

A Scalable Multi-Resolution Spatio-Temporal Model for Brain Activation and Connectivity in fMRI Data

Stefano Castruccio^{1,*}, Hernando Ombao² and Marc G. Genton²

Department of Applied and Computational Mathematics and Statistics, University of Notre Dame,
153 Hurley Hall, Notre Dame, Indiana 46556, U.S.A.
Statistics Program, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900,
Saudi Arabia
**email*: scastruc@nd.edu

SUMMARY. Functional Magnetic Resonance Imaging (fMRI) is a primary modality for studying brain activity. Modeling spatial dependence of imaging data at different spatial scales is one of the main challenges of contemporary neuroimaging, and it could allow for accurate testing for significance in neural activity. The high dimensionality of this type of data (on the order of hundreds of thousands of voxels) poses serious modeling challenges and considerable computational constraints. For the sake of feasibility, standard models typically reduce dimensionality by modeling covariance among regions of interest (ROIs)—coarser or larger spatial units—rather than among voxels. However, ignoring spatial dependence at different scales could drastically reduce our ability to detect activation patterns in the brain and hence produce misleading results. We introduce a multi-resolution spatio-temporal model and a computationally efficient methodology to estimate cognitive control related activation and whole-brain connectivity. The proposed model allows for testing voxel-specific activation while accounting for non-stationary local spatial dependence within anatomically defined ROIs, as well as regional dependence (between-ROIs). The model is used in a motor-task fMRI study to investigate brain activation and connectivity patterns aimed at identifying associations between these patterns and regaining motor functionality following a stroke.

KEY WORDS: Big data; Brain imaging; Functional magnetic resonance image; Gaussian processes; Multi-resolution model; Space-time statistics.

1. Introduction

Detecting and understanding significant patterns of brain activity is among the most important challenges of contemporary science. To this end, functional Magnetic Resonance Imaging (fMRI) data has been the primary tool of neuroscience investigations for the last 20 years. It measures brain activity by detecting changes in neural activity associated with blood flow using the contrast between deoxygenated hemoglobin (which is paramagnetic) to oxygenated hemoglobin (which is diamagnetic) in localized spatial volumes called *voxels*. An fMRI scan produces a highly spatially resolved brain imaging data set. The challenges for statisticians are to develop a model that is able to detect activated voxels and regions of interest (ROIs) during a performed task; to understand how each ROI functions as a unit and how they all work together to execute tasks (there may be as many as 150,000 voxels over the entire brain volume for each scan); and to describe potentially complex spatial dependence structures while still allowing for scalable inference.

In this article, we report the results of a collaboration with the stroke rehabilitation center at UC Irvine (PI: Dr. Steven C. Cramer) on a project that aims to identify brain

activation and connectivity patterns during the execution of a motor task (e.g., hand grasping). Throughout this article, connectivity will be defined as the conditional dependence across ROIs. The ultimate goal is to find associations between these patterns and the ability of an affected stroke patient to regain motor functionality.

The simplest statistical approach in analyzing fMRI data is to fit a linear model (often termed *general linear model* in brain imaging literature) for voxel-specific (or ROI-specific) time series. However, this approach does not account for spatial dependence across voxels or between ROIs, and hence could potentially result in misleading conclusions. It is well known that if spatial dependence is ignored, then the uncertainty of the estimators is assessed incorrectly, thus inducing inflated Type I error rates (Dubin, 1988) when testing for significance. While it is possible to partially adjust the analysis for spatial correlation (e.g., by post-processing the data via spatial smoothing), our goal in this article is to develop a comprehensive model that defines the spatio-temporal dependence of a complete fMRI data set which is necessary to fully and correctly account for these effects. Although most researchers acknowledge that taking spatial dependence into account is important, the key obstacle has been the seemingly insurmountable computational cost. Here, we develop a computationally efficient algorithm and a step-wise inference approach that overcomes this major limitation in spatio-temporal models for fMRI data.

[Corrections added on August 8, 2018, after first online publication: Grammar correction within section 1. Introduction, Figure 1 caption updated, equation corrections within sections 3.2. Step 2: Estimating Local Effects and 4. Simulation Studies, and spelling correction within section 7. Supplementary Materials]

The earliest approaches to modeling fMRI data (Worsley and Friston, 1995; Locascio et al., 1997; Bullmore et al., 2001) focused only on within-voxel temporal correlation, either ignoring spatial correlation or imposing unrealistic constraints to reduce the computational burden associated with the high spatial dimensionality of fMRI data. Seminal works on statistical methods for fMRI data (the random field theory developed in Worsley et al. (1992) and Worsley and Friston (1995)) indirectly accounts for spatial correlation by assuming that the voxel-specific test statistics (e.g., the t -statistic or F -statistic) are realizations of some random field. This approach is unrealistic because it does not directly model spatial covariance via the fMRI time series at different voxels and is prone to subjective interpretation of the clustering of activation patterns.

Bowman (2005) proposed a two-stage hierarchical Bayesian approach to first estimate activation patterns under the assumption of spatial independence and then to model the spatial dependence of the mean within regions. Bowman (2007) and Bowman et al. (2008) extended this work by allowing correlation for each voxel within a region, and Derado et al. (2010) proposed a model to also account for temporal correlation between multiple experimental effects. The two-stage approach was the first rigorous framework to acknowledge spatial correlation and has given rise to a large body of literature on Bayesian models for fMRI data (see Zhang et al. (2015) for a comprehensive review). Although this framework has been demonstrated to be flexible and to produce useful information for practitioners (e.g., posterior probability maps for activation), two main factors still present limitations to the development of spatio-temporal models for fMRI data. Firstly, dependence of activation patterns (or some differencing of them) assumes a Gaussian Markov random field, which is a natural choice for the gridded geometry of fMRI, but implies stationarity on the local scale, which is overly simplistic, as we show in this article. Secondly, inference is often limited to subsamples such as two-dimensional slices to reduce the dimensionality and consequently the computational time. Zhang et al. (2015) claim that “the large dimensionality of the data makes it impossible to model the entire 3D maps of the data at once.” Here, we demonstrate that it is possible to provide activation maps for the entire brain with nontrivial spatio-temporal models, provided that appropriate computational power is available and a suitable inferential scheme that fully exploits distributed computing is implemented.

An alternative approach to modeling spatial dependence was proposed by Kang et al. (2012), where a spatio-spectral mixed-effects model that captures multi-scale spatial correlation (among ROIs and within ROIs) was defined. By defining the model on the spectral domain (i.e., modeling the Fourier coefficients rather than the fMRI time series), their approach holds promise for scalability because the Fourier coefficients are approximately uncorrelated across different frequencies under a temporally stationary assumption. By incorporating voxel-specific and ROI-specific random effects, the model captures the spatial covariance structure both on a local level (where the local correlation between voxels depends on their distance) and on a regional level (where the correlation between regions is not forced to depend on distance) without reducing the analysis to 2D slices. However, the activation

was assumed at the regional level, not at the voxel level, and the assumption of isotropy within a region, as we show in this article, is not appropriate. With more powerful computing resources available and a more sophisticated methodology to handle nonstationary spatial data, this assumption can and should be relaxed.

Degras and Lindquist (2014) developed a hierarchical model for voxel-specific and condition-specific activation and inference in a multi-subject setting. The problem of simultaneously estimating the Hemodynamic Response Function (HRF) and voxel activation was previously discussed in Makni et al. (2005). The approach in Degras and Lindquist (2014) uses a set of B-spline basis functions to represent the HRF. The coefficients of these functions are allowed to vary between experimental conditions, across voxels in space and over all subjects. Similarly, Zhang et al. (2012, 2013, 2014) proposed different estimation strategies of the HRF, both in the parametric and semi-parametric setting for multiple subjects. The proposed models are flexible and the estimation-inference procedure is rigorously developed; however, all of them still lack a sound formulation of the spatial covariance.

More recently, Zhu et al. (2014) and Hyun et al. (2014) proposed mixed effect models with spatially varying coefficients that allow for spatial discontinuities in Blood Oxygenation Level Dependent (BOLD) activation and spatial dependence. While these models allow for a rich and flexible structure in the mean function, they also require an explicit and interpretable definition of a functional structure allowing jumps, which is difficult to implement with a very large number of voxels. Also, the covariance structure was assumed to have either a low rank representation (Cressie and Johannesson, 2008) or to rely on Gaussian predictive processes (Banerjee et al., 2008) which do not explicitly model nonstationarity and lead to loss of information when the spatial correlation is moderate or strong (Stein, 2014).

In this article, we propose a new model that overcomes the aforementioned limitations by capturing both nonstationary local and regional dependence. This model is used to test for significant voxel-specific activation and connectivity. Inference can be achieved within the context of our data example, comprising of more than 22 million observations, via a multi-resolution approach and an intensive use of distributed computing, without resorting to subsampling either with 2D slices or by assuming a coarse-level activation structure. The isotropy assumption for dependence is also relaxed assuming instead a locally anisotropic model (Fuentes, 2001) for each ROI. Subsequently, generalized shrinkage (Fiecas and Ombao, 2011) with the empirical covariance matrix is performed to allow sufficient flexibility in capturing high spatial frequencies. The regional dependence structure is estimated with a sparse inverse structure via graphical Least Absolute Shrinkage and Selection Operator (LASSO) (Friedman et al., 2008) which allows for visualization of connectivity patterns across ROIs, exploiting its interpretability as a graphical model (Meinshausen and Bühlmann, 2006).

The Gaussian model that we propose allows both a local and a regional structure so that it is possible to realistically capture the complex multi-scale spatio-temporal dependence across voxel-specific fMRI time series. Our model has two advantages. First, a more realistic description of spatial

dependence allows for an improved inference when testing for activation with less false positives than a model that assumes independence or an overly simplistic dependence. Thus, our proposed model produces more accurate detection of activated ROIs when the patient is performing the motor-task (although this approach can be generalized to other types of clinical trials). Second, by using a realistic model for regional dependence with a sparsity structure, one can deduce an interpretable functional connectivity graph that provides information on the connectivity structure of these ROIs during the task.

The remainder of the article is organized as follows: Section 2 introduces the temporal, local, and regional spatial structure of the spatio-temporal statistical model; Section 3 gives a multi-resolution inference scheme and discusses how distributed computing is instrumental in fitting the model to a data set of such scale; Section 4 presents two simulation studies to highlight the need for nonstationary models on the local scale, and Section 5 demonstrates that the proposed spatio-temporal model has the ability to detect more active voxels when compared to a model with spatial independence and how these voxels correspond to interpretable patterns related to the task performed in the motor experiment. Section 6 concludes with a discussion and future directions of investigation. In the Supplementary Material, we included a description on the experimental design and the fMRI data.

2. The Statistical Model for fMRI Data

Throughout this article, $Y_v(t)$ denotes the fMRI time series at voxel v and r_v denotes the ROI to which voxel v belongs. At each time point, the condition is either “rest” (no external stimulus applied) or “active” (when a stimulus was applied). Let $S_1(t)$ be the indicator function during “task” in the block design in Figure S1 and $S_2(t) = 1 - S_1(t)$ the indicator for “rest,” so that exactly one condition is present for each time t .

We further assume that the HRF, denoted $h(t)$, is known and is common across all voxels. Voxel- and ROI-specific HRFs have been developed (see, e.g., Degras and Lindquist (2014)) by assuming that the HRF is the difference between two gamma functions or has a Poisson distribution (Zhang et al., 2015), and allowing spatially varying parameters. We use the canonical HRF from the Statistical Parametric Mapping software to produce a BOLD response associated with each of the two stimuli. The BOLD response for the active condition, denoted $X_1(t)$, is the convolution of the HRF, $h(t)$, with the active stimulus indicator, $S_1(t)$: $X_1(t) = (h * S_1)(t) = \int_{-\infty}^{\infty} h(t)S_1(t - \tau)d\tau$. Similarly, the BOLD response for the rest condition is $X_2(t) = (h * S_2)(t)$.

2.1. Voxel-Wise Activation

Let $\mathbf{Y}(t) = \{Y_1(t), \dots, Y_V(t)\}^\top$ be the fMRI intensity at time $t = 1, \dots, T$ for all voxels $v = 1, \dots, V$. We assume the following standard model for fMRI:

$$\mathbf{Y}(t) = \boldsymbol{\beta}_0 + \sum_{j=1}^J \boldsymbol{\beta}_j I(t \in C_j) + \boldsymbol{\beta}_{J+1}(t \bmod T/3) + \boldsymbol{\beta}_{J+2} X_1(t) + \boldsymbol{\beta}_{J+3} X_2(t) + \boldsymbol{\varepsilon}(t), \quad (1)$$

where C_j indicates the j th session; $\boldsymbol{\beta}_j = (\beta_{j,1}, \dots, \beta_{j,V})^\top$ for $j = 0, \dots, J + 3$ are the covariates that are allowed to change at each voxel. Specifically, $\boldsymbol{\beta}_0$ is the intercept, $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_J$ allow for a changing mean for the sessions (in our case $J = 2$, the third session mean is equal to zero for identifiability) while $\boldsymbol{\beta}_{J+1}$ accounts for a temporal effect. Here, $\boldsymbol{\beta}_{J+2}$ represents the linear contribution of the BOLD response, $X_1(t)$, while $\boldsymbol{\beta}_{J+3}$ accounts for $X_2(t)$. The proposed model could also be extended to M stimuli, with terms $\boldsymbol{\beta}_{J+m+1} X_m(t)$ for $m = 1, \dots, M$ and $X_m(t) = (h * S_m)(t)$ with $S_m(t)$ being the indicator for the m th stimulus.

The noise $\boldsymbol{\varepsilon}(t)$ is modeled as a vector autoregressive process of order 2 (VAR(2)). Note that the most common analyses use the VAR of order 1 or 2, and even such low VAR orders already have $2V^2$ or $3V^2$ unknown parameters, respectively. Higher orders may be possible with a small number of ROIs but could be potentially computationally challenging at the voxel level (Bowman, 2007; Degras and Lindquist, 2014). Thus, $\boldsymbol{\varepsilon}(t)$ is written as

$$\boldsymbol{\varepsilon}(t) = \boldsymbol{\Phi}_1 \boldsymbol{\varepsilon}(t - 1) + \boldsymbol{\Phi}_2 \boldsymbol{\varepsilon}(t - 2) + \mathbf{S}(\boldsymbol{\Omega} \mathbf{H}_1(t) + (\mathbf{I}_V - \boldsymbol{\Omega}) \mathbf{H}_2(t)), \quad (2)$$

where \mathbf{I}_V is the identity matrix of size $V \times V$ and $\boldsymbol{\Phi}_1 = \{\phi_{1,v}\}$, $\boldsymbol{\Phi}_2 = \{\phi_{2,v}\}$ are $V \times V$ diagonal matrices with coefficients representing the autoregressive components of $\boldsymbol{\varepsilon}(t)$. Here, $\mathbf{S} = \{\sigma_v\}$ is a diagonal matrix with voxel-wise standard deviations, and $\boldsymbol{\Omega} \mathbf{H}_1(t) + (\mathbf{I}_V - \boldsymbol{\Omega}) \mathbf{H}_2(t)$ is the vector of normalized innovations. We assume that $\mathbf{H}_u(t) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_u)$ for $u = 1, 2$ where $\boldsymbol{\Sigma}_u$ are correlation matrices and that $\mathbf{H}_1(t)$ is independent from $\mathbf{H}_2(t)$. The vector $\mathbf{H}_1(t)$ controls the local (voxel-specific) scale dependence: its covariance $\boldsymbol{\Sigma}_1$ is a block diagonal matrix, where each block $\boldsymbol{\Sigma}_{1,r}$ corresponds to the dependence within ROI r . The vector $\mathbf{H}_2(t)$ controls the regional scale dependence, representing the ROI specific effect with correlation $\boldsymbol{\Sigma}_2$. $\boldsymbol{\Omega}$ is a $V \times V$ diagonal matrix with diagonal elements $\omega_r \in [0, 1]$ for each ROI r , which represents the relative contribution of the local random effect $\mathbf{H}_1(t)$ compared to the regional random effect $\mathbf{H}_2(t)$. We denote by $\mathbf{H}_1(v)$ and by $\mathbf{H}_2(r_v)$ the value of \mathbf{H}_1 and \mathbf{H}_2 at voxel v , respectively (the time index is removed for simplicity).

2.2. Modeling Intra-ROI Dependence

We assume that $\text{cov}\{\mathbf{H}_1(v), \mathbf{H}_1(v')\} = 0$ if $r_v \neq r_{v'}$; that is, voxels in different ROIs have no dependence through \mathbf{H}_1 . If the voxels belong to the same ROI, that is, if $r_v = r_{v'} = r$, a model for the spatial dependence must be defined. Previous works (Kang et al., 2012, 2013) have proposed a nonparametric isotropic model based on voxel distance (within the same ROI). The simulation study in Section 4, however, shows that such an assumption is overly simplistic and a nonstationary model is necessary to adequately capture the complex spatial structure in fMRI data. A wide variety of approaches have been proposed to construct nonstationary spatial processes, see Sampson (2010) for a comprehensive review. In this article, we adopt the construction of Fuentes (2001), who proposed that $\mathbf{H}_1(v)$ is a linear combination of L

independent locally geometrically anisotropic processes; that is

$$\mathbf{H}_1(v) = \sum_{l=1}^L \mathbf{H}_1^l(v) w_l(v), \quad (3)$$

where $\mathbf{H}_1^l(v)$ are independent, mean zero Gaussian processes across l with $\text{cov}\{\mathbf{H}_1^l(v), \mathbf{H}_1^l(v')\} = \frac{1}{\Gamma(\nu_{l,r}) 2^{\nu_{l,r}-1}} \left[\frac{[\mathbf{R}(v-v')]^\top \mathbf{R}(v-v')}{\theta_{l,r}} \right]^{\nu_{l,r}} K_{\nu_{l,r}} \left[\frac{[\mathbf{R}(v-v')]^\top \mathbf{R}(v-v')}{\theta_{l,r}} \right]$, where K_ν is a Bessel function of the third kind, \mathbf{v} and \mathbf{v}' are the coordinates in 3D space of v and v' , respectively. In other words, $\mathbf{H}_1^l(v)$ has a Matérn covariance with scale $\theta_{l,r}$, smoothness $\nu_{l,r}$ and a distance matrix \mathbf{R} such that the isocovariance curves are ellipsoids with semi-principal axes of length $\{\ell_{1;l,r}, \ell_{2;l,r}, \ell_{3;l,r}\}$ and with a rotation of angle $\xi_{1;l,r}$ with respect to the $x-y$ plane and $\xi_{2;l,r}$ with respect to the $x-z$ plane. The weights $w_l(v)$ in (3) are the inverse distance of v from the centroid of region l , normalized to yield unit variance. The nonstationary model in (3) allows for a different spatial dependence over the L subregions of each ROI and thus accounts for the different local structure. The added flexibility comes at the price of an increased computational burden in the inference, as shown in Section 3.2.

2.3. Modeling Inter-ROI Dependence

From this part, for simplicity, $\mathbf{H}_2(t)$ denotes the $R \times 1$ vector of the region-specific effect instead of the $V \times 1$ vector as described in Section 2.

\mathbf{H}_2 is then a ROI-specific random effect with correlation matrix $\boldsymbol{\Sigma}_2$, such that

$$\text{cov}\{\mathbf{H}_2(r), \mathbf{H}_2(r')\} = \begin{cases} (\boldsymbol{\Sigma}_2)_{r,r'} = \rho(r, r') & \text{if } r \neq r', \\ (\boldsymbol{\Sigma}_2)_{r,r} = 1 & \text{if } r = r', \end{cases} \quad (4)$$

which does not depend on the Euclidean distance between ROIs. Here, $\rho(r, r')$ is a symmetric function such that the resulting covariance is positive definite.

3. Inference on fMRI Activation and Connectivity

Inference for the model defined in Section 2 needs to be performed for a data set with more than 22 million data points. Due to the extremely high dimensionality of the data, it is necessary to develop some approximations to the likelihood in order to perform inference. We propose a three-step likelihood approximation, where the second and third steps are performed conditional on the first, and assume independence across subsets of data. In the first step, we fit a profile likelihood for each individual voxel to extract temporal dependence. The second step estimates $\mathbf{H}_1(t)$ separately for each ROI. The last step estimates the regional effect $\mathbf{H}_2(t)$ with the entire data set. The diagram in Figure 1 shows the last two inference steps. While this approach does not yield a global maximum likelihood (which is practically impossible to achieve), it has shown near-optimal results in terms of both bias and error propagation over a wide range of applications in environmental statistics; see, for example, Castruccio and Stein (2013).

3.1. Step 1: Voxel-Specific Profile Likelihood

We initially consider models (1) and (2), assuming that there is no spatial dependence. Thus, the fit of $\mathbf{Y}_v = \{\mathbf{Y}_v^\top(1), \dots, \mathbf{Y}_v^\top(T)\}^\top$ can be performed independently for each voxel via profile likelihood. If we denote the $T \times 6$ design matrix induced by (1) for every voxel as $\tilde{\mathbf{X}}$, the vector of parameters for voxel v as $\boldsymbol{\theta}_v = (\phi_{1;v}, \phi_{2;v}, \sigma_v^2)^\top$, and the temporal covariance matrix induced by the AR(2) structure in (2) as $\mathbf{K}(\boldsymbol{\theta}_v)$, then the profile likelihood can be written as (Stein, 1999)

$$\ell(\boldsymbol{\theta}_v; \mathbf{Y}_v) = -\frac{T}{2} - \frac{1}{2} \log |\mathbf{K}(\boldsymbol{\theta}_v)| - \frac{1}{2} \mathbf{Y}_v^\top \{\mathbf{K}(\boldsymbol{\theta}_v)\}^{-1} - \mathbf{K}(\boldsymbol{\theta}_v)^{-1} \tilde{\mathbf{X}} \mathbf{W}(\boldsymbol{\theta}_v)^{-1} \tilde{\mathbf{X}}^\top \mathbf{K}(\boldsymbol{\theta}_v)^{-1} \mathbf{Y}_v,$$

where $\mathbf{W}(\boldsymbol{\theta}_v) = \tilde{\mathbf{X}}^\top \mathbf{K}(\boldsymbol{\theta}_v)^{-1} \tilde{\mathbf{X}}$, and the mean vector $\boldsymbol{\beta}_v$ can be obtained via generalized least squares:

$$\hat{\boldsymbol{\beta}}_v(\boldsymbol{\theta}_v) = \mathbf{W}(\boldsymbol{\theta}_v)^{-1} \tilde{\mathbf{X}}^\top \mathbf{K}(\boldsymbol{\theta}_v)^{-1} \mathbf{Y}_v.$$

Figure 2 shows the results of the fit for four randomly chosen voxels. It is apparent how the linear model (1) is able to adequately capture both the mean and the uncertainty for all voxels, including the ones that show a discontinuous change in the mean for different sessions in the experiment.

3.2. Step 2: Estimating Local Effects

We now consider (1) and (2) assuming that $\mathbf{H}_1(t)$ has the spatial structure defined in Subsection 2.2, while for \mathbf{H}_2 , we assume $\rho(r, r') = 0$ for every $r \neq r'$ in (4) so that $\boldsymbol{\Sigma}_2 = \mathbf{I}_R$. In other words, the process assumes that the voxels within the same ROI are correlated but voxels from different ROIs are not correlated.

Denote by

$$\mathbf{e}(t) = \hat{\mathbf{S}}^{-1} \{\hat{\mathbf{Y}}(t) - \hat{\boldsymbol{\Phi}}_1 \hat{\mathbf{Y}}(t-1) - \hat{\boldsymbol{\Phi}}_2 \hat{\mathbf{Y}}(t-2)\},$$

where $(\hat{\boldsymbol{\Phi}}_1, \hat{\boldsymbol{\Phi}}_2, \hat{\mathbf{S}})$ and $\hat{\mathbf{Y}}(t) = \mathbf{Y}(t) - (\mathbf{I}_V \otimes \tilde{\mathbf{X}}) \hat{\boldsymbol{\beta}}$ are estimated from the previous step. These residuals can then be used to estimate $\text{cov}\{\boldsymbol{\Omega} \mathbf{H}_1(t) + (\mathbf{I}_V - \boldsymbol{\Omega}) \mathbf{H}_2(t)\} = \boldsymbol{\Omega} \boldsymbol{\Sigma}_1 \boldsymbol{\Omega}^\top + (\mathbf{I}_V - \boldsymbol{\Omega})(\mathbf{I}_V - \boldsymbol{\Omega})^\top$, and since this matrix has a block diagonal structure, the fit for each ROI can be performed independently. We thus focus on $\mathbf{e}_r(t)$, the collection of all values of $\mathbf{e}(t)$ in ROI r , and fit a zero mean Gaussian process with locally anisotropic covariance function (3) to estimate $\boldsymbol{\Sigma}_1$ (a plot of the entries of $\boldsymbol{\Omega}$, that is, the relative contributions of the local versus regional covariance, can be found in the Supplementary Material).

The choice of the number L and shape of the regions in (3) is different for every ROI and is performed according to a model selection procedure. Each ROI is divided into $L = L_x \times L_y \times L_z$ sub-regions, where the x (y , z) axis is divided into L_x (L_y , L_z) equally spaced intervals with regions with less than 36 points merged with the largest neighboring region (Fuentes, 2001). For each parallelepiped i , the Maximum Likelihood Estimator is computed for a geometrically anisotropic Matérn model assuming $\xi_{1;l,r} = \xi_{2;l,r} = 0$ (a model selection with further estimation of the angles was computationally infeasible). The global covariance function is

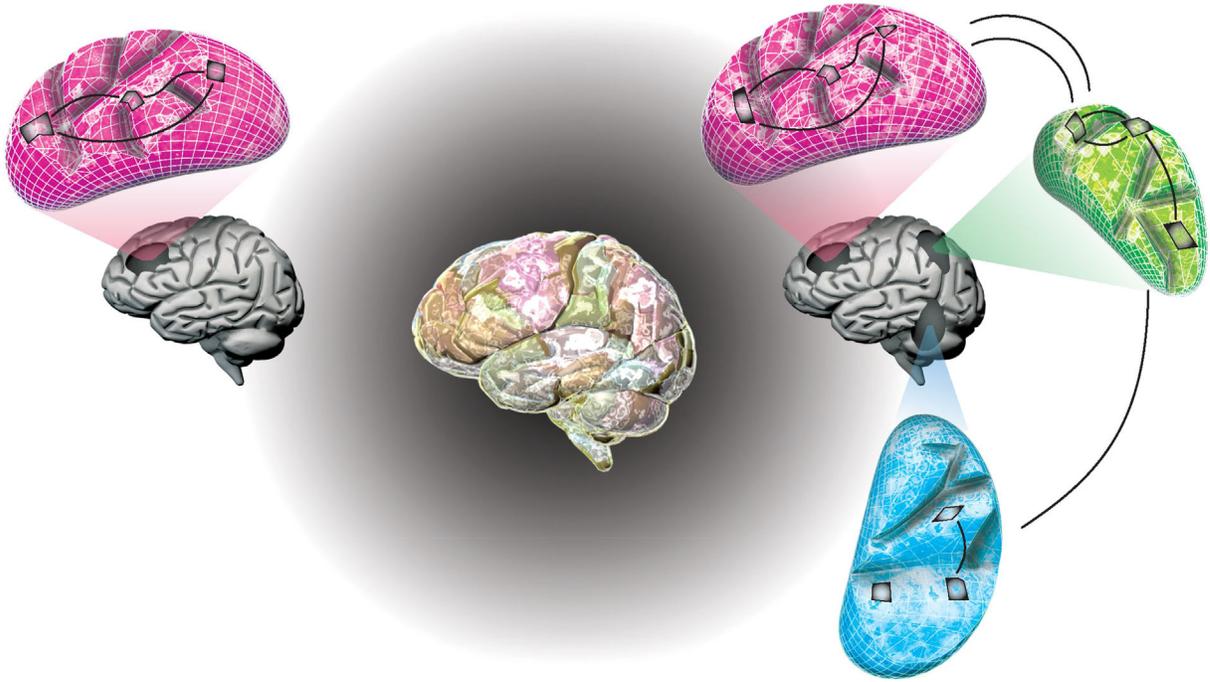


Figure 1. Concept figure of the last two inference steps. The leftmost figure represents the spatial dependence within each ROI assuming independence across them, and the rightmost considers the connectivity across ROIs. The middle picture represents possible activation patterns across ROIs. Figure 1 was produced by Heno Hwang, scientific illustrator at King Abdullah University of Science and Technology (KAUST). This figure appears in color in the electronic version of this article.

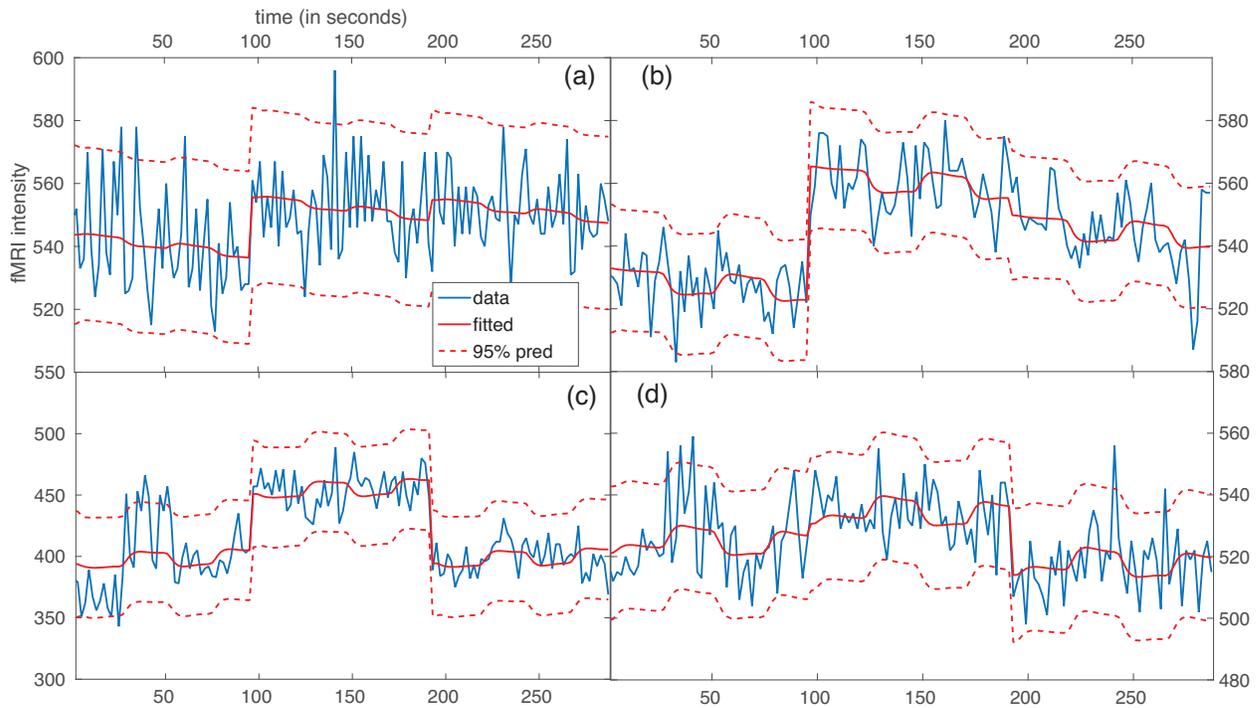


Figure 2. Examples of fMRI intensity fit for four randomly selected voxels. The fitted value according to (1) (solid) with its associated 95% prediction intervals (dashed) follow the data (solid). The four voxels belong to (a) middle frontal gyrus, right orbital lobe, (b) left precuneus, (c) right inferior temporal gyrus, and (d) right fusiform gyrus. This figure appears in color in the electronic version of this article.

estimated as a weighted sum of the local covariance functions and the Bayesian Information Criterion (Schwarz, 1978) is computed. The optimal L is chosen according to the following steps:

1. Start with a single region, that is with a geometrically anisotropic model with $L_x = L_y = L_z = 1$.
2. Evaluate the BIC on all the neighboring configurations of (L_x, L_y, L_z) , that is $(L_x + 1, L_y, L_z)$, $(L_x - 1, L_y, L_z)$, \dots . Also, evaluate the BIC at 25 randomly drawn locations (to avoid local minima).
3. If there is a configuration with smaller BIC, redo point 2, otherwise stop.
4. Once the optimal configuration is obtained, re-estimate \mathbf{R} with the rotation angles.

A comparison of all ROIs in terms of BIC for the isotropic and locally geometrically anisotropic model, as shown in Figure 3, demonstrated that the proposed model is substantially more suitable for fMRI data within the same ROI: the locally anisotropic model is on average approximately 28,000 BIC unit better than the isotropic one.

Despite its flexibility, a preliminary analysis has shown that an estimated covariance function with (3), even after model selection, still leaves a considerable margin for improvement. We thus compute a new estimate of the covariance via generalized shrinkage (Friedman et al., 2008), which is a linear combination of the empirical covariance and an estimate from a parametric model (Fiecas and von Sachs, 2014). We denote the estimated covariance resulting from the Maximum Likelihood Estimate of the r th block of $\hat{\Sigma}_1$ according to (3) as $\hat{\Sigma}_{1,r}^{\text{MLE}}$, and $\hat{\Sigma}_{1,r} = \frac{1}{T} \sum_{t=1}^T \mathbf{e}_r(t) \mathbf{e}_r(t)^\top$ is the sample covariance matrix of $\mathbf{e}_r(t)$ computed from the temporal replicates. The shrunk covariance is defined as: $\hat{\Sigma}_{1,r}(\delta_r) = (1 - \delta_r) \hat{\Sigma}_{1,r} + \delta_r \hat{\Sigma}_{1,r}^{\text{MLE}}$, where $\delta_r \in (0, 1)$ is a suitable constant, chosen so that dependence at high spatial frequencies (represented by the contrasts) matches that of $\Sigma_{1,r}$ within some tolerance (see Figure S3). This is obtained with the following steps:

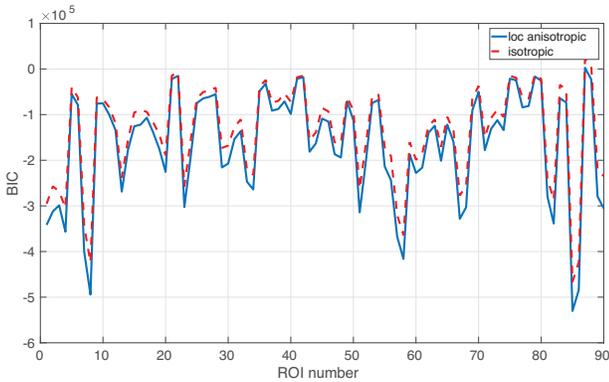


Figure 3. BIC for all 90 ROIs for the isotropic model and the locally anisotropic model; the y -axis is on the 10^5 scale. This figure appears in color in the electronic version of this article.

- The contrast

$$c_x = \frac{1}{|N_x|} \sum_{(y,z) \in \mathcal{S}_x} \{ (\hat{\Sigma}_{1,r})_{(x,y,z),(x,y,z)} + (\hat{\Sigma}_{1,r})_{(x+1,y,z),(x+1,y,z)} - 2(\hat{\Sigma}_{1,r})_{(x+1,y,z),(x,y,z)} \}$$

is computed, where N_x is the set of all pairs (y, z) such that the points (x, y, z) and $(x + 1, y, z)$ belong to the grid. Similarly, we define c_y and c_z .

- A smoothing spline with penalization parameter p is fit to c_x , c_y , and c_z , which we call c_x^p , c_y^p , and c_z^p , respectively.
- The parameter δ_r generates $\Sigma_{1,r}(\delta_r)$ such that

$$c_x(\delta_r) = \frac{1}{|S_x|} \sum_{(y,z) \in \mathcal{S}_x} \{ (\hat{\Sigma}_{1,r}(\delta_r))_{(x,y,z),(x,y,z)} + (\hat{\Sigma}_{1,r}(\delta_r))_{(x+1,y,z),(x+1,y,z)} - 2(\hat{\Sigma}_{1,r}(\delta_r))_{(x+1,y,z),(x,y,z)} \}.$$

Similarly, we define $c_y(\delta_r)$ and $c_z(\delta_r)$. We choose δ_r , such that $\|c_x^p - c_x(\delta_r)\|_2^2 + \|c_y^p - c_y(\delta_r)\|_2^2 + \|c_z^p - c_z(\delta_r)\|_2^2$ is minimized.

Thus, it is possible to control the spatial structure via the penalization parameter p . Figure S3 shows an example of the shrinkage to the empirical covariance for a fixed ROI and different values of δ_r in terms of their spatial contrasts. For our analysis, we choose a penalty term of $p = 0.3$, because it allows for some flexibility in the pattern structure, and because values of p in this neighborhood have yielded qualitatively indistinguishable results.

3.3. Step 3: Estimating the Regional Effects

Conditional on the estimated covariance structure in the previous steps, we estimate the covariance structure of $\mathbf{H}_2(t)$, that is, the entries of Σ_2 . If we denote as $\bar{\mathbf{e}}(t) = \{\bar{\mathbf{e}}_1^\top(t), \dots, \bar{\mathbf{e}}_R^\top(t)\}^\top$, and $\bar{\mathbf{e}}_r(t)$ the average of $\mathbf{e}(t)$ for ROI r , $\mathbf{A} = \frac{1}{T} \sum_{t=1}^T \bar{\mathbf{e}}(t) \bar{\mathbf{e}}(t)^\top$ provides a nonparametric estimation of Σ_2 , the correlation matrix for the ROI specific effects. However, this would require estimating $R(R - 1)/2$ entries, with no insight on the connectivity patterns induced by the experiment, that is, on which ROIs are significantly connected. We choose to estimate the $R \times R$ inverse covariance $\mathbf{W}_{\text{brain}}$ of $\mathbf{H}_2(t)$, by imposing an ℓ^1 constraint on the number of nonzero entries of the inverse correlation matrix. This penalized likelihood approach on the inverse covariance was proposed by Yuan and Lin (2007) as an application of the maxdet problem (Vandenberghe et al., 1998); a faster approach for high-dimensional covariance estimation, the graphical LASSO (Friedman et al., 2008), has become widely popular in recent years, also in the context of neurological data (Varoquaux et al., 2010; Cribben et al., 2012). The main idea is to rewrite the loglikelihood for $\bar{\mathbf{e}}(t)$, that is $-\frac{T}{2} \log(2\pi) + \frac{1}{2} \log |\mathbf{W}_{\text{brain}}| - \frac{1}{T} \sum_{t=1}^T \bar{\mathbf{e}}(t)^\top \mathbf{W}_{\text{brain}} \bar{\mathbf{e}}(t)$ as a disciplined convex problem, and to further impose an ℓ^1 penalty on the

number of nonzero entries of $\mathbf{W}_{\text{brain}}$:

$$-\frac{T}{2} \log(2\pi) + \frac{1}{2} \log |\mathbf{W}_{\text{brain}}| + \text{tr}(\mathbf{W}_{\text{brain}}\mathbf{A}) + \lambda \sum_{r \neq r'} |(\mathbf{W}_{\text{brain}})_{r,r'}|. \quad (5)$$

This allows us to obtain an estimate of the inverse correlation that is sparse, and consequently interpretable results for ROI connectivity. The optimal $\hat{\lambda}$ is chosen by cross-validation (Friedman et al., 2008): the inverse correlation is evaluated with (5) for 90% of the data. Then, each ROI is predicted on the remaining 10% of the data by leave-one-out cross-validation. Figure 4a shows a plot of the error sum of squares against λ . A small penalty results in severe overparametrization and as the penalty increases, the estimated $\mathbf{W}_{\text{brain}}$ is closer to the diagonal matrix and results in a suboptimal fit. In the supplement, the file `movie_glasso.avi` shows how the sparsity increases and where it occurs as λ increases. Figure 4b shows all the nonzero elements (53% of the total number of entries) of $\mathbf{W}_{\text{brain}}$ for the optimal $\hat{\lambda} = 8 \times 10^{-4}$.

Once the optimal sparse inverse correlation is computed, it is possible to infer the structure of brain connectivity from the structure of the graph, as will be discussed in Section 5.

3.4. Computational Considerations

Besides providing a flexible framework to allow for inference from a complex model with more than 22 million data points, this multi-resolution scheme also allows for an extensive use of distributed computing to achieve a scalable inference. In step 1, the code is parallelized so that inference is performed independently for every voxel time series. In step 2, model (3) (and similarly the simpler models in the simulation studies) can be estimated independently for each core in a cluster,

allowing as many as R independent simultaneous estimations. It is also possible to focus on a single ROI and parallelize the model selection algorithm detailed in Subsection 3.2, but this approach yielded a suboptimal performance, as the communication overhead across processors generated significant latency especially with small-sized ROIs.

3.5. Testing for Voxel-Specific Activation

Activation for each voxel can be tested by combining a fine-scale estimation of the dependence structure with the definition of a voxel-specific mean structure. We assume that $\beta_0, \dots, \beta_{J+1}$ are fixed and equal to the estimated values obtained from the fit of the temporal and mean parameter in (1) assuming spatial independence, while we re-estimate β_{J+2} and β_{J+3} , for which we want to test $H_0: \beta_{J+2} - \beta_{J+3} = 0$ versus $H_1: \beta_{J+2} - \beta_{J+3} \neq 0$. We allow β_{J+2} to be smoothly varying in space, thus allowing to borrow strength across neighboring locations and yielding a more accurate test for activation. We assume that the spatial variation is parametrized by Fourier coefficients: for each ROI r , we have

$$\beta_{J+2}^r(x, y, z) = \sum_{n=1}^N \sum_{b \in \{x,y,z\}} a_{b,n,1}^r \cos(2\pi bn/d_b^r) + a_{b,n,2}^r \sin(2\pi bn/d_b^r), \quad (6)$$

where x, y , and z are the (normalized) coordinates over the three axes, and d_x^r, d_y^r , and d_z^r are the maximum lengths across the three axes. To avoid identification issues, we set β_{J+3} constant across the ROI. We thus need to estimate $\{a_{b,n,1}^r, a_{b,n,2}^r\}$ for each harmonic $n = 1, \dots, N$, for each coordinate $b \in \{x, y, z\}$ and for each ROI $r = 1, \dots, R$. As in the previous sections, parameters in each ROI can be estimated independently and the code can be easily parallelized, but the computational and memory demand for each likelihood evaluation allowed only to fit the model for $N = 1$. Figure S4 shows the activated voxels for each ROI with False Discovery

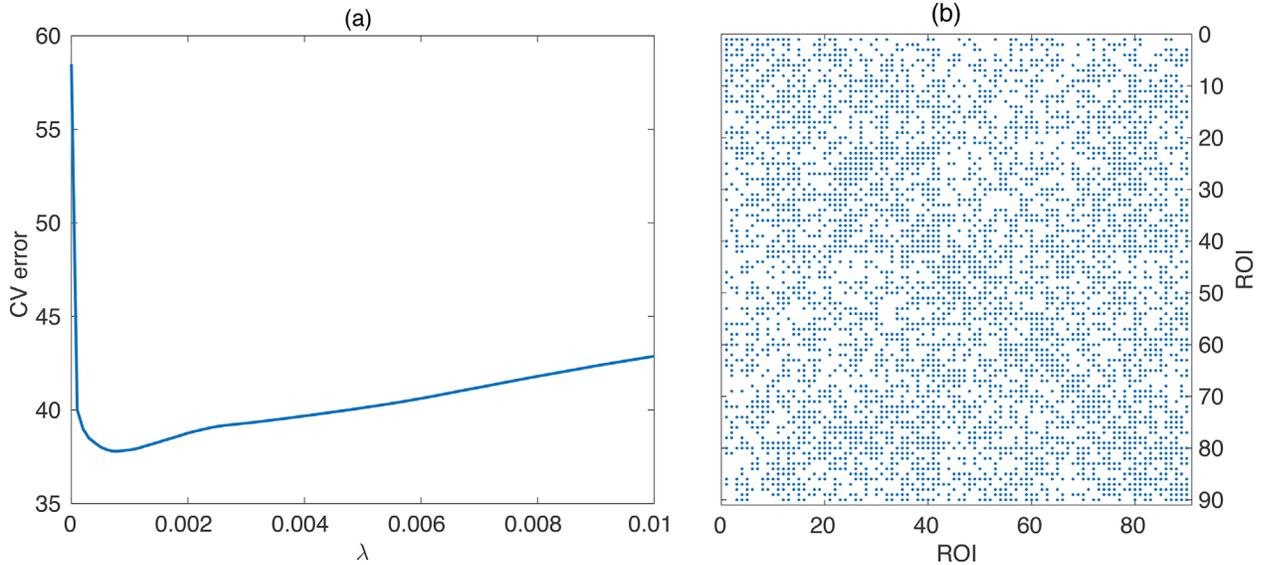


Figure 4. Graphical LASSO for brain connectivity. (a) Sum of squares cross-validation error versus different choices of penalty in (5), (b) plot of all nonzero elements in $\mathbf{W}_{\text{brain}}$ for the optimal $\hat{\lambda}$. This figure appears in color in the electronic version of this article.

Rate (Benjamini and Hochberg, 1995) at 5% assuming independence and spatial dependence. The results for particular areas will be discussed in Section 5.

Alternatively, the spatial variation could be also modeled hierarchically, that is, by assuming $\beta_{J+2} \sim \mathcal{N}(\mathbf{m}, (\lambda \mathbf{Q})^{-1})$ as a latent process with mean \mathbf{m} , precision matrix $\lambda \mathbf{Q}$ with (fixed or random) parameter λ . However, the computational burden of evaluating the likelihood for the entire ROI would make the inference considerably more challenging.

4. Simulation Studies

To further support our choice of a locally anisotropic model over the simpler existing alternatives described in Subsection 3.2, we performed two simulation studies: one focused on an activation test and another on spatial interpolation.

In the first study, we focused on the improved performance of a locally anisotropic model (*l-aniso*) against a general linear model (denoted *glm*), an isotropic model (*iso*) and an anisotropic model (*aniso*) for a single ROI. For each ROI, we performed 100 simulations, where the true spatial covariance is the sample covariance obtained from the fMRI data of the subject in the case study after detrending with Ordinary Least Squares in time. This approach ensures that the performance of each method is compared against a true model that is as close as possible to the original data set. For computational reasons, whenever an ROI has more than 1000 voxels, a random sample of this size (the same for every simulation) is considered for the analysis. We assume a common mean across the ROI with no time or session effect and only the hemodynamic response terms in the mean structure, which we denote as β_1 and β_2 . The common mean does not allow for a straightforward comparison with the results in Subsection 3.5, but it does significantly reduce the considerable computational burden and allows for a closer comparison with the similar simulation studies in Kang et al. (2012). For each simulation, $\hat{\beta}_1 - \hat{\beta}_2$ and $\text{var}(\hat{\beta}_1 - \hat{\beta}_2)$ are computed according to the four models and a test is performed to determine the presence of activation, that is, if $\beta_1 \neq \beta_2$. Model selection for the locally anisotropic model is avoided because it would require several days per simulation on a large computational facility; the geometry of the anisotropic subsets of the brain is obtained from the model selection step in Subsection 3.2.

In the first study, we assumed that the ROI is not active, that is, $\beta_1 = \beta_2$, and we compared the false positives at a 5% confidence level out of the 100 simulations. Table 1 shows the results, where the first three rows represent the number of false positives for three randomly drawn ROIs and the last row computes an average across all ROIs. The assumption

of no spatial dependence in *glm* implies a very high number of false positives, as was previously reported by Kang et al. (2012). Even with the simple isotropic model *iso*, the assumption of spatial dependence brings dramatic improvements in the accuracy of the test. The models *aniso* and *l-aniso* bring a further improvement, although it comes at the expense of a longer computational time. It is also remarkable how these results are markedly larger than the nominal 5% rate of false positive, thus indicating how the nonstationarity dependence within ROI is very complex and more sophisticated models could further improve the results. A fully nonparametric approach of estimating the empirical covariance (results not shown) proved considerably worse, yielding nearly 100% false positives, indicating that some parametric description of the nonstationarity is needed. In the Supplementary Material, we show the power curves for this study.

In the second study, we compared the effect of the four models on interpolation. In the setting of the previous studies, for each simulation, we remove 50 random voxels (the same across all simulations for the same ROI), we interpolate their values with kriging and compute the Root Mean Squared Error (RMSE) with the true value. The results in parenthesis of Table 1 show how the independence assumption is largely inadequate, and how *iso*, *aniso* and *l-aniso* perform incrementally better as they yield interpolated values closer to the true simulated data. As in the first simulation study, a smaller RMSE for *aniso* and *l-aniso* comes at the cost of a much more computationally challenging estimation. Although the difference in performance between *aniso* and *l-aniso* seems small, note that the model selection was not performed for every simulation because it would have been too computationally demanding. Thus, further improvements could be expected if the regions of anisotropy were not predefined from the real data set.

5. Discussion of the Results

The model selection procedure suggests that a locally anisotropic model outperforms the isotropic model uniformly across ROIs in terms of lower false positive rates and lower root mean squared error (see Figure 3 and Table 1). Due to the computational complexity, previous models (e.g., Kang et al., 2012, 2013) have simply assumed an isotropic structure within each ROI. However, our method suggests non-stationary behavior—even within an ROI—which means complexity of the spatial covariance structure that requires more sophisticated modeling. In particular, the degree of correlation between voxels may vary as a function of their Euclidean distance, although this correlation may differ depending on the

Table 1

Percentage of false positives assuming no activation for three randomly chosen ROIs (first three rows, see Table S1 for the abbreviation) and the mean across ROIs for the four models (last row). In parenthesis, $\text{RMSE} \times 10^3$ for 50 randomly removed points.

ROI	General linear model	Isotropic	Anisotropic	Locally anisotropic
Superior occipital gyrus, right	80 (6.97)	26 (1.09)	28 (1.05)	13 (1.04)
Parahippocampal gyrus, right	72 (6.89)	14 (1.05)	12 (0.94)	11 (0.91)
Orbital superior frontal gyrus, left	64 (6.93)	35 (0.99)	31 (0.99)	27 (0.95)
Mean	78.7 (6.66)	31.9 (1.10)	28.3 (1.01)	26.3 (0.99)

exact location of these voxels in an ROI. This difference could have a significant impact on inference (e.g., testing for activation) and hence must be properly specified when computing the test statistic.

One of the aims of our analysis is to identify the voxels that are activated during the hand-grasping task. In Figure S4, we show the contrasts for the analysis that ignores spatial correlation (in black) and the analysis that incorporates local anisotropy (in blue) for the left supplementary motor area (SMA-L), while in the Supplementary Material, we report (Table S2) the percentage of all activated voxels according to the two models for all ROIs.

We note that the activation patterns differ in many respects. First, estimation of contrasts for the spatio-temporal model is considerably more conservative (as shown in Figure S4), and on average there is a smaller number of voxels with significantly different BOLD activation between the hand-grasping and rest conditions (92.8% for the general linear model and 84.8% for the spatio-temporal model, from Table S2). Here, it is very likely that the independence model flagged a number of voxels as significant that in fact do not display a differential BOLD response level. Secondly, among all ROIs, SMA-L had the highest percentage of voxels with significantly greater amplitude of the BOLD response during the hand-grasping task compared to the resting state (see Table S2) under the spatio-temporal model. This result was expected, since the subject was performing a motor task using his right (dominant) hand. Besides, the corresponding ROI in the right hemisphere, the right supplementary motor area (SMA-R) has the third largest proportion of activated voxels at 96%. Under

the independence model, SMA-L and SMA-R ranked 10th and 6th, respectively.

We also identified an unexpected result using the independence model that carries no neurophysiological justification: right pre-central area was flagged as the most significant region. This result was quite unexpected since the right area is indicated as more active, as it should be in a left-handed subject, and the SMA-L is only ranked 10th in terms of active voxels. Also, the cuneus left (CUN-L) had the second largest percentage of voxels with higher BOLD amplitude during the motor task, and this region is not located in the primary motor cortex. This underscores the fact that ignoring spatial covariance in the analysis could produce unexpected results that are likely incorrect since the task is purely motor and would not require higher cognitive processing.

In addition to examining activation, the proposed procedure can also be used to study potentially complex dependence between voxels in an ROI through the localized spatial covariance. As noted, the BIC model selection metric chose the locally anisotropic model over the isotropic model. This concept of local anisotropy is not just a theoretical construct. First, we need the spatial covariance to be properly specified to give us confidence about the level and power of our testing procedure. Next, the model provides us with information on the strength of correlation between pairs of voxels and how this could vary depending on whether the pair is located on the anterior or posterior neighborhood of the ROI. For example, in Figure 5c, the anterior part of SMA-L, has a correlation structure that is more circular than ellipsoidal

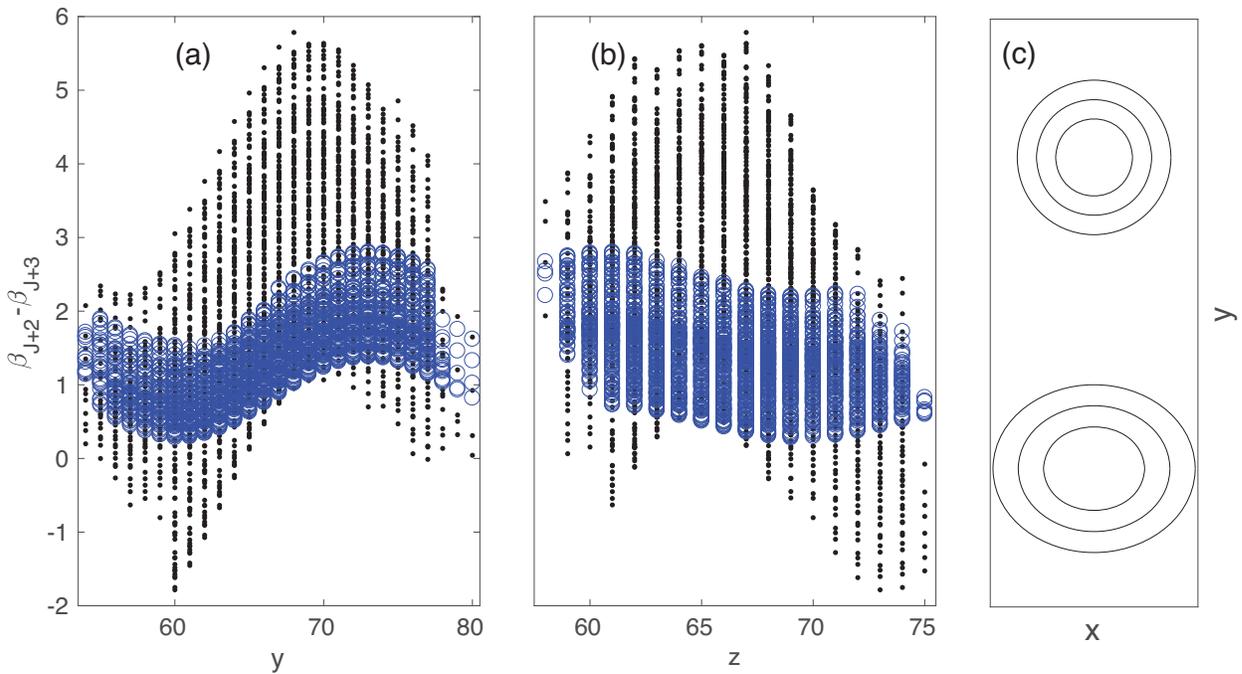


Figure 5. Activation contrasts against the y (a) and z-axis (b) for SMA-L. The dots represent the contrast according to the model with spatial independence, while the circles represent the spatio-temporal model. (c) The contour of the estimated covariance function for two points in the anterior (top) and posterior (bottom) part of the ROI. This figure appears in color in the electronic version of this article.

(meaning there is no preferred direction in the dependence structure), while the opposite is true of the anterior right part of SMA-L. This finding is very interesting for our stroke collaborators who will look into confirming these findings using other types of data modalities and modeling techniques (e.g., probabilistic tractography and fractional anisotropy/mean diffusivity in diffusion tensor imaging). However, this is beyond the current scope of this article.

The task of modeling and understanding connectivity is intertwined with activation and there is particular interest in stroke patients because neurons from other regions activate at a higher than normal level to compensate for reduced activation in regions with injury (Tombari et al., 2004). The estimated inverse covariance matrix revealed a number of interesting connections between a few pairs of ROIs. Using a high threshold at $\lambda = 0.01$, the pair of regions that survived the stringent threshold, indicating the strongest pairwise direct dependence, is left frontal operculum and left rolandic operculum. The left frontal operculum refers to the small region in the frontal lobe that overlies the rostradorsal portion of the insula in primates. Alexander et al. (1990) found nine cases of aphasia (broadly defined as a difficulty with speech and writing) following lesions in the region of the left frontal operculum. Moreover, in Tonkonogy and Goodglass (1981), two cases of articulatory difficulty were associated with lesions in the rolandic operculum. It is interesting that these two ROIs are almost adjacent, which partly explains how damages to these regions result in similar symptoms of aphasia. This suggests that the strongest direct link between ROIs may not be due to these regions having shared functional dependence because these ROIs are not shown to be implicated in motor task activity, however, this can be explained by the actual anatomical proximity.

6. Conclusions

This article addressed the issue of enhancing the detection of activation and connectivity in fMRI data by explicitly modeling spatial dependence. Motivated by the need to develop flexible models for enhancing detection of neurological patterns in the recovery of a patient suffering stroke, this article provides a general framework for whole-brain modeling.

The multi-resolution model introduced allows for spatially varying coefficients and local nonstationarity within ROIs in the error structure. We have demonstrated with numerical studies how isotropy within an ROI, a widespread simplifying assumption, is largely inappropriate for fMRI data even for small ROIs, as it leads to suboptimal tests for activation. Compared to the current methodology, our proposed model showed clear improvements. In particular, numerical studies suggest the nonstationarity is complex and a locally anisotropic model is not fully adequate. Future work will investigate the use of other constructions for nonstationary processes such as the one proposed in Paciorek and Schervish (2006), multiresolution models with random coefficients with sparse dependence (Nychka et al., 2015) and continuous time/discrete space models (Quick et al., 2013) that would allow inference on a temporal scale smaller than the measurement frequency, while still borrowing strength across locations with a sparse precision matrix.

This model shares common features from the two-stage Bayesian hierarchical approach introduced in Bowman (2007) and the spatio-spectral mixed model in Kang et al. (2012) in that it aims at modeling spatial dependence directly instead of mitigating its effect via spatial smoothing. However, it provides a finer spatial scale information on activation, which is attained at the price of a substantially increased computational burden, requiring a tailored inference scheme that fully exploits parallelization. In addition to the advantage of describing finer scale information, voxel-specific (rather than an ROI-specific) activation bypasses the problem of determining how many active voxels comprise an active ROI (Kang et al., 2012). The connectivity is captured via the ROI-specific random effect as in Kang et al. (2012), but our proposed model is more appealing, as graphical LASSO yields undirected graphs that are interpretable and partially avoids the overparametrization of the empirical covariance.

7. Supplementary Materials

Web Appendices, Tables, and Figures referenced in Sections 1, 3, 4, and 5, as well as the code for the fit of the steps 1 and 2, and of the model are available with this article at the *Biometrics* website on Wiley Online Library.

REFERENCES

- Alexander, M., Naeser, M., and Palumbo, C. (1990). Broca's area aphasias: Aphasia after lesions including the frontal operculum. *Neurology* **40**, 353–362.
- Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008). Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society, Series B* **70**, 825–848.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B* **57**, 289–300.
- Bowman, D. (2007). Spatiotemporal models for region of interest analyses of functional neuroimaging data. *Journal of the American Statistical Association* **102**, 442–453.
- Bowman, F. (2005). Spatio-temporal modeling of brain activity. *Biostatistics* **6**, 558–575.
- Bowman, F. D., Caffo, B., Bassett, S., and Kilts, C. (2008). A bayesian hierarchical framework for spatial modeling of fmri data. *Neuroimage* **39**, 146–156.
- Bullmore, E., Long, C., Suckling, J., Fadili, J., Calvert, G., Zelaya, F., et al. (2001). Colored noise and computational inference in neurophysiological (fmri) time series analysis: Resampling methods in time and wavelet domains. *Human Brain Mapping* **12**, 61–78.
- Castruccio, S. and Stein, M. L. (2013). Global space-time models for climate ensembles. *Annals of Applied Statistics* **7**, 1593–1611.
- Cressie, N. and Johannesson, G. (2008). Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society, Series B* **70**, 209–226.
- Cribben, I., Haraldsdottir, R., Atlas, L. Y., Wager, T. D., and Lindquist, M. A. (2012). Dynamic connectivity regression: Determining state-related changes in brain connectivity. *Neuroimage* **61**, 907–920.
- Degras, D. and Lindquist, M. (2014). A hierarchical model for simultaneous detection and estimation in multi-subject fmri studies. *NeuroImage* **98**, 61–72.

- Derado, G., Bowman, F., and Kilts, C. (2010). Modeling the spatial and temporal dependence in fmri data. *Biometrics* **66**, 949–957.
- Dubin, R. A. (1988). Estimation of regression coefficients in the presence of spatially autocorrelated error terms. *The Review of Economics and Statistics* **70**, 466–474.
- Fiecas, M. and Ombao, H. (2011). The generalized shrinkage estimator for the analysis of functional connectivity of brain signals. *Annals of Applied Statistics* **5**, 1102–1125.
- Fiecas, M. and von Sachs, R. (2014). Data-driven shrinkage of the spectral density matrix of a high-dimensional time series. *Electronic Journal of Statistics* **8**, 2975–3003.
- Friedman, J., Hastie, T., and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9**, 432–441.
- Fuentes, M. (2001). A high frequency kriging approach for non-stationary environmental processes. *Environmetrics* **12**, 469–483.
- Hyun, J. W., Li, Y., Gilmore, J. H., Lu, Z., Styner, M., and Zhu, H. (2014). Sgpp: Spatial gaussian predictive process models for neuroimaging data. *NeuroImage* **89**, 70–80.
- Kang, H., Yang, X., Bryan, F. W., Tripp, C. M., and Landman, B. A. (2013). Whole brain functional connectivity using multi-scale spatio-spectral random effects model. In *Multimodal Brain Image Analysis. Lecture Notes in Computer Science, vol 8159*, L. Shen, T. Liu, P. T. Yap, H. Huang, D. Shen, and C. F. Westin (eds), Cham: Springer.
- Kang, H., Ombao, H., Linkletter, C., Long, N., and Badre, D. (2012). Spatio-spectral mixed-effects model for functional magnetic resonance imaging data. *The Journal of the Acoustical Society of America* **107**, 568–577.
- Locascio, J., Jennings, P., Moore, C., and Corkin, S. (1997). Time series analysis in the time domain and resampling methods for studies of functional magnetic resonance brain imaging. *Statistics in Medicine* **5**, 168–193.
- Makni, S., Idier, J., Vincent, T., Thirion, B., Dehaene-Lambertz, G., and Ciuciu, P. (2005). A fully bayesian approach to the parcel-based detection-estimation of brain activity in fmri. *Human Brain Mapping* **49**, 941–969.
- Meinshausen, N. and Bühlmann, P. (2006). High dimensional graph and variable selection with the lasso. *Annals of Applied Statistics* **34**, 1436–1462.
- Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., and Sain, S. (2015). A multiresolution gaussian process model for the analysis of large spatial datasets. *Journal of Computational and Graphical Statistics* **24**, 579–599.
- Paciorek, C. J. and Schervish, M. J. (2006). Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics* **17**, 483–506.
- Quick, H., Banerjee, S., and Carlin, B. P. (2013). Modeling temporal gradients in regionally aggregated california asthma hospitalization data. *Annals of Applied Statistics* **7**, 154–176.
- Sampson, P. D. (2010). Constructions for nonstationary processes. *The Handbook of Spatial Statistics*. Boca Raton, FL: Chapman & Hall/CRC.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics* **6**, 461–464.
- Stein, M. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. NY: Springer.
- Stein, M. L. (2014). Limitations on low rank approximations for covariance matrices of spatial data. *Spatial Statistics* **8**, 1–19.
- Tombari, D., Loubinoux, I., Pariente, J., Gerdelat, A., Albucher, J. F., Tardy, J., et al. (2004). A longitudinal fmri study: In recovering and then in clinically stable sub-cortical stroke patients. *NeuroImage* **23**, 827–839.
- Tonkonogy, J. and Goodglass, H. (1981). Language function, foot of the third frontal gyrus, and rolandic operculum. *Archives of Neurology* **38**, 486–490.
- Vandenberghe, L., Boyd, S., and Wu, S.-P. (1998). Determinant maximization with linear matrix inequality constraints. *SIAM Journal on Matrix Analysis and Applications* **19**, 499–533.
- Varoquaux, G., Gramfort, A., Jean-Baptiste, P., and Thirion, B. (2010). Brain covariance: Better individual functional connectivity models using population prior. In *Advances in Neural Information Processing Systems 23*, J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta (eds), 2334–2342. NY: Curran Associates, Inc. <http://papers.nips.cc/paper/4080-brain-covariance-selection-better-individual-functional-connectivity-models-using-population-prior.pdf>
- Worsley, K., Evans, A., Marrett, S., and Neelin, P. (1992). 3d statistical analysis for rcbf activation studies in the human brain. *Journal of Cerebral Blood Flow & Metabolism* **12**, 900–918.
- Worsley, K. and Friston, K. (1995). Analysis of fmri time series revisited-again. *NeuroImage* **40**, 173–181.
- Yuan, M. and Lin, Y. (2007). Model selection and estimation in the gaussian graphical model. *Biometrika* **94**, 19–35.
- Zhang, L., Guindani, M., and Vannucci, M. (2015). Bayesian models for functional magnetic resonance imaging data analysis. *Wiley Interdisciplinary Reviews Computational Statistics* **7**, 21–41.
- Zhang, T., Li, F., Beckes, L., Brown, C., and Coan, J. A. (2012). Nonparametric inference of hemodynamic response for multi-subject fmri data. *NeuroImage* **63**, 1754–1765.
- Zhang, T., Li, F., Beckes, L., and Coan, J. A. (2013). A semi-parametric model of the hemodynamic response for multi-subject fmri data. *NeuroImage* **75**, 136–145.
- Zhang, T., Li, F., Gonzalez, M. Z., Maresh, E. L., and Coan, J. A. (2014). A semi-parametric nonlinear model for event-related fmri. *NeuroImage* **97**, 178–187.
- Zhu, H. T., Fan, J. Q., and Kong, L. L. (2014). Spatially varying coefficient models with applications in neuroimaging data with jumping discontinuity. *The Journal of the Acoustical Society of America* **109**, 1084–1098.

Received June 2017. Revised October 2017.

Accepted November 2017.