

Parametric Response Surface Models for Analysis of Multi-Site fMRI Data

Seyoung Kim¹, Padhraic Smyth¹, Hal Stern¹, Jessica Turner², and FIRST BIRN

¹ Bren School of Information and Computer Sciences, University of California, Irvine
{sykim, smyth}@ics.uci.edu, sternh@uci.edu

² Department of Psychiatry and Human Behavior, University of California, Irvine
turnerj@uci.edu

Abstract. Analyses of fMRI brain data are often based on statistical tests applied to each voxel or use summary statistics within a region of interest (such as mean or peak activation). These approaches do not explicitly take into account spatial patterns in the activation signal. In this paper, we develop a response surface model with parameters that directly describe the spatial shapes of activation patterns. We present a stochastic search algorithm for parameter estimation. We apply our method to data from a multi-site fMRI study, and show how the estimated parameters can be used to analyze different sources of variability in image generation, both qualitatively and quantitatively, based on spatial activation patterns.

1 Introduction

Functional magnetic resonance imaging (fMRI) is one of the primary techniques for studying how the brain functions when triggered by external stimuli. Neural activity in local areas of the brain changes the oxygenated blood level, changes that in turn can be detected by MRI scanners. The data collected in an fMRI session consist of a time-series of voxel images, containing both temporal and spatial information about brain activation patterns. The temporal aspect of the data is often aggregated by fitting a linear regression model that relates the activation data at each voxel to the stimulus signal and perhaps other predictors as well. The regression coefficients of the stimulus signal at each voxel (often denoted as β) or other regression summary (e.g., t -statistic) can then be used as a statistical parametric map of activity [1]. There is a significant body of prior work using techniques such as hypothesis testing and thresholding to search these statistical voxel images for activated brain regions. Test statistics used in these techniques include the maximum activation intensity within a selected region of interest, or the spatial extent of the largest set of contiguous voxels above a threshold [2, 3].

An important aspect of the data that is typically not considered in these studies is the spatial pattern of activation across voxels. Voxels with a high activation level tend to appear in spatial clusters in the β -maps, forming what looks like a bell shape for each activation cluster. These spatial clusters can

be explained to first order by assuming that the underlying true activations are point sources that are then spatially smoothed (convolved) by Gaussian-like filters.

In this paper we characterize fMRI activation patterns in the brain by modeling the shape of each “activation cluster” by (a) its height or peak value, representing the amount of activation, and (b) the location of the cluster, modeling the center of activation in voxel-space. We represent these features using a parametric model of the functional form of the surface with free parameters for the heights and center locations of the “bumps.” These parameters can be estimated in a statistical manner from the β activation maps.

In terms of relevant prior work, Hartvig [5] also used a similar parametric function for modeling spatial activation patterns. However, the focus in this earlier work was on extracting activated voxels by thresholding after the model parameters were estimated. In the work described in this paper we are interested in the shape features themselves, and analyze the activation pattern at the feature level rather than at a voxel level. In our experiments we show how the estimated features of the local activation clusters can be used in a multi-site study of fMRI to analyze sources of variability. Evidence in this direction is suggested by Fig. 1, where the locations of the peaks of activation are more consistent in the four runs of the same subject within the same site (MRI scanner), than between the two sites.

Multi-site fMRI studies are not currently common, but are becoming more so as the need is growing to combine imaging data across sites to create larger datasets than would be possible at a single site. These datasets allow imaging experiments in rare diseases where few subjects would be available at any given site, or with very complex hypotheses where a single site study would be underpowered. In particular, in one multi-site study of signal reliability across sites [4] the activation levels and locations were similar, but the pattern of activation across sites was very different. The ability to assess such patterns of activation provides the primary motivation for the method proposed in this paper.

The paper is organized as follows. In Sect. 2, we briefly describe the data collection process and preprocessing steps used to produce the activation maps analyzed in this paper. In Sect. 3, we introduce a spatial model for activation patterns, describe inference procedures for this model and present experimental results based on fitting this model to multi-site fMRI data. In Sect. 4, we show how the shape features estimated from the surface modeling can be used to investigate multi-site variability. Sect. 5 concludes with a brief discussion on future work.

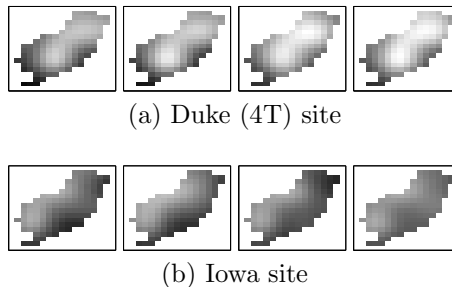


Fig. 1. Raw data (β maps) for a cross section at $z = 53$ of the right precentral gyrus of four runs within visit 2 of subject 3

2 Multi-Site Data Collection and Preprocessing

fMRI scans for the same five control subjects were collected from 10 different scanners (UCSD, UCI, Stanford, Duke (1.5T), Duke (4T), New Mexico, Minnesota, Iowa, BWH, MGH) as part of a multi-site study of functional brain images, known as FIRST BIRN (Functional Imaging Research on Schizophrenia Test-bed Biomedical Informatics Research Network), also known as Function BIRN or fBIRN. For each subject there were two visits to each site, and at each visit fMRI data were collected for four runs of a sensorimotor task and two runs of breathholding, resting, and two cognitive tasks. A primary goal of this data collection is to better understand the variability of fMRI response patterns across runs, visits, scanners (sites) and subjects, so that future data collected across sites and subjects can be analyzed collectively and consistently. In this paper, we use the data from the sensorimotor task, and we focus on activation within specific regions of interest such as the right precentral gyrus and the left superior temporal gyrus that are relevant to this task.

Each run of the sensorimotor task produces a series of 85 scans that can be thought of as a large time-series of voxel images. The set of scans for each run is preprocessed in a standard manner using SPM99, with the default settings. The preprocessing steps include correction of head motion, normalization to a common brain shape (MNI template), and spatial smoothing.

A general linear model is then fit to the time-series data for each voxel. A β -map is a voxel image of the regression coefficients (β 's) that summarizes the activation across time as an activation map. Binary masks for regions of interest from the normalized atlas were then used to extract the β values for all voxels within the region.

3 Activation Surface Modeling

3.1 Model

We model the spatial activation pattern (β -map) for a region of interest as a superposition of multiple Gaussian-shaped surfaces, with a constant background term for non-activation, and with additive Gaussian measurement noise. Below we develop the model for the case of 2-dimensional slices of pixels—the 3-dimensional case can be derived directly as an extension of the 2-dimensional case, but is not pursued in this paper.

Assuming the number M of Gaussian surfaces for a 2-dimensional slice of a region is known and fixed, the model for the activation value β at position $\mathbf{x} = (x_1, x_2)'$ is defined as

$$\beta = \mu + \sum_{m=1}^M k_m \exp(-(\mathbf{x} - \mathbf{b}_m)'(\mathbf{x} - \mathbf{b}_m)/\sigma_m) + \varepsilon, \quad (1)$$

where ε is distributed as $N(0, \sigma^2)$, and μ is the background level. $\boldsymbol{\theta}_m = \{k_m, \mathbf{b}_m, \sigma_m\}$, $m = 1, \dots, M$, is the set of parameters for the m th Gaussian surface,

corresponding to activation centered around $\mathbf{b}_m = (b_{1m}, b_{2m})'$ with height of k_m , and σ_m controlling the volume under the surface. By using σ_m instead of a 2×2 covariance matrix, we assume a spherical shape with no correlation between x_1 and x_2 directions.

In the Bayesian estimation setting, we are interested in the posterior distribution of the unknown parameters $p(\mu, \sigma, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M | \boldsymbol{\beta})$. Using Bayes' rule, we write $p(\mu, \sigma, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M | \boldsymbol{\beta}) = p(\boldsymbol{\beta} | \mu, \sigma, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M) p(\mu, \sigma, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M) / p(\boldsymbol{\beta})$, where $p(\boldsymbol{\beta} | \mu, \sigma, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M)$ is the likelihood of the $\boldsymbol{\beta}$ values $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_N\}$ for a region with N voxels given the model in (1). Assuming independent priors for (μ, σ^2) , and $\boldsymbol{\theta}_m$'s, we choose a uniform prior $p(\mu, \log \sigma) \propto 1$ for $(\mu, \log \sigma)$, $p(k_m) \propto 1$ for k_m 's, where $k_m > 0$ for positive activations, and assign a Gamma(a, b) prior on σ_m with hyperparameters a and b . Given that it is reasonable to believe that the center of activation \mathbf{b}_m is located inside (or close to) the region of brain under consideration, we assume $p(\mathbf{b}_m)$ is uniform inside or within one pixel of the region. It is straightforward to verify that this choice of prior distributions yields a proper posterior distribution.

3.2 Parameter Estimation Using Stochastic Search

Because of the nonlinearity of the model described in Sect. 3.1, it is not possible to determine the posterior distribution of the parameters analytically, and we need to rely on simulation methods to obtain samples from the posterior. In this work, we use a combination of the Gibbs sampler and the Metropolis algorithm.

In the Gibbs sampler, starting from initial values for the unknown parameters, at each iteration we update each parameter in turn by replacing it with a sample from its conditional posterior distribution given the data and the most recent samples of the other parameters. By repeating this update, the Gibbs sampler converges (under mild conditions) to a state where the samples can be presumed to be coming from the joint posterior distribution of interest. Direct sampling from the conditional posterior distribution is possible for μ and σ^2 , but not for the $\boldsymbol{\theta}_m$'s. The Metropolis algorithm is used to sample $\boldsymbol{\theta}_m$'s inside the Gibbs sampler. We describe the details of the sampling algorithm below.

1. Sample a new value for μ from its conditional posterior distribution $p(\mu | \sigma, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M, \boldsymbol{\beta}) = N(\hat{\mu}, \sigma^2/N)$, where $\hat{\mu} = (\sum_{i=1}^N (\beta_i - \sum_{m=1}^M k_m \exp(-(\mathbf{x}_i - \mathbf{b}_m)'(\mathbf{x}_i - \mathbf{b}_m)/\sigma_m)))/N$.
2. Sample a new value for σ^2 from its conditional posterior distribution. Given the prior $p(\log \sigma) \propto 1$, which is equivalent to $p(\sigma^2) \propto \sigma^{-2}$, the conditional posterior distribution for σ^2 can be found as $p(\sigma^2 | \mu, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M, \boldsymbol{\beta}) = \text{Inv-gamma}(\nu/2, (\nu/2)s^2)$, where the degree of freedom $\nu = N$ and $s^2 = (\sum_{i=1}^N (\beta_i - \mu - \sum_{m=1}^M k_m \exp(-(\mathbf{x}_i - \mathbf{b}_m)'(\mathbf{x}_i - \mathbf{b}_m)/\sigma_m))^2)/N$.
3. For $m = 1, \dots, M$,
 - (a) Sample k_m^* from the jumping distribution $N(k_m, (\tau_k)^2)$ and compute $r = (p(\boldsymbol{\beta} | k_m^*, \boldsymbol{\theta}_{-k_m}) p(k_m^*)) / (p(\boldsymbol{\beta} | k_m, \boldsymbol{\theta}_{-k_m}) p(k_m))$, where $\boldsymbol{\theta}_{-k_m}$ represents all parameters other than k_m . If $r > 1$, accept k_m^* as a new value for k_m , otherwise, accept it with the probability r .

- (b) Update \mathbf{b}_m as in 3a, but using the jumping distribution $N_2(\mathbf{b}_m, (\tau_{\mathbf{b}})^2 I)$.
- (c) Update σ_m as in 3a but using the jumping distribution $N(\sigma_m, (\tau_{\sigma})^2)$.

Before the start of the sampler, the values for $(\tau_k)^2$, $(\tau_{\mathbf{b}})^2$, and $(\tau_{\sigma})^2$ in Step 3 are adjusted to values that keep the acceptance rate at roughly 30-50%.

In the results described below we summarize the posterior distributions via “point estimates” (specific values of the parameters), corresponding to the mean values for each parameter over the samples (i.e., the estimated posterior mean).

3.3 Results from Surface Fitting

In this section, we discuss results of our modeling procedure for two cross sections, one at $z = 53$ of the right precentral gyrus and another at $z = 33$ of the left superior temporal gyrus, for subjects 1 and 3 in our study, for the sensorimotor task at all sites.

From visual inspection of the images, we set the number of Gaussians in the model to $M = 1$ for subject 1, and $M = 2$ for subject 3 in the case of the right precentral gyrus, and $M = 2$ for both subjects in the case of the left superior temporal gyrus.

From an exploratory analysis of the data, the hyperparameters a and b of the gamma prior distributions of the σ 's were set so that the mode and variance are approximately 15 and 160 ($a = 3.0$ and $b = 0.1365$). With this setting of the model, we ran the Markov chain simulation for 20,000 iterations for each image until convergence, and output the mean of the samples from the last 10,000 iterations to provide point estimates of the parameters.

A typical example (for a particular run, visit, and subject) is shown in Fig. 2, comparing the raw β maps with the activation maps for the fitted model based on point estimates of the parameters from the same raw data. The shapes of the activation pattern in the estimated surfaces such as the locations of the peaks are consistent with what we see in the raw data. Table 1 shows an average across multiple images of the sum of squared errors (across voxels) between the fitted model and the data, divided by the voxel variance in each image and expressed as a percentage—in effect, the average amount of variability in the data not explained by the model. Both Fig. 2 and Table 1 suggest that while the functional form of the model is a result of various simplifying assumptions it nevertheless provides a good fit to the observed spatial activation patterns.




Fig. 2. Raw data vs. estimated surfaces from run 3, visit 2, and subject 3 at Iowa site.

Table 1. Percentage of variance not explained by the models.

	subject 1	subject 3
Right precentral gyrus ($z = 53$)	15.4%	7.1%
Left superior temporal gyrus ($z = 33$)	16.1%	5.6%

4 Modeling Sources of Variability

We can think of the estimated parameters of the model as features that summarize the activation pattern, and analyze the statistical variation of these features rather than that of the raw data. Fig. 3 visually shows the variabilities in the estimated location parameters of the local activation centers at each of run, visit, and site levels for the cross section $z = 53$ of the right precentral gyrus of subject 3. The estimated locations are represented as o's for visit 1 and +'s for visit 2. The two bumps are connected with a line if they come from the same run. Across certain groups of sites (e.g., UCI and UCSD) there appears to be little cross-site variability in the estimated spatial locations, but between other pairs (e.g., MGH and UCI) the cross-site variation is much larger than the within-site variation. Qualitatively similar results are obtained for other cross-sections.

In this section, we attempt to quantify the contributions of site, visit, and run effects to variation in features of the local activations within each subject. In the case of $M = 2$, we first find the correspondence of the two bumps across images from different sites, visits and runs based on the estimated location parameters using a simple heuristic algorithm, and consider each bump separately in the variance component analysis. In Sect. 4.1 we discuss a Bayesian hierarchical model for a variance component analysis. In Sect. 4.2 we present the results from the model.

4.1 Bayesian Variance Component Analysis

A traditional variance components analysis can be used to decompose the observed variation in signal into portions attributable to site, visit, and run variabilities. One such model that we choose to use in our analysis is $y_{ijk} = \mu + s_i + v_{ij} + r_{ijk}$, where y_{ijk} is the response measure, μ is an overall mean, s_i is a site effect, v_{ij} is a visit effect and r_{ijk} is a run effect (essentially all sources of variation

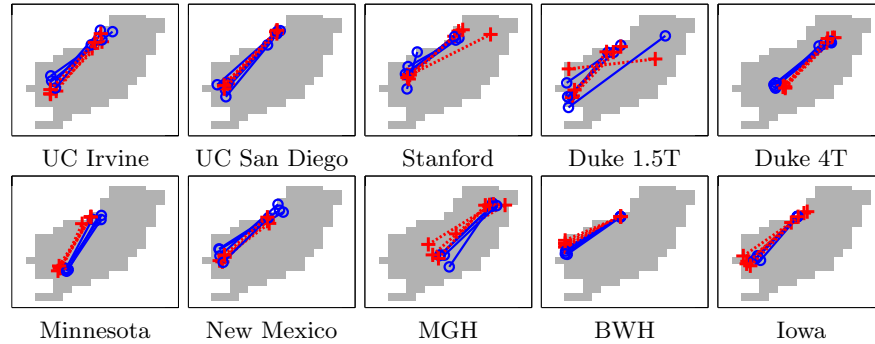


Fig. 3. Variability in locations parameters estimated from right precentral gyrus at $z = 53$ of subject 3. o's are visit 1, and +'s are visit 2. Two locations are connected with a line if they come from the same image.

not included in previous terms). The responses y_{ijk} 's considered in this analysis are the estimated locations and heights of each of the Gaussian components. Each of the effects (site, visit, run) is further modeled as a Gaussian random variable with mean zero and variance/covariance parameter that characterizes the contribution to overall response variability of the particular source. It is the variance parameters, or the variance components, that are of interest.

Parameters are estimated as means of the posterior distribution of the variance components from a Bayesian analysis with weak prior distributions on the variance components. The prior distributions for the analyses reported here are $N(0, 10^{10})$ for μ (a common vague prior for the mean of a Gaussian distribution), gamma with mean 1 and variance 100 on the one dimensional precision parameters (reciprocals of the variances), and Wishart distribution with degree of freedom 2 and scale parameters estimated empirically from the y_{ijk} 's for the covariance matrices. Samples of the variance parameters from their posterior distribution given the data are obtained using WinBugs, a freely available software implementation of the Gibbs sampler. The Gibbs sampler was run for 1,000,000 iterations by which time the draws were determined to be representative of the posterior distribution. The means of the last 200,000 draws are reported as estimates for the variance components.

4.2 Results

We compute the variance components for sites, visits, and runs of location and height parameters, and show the proportions of variance with respect to the total variation in Tables 2 and 3. For the 2×2 covariance matrices of location parameters, we use the absolute value of the determinant of the matrix as a summary of the amount of variation instead of the full matrix.

In the cross section of $z = 53$ of the right precentral gyrus the effect size due to sites is consistently larger than visit and run effects across different bumps and subjects in both locations and heights. The visit effects account for the smallest proportion of the total variabilities. Note also that the relative sizes of variance components (for sites, visits, runs) are similar for both subjects.

A slightly different pattern is observed in the cross section of $z = 33$ of the left superior temporal gyrus. The contribution from visit effects is still the smallest, but the run variability is often larger than the site variability. Overall we see more consistent patterns of variance components in the right precentral gyrus than in the left superior

temporal gyrus. The right precentral gyrus is considerably smaller in size, therefore, is less affected by the limitation that we are looking at a 2 dimen-

Table 2. Proportions of variance components for right precentral gyrus at $z = 53$

	Subject 1		Subject 3			
	Height	Location	Height		Location	
			Bump 1	Bump 2	Bump 1	Bump 2
Site	0.51	0.86	0.50	0.58	0.67	0.90
Visit	0.22	0.07	0.13	0.18	0.02	0.03
Run	0.27	0.07	0.37	0.24	0.31	0.07

Table 3. Proportions of variance components for left superior temporal gyrus at $z = 33$

	Subject 1				Subject 3			
	Height		Location		Height		Location	
	Bump 1	Bump 2	Bump 1	Bump 2	Bump 1	Bump 2	Bump 1	Bump 2
Site	0.25	0.52	0.05	0.06	0.49	0.59	0.40	0.45
Visit	0.42	0.25	0.18	0.03	0.12	0.05	0.04	0.02
Run	0.33	0.23	0.77	0.91	0.39	0.36	0.56	0.53

sional cross section instead of taking into account the variability in all of the three dimensions.

An alternative non-Bayesian method of moments approach provides similar results with the exception of some negative estimates of variances (a known disadvantage of this approach).

5 Conclusions

We have shown that spatial modeling of fMRI activation patterns in local brain regions can extract reliable and useful information providing (for example) a basis for statistical analysis of variability. Future work includes analyzing more regions of interest with all 5 subjects in the study, extending the model to 3-dimensional voxel data, and developing techniques for automatically selecting the number of bumps M .

Acknowledgements. This material was supported in part by the National Science Foundation under grants No. SCI-0225642 and IIS-0431085. The authors also acknowledge the support of the First BIRN (M01RR00827), the Transdisciplinary Imaging Genetics Center (P20RR020837-01) and the National Alliance for Medical Image Computing (Grant U54 EB005149), funded by NIH/NCRR.

References

1. Bullmore, E., Brammer, M., Williams, S.C., Babe-Hesketh, S., Janot, N., David, A., Mellers, J., Howard, R., & Sham, P. (1996). Statistical methods of estimation and inference for functional MR image analysis. *Magnetic Resonance in Medicine*, 35(2):261–277.
2. Cao, J., & Worsley, K. (2001). Applications of random fields in human brain mapping. In *M. Moore (Ed.) Spatial Statistics: Methodological Aspects and Applications*, Springer Lecture Notes in Statistics, 159:169–182.
3. Friston, K.J., Holmes, A.P., Price, C.j., Buchel, C., & Worsley, K.J. (1999). Multi-subject fMRI studies and conjunction analysis. *NeuroImage*, 10:385-396.
4. Casey, B. J., Cohen, J. D., O’Craven, K., Davidson, R. J., Irwin, W., Nelson, C. A. et al. (1998). Reproducibility of fMRI results across four institutions using a spatial working memory task. *Neuroimage*, 8:249-261.
5. Hartvig, N. (1999). A stochastic geometry model for fMRI data. Research Report 410, Department of Theoretical Statistics, University of Aarhus.