

cingulate cortex (PCC), and the intraparietal sulcus (IPS). Only voxels whose correlation with the seed time course exceeded 0.4 were included in the system for each subject. We used the k-means implementation of the mixture model and applied k-means and spectral clustering to partition the brain into five systems.

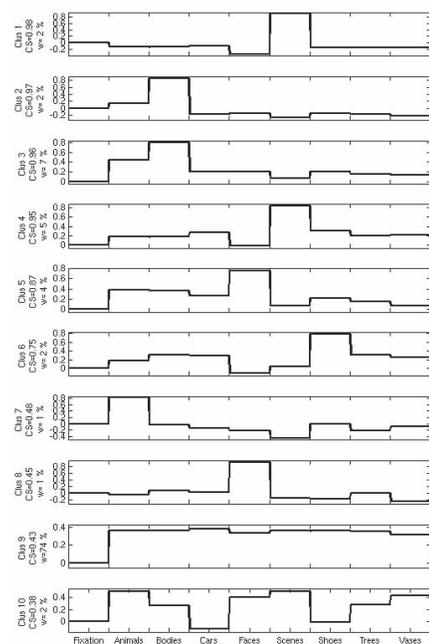
Using random subsamples of the data, we investigated the stability of the Nyström approximation for spectral clustering. We observed that for Nyström sets of 2,000 voxels and above, the resulting partitions varied little across runs. Specifically, the average difference in the spatial map for each cluster was below 1% of the cluster size. A set of 2,000 voxels represents 1% of the whole brain in this data set, offering a dramatic reduction in computational complexity of the spectral clustering algorithm.

Figure 2 illustrates the clusters identified by the three methods. Only voxels assigned to the respective system in at least half of the participants are shown for each method. Figure 2 shows clearly that both spectral clustering and mixture model can identify well-known structures such as the default network (a-f), the visual cortex (g-i), the motor cortex (j-l), and the dorsal attention system (m-o). Spectral clustering and mixture model also identified white matter (p-q). In general, one would not attempt to delineate this region using seed-based analysis. Since we regress out the white matter signal during the preprocessing, we would not expect seed-based analysis to isolate white matter as a system. In our experiments, spectral clustering and mixture modeling achieve similar clustering results across participants. Furthermore, both methods identify the same functional systems as seed-based analysis without requiring *a priori* knowledge about the brain.

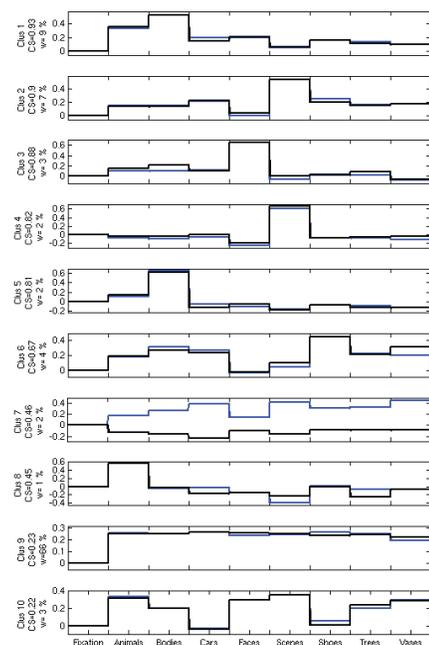
### C. Activation Profiles in Visual Experiments

We use a visual fMRI study to demonstrate the application of our mixture modeling in the space of activation profiles. In this experiment, six subjects viewed images from eight categories of visual stimuli. To construct the space of activation profiles, we estimated the regression coefficients for all eight categories based on the General Linear Model [14]. We excluded from the analysis all voxels that did not pass a stimulus-versus-fixation contrast at a significance level  $\alpha = 10^{-4}$ .

Figure 3a illustrates an application of our method to this data for  $K = 10$ . The plots show the relative strength of response to each category for the activation profiles detected by the algorithm, sorted in the decreasing order of consistency across subjects. These results clearly contain previously identified category-selective profiles. For example, profiles #1 and #4 correspond to the place-selective voxels (PPA), while the profiles #5 and #8 correspond to the face-selective voxels (FFA). To study the sensitivity of the method to the number of distinct categories and the



(a) 1 repetition



(b) 2 repetitions

Fig. 3. Ten representative activation profiles (cluster means) detected by the algorithm. The profiles are sorted in the decreasing order of functional consistency across subjects. For each profile, the plot shows the relative strength of activation for each category, the consistency score (CS) and the weight ( $\lambda$ ). (a) All stimuli from the same category were used with the same category label in the estimation procedure. (b) The stimuli within each category were divided into two artificial categories. Blue and black show the level of activation for each “virtual” category.

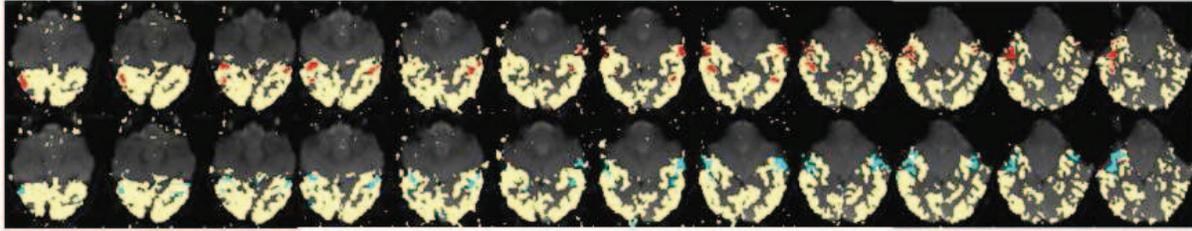


Fig. 4. Spatial maps of the face selective regions found by the statistical test (red) and the mixture model (blue). The mask of visually responsive voxels is shown in yellow in the background to help the visual comparison.

number of presentations for each category, we separated presentations for each original category into two new “virtual” categories, creating a total of 16 categories. On this data set, the algorithm was forced to identify the most stable activation profiles of length 16. Figure 3b displays the detected activation profiles by overlaying the activation level for matching “virtual” categories in two different colors. The top activation profiles clearly exhibit nearly identical values for the two category duplicates. We also observe that as the functional consistency across subjects decreases, so does the consistency across category duplicates.

Figure 4 visually compares the spatial map of face-selective locations identified by our method and the face-selective FFA area identified through the traditional hypothesis-driven method for one of the subjects in the study, demonstrating fairly good agreement. To create the traditional FFA map, we applied the  $t$ -test comparing the response for faces to the response for objects (shoes, tools, vases), at  $\alpha = 10^{-4}$ . For our algorithm’s results, we first identified all profiles whose component for faces was at least 1.5 times higher than all other components. We then assigned each voxel to its corresponding MAP cluster label to construct a binary map. We also observed similar agreement for other well known areas, such as PPA and EBA, in all subjects in the study.

In this experiment, increasing the number of clusters resulted only in the split of previously detected clusters, and did not significantly alter the pattern of the discovered profiles.

#### IV. CONCLUSIONS

In this work, we demonstrated that clustering can be effectively used for finding structure in fMRI data. Specifically, we employed mixture modeling and spectral clustering to identify spatial patterns of co-activation on the cortex. We also used mixture modeling to detect and characterize stable activation profiles in experiments with multiple stimulus categories. Group-wise analysis promises to yield comprehensive models of functional organization of the brain across subjects, and represents an interesting direction for future research.

#### REFERENCES

- [1] Baumgartner, R., *et al.* Fuzzy clustering of gradient-echo functional MRI in the human visual cortex. *J Magnetic Resonance Imaging*, 7(6):1094-101, 1997.
- [2] Beckmann, C.F., and Smith, S.M. Tensorial Extensions of Independent Component Analysis for Group FMRI Data Analysis. *NeuroImage*, 25(1):294-311, 2005.
- [3] Bell, A.J., and Sejnowski, T.J. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7:1129-159, 1995.
- [4] Biswal, B., *et al.* Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34:537-541, 1995.
- [5] Buckner, R.L., and Vincent, J.L. Unrest at rest: Default activity and spontaneous network correlations. *NeuroImage*, 37:1091-1096, 2007.
- [6] Cordes, D., *et al.* Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magnetic Resonance Imaging*, 20(4):305-317, 2002.
- [7] Dempster, A., Laird, N., and Rubin, D. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Society, Series B*, 39(1):1-38, 1977.
- [8] Fadili, M.J., *et al.* A multistep Unsupervised Fuzzy Clustering Analysis of fMRI time series. *Human Brain Mapping*, 10(4):160-178, 2000.
- [9] Fox, M.D. and Raichle, M.E. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature*, 8:700-711, 2007.
- [10] Filzmoser, P., Baumgartner, R., and Moser, E. A hierarchical clustering method for analyzing functional MR images. *Magnetic Resonance Imaging*, 10:817-826, 1999.
- [11] Fowlkes, C., *et al.* Spectral grouping using the nystrom method. *IEEE PAMI*, 26:214-225, 2004.
- [12] Friston, K.J., *et al.* Functional connectivity: the principle component analysis of large (PET) data sets. *J. Cereb. Blood Flow Metab.*, 13:5-14, 1993.
- [13] Friston, K.J. Functional and effective connectivity: a synthesis. *Human Brain Mapping*, 2:56-78, 1994.
- [14] Friston, K.J., *et al.* Statistical parametric maps in functional imaging: a general linear approach. *Human*

- Brain Mapping*, 2:189-210, 1995.
- [15] Golland, Y., *et al.* Extrinsic and intrinsic systems in the posterior cortex of the human brain revealed during natural sensory stimulation. *Cereb. Cortex*, 17:766-777, 2007.
  - [16] Golland, P., *et al.* Detection of Spatial Activation Patterns As Unsupervised Segmentation of fMRI Data. *In Proc. MICCAI*, LNCS 4791:110-118, 2007.
  - [17] Goutte, C., *et al.* On clustering fMRI time series. *Neuroimage*, 9(3):298-310, 1999.
  - [18] Goutte, C., *et al.* Feature-space clustering for fMRI meta-analysis. *Human Brain Mapping*, 13:165-183, 2001.
  - [19] Grill-Spector, K., *et al.* A sequence of early object processing stages revealed by fMRI in human occipital lobe. *Human Brain Mapping*, 6(4):316-328, 1998.
  - [20] Kahn, I., *et al.* Distinct cortical anatomy linked to subregions of the medial temporal lobe revealed by intrinsic functional connectivity. *J. of Neurophysiology*, 100:129-139, 2008.
  - [21] Kanwisher, N.G. The ventral visual object pathway in humans: evidence from fMRI. *In The Visual Neurosciences*, Eds. L. Chalupa and J. Werner, 1179-1189, 2003.
  - [22] Lashkari, D., *et al.* Discovering Structure in the Space of Activation Profiles in fMRI. *In Proc. MICCAI*, LNCS 5241:1015-1024, 2008.
  - [23] Mardia, K.V. Statistics of directional data. *J. R. Statist. Soc. Series B*, 37:349-393, 1975.
  - [24] McKeown, M.J., *et al.* Analysis of fMRI data by blind separation into spatial independent components. *Human Brain Mapping*, 6:160-188, 1998.
  - [25] McLachlan, G.J., and Basford, K.E. Mixture models. Inference and applications to clustering. Dekker, 1988.
  - [26] Moser, E., *et al.* Fuzzy clustering of gradient-echo functional MRI in the human visual cortex. Part II: quantification. *J. Magnetic Resonance Imaging*, 7(6):1102-8, 1997.
  - [27] Nadler, B., *et al.* Diffusion maps, spectral clustering and reaction coordinates of dynamical systems. *ACHA*, 21:113-127, 2006.
  - [28] Shi, J., and Malik, J. Normalized cuts and image segmentation. *IEEE PAMI*, 22:888-905, 2000.
  - [29] Thirion, B., and Faugeras, O. Feature Detection in fMRI Data: The Information Bottleneck Approach. *Medical Image Analysis*, 8:403-419, 2004.
  - [30] Thirion B., Dodel S., and Poline J.B. Detection of signal synchronizations in resting-state fMRI datasets. *NeuroImage*, 29:321-327, 2006.
  - [31] Venkataraman, A., *et al.* Exploring functional connectivity in fMRI via clustering. *Submitted to ICASSP'09*.