SOFTWARE HIGHLIGHT

NPBayes-fMRI: Nonparametric Bayesian General Linear Models for Single- and Multi-Subject fMRI Data

Jeong Hwan Kook, Marina Vannucci
Rice University
jk5@rice.edu, marina@rice.edu

Functional magnetic resonance imaging (fMRI) is a noninvasive neuroimaging technique which measures the blood oxygenation level dependent (BOLD) contrast, i.e. the difference in magnetization between oxygenated and deoxygenated blood arising from changes in regional cerebral blood flow. In a typical task-related fMRI experiment, a subject is presented a set of stimuli while the whole brain is scanned at multiple time points. Each scan is arranged as a 3D array of volume elements (or "voxels"), and the experiment produces time series of BOLD responses acquired at each voxel.

Common modeling approaches for the analysis of task-related fMRI data rely on the general linear model formulation that was first proposed by [8] and subsequently investigated by many other authors, particularly for single-subject data, see for example [7, 23, 21, 13, 18, 11, 25], among many others. For multi-subject studies, two-stage "group analysis" approaches are often adopted as computationally attractive methods where summary estimates of model parameters are obtained at the individual level and then used in a second stage model at the group/population level, see for example [1, 20, 19, 12]. Also, newer data-driven methods for analyzing fMRI, for example those that use model-free methods such as independent component analysis (ICA) and tensor-product ICA (T-PICA), have been developed to detect the presence of subgroups of participants within a population as in [3], but these approaches still involve multiple estimation steps, and therefore do not properly take into account variability and heterogeneity in the data.

NPBayes-fMRI [10] is a user-friendly MATLAB [14] GUI that practitioners can use for both model fitting and visualization of the results. For inference, the software extends a unified, single-stage Bayesian approach for the analysis of task-related brain activity proposed by [24]. This model formulation considers a spatio-temporal linear regression model that specifically accounts for between-subject heterogeneity in neuronal activity via a spatially informed multi-subject nonparametric variable selection prior.

1 Bayesian spatio-temporal models for fMRI data

For subject $i = 1, \ldots, N$, let $Y_{i\nu} = (Y_{i\nu 1}, \ldots, Y_{i\nu T})^T$ be the vector of the BOLD response data at voxel ν, with $\nu = 1, \ldots, V$. We model the data as

$$Y_{i\nu} = X_{i\nu}\beta_{i\nu} + \varepsilon_{i\nu},\ \varepsilon_{i\nu} \sim N_T(0, \Sigma_{i\nu}), \qquad (1)$$

where $X_{i\nu}$ is a known $T \times p$ covariate matrix and $\beta_{i\nu} = (\beta_{i\nu 1}, \ldots, \beta_{i\nu p})^T$ is a $p \times 1$ vector of regression coefficients. Without loss of generality, we center the data and thus do not include the intercept term in the model. Let $X_{i\nu j}$ be the jth column of $X_{i\nu}$. Then $X_{i\nu j}$ is modeled as the convolution of the j-th stimulus pattern with a hemodynamic response function (HRF) [2], that is

$$X_{i\nu j}(t) = \int_0^t x_j(s) h_{\lambda_{i\nu j}}(t-s)ds, \qquad (2)$$

where $x_j(s)$ represents the stimulus pattern. One common choice is a Poisson HRF, that is $h_{\lambda_{i\nu j}} = \exp(-\lambda_{i\nu j})\lambda_{i\nu j}^t / t!$. The parameter $\lambda_{i\nu j}$ can be interpreted as the delay of the response with respect to the stimulus onset and it is often modeled as an unknown voxel-dependent parameter.

The error term in equation (1) is modeled as auto-correlated, specifically long memory and Discrete wavelet transforms (DWT) are employed as a way to decorrelate the data. This is a common approach in the fMRI literature [4, 15, 19, 25]. After applying the DWT to equation (1) the model in the wavelet domain can be written as

$$Y^*_{i\nu} = \sum_{j=1}^{p} X^*_{i\nu j} \circ \beta_{i\nu j} + \varepsilon^*_{i\nu},\ \varepsilon^*_{i\nu} \sim N_T(0, \Sigma^*_{i\nu}), \quad (3)$$

with $\circ$ the element-by-element (Hadamard) product, and where W is a $T \times T$ matrix corresponding to the wavelet transform, $Y^*_{i\nu} = WY_{i\nu}, X^*_{i\nu} = WX_{i\nu}$, and $\varepsilon^*_{i\nu} = W\varepsilon_{i\nu}$, and with the covariance matrix $\Sigma^*_{i\nu}$ approximately diagonal with elements $\psi_{i\nu}\sigma^2_{imn}$ indicating the variance of the nth

wavelet coefficient at the mth scale. We follow the variance progression method of [22] for the wavelet coefficients,

$$\psi_{i\nu}\sigma^2_{imn} = \psi_{i\nu}(2^{\alpha_{i\nu}})^{-m}, \quad (4)$$

with $\psi_{i\nu}$ the innovation variance and $\alpha_{i\nu} \in (0,1)$ the long memory parameter. This structure encompasses the general fractal process, which includes long memory.

Detecting voxels that activate in response to a stimulus is equivalent to identifying the non-zero regression coefficient $\beta_{i\nu j}$ in model (3). In formulas, let $\gamma_{i\nu j}$ be a binary indicator of whether a given voxel is activated or not, that is, $\gamma_{i\nu j} = 0$ if $\beta_{i\nu j} = 0$ and $\gamma_{i\nu j} = 1$ otherwise. A spiked nonparametric prior is imposed on the coefficients

$$\beta_{i\nu j}|\gamma_{i\nu j}, G_i \quad \sim \quad \gamma_{i\nu j}G_{ij} + (1-\gamma_{i\nu j})\delta_0, \quad (5)$$

where δ_0 is a point mass at zero and G denotes a known distribution. With multiple subjects, a hierarchical Dirichlet Process (HDP) prior can be specified as the nonparametric slab, inducing clustering among voxels within a subject on one level and between subjects on the second level. This construction enables the model to borrow information from subjects exhibiting similar activation patterns in estimating parameters of interest and also capture spatial correlation among distant voxels. For single-subject analysis, the HDP reduces to a Dirichlet process (DP) prior.

In addition to the prior construction above, spatial correlation among neighboring voxels within a subject is modeled via a Markov Random Field (MRF) prior imposed on $\gamma_{i\nu j}$,

$$P(\gamma_{i\nu j}|d, e, \gamma_{ikj}) \sim \exp(\gamma_{i\nu j}(d + e \sum_{k\in N_{i\nu}} \gamma_{ikj})),$$

with $N_{i\nu}$ the set of neighboring voxels of voxel ν for subject i, and $p(\gamma_{i\nu}) = \prod_{j=1}^{p} p(\gamma_{i\nu j})$. The sparsity parameter $d \in (-\infty, \infty)$ represents the expected prior number of activated voxels, while the smoothness parameter $e > 0$ controls the probability of identifying a voxel as active based on the activation of the neighboring voxels. The prior model is completed by considering a uniform prior distribution on the delay parameter, $\lambda_{i\nu j} \sim \text{U}(u_1, u_2)$, an Inverse Gamma (IG) prior on the innovation variance parameter, $\psi_{i\nu} \sim \text{IG}(a_0, b_0)$, and a Beta distribution on the long memory parameter, $\alpha_{i\nu} \sim \text{Beta}(a_1, b_1)$.

For posterior inference, [24] use Variational Bayes (VB) algorithms which, unlike MCMC methods, does not rely on numerical integration. VB methods have been employed successfully in Bayesian models for single-subject fMRI data [17, 6, 9]. These methods find an optimal approximation to the posterior that minimizes the Kullback-Leibler (KL) divergence. Typically, VB approaches provide good estimates of means, although they tend to underestimate posterior variances and also to poorly estimate the correlation structure of the data. This can still be an acceptable trade-off for our inferential purposes, as we are only interested in the identification of broad areas of activations. When analytically tractable updates for some of the parameters are not available, the VB algorithm can be combined with importance sampling. Table 1 reports a schematic representation of the algorithm used in [24].

The primary interest of the inference is in the estimation of the selection parameters, γ, and the regression coefficients, β. These can be used to obtain activation maps, by subject and by stimulus. Using the output from the VB algorithm, posterior probabilities of inclusion (PPIs) for stimulus $j, p(\gamma_{i\nu j} = 1)$, for $j = 1, \dots P$, are approximated as weighted averages of the variational distribution values $q(\gamma_{i\nu j} = 1)$ estimated across the iterations of the outer loop of the algorithm (see Table 1). Activation maps can then be obtained by thresholding the PPIs using a threshold value to ensure a pre-defined Bayesian false discovery rate (FDR) [16]. This produces a spatial mapping of the activated brain regions, for each subject. Corresponding posterior β-maps can be calculated by estimating the β coefficients via weighted averages of the variational distribution values, on active voxels. An additional feature of the modeling approach of [24] is that the use of the nonparametric HDP prior construction (5) can be exploited to obtain a clustering of the subjects for possible discovery of differential activations. For an individual stimulus, and given a pre-specified threshold (or FDR) value on the PPIs, a dissimilarity matrix can be calculated based on the squared Euclidean distances between each pair of subjects and transformed into a tree via hierarchical clustering, then a dendrogram can be obtained using the linkage method with Ward's minimum variance. An optimal number of clusters can finally be selected by visual inspection of the dendrogram and group-level β-maps can be calculated by averaging the posterior maps of the non-zero β coefficients in each cluster. Finally, when analyzing experimental data with multiple stimuli, contrast maps can be produced to compare the effects of different

Algorithm 1 VB Algorithm (with Poisson HRF)

for $l = 1 : L$ iterations **do**
  Update α_{iv}^l and λ_{iv}^l $i = 1, \ldots, N, v = 1, \ldots, V, j = 1, \ldots, P$, via importance sampling.
  for $m = 1 : M$ iterations **do**
    Using the VB method,
    Update ψ_{iv}^m as the mean of its variational distribution $q(\psi_{iv})$.
    Update γ_{ivj}^m from its variational distribution $q(\gamma_{ivj})$.
    This update takes the neighboring structure of voxels into account via the MRF prior.
    Update β_{ivj}^m when $\gamma_{ivj}^m = 1$. Otherwise, set $\beta_{ivj}^m = 0$.
    Store the last update at $m = M$ as final update.
  end for
end for
Compute the importance sampling weights w_{ivl} and normalize them to $\hat{w}_{ivl}$.
Estimate the model parameters as weighted averages.

treatments, by subject, by estimating probability maps of the type $p(\beta_j - \beta_{j'} > \kappa)$, with j and j' a pair of stimuli and κ a pre-defined hypothesized value.

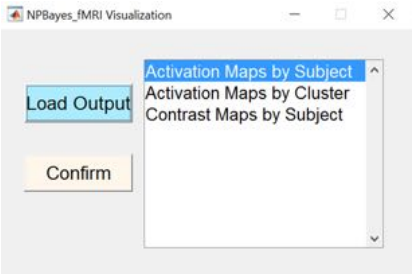

Figure 2: Main interface for Visualization.

2 Software design and implementation

In this section, we discuss the interface of `NPBayes-fMRI`. More details regarding parameter setting and input arguments can be found on [10]. `NPBayes-fMRI` comprises of two main interfaces, one for model fitting and one for the visualization of the results. They are organized as shown on Figures 1 and 2.

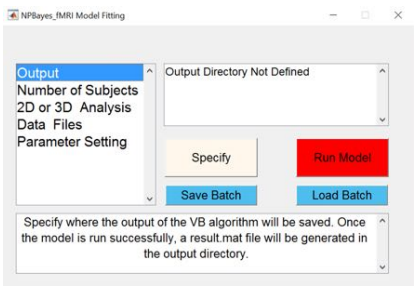

Figure 1: Main interface for Model Fitting

Model Fitting – Inference Stage. For model fitting, a set of parameters must be defined using the *NPBayes-fMRI: Model Fitting* shown in Figure 1, by first selecting the object in the listbox and then clicking the `Specify` button:

- **Output**: The user specifies where to save the final output of the VB algorithm. Once the model is run successfully, a `result.mat` file will be generated in the output directory.
- **Number of Subjects**: The user inputs the number of subjects used for the analysis. When this variable is set to 1, a DP will be used for the slab distribution in equation (5), while a HDP is used otherwise.
- **2D or 3D Analysis**: This option allows to specify the type of analysis, that is, whether it is performed on a single 2D slice or on a 3D whole-brain parcellation. Based on this, the user will be prompted to either define or load additional files. These arguments will be used later for visualization of the results.
- **Data Files**: The user needs to load a .mat file that consists of two matrices: `xtdat`, a $T \times P$ binary design matrix, with T the number of time points and P the number of stimuli, and `y_dat`, a $T \times (N \times V)$ ma-

trix of BOLD signals, with N the number of subjects and V the number of voxels (for 2D analyses) or ROIs (for 3D analyses). For both 2D and 3D analyses, the percent signal change normalization and the DWT are applied as part of the model fitting stage. For DWT, Daubechies minimum phase wavelets with 4 vanishing moments are used.

- **Parameter Setting**: The user can choose to run the model with a default parameter setting or to manually set the parameters.

Visualization – Visualization Stage. This interface is used to visualize the results using the $result.mat$ file obtained from running the algorithm with the *NPBayes-fMRI: Model Fitting* interface. It comprises of three components, described below. For simplicity, we consider 3D data.

- **Activation Maps by Subject**: This function allows the user to view the activation maps, the posterior β-maps and the HRF maps for a single subject. Clicking on `Map Type` allows the user to select either `Probability Map`, which allows to view PPI activation maps, or `Activation Map`, to view the posterior β-maps, or `HRF Map`, to view posterior maps for the HRF parameters. Matlab's built-in colormaps can be selected via the `Color Map` pop-up menu. `Range` lets the user define the axes limits. Two sliders appear on the right-hand side of the interface. The bigger slider can be used to adjust the PPI threshold and the FDR value. These values can also be set manually by the user. The smaller slider, circled in red in Figure 3, appears only when a Brain Template Image has been uploaded for 3D Analysis. This slider allows the user to control the transparency of the activation map that will be overlayed on top of the Reference Image. The X, Y and Z sliders are used to define the coordinates of the 3D NIFTI brain image in sagittal, coronal and axial orientation. If the user desires to view multiple slices of the brain in one particular orientation for a given stimulus, the `Multi-Slice` option can be used instead. The `Viewing Options` tab can be used to view either all stimuli at once or a single stimulus at a time, as shown in Figure 3.

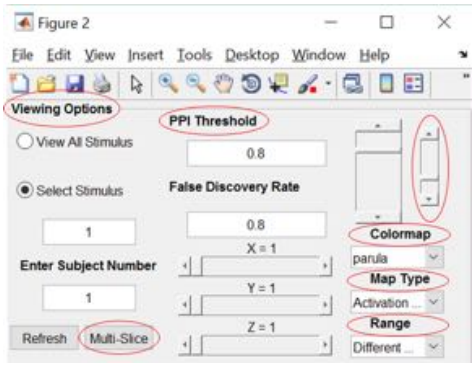

Figure 3: NPBayes-fMRI: Visualization. Interface for viewing activation maps by subject. The `Viewing Options` tab allows the user to view all stimuli at once or one at a time. By clicking on the `Map Type`, `Range` and `Colormap` pop-up menus, the user can define the type of maps to visualize, set the axes ranges and the desired colormap setting.

- **Activation Maps by Cluster**: This function is used to view cluster-level activation maps, for a given stimulus and PPI (or FDR) threshold. Clusters are defined based on a dendrogram obtained by applying hierarchical clustering with Ward's linkage method to a dissimilarity matrix defined based on the posterior mean estimates of the non-zero β coefficients. By clicking on `Load Cluster Defined From Dendrogram` as shown in Figure 4 the user can insert the number of clusters by which the subjects will be grouped.

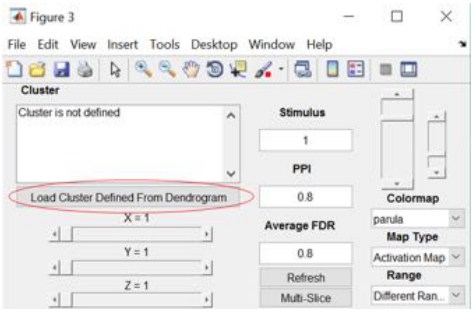

Figure 4: NPBayes-fMRI: Visualization. Interface for viewing activation maps by cluster, for a given stimulus. The user selects the stimulus and the PPI (or FDR) threshold. The corresponding dendrogram will be displayed, and the user can then specify the number of clusters and click on the `Load Cluster Defined from Dendrogram` tab. When confirmed, the cluster indices will be displayed in the `Cluster` tab.

- **Contrast Maps by Subject**: For multiple stimuli, this function lets the user define a

contrast by subject by defining a `Contrast Vector` and `Hypothesis Value` using the `Define Contrast` option. The length of the `Contrast Vector` must not be greater than the number of stimuli and the entries must sum to 0. Once a contrast has been defined, the user can use the slider to adjust the `Threshold Probability` and view different subjects by entering the subject numbers.The interface is shown Figure 5.

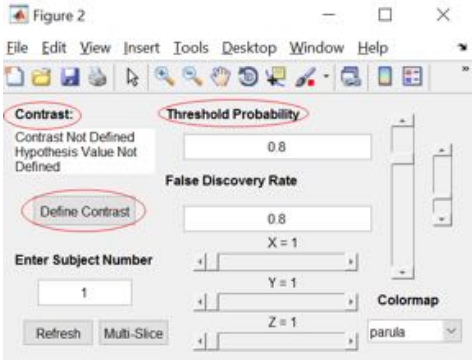

Figure 5: NPBayes-fMRI: Visualization. Interface for viewing contrast maps by subject (for multiple stimuli). The user must use the `Define Contrast` tab to insert the contrast vector and hypothesized value. The slider bar can be used to adjust the PPI (or FDR) value.

3 Inference and Visualization on 3D data

`NPBayes-fMRI` includes data of 30 subjects performing an experiment with three different stimuli. The dataset is part of a pilot study on variability in the cognitive and neural processes involved in reading, conducted at Rice University [5]. A 3D parcellation of the data was performed using the *MarsBaR* toolbox in SPM 12. The Automatic Anatomical Labeling (AAL) brain atlas was used to obtain the parcellation, resulting in 90 ROIs, excluding the regions associated with the cerebellum. Instructions on how to upload the data into the toolbox are given in Table 2. A neighboring matrix is also available, based on thresholded Euclidean distances between pairs of ROIs, calculated using the coordinates defined in the Montreal Neurological Institute (MNI) space. The threshold was chosen so that ROIs would have five neighbors on average. This matrix can be used to define the neighboring structures among ROIs for the specification of the MRF prior.

Figure 6 shows the posterior β-maps for one of the subjects, for stimulus 2, obtained at a PPI threshold of 0.9. A multi-slice sagittal view is shown in the middle panel. The smaller slider can be scrolled down if one wishes to see more of the brain structure through the overlayed activation map. The bottom subplot displays the activation map at coordinates $X = 92, Y = 115, Z = 111$. If `View all Stimulus` is selected under `Viewing Options`, then a 3×3 plot of activation maps will be displayed. Different locations of the brain can be examined by using the three sliders to control the X, Y, Z coordinates.

For stimulus 2 and a PPI threshold of .9, Figure 7 shows the dendrogram (middle), obtained by clustering the posterior β estimates of all 30 subjects, and the cluster-level β-maps (bottom) when 3 clusters are selected. The subject numbers corresponding to each cluster are displayed on the interface that controls the dendrogram and activation maps (top).

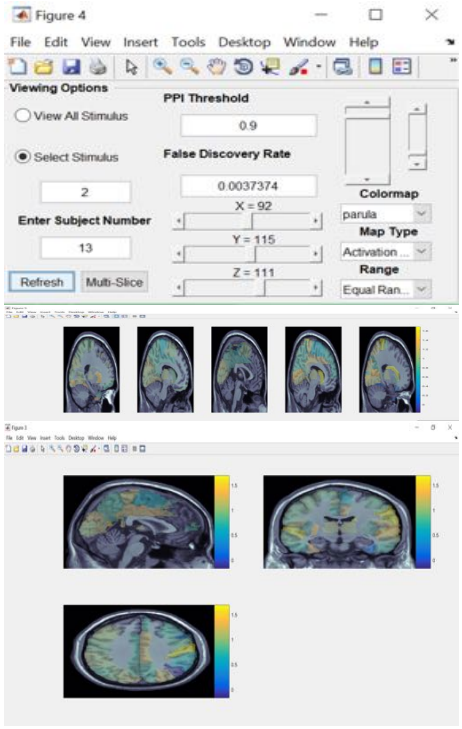

Figure 6: **3D Analysis:** *Example of activation β-maps, with `Range` set to `Equal Range`, for stimulus 2 and PPI threshold of .9. The middle subplot displays a multi-slice sagittal view, the bottom subplot displays the activation map at coordinates $X = 92, Y = 115, Z = 111$.*

Algorithm 2 Instructions for running the example dataset

Once you run *NPBayes_fMRI*, select **Model Fitting** and follow these instructions:

1. Define the output directory by clicking on "OUTPUT"
2. Insert 30 for `Number of Subjects`
3. For `2D or 3D Analysis`, select "3-Dimensional" and insert the following:
   - `Matrix of ROI names`: Select *ROI_names.mat* from the Example_ROIs subfolder
   - `Neighbor Matrix`: Select *nei_vec.mat* from the Example folder
   - `ROI NIFTI Directory`: Select the Example_ROIs folder
   - `Brain Template Image`: Select 'Load Nifti Brain Template Image' and load the *ch2.nii* file found in the Example folder
4. For "Data Files", select:
   multi_data.mat file found on the Example folder
5. For `Parameter Setting` select the default setting by clicking "yes"
6. Initiate model fitting by pressing "Run Model"

To visualize the results, select Visualization and load the *result.mat* file located in the output directory using Load Output.

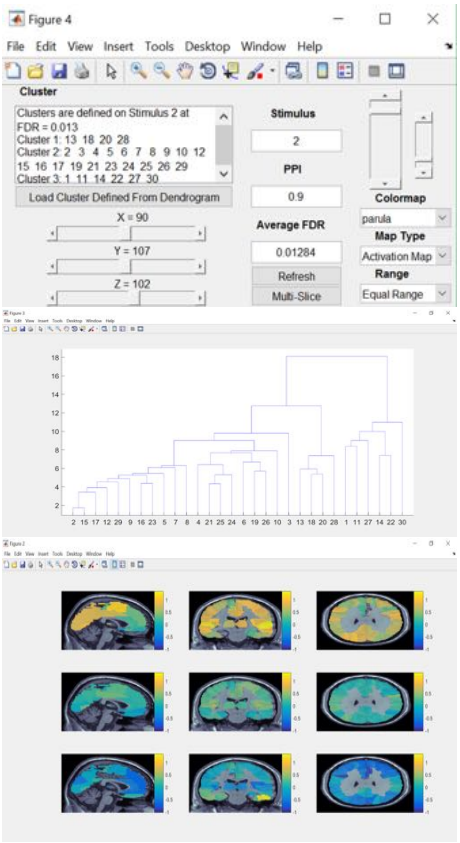

Figure 7: **3D Analysis:** *Example of dendrogram (middle), for stimulus 2 and a PPI threshold of .9, and cluster-level β-maps (bottom), obtained with three clusters. The subject cluster memberships are displayed in the* `Cluster` *tab of the interface (top).*

4 More information

The `NPBayes-fMRI` software is available for download at: https://github.com/rimehi/NPBayes_fMRI and http://www.stat.rice.edu/~marina/software.html

References

[1] F. Bowman, B. Caffo, S. Bassett, and C. Kilts. A Bayesian hierarchical framework for spatial modeling of fMRI data. *NeuroImage*, 39(1):146–156, 2008.

[2] R. Buxton and L. Frank. A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation. *Journal of Cerebral Blood Flow & Metabolism*, 17(1):64–72, 1997.

[3] L. Cerliani, M. Thomas, D. Aquino, V. Contarino, and A. Bizzi. Disentangling subgroups of participants recruiting shared as well as different brain regions for the execution of the verb generation task: A data-driven fMRI data. *Cortex*, 86:247–259, 2017.

[4] M.J. Fadili and E.T. Bullmore. Wavelet-generalised least squares: A new BLU estimator of linear regression models with 1/*f* errors. *NeuroImage*, 15:217–232, 2002.

[5] S. Fischer-Baum, J.H. Kook, Y. Lee, A. Ramos-Nunez, and M. Vannucci. Hetero-

geneity in in the neural and cognitive mechanisms of single word reading. *Frontiers in Human Neuroscience*, in revision, 2017.

[6] G. Flandin and W.D. Penny. Bayesian fMRI data analysis with sparse spatial basis function priors. *NeuroImage*, 34(3):1108–1125, 2007.

[7] K.J. Friston, A.P. Holmes, J.B. Poline, P.J. Grasby, S.C.R. Williams, R.S.J. Frackowiak, and R. Turner. Analysis of fMRI time-series revisited. *NeuroImage*, 2(1):45–53, 1995.

[8] K.J. Friston, P. Jezzard, and R. Turner. Analysis of functional MRI time-series. *Human Brain Mapping*, 1(2):153–171, 1994.

[9] L.M. Harrison and G.G.R. Green. A Bayesian spatiotemporal model for very large data sets. *NeuroImage*, 50(3):1126–1141, 2010.

[10] J.H. Kook, M. Guindani, L. Zhang, and M. Vannucci. NPBayes-fMRI: Nonparametric Bayesian general linear models for single- and multi-subject fMRI data. *Statistics in Biosciences*, pages 1–19, 2017.

[11] K. Lee, G.L. Jones, B.S. Caffo, and S.S. Bassett. Spatial Bayesian variable selection models on functional magnetic resonance imaging time-series data. *Bayesian Analysis*, 9(3):699–732, 2014.

[12] F. Li, T. Zhang, Q. Wang, Gonzalez M.Z., Maresh E.L., and J. Coan. Spatial Bayesian variable selection and grouping in high-dimensional scalar-on-image regressions. *Annals of Applied Statistics*, 9(2):687–713, 2015.

[13] M. Lindquist. The statistical analysis of fMRI data. *Statistical Science*, 23(4):439– 464, 2008.

[14] MATLAB. *version 9.0.0.341360 (R2016a)*. The MathWorks Inc., Natick, Massachusetts, 2016.

[15] F.G. Meyer. Wavelet-based estimation of a semiparametric generalized linear model of fMRI time-series. *IEEE transactions on Medical Imaging*, 22(3):315–322, 2003.

[16] M.A. Newton, A. Noueiry, D. Sarkar, and P. Ahlquist. Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics*, 5(2):155–176, 2004.

[17] W. Penny, S. Kiebel, and K.J. Friston. Variational bayesian inference for fmri time series. *NeuroImage*, 19(3):727–741, 2003.

[18] A. Quirós, R.M. Diez, and D. Gamerman. Bayesian spatiotemporal model of fMRI data. *NeuroImage*, 49(1):442–456, 2010.

[19] N. Sanyal and M.A. Ferreira. Bayesian hierarchical multi-subject multiscale analysis of functional MRI data. *NeuroImage*, 63(3):1519–1531, 2012.

[20] S. Su, B. Caffo, E. Garrett-Mayer, and S. Bassett. Modified test statistics by inter-voxel variance shrinkage with an application to fMRI. *Biostatistics*, 10(2):219–227, 2009.

[21] M.W. Woolrich, M. Jenkinson, J.M. Brady, and S.M. Smith. Fully Bayesian spatio-temporal modeling of fMRI data. *IEEE Transactions on Medical Imaging*, 23(2):213–231, 2004.

[22] G.W. Wornell and A.V. Oppenheim. Estimation of fractal signals from noisy measurements using wavelets. *IEEE transactions on Signal Processing*, 40(3):611–623, 1992.

[23] K.J. Worsley and K.J. Friston. Analysis of fMRI time-series revisited-again. *NeuroImage*, 2:173–181, 1995.

[24] L. Zhang, M. Guindani, F. Versace, J. Englemann, and M. Vannucci. A spatiotemporal nonparametric Bayesian model of multi-subject fMRI data. *Annals of Applied Statistics*, 10(2):638–666, 2016.

[25] L. Zhang, M. Guindani, F. Versace, and M. Vannucci. A spatio-temporal nonparametric Bayesian variable selection model of fMRI data for clustering correlated time courses. *NeuroImage*, 95:162–175, 2014.