

***NPBayes-fMRI*: Non-parametric Bayesian General Linear Models for Single- and Multi-Subject fMRI Data**

**Jeong Hwan Kook¹ · Michele Guindani² ·
Linlin Zhang³ · Marina Vannucci¹**

Received: 27 February 2017 / Accepted: 10 October 2017 / Published online: 31 October 2017
© International Chinese Statistical Association 2017

Abstract In this paper, we introduce *NPBayes-fMRI*, a user-friendly MATLAB GUI that implements a unified, probabilistically coherent non-parametric Bayesian framework for the analysis of task-related fMRI data from multi-subject experiments. The modeling approach is based on a spatio-temporal linear regression model that specifically accounts for the between-subjects heterogeneity in neuronal activity via a spatially informed multi-subject non-parametric variable selection prior. A characteristic feature of the approach is that it results in a clustering of the subjects into subgroups characterized by similar brain responses, while simultaneously producing group-level as well as subject-level activation maps. This is distinct from two-stage “group analysis” approaches traditionally considered in the fMRI literature, which separate the inference on the individual fMRI time courses from the inference at the population level. Here, we first describe the models and a Variational Bayes algorithm

This work was partially supported by NSF-SES 1659921 and NSF-SES-1659925.

✉ Marina Vannucci
marina@rice.edu

Jeong Hwan Kook
jk5@rice.edu

Michele Guindani
michele.guindani@UCI.edu

Linlin Zhang
linlin.zhangrice@gmail.com

¹ Rice University, Houston, TX, USA

² UC Irvine, Irvine, CA, USA

³ Schlumberger, Houston, TX, USA

for posterior inference. Next, we introduce the toolbox and illustrate its features via an example.

Keywords General linear model · MATLAB · Multi-subject fMRI data · Non-parametric variable selection priors · Spatio-temporal modeling · Variational Bayes

1 Introduction

Functional magnetic resonance imaging (fMRI) is a non-invasive 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 Friston et al. [10] and subsequently investigated by many other authors, particularly for single-subject data, see for example [11, 12, 17, 19, 25, 33, 36, 38], among many others. Many of these models incorporate the complex spatial and temporal correlation structure of the fMRI data. Bayesian approaches, in particular, allow flexible modeling of spatial and temporal correlations via suitable prior models and can achieve increased signal detection and fewer false-positive counts with respect to simpler approaches that do not appropriately account for the spatio-temporal variability of the data. See for example [39] for a review of recent Bayesian models.

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 [2, 15, 18, 27, 29]. 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 Cerliani et al. [4], but these approaches still involve multiple estimation steps, and therefore do not properly take into account variability and heterogeneity in the data.

In this paper, we review and extend a unified, single-stage Bayesian approach for the analysis of task-related brain activity proposed by Zhang et al. [40]. 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 non-parametric variable selection prior. This effectively captures correlation among time-series voxels within and across subjects, by inducing clustering among voxels within a subject at one level of the hierarchy and between subjects at the second level. In the fMRI literature, capturing statistical dependence among possibly remote neurophysiological events is often viewed as an aspect of “functional” connectivity [9, 13]. The approach of Zhang et al. [40] further takes into account the spatial prox-

imity of potential activations within a subject by employing a Markov random field (MRF) prior on the selection indicators of the spike-and-slab distribution. Posterior inference is performed via fast Variational Bayes (VB) algorithms. We extend the modeling approach of Zhang et al. [40] to allow for multiple stimuli and different choices of the hemodynamic response function. Furthermore, we consider formulations of the model that can be used to analyze either voxel-based 2D slices or 3D data in the form of brain parcellations. We also discuss the non-parametric prior formulation in the case of a single-subject analysis. Finally, we show how to derive contrast maps based on the VB output.

A well-recognized challenge in the use of Bayesian models with complex methods is the lack of user-friendly software that can be used by practitioners to apply the methods to their experimental data. In an attempt to narrow the gap, we introduce *NPBayes-fMRI*, a MATLAB GUI that implements the non-parametric spatio-temporal models described in the paper. The GUI comprises two components, one for model fitting and another one for visualization of the results. Within the model fitting interface, the user can define the type of analysis (voxel-based or whole-brain parcellation into regions of interest, i.e., ROIs) and the model parameters. Users have the option of a pre-defined default setting for all parameters that also allows customized choices. The GUI also accommodates single-subject analyses. The VB algorithm can be run within the GUI. After running the algorithm, the output file can be uploaded via the visualization interface and used to plot subject-level activation maps, contrasts maps, and cluster-defined averaged β -maps. No additional MATLAB toolboxes are required to run *NPBayes-fMRI*.

The rest of the paper is organized as follows: Section 2 introduces the spatio-temporal model, the non-parametric variable selection prior, and the methods for posterior inference, for both multiple and single-subject data. Section 3 describes the *NPBayes-fMRI* MATLAB GUI. Section 4 provides illustrations of the methods using the MATLAB GUI. Section 5 concludes the paper.

2 Methods

In this section, we review the general framework of the model as proposed by Zhang et al. [40]. We first consider multiple subjects and then describe the simplified model for single-subject analysis. In both cases, we provide the formulation of the model for the general case of multiple tasks.

2.1 Multi-Subject Spatio-Temporal Model

For subject $i = 1, \dots, N$, let $Y_{iv} = (Y_{iv1}, \dots, Y_{ivT})^T$ be the vector of the BOLD response data at voxel v , with $v = 1, \dots, V$. We model the data as

$$Y_{iv} = X_{iv}\beta_{iv} + \varepsilon_{iv}, \quad \varepsilon_{iv} \sim N_T(0, \Sigma_{iv}), \quad (1)$$

where X_{iv} is a known $T \times p$ covariate matrix and $\beta_{iv} = (\beta_{iv1}, \dots, \beta_{ivp})^T$ 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. Additionally, as typical with multiple subjects data, to make the BOLD signal levels consistent across subjects, we consider transformed the data by percent signal change normalization, i.e., $y_{ivt}^* = y_{ivt}/\bar{y}_{iv} \times 100$, where y_{ivt} denotes the BOLD signal of subject i in voxel v at time point t , and \bar{y}_{iv} the mean signal level across the T time points. Let X_{ivj} be the j th column of X_{iv} . Then X_{ivj} is modeled as the convolution of the j -th stimulus pattern with a hemodynamic response function (HRF) [3], that is,

$$X_{ivj}(t) = \int_0^t x_j(s)h_{\lambda_{ivj}}(t-s)ds, \quad (2)$$

where $x_j(s)$ represents the stimulus pattern. One common choice is a Poisson HRF, that is $h_{\lambda_{ivj}} = \exp(-\lambda_{ivj})\lambda_{ivj}^t/t$. The parameter λ_{ivj} 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. Other popular choices are a canonical HRF, that is, $h_{A_{ivj}} = A_{ivj} \left(\frac{t^{\alpha_1-1}\beta_1^{\alpha_1}}{\Gamma(\alpha_1)} - c \frac{t^{\alpha_2-1}\beta_2^{\alpha_2}}{\Gamma(\alpha_2)} \right)$, where $\alpha_1 = 6$, $\alpha_2 = 16$, $\beta_1 = \beta_2 = 1$, $c = 1/6$, and a gamma HRF, that is, $h_{a_{ivj}, b_{ivj}} = \frac{b_{ivj}^{-a_{ivj}}}{\Gamma(a_{ivj})} t^{a_{ivj}-1} \exp(-t/b_{ivj})$ [19].

The error term in equation (1) is modeled as a long memory process. Specifically, the covariance matrix is written as $\sum_{iv}(t, s) = [\gamma(|t-s|)]$, with the auto-covariance function $\gamma(h)$ defined as

$$\gamma(h) \sim Ch^{-\alpha}, \quad (3)$$

with $C > 0$, $0 < \alpha < 1$, and h large. This choice accounts for low-frequency noise which induces slow changes in voxel intensity over time, such as scanner drift, and for physiological noise, due to patient motion, respiration, and heartbeat causing fluctuations in signal across both space and time. In an analysis of single-subject fMRI data, [38] show that such modeling strategy improves the deconvolution of the signal and the noise, leading to the detection of more localized, fewer false-positive, and sparser activations with respect to using auto-regressive error structures.

Discrete wavelet transforms (DWT) are often employed in the fMRI literature as a way to decorrelate the data [6, 16, 20, 27, 38]. After applying the DWT to equation (1) the model in the wavelet domain can be written as

$$Y_{iv}^* = \sum_{j=1}^p X_{ivj}^* \circ \beta_{ivj} + \varepsilon_{iv}^*, \quad \varepsilon_{iv}^* \sim N_T(0, \Sigma_{iv}^*), \quad (4)$$

with \circ the element-by-element (Hadamard) product, and where W is a $T \times T$ matrix corresponding to the wavelet transform, $Y_{iv}^* = WY_{iv}$, $X_{iv}^* = WX_{iv}$, and $\varepsilon_{iv}^* = W\varepsilon_{iv}$, and with the covariance matrix Σ_{iv}^* approximately diagonal with elements $\psi_{iv}\sigma_{imn}^2$ indicating the variance of the n th wavelet coefficient at the m th scale. We follow the variance progression method of Wornell and Oppenheim [35] for the wavelet coefficients,

$$\psi_{iv}\sigma_{imn}^2 = \psi_{iv}(2^{\alpha_{iv}})^{-m}, \quad (5)$$

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

2.2 Non-parametric Variable Selection Prior

Detecting voxels that activate in response to a stimulus is equivalent to identifying the non-zero regression coefficient β_{ivj} in model (4). Zhang et al. [40] embed the selection into a clustering framework, effectively defining a multi-subject non-parametric variable selection prior with spatially informed selection within each subject. More specifically, they employ a hierarchical Dirichlet Process (HDP) prior [31], which implies that the non-zero β 's within subject i are drawn from a mixture model and possibly shared between subjects. The HDP prior construction effectively captures correlation among time-series voxels within and across subjects, by inducing clustering among voxels within a subject at one level of the hierarchy and between subjects at the second level. This allows, in particular, to capture spatial correlation among potential activations of distant voxels, within a subject, while simultaneously borrowing strength in the estimation of the parameters from subjects with similar activation patterns. For the multi-stimuli formulation of the model, let γ_{ivj} be a binary indicator of whether a given voxel is activated or not under stimulus j , that is, $\gamma_{ivj} = 0$ if $\beta_{ivj} = 0$ and $\gamma_{ivj} = 1$ otherwise. A spiked non-parametric prior is imposed on the coefficients β_{ivj} , i.e., a spike-and-slab prior where the slab distribution is modeled as a non-parametric prior, as

$$\beta_{ivj} | \gamma_{ivj}, G_i \sim \gamma_{ivj} G_{ij} + (1 - \gamma_{ivj}) \delta_0, \quad (6)$$

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 non-parametric slab,

$$\begin{aligned} G_{ij} | \eta_1, G_0 &\sim DP(\eta_1, G_0) \\ G_0 | \eta_2, P_0 &\sim DP(\eta_2, P_0) \\ P_0 &= N(0, \tau), \end{aligned} \quad (7)$$

where τ, η_1, η_2 are fixed parameters and P_0 is the base distribution. Parameters η_1, η_2 control the variability of the coefficients at the subject and population level, respectively. The HDP prior consists of two levels of hierarchy, which induce 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. Using both simulated and real data, [40] show increased detection power and lower numbers of false-positive calls with respect to common two-stage estimation approaches which separate the inference on the individual fMRI time courses from the inference at the population level.

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 γ_{ivj} ,

$$P(\gamma_{ivj}|d, e, \gamma_{ikj}, k \in N_{iv}) \sim \exp \left(\gamma_{ivj} \left(d + e \sum_{k \in N_{iv}} \gamma_{ikj} \right) \right), \quad (8)$$

with N_{iv} the set of neighboring voxels of voxel v for subject i , and $p(\gamma_{iv}) = \prod_{j=1}^p p(\gamma_{ivj})$. This prior reduces to an independent Bernoulli with parameter $\exp(d)/[1 + \exp(d)]$ if a voxel does not have any neighbors. The sparsity parameter $d \in (-\infty, \infty)$ in (8) 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 use of MRF priors has become quite popular in recent years in the Bayesian modeling of fMRI data [17,28,37,38].

The prior model is completed by considering a uniform prior distribution on the delay parameter, $\lambda_{ivj} \sim U(u_1, u_2)$, for a Poisson HRF, or the amplitude parameter, $A_{ivj} \sim U(u_1, u_2)$, for a Canonical HRF, or the shape and scale parameters, $a_{ivj} \sim U(u_1, u_2)$, $b_{ivj} \sim U(u_3, u_4)$, for a gamma HRF. Also, an Inverse Gamma (IG) prior is imposed on the innovation variance parameter, $\psi_{iv} \sim \text{IG}(a_0, b_0)$, and a Beta distribution on the long memory parameter, $\alpha_{iv} \sim \text{Beta}(a_1, b_1)$.

2.3 Single-Subject Modeling

For a single subject, the non-parametric prior reduces to a Dirichlet process (DP) prior of the type

$$\begin{aligned} G_0|\eta, P_0 &\sim DP(\eta, P_0) \\ P_0 &= N(0, \tau), \end{aligned} \quad (9)$$

where τ, η are fixed and P_0 is the base distribution. The mass parameter η is used to regulate the variability of the coefficients at the subject level. This prior allows, in particular, to capture spatial correlation among activations of distant voxels. In the fMRI literature, capturing statistical dependence among possibly remote neurophysiological events is often viewed as an aspect of “functional” connectivity [9,13].

2.4 Prior Specification

We provide here some general guidelines on the choice of the prior hyperparameters. Zhang et al. [40] comment on the sensitivity of their results to different choices of the priors when using simulated data. The authors notice that, in general, modest changes of the values of the variance parameter τ in the base measure of the HDP prior and of the hyperparameters a_0, b_0, a_1, b_1 , of the priors on the variance parameters ψ 's and

the long memory parameters α 's, did not affect the accuracy of the estimation results. Consequently, these parameters can be set to default values that correspond to vague or uninformative priors. Employing non-informative or vague priors is a common choice in Bayesian statistics in the absence of prior knowledge about the unknown parameters. Here, in particular, the hyperparameter τ can be set to a large value (e.g., > 50), non-informative priors can be specified on the long memory parameters by setting $a_1 = b_1 = 1$ and vague priors on the innovation variance parameters by setting $a_0 = 3, b_0 = 2$. As for the concentration parameters η_1 and η_2 of the HDP prior, larger values of these parameters tend to generate larger numbers of components across and within-subjects when fitting the model. We recommend using a non-informative setting by setting $\eta_1 = \eta_2 = 1$. Also, vague specifications can be adopted on the parameters of the HRF, for example, by specifying a uniform prior with, say, $u_1 = 0, u_2 = 8$, in the case of a Poisson HRF, and similarly for the other HRFs.

Some sensitivity should be expected in regards to the MRF parameters. In particular, as noted by [40], larger values of d or e lead to lower FNRs, at the expense of higher FPRs and lower precisions. We suggest fixing d to reflect a prior belief in a sparse model. For example, $d = -2.5$ implies that the prior probability of activation is less than 10% when a voxel has no neighbors. Also, a value of e in the range (0.3–0.5) generally results in values below the phase transition point, which can be estimated using the algorithm proposed by Propp and Wilson [24].

2.5 Model Fitting by Variational Bayes

Variational Bayes (VB) algorithms are an alternative method for posterior inference that, unlike MCMC methods, does not rely on numerical integration. Variational Bayes methods have been employed successfully in Bayesian models for single-subject fMRI data [8, 14, 22, 23, 34]. 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 [1, 26]. 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. Indeed, [40] perform a thorough comparison between MCMC and VB on simulated data, showing very good performance of the VB algorithm in the estimation of the model parameters. They also notice a remarkable improvement in computing time, with 1000 MCMC iterations taking approximately 7 h, on a double core @Intel @Xeon processor with 16 GB of memory, 2.2 GHz., while a VB with 50 inner loop iterations and 100 outer loop iteration took approximately 34 min.

When using VB methods within HDP frameworks, such as the spiked HDP prior distribution (7) on the β 's parameters, it is beneficial to employ the truncated stick-breaking construction, to exploit conjugacy and allow for analytically tractable updates of the parameters [32]. In our model formulation, the parameters of the HRF appear through convolution (2) and the α 's via the variance progression formula (3). This makes it impossible to derive analytically tractable updates for these parameters. Zhang et al. [40] address the problem by combining the VB algorithm with an importance

Table 1 VB Algorithm (with Poisson HRF)

for $l = 1 : L$
 Update α_{iv}^l and λ_{ivj}^l $i = 1, \dots, N, v = 1, \dots, V, j = 1, \dots, P$, via importance sampling.
for $m = 1 : M$
 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.

sampling procedure. The resulting algorithm has two major components. The first component (inner loop) approximates the posterior distribution of the regression coefficients, the selection parameters, and the innovation variance parameters via mean field variational inference with a coordinate ascent algorithm. The second component (outer loop) estimates the parameters of the HRF and the α 's via importance sampling, with the importance sampling weights calculated based on the optimal solution from the first component. A schematic representation of the algorithm is given in Table 1.

2.6 Posterior Inference

For posterior inference, primary interest 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 , $p(\gamma_{ivj} = 1)$, for $j = 1, \dots, P$, are approximated as weighted averages of the variational distribution values $q(\gamma_{ivj} = 1)$ 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) [5, 21, 30]. For subject i and stimulus j , the Bayesian FDR is defined as

$$FDR_{ij}(\kappa_i j) = \frac{\sum_{v=1}^V (1 - PPI_{ivj}) I_{(PPI_{ivj} > \kappa_{ij})}}{\sum_{v=1}^V I_{(PPI_{iv} > \kappa_i)}}, \quad (10)$$

where PPI_{ivj} is the PPI for subject i at voxel v and stimulus j , and $I_{(PPI_{ivj} > \kappa_{ij})}$ is the indicator function such that $I_{(PPI_{ivj} > \kappa_{ij})} = 1$ if $PPI_{ivj} > \kappa_{ij}$ and 0 otherwise, with κ_i a threshold value. In the data analyses, one can set the FDR to a pre-specified value, typically 0.05 or 0.1, and then choose κ_i accordingly. 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 our modeling approach is that the use of the non-parametric HDP prior construction (6) 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 as

$$d_{i,i'} = \sqrt{(\hat{B}_i - \hat{B}_{i'})^T (\hat{B}_i - \hat{B}_{i'})},$$

with \hat{B}_i denoting the posterior estimate of $B_i = (\beta_{i1j}, \dots, \beta_{iVj})^T$. The dissimilarity matrix can then be transformed into a tree via hierarchical clustering and 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 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. Within the VB framework, a contrast map can be obtained by thresholding the probabilities

$$p_v = \sum_{l=1}^L \hat{w}_{vl} I_{((\sum_{j=1}^J \pi_j \cdot \beta_{vjl}) > \kappa)}, \tag{11}$$

with $\pi = \pi_1, \dots, \pi_J$ a contrast weight vector summing to 0, L the number of outer loop VB iterations, B_{vjl} the $V \times J \times L$ matrix storing the updated β value for all voxels and stimuli, across the L iterations, and \hat{w}_{vl} the normalized importance weights. For each subject i and outer loop iterations $l = 1, \dots, L$, the importance weight is computed as

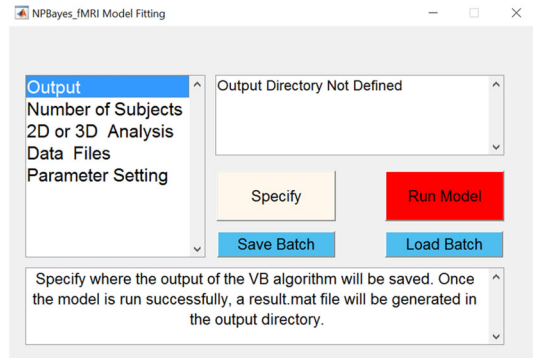
$$w_{vl} = \frac{q(\alpha_{iv}^{(l)}, \lambda_{iv}^{(l)})}{\tilde{p}(\alpha_{iv}^{(l)}, \lambda_{iv}^{(l)})}, \tag{12}$$

with $q(\alpha_{iv}^{(l)}, \lambda_{iv}^{(l)})$ the variational distribution of $(\alpha_{iv}^{(l)}, \lambda_{iv}^{(l)})$ and $\tilde{p}(\alpha_{iv}^{(l)}, \lambda_{iv}^{(l)})$ the importance sampling density of $(\alpha_{iv}^{(l)}, \lambda_{iv}^{(l)})$. Once all outer loop iterations terminate, the importance weight is normalized to obtain \hat{w}_{vl} .

3 The *NPBayes-fMRI* GUI

We now provide a detailed description of the *NPBayes-fMRI* MATLAB GUI that implements the spatio-temporal general linear regression models described in the

Fig. 1 *NPBayes-fMRI*: Main interface of the *Model Fitting* GUI



previous Section. The GUI comprises two main interfaces, one for model fitting and one for the visualization of the results.

3.1 Model Fitting

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

Output: The user is asked to specify the directory where the output of the VB algorithm will be saved. Once the model is run successfully, a `result.mat` file will be generated in the output directory.

Number of Subjects: The user is asked to specify the number of subjects that are being used for the analysis. When this variable is set to 1, a DP will be used for the slab distribution in equation (6), 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. If 2D is selected, then the user is prompted to define the threshold for 2D image and the dimension of the 2D slice. The threshold is a value used to define the gray matter mask on the fMRI data (see Data Files below). The parameter dimension of the 2D slice is defined by the number of rows and columns of the fMRI slice. The number of rows and columns must be such that their product is equal to V . These arguments will be used later for visualization of the results. If 3D is selected, the following arguments must be defined:

Matrix of ROI names: This is a `.mat` file containing a variable named `ROI_names`, which is a $1 \times V$ cell, with each cell entry containing the brain regions names as defined by the parcellation, with the `.nii` extension (e.g., `Amygdala_L.nii`).

Neighbor matrix: This is a `.mat` file containing a variable named `nei_vec`, whose first row contains the indices of those regions that are neighbors to the region whose index appears on the second row. For example, if region 5 has regions 6, 10, 15, 20 as its neighbors, then `nei_vec` will have entries $\begin{bmatrix} 6 & 10 & 15 & 20 \\ 5 & 5 & 5 & 5 \end{bmatrix}$.

This information is also used to specify the MRF prior (8) by calculating Euclidean distances between the centroids of the ROIs.

ROI NIFTI Directory: This is the directory where all NIFTI files for the ROIs are stored. The names of these files should be equal to the names provided by the *Matrix of ROI names*. These files will be used to map back the inference results into 3-Dimension for visualization.

Brain Template Image: For better visualization of the inference results, the user has the option to upload a brain template NIFTI file. When available, the brain template will be used as a background image, and the resulted image from running the model will be overlaid on top. This will only work when the dimension of the brain template file is equivalent to the dimensions of the ROI NIFTI images. If a template image is not provided, the visualizations will take place without a background image.

Data Files: The user needs to load a *.mat* file that consists of two matrices: *xt_dat*, a $T \times P$ binary design matrix, with T the number of time points and P the number of stimuli, and *y_dat*, a $T \times (N \times V)$ matrix of BOLD signals, with N the number of subjects and V the number of voxels (for 2D analyses) or ROIs (for 3D analyses). So if we let y_i be a $N \times V$ BOLD signals for subject i , $y_dat = (y_1, y_2, \dots, y_N)$. For 2D analyses, a gray matter mask is applied to the fMRI data and inference is performed based on those voxels where *y_dat* is greater than the threshold specified by the user. For 3D analyses, *y_dat* should contain the voxel time-series data averaged by ROI, listed in the same order provided by the *Matrix of ROI names .mat* file. 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 software includes a pre-defined default setting for all hyperparameters and VB parameters. The user also has the option of setting some of the model parameters manually, including selecting the type of HRF distribution, the prior setting for the HRF parameters and the MRF parameters, and the number of VB iterations. In the default parameter setting, the Poisson distribution is automatically selected as the HRF distribution, with its default prior setting.

Once all the variables have been defined, the *Run Model* button in the GUI will turn green from red, and pressing the button will start the algorithm. If one chooses to run the model later, it is also possible to press the *Save Batch* button to store the model specification. In this case all the parameter settings are saved in a *.mat* file, which can later be loaded by clicking on the *Load Batch* button.

3.2 Visualization

The *NPBayes-fMRI: Visualization* GUI, shown in Fig. 2, is used to visualize the results using the *result.mat* file obtained from running the algorithm using *NPBayes-fMRI: Model Fitting* interface. The GUI comprises the three components described below. For simplicity, we consider the case of 3D data only.

Fig. 2 *NPBayes-fMRI*: Main interface of the *Visualization* GUI

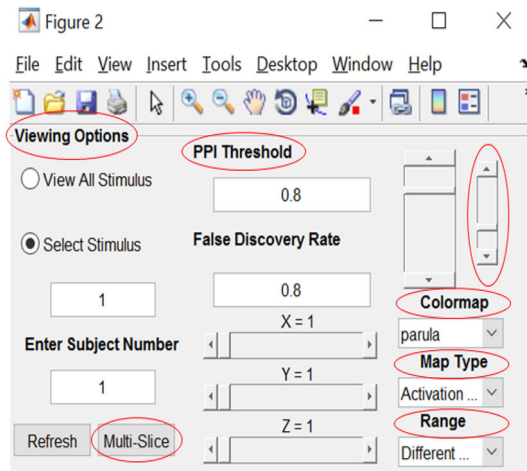
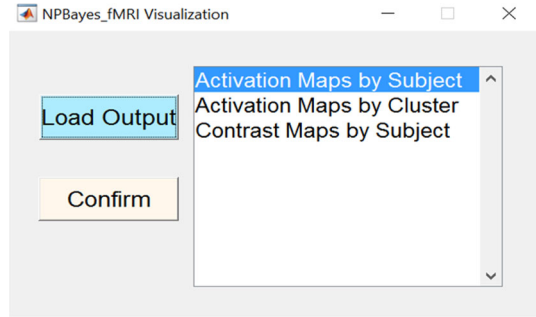


Fig. 3 *NPBayes-fMRI: Visualization*. Interface for viewing activation maps by subject. The *Viewing Options* tab allows the user to view all stimuli at once or one at a time. By clicking on the *Map Type*, *Range* and *Colormap* pop-up menus, the user can define the type of maps to visualize, set the axes ranges and the desired colormap setting. The bigger slider adjusts the PPI threshold and FDR value, the smaller one controls the transparency of the activation map when a Brain Template Image has been uploaded for 3D Analysis. The *Multi-Slice* option can be used to view multiple slices of the brain in one particular orientation for a given stimulus

Activation Maps by Subject (Fig. 3): 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. Depending on the HRF distribution, the HRF map will display values for the λ_{ivj} , A_{ivj} , or $a_{ivj} \cdot b_{ivj}$ (as the mean of a gamma distribution) parameters, for the Poisson, Canonical, and Gamma HRFs, respectively. MATLAB's built-in colormaps can be selected via the *Color Map* pop-up menu. *Range* lets the user define the axes limits. When set to *Equal_Range*, all figures will be defined on the same range. This may be useful when one is comparing posterior β -maps across all stimuli. If, however, the β values

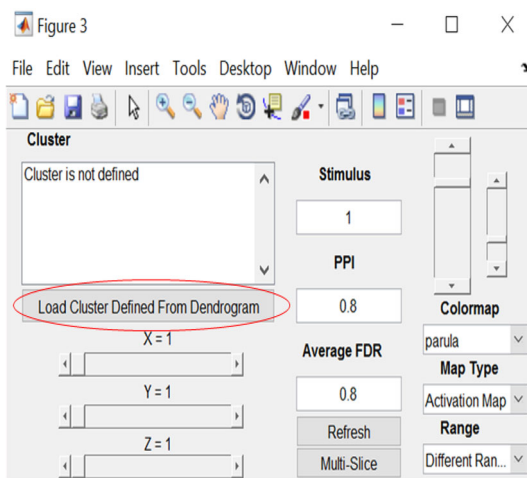
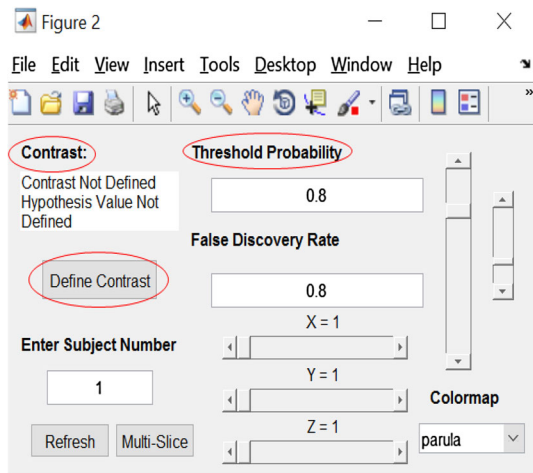


Fig. 4 NPBayes-fMRI: Visualization. Interface for viewing activation maps by cluster, for a given stimulus. The user selects the stimulus and the PPI (or FDR) threshold. The corresponding dendrogram will be displayed, and the user can then specify the number of clusters and click on the Load Cluster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

Fig. 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



are significantly smaller in one stimulus than another, then the Different_Range option is preferable when inspecting the activation maps. 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 Fig. 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,

Table 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:
 - a) *data.mat* file found on the Example folder
- For long time points select:
- a) *multi_data.mat* file found on the Example folder for the multiple stimulus case
 - b) *single_data.mat* file found on the Example folder for the single stimulus case
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.

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.

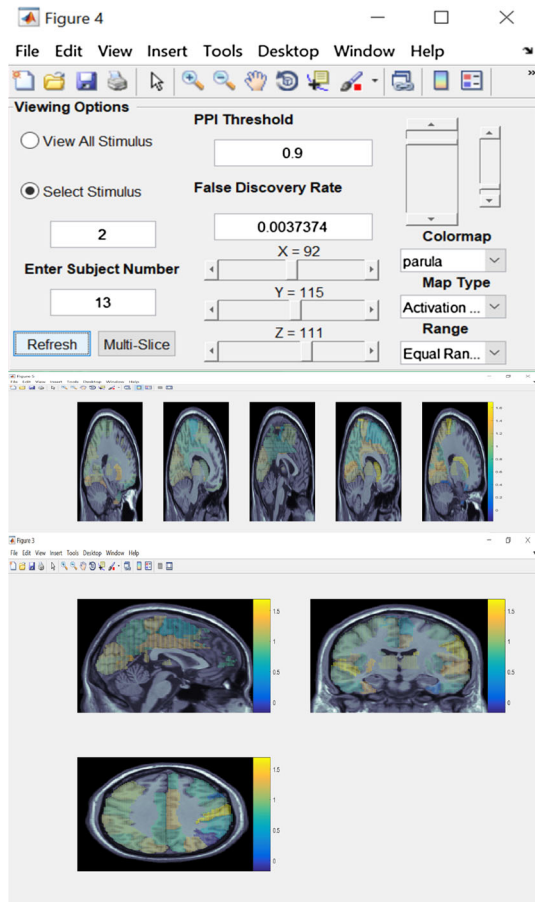
Activation Maps by Cluster (Fig. 4): This function is used to view cluster-level activation maps, for a given stimulus and PPI (or FDR) threshold. Clusters are defined based on a dendrogram obtained by applying hierarchical clustering with Ward’s linkage method to a dissimilarity matrix defined based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

Contrast Maps by Subject (Fig. 5): 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.

4 Illustration with 3D Data

In order to better illustrate the features of the *NPBayes_fmri: Visualization* GUI, we show some selected output from an analysis of a dataset consisting of fMRI data on 30

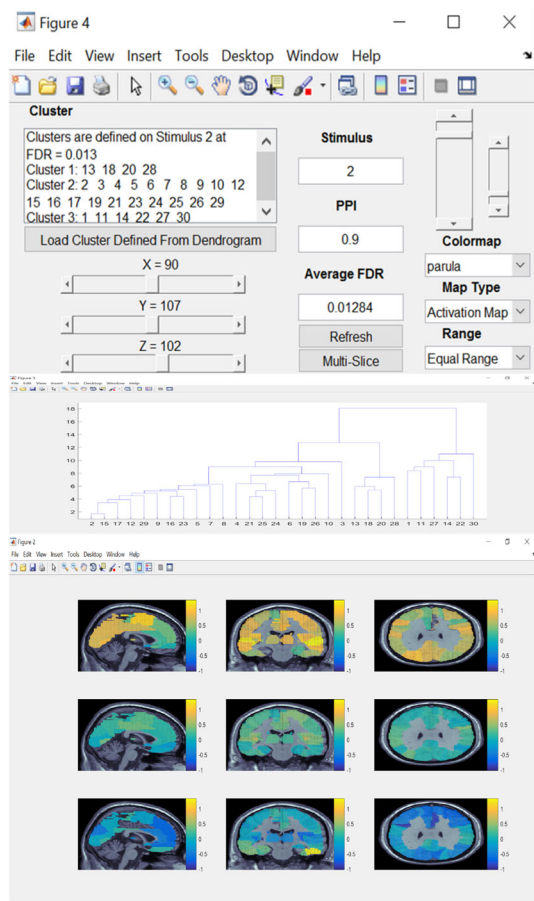
Fig. 6 3D Analysis: Example of activation β -maps, with Range set to Equal Range, for stimulus 2 and PPI threshold of 0.9. The middle subplot displays a multi-slice sagittal view, the bottom subplot displays the activation map at coordinates $X = 92$, $Y = 115$, $Z = 111$



subjects performing an experiment with three stimuli. The dataset is part of a pilot study on variability in the cognitive and neural processes involved in reading, conducted at Rice University [7]. 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. Euclidean distances between pairs of ROIs were calculated using the coordinates defined in the Montreal Neurological Institute (MNI) space, and a neighboring matrix was calculated by thresholding the distances. The threshold was chosen so that ROIs would have five neighbors on average. This matrix was then used to define the neighboring structures among ROIs for the specification of the MRF prior given in equation (8). Instructions on how to upload the data into the toolbox are given in Table 2.

Results shown here were obtained by running the *NPBayes-fMRI: Model Fitting* interface with the default hyperparameter setting. Fig. 6 shows the posterior β -maps for one of the subjects, for stimulus 2, obtained at a PPI threshold of 0.9. The mid-

Fig. 7 3D Analysis: Example of dendrogram (middle), for stimulus 2 and a PPI threshold of 0.9, and cluster-level β -maps (bottom), obtained with three clusters. The subject cluster memberships are displayed in the Cluster tab of the interface (top)



dle subplot displays a multi-slice sagittal view at $X = [70, 80, 90, 100, 110]$. The smaller slider can be scrolled down if one wishes to see more of the brain structure through the overlaid 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 0.9, Fig. 7 shows the dendrogram (middle) obtained by clustering the posterior β estimates and the cluster-level β -maps (bottom) when three clusters are selected. The subject numbers corresponding to each cluster are displayed on the interface that controls the dendrogram and activation maps (top). Finally, Fig. 8 displays the estimated contrast probability map $p(\beta_2 - \beta_3 > 0) = p(\beta_2 > \beta_3)$, for one of the subjects, for a threshold probability of 0.9, corresponding to a FDR value of 0.0096107.

Fig. 8 3D Analysis: Example of contrast map, for subject 10 and a threshold probability of 0.9, corresponding to an FDR of 0.0096107



5 Conclusions

In this paper, we have described a unified, probabilistically coherent framework for the analysis of task-related brain activity in multi-subject fMRI experiments. The model, proposed by Zhang et al. [40], builds upon the large literature on spatio-temporal linear regression models by specifically accounting for within- and between-subjects heterogeneity via a non-parametric Bayesian variable selection prior. Furthermore, posterior inference is carried out via a variational Bayes algorithm that allows scalability. Zhang et al. [40] demonstrate that this probabilistically coherent modeling approach can improve estimation performance with respect to two-stage approaches. They also show, on real data, that a multi-subject modeling strategy leads to a more accurate detection of the activated areas than single-subject models. This is an important stepping stone in the development of reliable detection methods that can be applied to full brain datasets and complex experimental designs.

Here we have also introduced *NPBayes-fMRI*, a MATLAB GUI that implements the non-parametric models proposed by Zhang et al. [40]. The GUI comprises two components: one for model fitting and another for visualization of the results. Within the model fitting interface, the user can define the type of analysis (voxel-based or whole-brain parcellation into regions of interest, i.e., ROIs) and the model parameters. Users have the option of a pre-defined default setting for all parameters that also allows customized choices. The GUI also accommodates single-subject analyses. The VB algorithm can be run within the GUI or in batch mode. After running the algorithm, the output file can be uploaded via the visualization interface and used to plot subject-level activation maps, contrasts maps, and cluster-defined averaged β -maps. The toolbox is available for download at <https://github.com/marinavannucci> and at <https://github.com/rimehi>. Detailed instructions on how to use the toolbox can be found in the Instructions text file. No additional MATLAB toolboxes are required to run *NPBayes-fMRI*.

References

1. Bishop C (2006) Pattern recognition and machine learning. Springer, Berlin
2. Bowman F, Caffo B, Bassett S, Kilts C (2008) A Bayesian hierarchical framework for spatial modeling of fMRI data. *NeuroImage* 39(1):146–156
3. Buxton R, Frank L (1997) A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation. *J Cereb Blood Flow Metab* 17(1):64–72
4. Cerliani L, Thomas M, Aquino D, Contarino V, Bizzi A (2017) 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
5. Efron B (2008) Microarrays, empirical Bayes and the two-groups model. *Stat Science* 23(1):1–22
6. Fadili M, Bullmore E (2002) Wavelet-generalised least squares: a new BLU estimator of linear regression models with 1/f errors. *NeuroImage* 15:217–232
7. Fischer-Baum S, Kook J, Lee Y, Ramos-Nunez A, Vannucci M (2017) Sight or sound? Heterogeneity in the neural and cognitive mechanisms of single word reading (submitted).
8. Flandin G, Penny W (2007) Bayesian fMRI data analysis with sparse spatial basis function priors. *NeuroImage* 34(3):1108–1125
9. Friston K (2011) Functional and effective connectivity: a review. *Brain Connect* 1(1):13–36
10. Friston K, Jezzard P, Turner R (1994) Analysis of functional MRI time-series. *Human Brain Mapp* 1(2):153–171
11. Friston K, Holmes A, Poline J, Grasby P, Williams S, Frackowiak R, Turner R (1995) Analysis of fMRI time-series revisited. *NeuroImage* 2(1):45–53
12. Friston K, Penny W, Phillips C, Kiebel S, Hinton G, Ashburner J (2002) Classical and Bayesian inference in neuroimaging: theory. *NeuroImage* 16:465–483
13. Friston KJ (1994) Functional and effective connectivity in neuroimaging: a synthesis. *Human Brain Mapp* 2:56–78
14. Harrison L, Green G (2010) A Bayesian spatiotemporal model for very large data sets. *NeuroImage* 50(3):1126–1141
15. Holmes A, Friston K (1998) Generalisability, random effects & population inference. *NeuroImage* 7:S754
16. Jeong J, Vannucci M, Ko K (2013) A wavelet-based Bayesian approach to regression models with long memory errors and its application to fMRI data. *Biometrics* 69:184–196
17. Lee K, Jones G, Caffo B, Bassett S (2014) Spatial Bayesian variable selection models on functional magnetic resonance imaging time-series data. *Bayesian Anal* 9(3):699–732
18. Li F, Zhang T, Wang Q, MZ G, EL M, Coan J (2015) Spatial Bayesian variable selection and grouping in high-dimensional scalar-on-image regressions. *Ann Appl Stat* 9(2):687–713
19. Lindquist M (2008) The statistical analysis of fMRI data. *Stat Sci* 23(4):439–464

20. Meyer F (2003) Wavelet-based estimation of a semiparametric generalized linear model of fMRI time-series. *IEEE Trans Med Imaging* 22(3):315–322
21. Newton M, Noueir A, Sarkar D, Ahlquist P (2004) Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics* 5(2):155–176
22. Penny W, Kiebel S, Friston K (2003) Variational bayesian inference for fmri time series. *NeuroImage* 19(3):727–741
23. Penny W, Trujillo-Barreto N, Friston K (2005) Bayesian fMRI time series analysis with spatial priors. *NeuroImage* 24(2):350–362
24. Propp J, Wilson D (1996) Exact sampling with coupled markov chains and applications to statistical mechanics. *Random Struct Algorithms* 9(1):223–252
25. Quirós A, Diez R, Gamerman D (2010) Bayesian spatiotemporal model of fMRI data. *NeuroImage* 49(1):442–456
26. Rue H, Maritno S, Chopin N (2009) Approximate Bayesian inference for latent Gaussian models by using integrate nested Laplace approximations. *J Royal Stat Soc Ser B* 71:319–392
27. Sanyal N, Ferreira M (2012) Bayesian hierarchical multi-subject multiscale analysis of functional MRI data. *NeuroImage* 63(3):1519–1531
28. Smith M, Fahrmeir L (2007) Spatial Bayesian variable selection with application to functional magnetic resonance imaging. *J Am Stat Assoc* 102(478):417–431
29. Su S, Caffo B, Garrett-Mayer E, Bassett S (2009) Modified test statistics by inter-voxel variance shrinkage with an application to fMRI. *Biostatistics* 10(2):219–227
30. Sun W, Reich B, Cai T, Guindani M, Schwartzman A (2015) False discovery control in large-scale spatial multiple testing. *J Royal Stat Soc* 77(1):59–83
31. Teh Y, Jordan M, Beal M, Blei D (2006) Hierarchical Dirichlet processes. *J Am Stat Assoc* 101(476):1–8
32. Wang C, Paisley J, Blei D (2011) Online variational inference for the hierarchical Dirichlet process. *International Conference on Artificial Intelligence and Statistics* pp 752–760
33. Woolrich M, Jenkinson M, Brady J, Smith S (2004) Fully Bayesian spatio-temporal modeling of fMRI data. *IEEE Trans Med Imaging* 23(2):213–231
34. Woolrich M, Behrens T, Smith S (2004b) Constrained linear basis sets for HRF modelling using variational Bayes. *NeuroImage* 21(4):1748–1761
35. Wornell G, Oppenheim A (1992) Estimation of fractal signals from noisy measurements using wavelets. *IEEE Trans Signal Process* 40(3):611–623
36. Worsley K, Friston K (1995) Analysis of fMRI time-series revisited-again. *NeuroImage* 2:173–181
37. Xia J, Liang F, Wang Y (2009) fMRI analysis through Bayesian variable selection with a spatial prior. *IEEE International Symposium on Biomedical Imaging* pp 714–717
38. Zhang L, Guindani M, Versace F, Vannucci M (2014) A spatio-temporal nonparametric Bayesian variable selection model of fMRI data for clustering correlated time courses. *NeuroImage* 95:162–175
39. Zhang L, Guindani M, Vannucci M (2015) Bayesian models for fMRI data analysis. *WIREs Comput Stat* 7(1):21–41
40. Zhang L, Guindani M, Versace F, Englemann J, Vannucci M (2016) A spatiotemporal nonparametric Bayesian model of multi-subject fMRI data. *Ann Appl Stat* 10(2):638–666