

UC Irvine

UC Irvine Previously Published Works

Title

Understanding the Impact of Stroke on Brain Motor Function: A Hierarchical Bayesian Approach

Permalink

<https://escholarship.org/uc/item/6bp0j3sz>

Journal

Journal of the American Statistical Association, 111(514)

ISSN

0162-1459

Authors

Yu, Zhe
Prado, Raquel
Quinlan, Erin B
[et al.](#)

Publication Date

2016-04-02

DOI

10.1080/01621459.2015.1133425

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



HHS Public Access

Author manuscript

J Am Stat Assoc. Author manuscript; available in PMC 2017 August 18.

Published in final edited form as:

J Am Stat Assoc. 2016 ; 111(514): 549–563. doi:10.1080/01621459.2015.1133425.

Understanding the Impact of Stroke on Brain Motor Function: A Hierarchical Bayesian Approach

Zhe Yu¹, Raquel Prado², Erin B. Quinlan³, Steven C. Cramer⁴, and Hernando Ombao¹

¹Department of Statistics, University of California at Irvine

²Department of Applied Mathematics & Statistics, University of California at Santa Cruz

³Department of Anatomy & Neurobiology, University of California at Irvine

⁴Department of Neurobiology, University of California at Irvine

Abstract

Stroke is a disturbance in blood supply to the brain resulting in the loss of brain functions, particularly motor function. A study was conducted by the UCI Neurorehabilitation Lab to investigate the impact of stroke on motor-related brain regions. Functional MRI (fMRI) data were collected from stroke patients and healthy controls while the subjects performed a simple motor task. In addition to affecting local neuronal activation strength, stroke might also alter communications (i.e., connectivity) between brain regions. We develop a hierarchical Bayesian modeling approach for the analysis of multi-subject fMRI data that allows us to explore brain changes due to stroke. Our approach simultaneously estimates activation and condition-specific connectivity at the group level, and provides estimates for region/subject-specific hemodynamic response functions. Moreover, our model uses spike and slab priors to allow for direct posterior inference on the connectivity network. Our results indicate that motor-control regions show greater activation in the unaffected hemisphere and the midline surface in stroke patients than those same regions in healthy controls during the simple motor task. We also note increased connectivity within secondary motor regions in stroke subjects. These findings provide insight into altered neural correlates of movement in subjects who suffered a stroke.

Keywords

fMRI; activation; connectivity; multi-subject

1. INTRODUCTION

Stroke is a disturbance in the blood supply to the brain that results in the death of neural tissue and subsequent behavioral deficits arising from the affected brain systems. One of the most commonly affected neural systems is the motor system, causing disability in people who suffered a stroke. In order to compensate for the loss of brain function in the affected

7. SUPPLEMENTARY MATERIAL

Supplementary document: the supplementary document contains details of the full conditional posterior distributions, the sampling algorithm, and the settings for the simulation study. (PDF file)

brain area, other areas within the motor system are often recruited—and, therefore, activated—to assist in the execution of movements in stroke-impaired subjects. Also, communications—termed as connectivity in the neuroscience literature—may be altered between regions within the motor network.

To better understand the impact of stroke in brain motor activation and connectivity, we developed a model for analyzing functional magnetic resonance imaging (fMRI) data from a multi-subject stroke study conducted at the University of California Irvine Neurorehabilitation Lab (PI: Cramer). The study recruited healthy subjects and stroke patients with residual motor deficit as it was focused on understanding how brain motor function is altered after stroke. Motor-task-related fMRI scans were acquired for all the participants.

As an imaging modality, fMRI is able to indirectly measure neuronal activity in the brain through the hemodynamic response: higher level of neuronal activity at a localized region requires a greater amount of oxygen, which results in higher level of the blood-oxygen-level-dependent (BOLD) contrast fMRI signal at that location. By analyzing fMRI data, one can study local neuronal activation as well as inter-regional connectivity. The hemodynamic response function (HRF) describes the shape of the hemodynamic response evoked by a point stimulus. The HRF is likely to vary across brain regions and subjects. Thus, it is very important to correctly estimate the HRF for each region and subject in order to correctly infer information regarding activation and connectivity, because neural activity is indirectly measured through the hemodynamic response. The goal of this Bayesian hierarchical modeling approach is to investigate the impact of stroke by comparing brain activation and connectivity of motor-related brain areas between stroke patients and healthy controls, while taking HRF variation into account.

1.1 Overview of the fMRI Stroke Study

The fMRI study involved patients who had suffered a stroke and had residual motor deficit on the right side of the body, and healthy subjects. During the fMRI experiment, each subject alternated between a right hand grasp-release movement (task condition, also called the stimulus condition) and rest (rest condition). Five brain regions of interest (ROI) known to be implicated in motor function are considered for this study (see Figure 1(a)): two primary motor regions including the left and right primary motor cortex (LM1 and RM1), three secondary motor regions including the left and right dorsal premotor cortex (LPMd and RPMd), and a midline supplementary motor area (SMA). Note that a stroke patient with injury on the left brain hemisphere will have motor function deficits on the right side of the body; and vice versa. Typically, when a healthy subject performs a simple right-hand movement task, it is expected that only LM1 (M1 in the left hemisphere) is likely to be activated. RM1 is expected to not activate, since it is primarily responsible for motor function of the left side of the body. The secondary motor regions (SMA and the two PMd regions) are primarily responsible for more advanced motor functions such as motor planning and bilateral coordination.

1.2 Proposed Approach

To analyze these multi-subject fMRI stroke data, we developed a Bayesian approach that jointly estimates activation and connectivity. Our approach provides estimates of subject and region-specific hemodynamic response functions (HRFs). It utilizes the general linear model (GLM) to describe activation, a Bayesian vector-autoregressive model (BVAR) to measure connectivity, and a constrained linear basis set to model the unknown HRFs. The proposed model provides a hierarchical framework that handles group-level activation and connectivity, as well as their variability among subjects. With the hierarchical structure, subject-specific estimates for activation and connectivity are obtained by pooling information from other subjects. Spike and slab priors are placed on the model parameters that describe the connectivity network. This allows us to explore the full posterior distribution of possible connectivity networks. Additionally, we allow condition-specific connectivity measures, thus making it possible to detect differences in connectivity between experimental conditions. Our goal is to study local activation and connectivity between brain regions, and compare the inferred patterns for stroke patients and healthy controls in order to explore the effects of stroke on brain motor function, while controlling for variation in the HRFs.

1.3 Current Statistical Methods and Their Limitations

In the past 20 years, there has been an increasing number of papers on statistical methods for fMRI data. Lindquist (2008) and Zhang et al. (2015) provide detailed reviews on existing methods. Methods for brain activation are proposed in Friston et al. (1994); Worsley and Friston (1995); Smith and Fahrmeir (2007); David et al. (2008); Guo and Pagnoni (2008); Xu et al. (2009); Degras and Lindquist (2014). Methods for brain connectivity are developed in Friston et al. (2003); Harrison et al. (2003); Bowman et al. (2008); Cribben et al. (2012); Kang et al. (2012); Gorrostiti et al. (2013); Luo (2014).

While methods for activation usually do not consider connectivity estimation, David et al. (2008) in fact considered connectivity estimation, but used a two-stage approach that estimates connectivity from the residuals of the activation analysis, and hence is not optimal for properly assessing the uncertainty associated to those estimates. Harrison et al. (2003) estimated connectivity using vector-autoregressive (VAR) models, while Cribben et al. (2012) developed dynamic connectivity regression which models changes in connectivity, but none of these approaches take into account the systematic changes in the BOLD response induced by external stimuli. Gorrostiti et al. (2013) considered the BOLD response evoked by stimuli and also the differential connectivity across different experimental conditions. However, the BOLD response shape was based on a fixed HRF shape for all brain regions and subjects, which may lead to erroneous conclusions when inferring activation and connectivity (Yu et al., 2015). Bowman et al. (2008) and Luo (2014) also took the experimental conditions into account, but modeled connectivity through a two-stage approach and used a pre-specified HRF. Kang et al. (2012) simultaneously analyzed activation and connectivity in the spectral domain, but also used a pre-specified HRF. Friston et al. (2003) used dynamic causal modeling (DCM) which addresses ROI-specific HRF and connectivity simultaneously. However, DCM heavily relies on the correct specification of

the biological assumptions, and needs a set of plausible connectivity networks a priori which are difficult to verify from the observed data.

There are many ways to parameterize HRFs. Non-linear parameterization approaches include gamma functions, cosine functions (Zarahn, 2002), or inverse logit functions (Lindquist et al., 2008). Linear basis approaches such as spline basis functions, finite impulse response bases (Ollinger et al., 2001) and semi-parametric models (Zhang et al., 2001) are also available. Non-linear parameterization takes more computational effort and is more prone to the local maxima problem and numerical stability issues. On the other hand, linear parameterization is easier from the inferential viewpoint, but often requires more parameters and constraints in order to provide enough flexibility in modeling HRF shapes. Woolrich et al. (2004) proposed constrained linear basis sets which are able to provide both flexible and reasonable HRF shapes with a relatively small number of HRF basis functions while keeping the computation cost low. Therefore, in this paper we adopt the constrained linear basis approach to parameterize the HRFs.

In a nutshell, we propose a model that has the following advantages over currently available methods: (1) it simultaneously estimates activation, connectivity and HRFs; (2) it provides ROI and subject-specific HRFs, as well as condition-specific connectivity measures; (3) it increases the power in detecting group differences by pooling information across subjects via hierarchical modeling; (4) it provides full posterior distributions of all the model parameters, including group-specific activation parameters and brain connectivity networks; (5) it can incorporate available scientifically relevant prior information. The remainder of the paper is organized as follows. In Section 2 we describe the stroke study and the fMRI data. In Section 3 we develop our approach including the model, the prior distributions and the inferential procedure. In Section 4 we present the simulation studies, and in Section 5 we present the analysis of the fMRI data from the stroke study. We summarize our approach and discuss future research directions in Section 6.

2. EXPERIMENTAL DESIGN AND FMRI DATA

Experiment

The study included 15 stroke patients and 12 healthy control subjects. All the subjects were right-handed. The stroke patients had ischemic stroke on the left brain hemisphere 11 – 26 weeks prior to the study assessments, and continued to have residual motor deficit on the right side of the body at the time of the experiment. During the task condition, the subjects used their right hand to perform the hand grasp-release movement task. For the stroke patients, the right hand was also the hand that was affected by stroke. The experiment was divided into three sessions for stroke patients and two sessions for healthy controls. Each session had 48 consecutive scans, alternating between task and rest conditions twice, but always starting with rest condition (see Figure 2).

FMRI data

The fMRI images were acquired using a T2*-weighted gradient-echo-planar imaging sequence with repetition time (TR) = 2 s. Functional data from all the sessions were

preprocessed using the SPM8 software (Wellcome Trust Center for Neuroimaging, UCL, 2009). Preprocessing steps included realignment to the first image, coregistration to the mean image, normalization to the standard MNI EPI template, and spatial smoothing (FWHM = 8 mm). The time series for each ROI were obtained by averaging the fMRI signals recorded across the voxels in the region. The mean time series were centered to zero and detrended with time as the covariate to remove the drift effect. Finally, the centered and detrended time series were scaled to have the same variance for all subjects and ROIs within each group.

Selected time series from the LM1 and RM1 regions in a healthy subject are shown in Figure 1(b). The block-shaped line at the bottom indicates the times when the subject performed the motor task. In LM1 (top plot) the fMRI signal generally follows the block-shaped wave with some lag. It rises soon after the task condition begins and drops soon after the task condition ends, which suggests that this region is implicated in the motor task. However, in RM1 (bottom plot) the time series does not appear to be obviously associated with the timing of the motor task.

3. METHODOLOGY

In this section we develop the Bayesian hierarchical modeling approach and related inferential procedures for investigating the effect of stroke on brain motor function, specifically on activation and connectivity of the five ROIs related to motor function. The proposed model will be utilized to infer activation and connectivity at the group level while taking into account subject/region-specific variability in the HRFs.

3.1 Hierarchical Model: Subject Level

To describe the fMRI signal for a given ROI and subject, we use the general linear model:

$$\begin{aligned} y_p^{(s,r)}(t) &= \beta_{p0}^{(s)} + x_{p1}^{(s,r)}(t)\beta_{p1}^{(s)} + x_{p2}^{(s,r)}(t)\beta_{p2}^{(s)} + u_p^{(s,r)}(t), \\ x_{pk}^{(s)}(t) &= [h_p^{(s)} * c_k](t), \end{aligned} \quad (1)$$

with ROI $p = 1, 2, \dots, P$; subject $s = 1, 2, \dots, S$; session $r = 1, 2, \dots, R$; and time (in a session) $t = 1, 2, \dots, T(\text{TR})$. The response variable $y_p^{(s,r)}(t)$ is the preprocessed fMRI signal; the covariate $x_{pk}^{(s,r)}(t)$ represents the expected shape of BOLD response, which is the convolution between the HRF, denoted as $h_p^{(s)}(t)$, and the known binary condition indicator with $c_k(t)$, with $c_k(t) = 1$ when the k th condition is on and zero otherwise (we use $k = 1$ for the task condition and $k = 2$ for the rest condition, see Figure 2 for the timing of the conditions); the regression coefficient $\beta_{pk}^{(s)}$ represents the amplitude of the BOLD response which indirectly reflects the strength of neuronal activity; and $u_p^{(s,r)}(t)$ is the noise that is not explained by the mean structure. The effect of the convolution is illustrated in Figure 3. The expected BOLD response is no longer block-shaped as the stimulus input, but is a smooth and delayed transformation of the input. In this model, we have imposed the following

assumptions: (1) the BOLD amplitude remains the same within a condition and across sessions; (2) the HRF shape within each ROI and subject remains constant across conditions and sessions.

On modeling activation—Neuroscientists are interested in estimating the difference between activity strength of two conditions as measured by a region-specific contrast, i.e.,

$\beta_{p,(c)}^{(s)} = \beta_{p1}^{(s)} - \beta_{p2}^{(s)}$. When ROI p shows stronger activity during the task than during the control condition, it suggests that this ROI is implicated in executing the task. Therefore, when $\beta_{p,(c)}^{(s)} > 0$, ROI p is labeled as active in subject s , and $\beta_{p,(c)}^{(s)}$ is used to measure the strength of the local activation.

On modeling connectivity—Following Gorrostieta et al. (2013), the noise $\mathbf{u}^{(s,r)}(t) \in \mathbb{R}^P$ is assumed to follow a Gaussian vector autoregressive model (VAR) with pre-determined order L and condition-dependent VAR coefficients, i.e.,

$$\mathbf{u}^{(s,r)}(t) = [u_1^{(s,r)}(t), \dots, u_p^{(s,r)}(t)]' = \sum_{\ell=1}^L \Phi^{*(s)}(\ell, t) \mathbf{u}^{(s,r)}(t-\ell) + \boldsymbol{\varepsilon}^{(s,r)}(t),$$

where $\boldsymbol{\varepsilon}^{(s)}(t) = [\varepsilon_1^{(s,r)}(t), \dots, \varepsilon_p^{(s,r)}(t)]' \stackrel{\text{ind}}{\sim} N(\mathbf{0}, \boldsymbol{\Sigma}^{(s)})$, and

$$\Phi^{*(s)}(\ell, t) = c_1(t-\ell)\Phi_1^{*(s)}(\ell) + c_2(t-\ell)\Phi_2^{*(s)}(\ell). \quad (2)$$

Denote the (p, q) -th element of $\Phi^{*(s)}(\ell, t)$ as $\Phi^{*(s)}(\ell, t)_{pq}$. This quantity describes the linear relation between ROIs p and q at time lag ℓ when $\Phi^{*(s)}(\ell, t)_{pq} > 0$, ROI q at time $t - \ell$ helps predict ROI p at the current time point t . The temporal precedence of ROI q in this relationship suggests the interpretation of $\Phi^{*(s)}(\ell, t)_{pq}$ as a measure of the directed influence from ROI q to ROI p . This type of influence is termed as “effective connectivity” in neuroscience. Therefore, $\Phi^{*(s)}(\ell, t)_{pq}$ can be used to measure effective connectivity from ROI q to ROI p . For simplicity, in the remainder of this paper we refer to this as “connectivity”. Also, it is possible that the manner in which an ROI influences another varies under different conditions. Hereby, we allow $\Phi^{*(s)}(\ell, t)_{pq}$ to depend on the condition at the corresponding past time point as specified in Equation (2).

We use a Bayesian variable selection approach (Mitchell and Beauchamp, 1988; George and McCulloch, 1993; Ishwaran and Rao, 2005) to determine which elements of $\Phi^{*(s)}(\ell, t)$ are different from zero, and therefore infer the connectivity network across ROIs. More specifically, we impose a “spike and slab” structure on $\Phi_k^{*(s)}$ given by

$$\Phi_k^{*(s)}(\ell)_{pq} = \Phi_k^{(s)}(\ell)_{pq} \boldsymbol{\xi}_{pq}^{(s)}(\ell), \quad \text{with } \Phi_k^{(s)}(\ell)_{pq} \in \mathbb{R}, \quad \boldsymbol{\xi}_{pq}^{(s)}(\ell) \in \{0, 1\}, \quad (3)$$

so that $\Phi_k^{*(s)}(\ell)_{pq}=0 \iff \xi_{pq}^{(s)}(\ell)=0$ (a. s.); $\Pr(\Phi_k^{*(s)}(\ell)_{pq} \neq 0)=\Pr(\xi_{pq}^{(s)}(\ell)=1)$, and $\Phi_k^{*(s)}(\ell)_{pq}=\Phi_k^{*(s)}(\ell)_{pq} \iff \xi_{pq}^{(s)}(\ell)=1$ (a. s.). When $\Phi_k^{*(s)}(\ell)_{pq}=0$, or equivalently when $\xi_{pq}^{(s)}(\ell)=0$ for all ℓ we say that there is no connectivity from ROI q to ROI p . Otherwise, there is connectivity from ROI q to ROI p .

On HRF modeling—A linear basis for the representation of the HRFs was obtained based on the approach of Woolrich et al. (2004). In particular, using this approach, the HRF of

subject s and region p can be written in terms of J basis vectors as $h_p^{(s)}(t)=\sum_{j=1}^J H(t, j)d_{pj}^{(s)}$, where H is the $T \times J$ basis matrix with each column being a basis vector, and $d_p^{(s)} \in \mathbb{R}^J$ is the vector of HRF basis coefficients for ROI p and subject s . For identifiability, we impose a normalization constraint on $d_p^{(s)}$ so that $\sum_{t=1}^T h_p^{(s)}(t)=1$. Further details on the basis representation and how to choose J are given in Section 3.3.

We use $A(i, :)$, $A(:, j)$, and $A(i, j)$ to denote, respectively, the i -th row, the j -th column, and the (i, j) -th element of any matrix A . Then, $h_p^{(s)}(t)$ can be written as $H(t, :)d_p^{(s)}$ with $d_p^{(s)} \in \mathbb{R}^J=[d_{p1}^{(s)}, \dots, d_{pJ}^{(s)}]$, and $x_{pk}^{(s)}(t)$ can be written as $x_{pk}^{(s)}(t)=[H * c_k](t, :)d_p^{(s)}:=\Lambda_k(t, :)d_p^{(s)}$, $t=1, 2, \dots, T$, where the matrix $\Lambda_k \in \mathbb{R}^{T \times J}$ consists of the convolutions between each linear basis vector $H(:, j)$, and the condition indicator $c_k(t)$. Therefore, combining this with Equation (1) leads to

$$y_p^{(s,r)}(t)=\beta_{p0}^{(s)}+\Lambda_1(t, :)d_p^{(s)}\beta_{p1}^{(s)}+\Lambda_2(t, :)d_p^{(s)}\beta_{p2}^{(s)}+u_p^{(s,r)}(t).$$

3.2 Hierarchical Model: Group Level

Since the signal to noise ratio is usually low in fMRI data (typically below 5%), it is beneficial to borrow information across subjects within a relatively homogeneous group. In order to do this, we use a Bayesian hierarchical modeling approach. We define

$$\begin{aligned} \beta_p^{(s)} &= [\beta_{p0}^{(s)}, \beta_{p1}^{(s)}, \beta_{p2}^{(s)}]', \beta^{(s)} \\ &= [(\beta_1^{(s)})', \dots, (\beta_p^{(s)})']'; \phi_k^{(s)} \\ &= [\text{vec}(\Phi_k^{(s)}(1))', \dots, \text{vec}(\Phi_k^{(s)}(L))'], \phi^{(s)} = [(\phi_1^{(s)})', (\phi_2^{(s)})']', \xi_{pq}^{(s)} = [\xi_{pq}^{(s)}(1), \dots, \xi_{pq}^{(s)}(L)] \end{aligned}$$

and $d^{(s)} = [(d_1^{(s)})', \dots, (d_p^{(s)})']'$. We then assume that each of the S subjects belongs to a single group, denoted by g_s , from a total of G groups (e.g., in our case we have stroke patients and healthy controls, so $G=2$), and consider the following distributions,

$$\begin{aligned} \beta^{(s)} &\stackrel{\text{ind}}{\sim} N(\boldsymbol{\mu}_\beta^{g_s}, \Sigma_\beta^{g_s}), & d^{(s)} &\stackrel{\text{ind}}{\sim} N(\boldsymbol{\mu}_d^{g_s}, \Sigma_d^{g_s}), & \phi^{(s)} &\stackrel{\text{ind}}{\sim} N(\boldsymbol{\mu}_\phi^{g_s}, \Sigma_\phi^{g_s}), \\ \xi_{pq}^{(s)} &\stackrel{\text{ind}}{\sim} \text{Multinomial}(1, \boldsymbol{\pi}_{pq}^{g_s}), & \Sigma^{(s)} &\stackrel{\text{ind}}{\sim} \text{InvWishart}(\Sigma^{g_s}, \nu_\Sigma). \end{aligned}$$

Here $\Pr(\xi_{pq}^{(s)} = \ell_j) = \pi_{pq}^{g_s}(j)$, where $\ell_j = [\underbrace{1, \dots, 1}_{j \text{ 1's}}, \underbrace{0, \dots, 0}_{L-j \text{ 0's}}]$, $j = 0, 1, \dots, L$. The j in $\xi_{pq}^{(s)} = \ell_j$ can be interpreted as the largest lag at which there is connectivity from ROI q to ROI p .

Therefore, all the subjects belonging to the same group, say group g , will have the same group parameters $\mu_\beta^g, \Sigma_\beta^g, \mu_d^g, \Sigma_d^g, \mu_\phi^g, \Sigma_\phi^g, \pi_{pq}^g$, and Σ^g . Note that μ_β^g are the mean BOLD amplitudes and intercepts for group g , μ_d^g are the mean HRF basis coefficients, μ_ϕ^g are the mean connectivity strengths, and $\pi_{pq}^g(\ell)$ is the overall probability that connectivity from ROI q to ROI p exists up to ℓ lags among the group. For simplicity, ν_2 is taken as a fixed constant across groups, and $\Sigma_\beta^{g_s}, \Sigma_d^{g_s}$ and $\Sigma_\phi^{g_s}$ are assumed to be diagonal.

For any g , the distributions for the group parameters are taken as

$$\mu_\beta^g \stackrel{ind}{\sim} N(\mu_{\mu_\beta,0}, \Sigma_{\mu_\beta,0}), \Sigma_\beta^g(i,i) \stackrel{ind}{\sim} InvGamma(a_\beta, b_\beta), \mu_d^g \stackrel{ind}{\sim} N(\mu_{\mu_d,0}, \Sigma_{\mu_d,0}), \Sigma_d^g(i,i) \stackrel{ind}{\sim} InvGamma(a_d,$$

, and $p(\Sigma^g) \propto |\Sigma^g|^{-\frac{1+P}{2}}$. Here $\mu_{\mu_\beta,0}, \Sigma_{\mu_\beta,0}, \mu_{\mu_\phi,0}, \Sigma_{\mu_\phi,0}, a_\beta, b_\beta, a_d, b_d, a_\phi, b_\phi$ and a_π are pre-determined constants. The choices of these constants will be discussed in Sections 4 and 5. In fMRI datasets with more than two groups, another level can be added to the hierarchy by imposing additional priors on the group parameters for activation, connectivity and those related to the HRFs.

The Bayesian hierarchical model is now fully specified. The quantities of interest are as follows.

Activation for ROI p — $\beta_{p1}^{(s)} - \beta_{p2}^{(s)}$ is the activation strength for subject s , and $\mu_{\beta,p1}^g - \mu_{\beta,p2}^g$ is the mean activation strength for group g (group-level activation strength). Here $\mu_{\beta,pk}^g$ denotes the element in μ_β^g corresponding to $\beta_{pk}^{(s)}$.

Connectivity from ROI q to ROI p — $\Pr(\xi_{pq}^{(s)}(\ell) = 1)$ is the probability that such connectivity is present at lag ℓ for subject s , for $\ell = 1, \dots, L$; $\Phi_k^{*(s)}(\ell)_{pq}$ is the measure of the connectivity strength under condition k at lag ℓ for subject s ; $\pi_{pq}^g(\ell)$ is the probability that the largest lag of such connectivity is ℓ for group g (group-level presence probability of connectivity), and $\mu_{\Phi_k(\ell),pq}^g$ is the mean connectivity for the group, (group-level connectivity strength).

HRF for ROI p — $h_p^{(s)} = Hd_p^{(s)}$ is the HRF of ROI p for subject s and $\mu_{h_p}^g = H\mu_{d_p}^g$ is the mean HRF of ROI p for group g .

In the particular case of the stroke study, based on a preliminary analysis of the data, we found that there is a large within-group variability in the shape of the HRFs across subjects (see Figure 4). The preliminary analysis was done by computing the posterior mode of the activation and HRF parameters using a conditional maximization algorithm, ignoring all the

dependence across time, ROIs and subjects. We note that averaging curves of very different shapes within a group is likely to produce meaningless overall group-specific shapes due to the non-linearity involved. Hence, it is prudent to not pool the highly heterogeneous shapes of the HRFs across subjects. Therefore, specifically for the analysis of the stroke data, we fix μ_d^g as a constant. This will be discussed in detail in the next subsection. However, we still combine information on activation and connectivity from other subjects.

3.3 Constrained Linear Basis for the HRFs

Here we briefly describe how to obtain the constrained linear basis for the HRFs. Following the general framework of Woolrich et al. (2004), one can simulate n discrete HRF vectors based on the half-cosine parameterization. The parameters involved are the temporal resolution δ , the duration parameters h_1, h_2, h_3, h_4 , and the depth parameters f_1 and f_2 (see Figures A.1(a) and A.1(b) in the Appendix). Using principal component analysis (PCA) one can obtain J linear basis vectors for the HRFs, and the corresponding basis coefficients (loadings of the top J components), denoted as d_i for the i -th simulated HRF, $i = 1, 2, \dots, n$. For the simulation study and the fMRI data analysis in this paper we used $n = 1000$, $\delta = 0.1$, $h_1 \sim \text{Unif}(0, 2)$, $h_2 \sim \text{Unif}(2, 7)$, $h_3 \sim \text{Unif}(2, 8)$, $h_4 \sim \text{Unif}(2, 12)$, $f_1 = 0$, $f_2 \sim \text{Unif}(0, 0.5)$, and $J = 5$. The reason for choosing $J = 5$ is that the top 5 principal components for the HRF basis set, i.e., the first $J = 5$ columns of the matrix H , explain about 99% of the total variability in the simulated HRFs (see Figure A.2(a) in the Appendix). Note that it is also the case that the first HRF basis vector is very similar to the canonical HRF (see Figure A.2(b) in the Appendix).

As a next step, we compute the empirical mean $\mu_d^* \in \mathbb{R}^J$ and variance $\sum_d^* \in \mathbb{R}^{J \times J}$ of the basis coefficients $\{d_i\}_{i=1}^n$. We then use the Gaussian distribution, $N(\mu_d^*, \sum_d^*)$, as the prior distribution for the HRF basis coefficients $d^{(s)}$, i.e., use $\mu_d^g = \mu_d^*$ and $\sum_d^g = \sum_d^*$ for all g . Such prior distribution on the constrained linear basis representation provides enough flexibility to capture the HRF shapes, e.g., time to peak, width of peak and time to undershoot, and simultaneously penalizes shapes that deviate from the space of reasonable shapes. Note that our prior beliefs about the HRFs are adopted through the choice of the duration parameters and the depth parameters. The normal distribution just provides a convenient approximation of such prior beliefs. In order to determine if the prior structure is reasonable, visual checks of the quality of the basis set can be performed by generating prior samples of the HRFs using the basis vectors and corresponding coefficients randomly sampled from the normal distribution described above. Figure A.2(c) displays examples of randomly generated HRFs from the HRF basis set used in our analysis. The HRFs sampled from the prior indicate that this HRF basis representation is sufficiently flexible in capturing time to peak, width of peak, time to undershoot, and lead to reasonable HRF shapes.

3.4 Posterior Inference via MCMC

The hierarchical structure of the model described in the previous section is such that information is pooled across subjects within a group and not across groups. This combined with the fact that the group-specific hyperparameters are fixed and known simplifies the

inference, as the full posterior distribution can be written as a product of group-specific components. This implies that sampling from the joint posterior distribution of all the parameters is equivalent to separately sampling from the group-specific joint posterior distributions. Therefore, we only discuss single group estimation and drop the group indicator in this section for convenience.

The subject level model can essentially be decomposed into three multivariate linear regression (MLR) components when conditioning on all the remaining components. We now describe each component. Define

$$\mathbf{y}^{(s,r)}(t) = [y_1^{(s,r)}(t), \dots, y_p^{(s,r)}(t)]', \quad \mathbf{x}_p^{(s)}(t) = [1, x_{p1}^{(s)}(t), x_{p2}^{(s)}(t)], \text{ and}$$

$$X^{(s)}(t) = \text{Bdiag}\{\mathbf{x}_1^{(s)}(t), \dots, \mathbf{x}_p^{(s)}(t)\}. \text{ Also define } \mathbf{y}_0^{(s,r)}(t) = \mathbf{y}^{(s,r)}(t) - [\beta_{10}^{(s)}, \dots, \beta_{p0}^{(s)}]', \text{ and}$$

$$Z^{(s)}(t) = \text{Bdiag}\{\Lambda_1(t, :)\beta_{11}^{(s)} + \Lambda_2(t, :)\beta_{12}^{(s)}, \dots, \Lambda_1(t, :)\beta_{p1}^{(s)} + \Lambda_2(t, :)\beta_{p2}^{(s)}\}. \text{ Then,}$$

$$\mathbf{y}^{(s,r)}(t) = X^{(s)}(t)\boldsymbol{\beta}^{(s)} + \mathbf{u}^{(s,r)}(t) \text{ and } \mathbf{y}_0^{(s,r)}(t) = Z^{(s)}(t)\mathbf{d}^{(s)} + \mathbf{u}^{(s,r)}(t). \quad (4)$$

Note that the noise terms $\mathbf{u}^{(s,r)}(t)$ are correlated over time due to the VAR structure, and therefore we need an extra step to calculate the full conditionals: “whitening” the data.

Define $\Phi_t^{(s)}(\mathbf{B}) = I_p - \sum_{\ell=1}^L \Phi^{*(s)}(\ell, t)\mathbf{B}^\ell$, where \mathbf{B} is the backshift operator, then we have $\boldsymbol{\varepsilon}^{(s,r)}(t) = \Phi_t^{(s)}(\mathbf{B})[\mathbf{u}^{(s,r)}(t)]$. Further define

$$\tilde{\mathbf{y}}^{(s,r)}(t) = \Phi_t^{(s)}(\mathbf{B})[\mathbf{y}^{(s,r)}(t)] = \mathbf{y}^{(s,r)}(t) - \sum_{\ell=1}^L \Phi^{*(s)}(\ell, t)\mathbf{y}^{(s,r)}(t-\ell), \quad (5)$$

$$\tilde{\mathbf{y}}_0^{(s,r)}(t) = \Phi_t^{(s)}(\mathbf{B})[\mathbf{y}_0^{(s,r)}(t)], \quad (6)$$

$$\tilde{X}^{(s)}(t) = \Phi_t^{(s)}(\mathbf{B})[X^{(s)}(t)], \quad (7)$$

$$\tilde{Z}^{(s)}(t) = \Phi_t^{(s)}(\mathbf{B})[Z^{(s)}(t)]. \quad (8)$$

Then, the equations in (4) can be written in terms of the temporally uncorrelated noise terms, leading to the following two MLR components, one for the BOLD amplitude parameters $\boldsymbol{\beta}^{(s)}$ and another for the HRF parameters $\mathbf{d}^{(s)}$:

$$\tilde{\mathbf{y}}^{(s,r)}(t) = \tilde{X}^{(s)}(t)\boldsymbol{\beta}^{(s)} + \boldsymbol{\varepsilon}^{(s,r)}(t), \quad (9)$$

$$\tilde{\mathbf{y}}_0^{(s,r)}(t) = \tilde{Z}^{(s)}(t)\mathbf{d}^{(s)} + \boldsymbol{\varepsilon}^{(s,r)}(t). \quad (10)$$

For computational efficiency, we work with the likelihood that is conditional on the initial L observations in each session, instead of the full likelihood.

The third MLR component comes from the VAR structure, when we use the conditional likelihood and condition on all other parameters. More specifically, define

$$\begin{aligned} \mathbf{w}_k^{(s,r)}(t) &= [c_k(t-1)\mathbf{u}^{(s,r)}(t-1)' \otimes I_P, \dots, c_k(t-L)\mathbf{u}^{(s,r)}(t-L)' \otimes I_P], \\ W^{(s,r)}(t) &= [\mathbf{w}_1^{(s,r)}(t), \mathbf{w}_2^{(s,r)}(t)], \quad \tilde{W}^{(s,r)}(t)_{p,i} = W^{(s,r)}(t)_{p,i} \boldsymbol{\xi}_{pq}^{(s)}(\ell), \\ \boldsymbol{\phi}^{*(s)}(i) &= \boldsymbol{\phi}^{(s)}(i) \boldsymbol{\xi}_{pq}^{(s)}(\ell), \quad \forall p, q, \ell, k, \\ \text{where } i &= J(p, q, \ell, k) := p + (q-1)P + (\ell-1)P^2 + (k-1)P^2L. \end{aligned} \quad (11)$$

Then, the third MLR component can be expressed as

$$\mathbf{u}^{(s,r)}(t) = \tilde{W}^{(s,r)}(t)\boldsymbol{\phi}^{*(s)} + \boldsymbol{\varepsilon}^{(s,r)}(t). \quad (12)$$

A MCMC algorithm was implemented to obtain samples from the joint posterior distribution of the model parameters. The fact that we have three MLR components simplifies the steps in the MCMC algorithm. In particular, combining Equations (9), (10) and (12) above with the Gaussian distributions for $\boldsymbol{\beta}^{(s)}$, $\mathbf{d}^{(s)}$ and $\boldsymbol{\phi}^{(s)}$, we have that the full conditional distributions for these parameters are also Gaussian, resulting in Gibbs steps. Similarly, the full conditional distributions for $\boldsymbol{\mu}_\beta^g$, $\boldsymbol{\mu}_d^g$ and $\boldsymbol{\mu}_\phi^g$ are also Gaussian. The diagonal entries of the covariance matrices, $\Sigma_{\boldsymbol{\beta}^g(i, i)}^{(g)}$, $\Sigma_{\boldsymbol{\phi}^g(i, i)}^{(g)}$ and $\Sigma_{\mathbf{d}^g(i, i)}^{(g)}$, are sampled from Inverse-Gamma distributions. The noise covariance matrices $\Sigma^{(s)}$, Σ^g are sampled from Inverse-Wishart distributions. Sampling for $\boldsymbol{\xi}_{pq}^{(s)}$ follows the approach of Koop and Korobilis (2009) for stochastic search variable selection. The probabilities $\pi_{pq}^{(g)}$ are sampled from Dirichlet distributions. The choice of conditionally conjugate priors for the model parameters largely simplifies the structure of the MCMC algorithm for posterior inference. The full conditional posterior distributions and all the steps of the MCMC algorithm are detailed in the supplementary material.

The point estimates for the quantities of interest presented here are based on posterior medians of the parameters obtained after MCMC convergence, the only exception being the

quantities $\Pr(\xi_{pq}^{(s)} = \ell_j | Y)$ for each p, q , for which we use the posterior means. More specifically, inference for each quantity of interest is as follows.

Activation—We calculate a $(1 - \alpha) \times 100\%$ posterior credible interval, referred to as CI, for the group level contrast $\mu_{\beta, p, (c)}^g = \mu_{\beta, p1}^g - \mu_{\beta, p2}^g$ for each ROI p . If the interval is entirely on the right side of zero, it indicates that ROI p is activated under the task condition for group g . For each specific subject, we calculate a $(1 - \alpha) \times 100\%$ posterior CI of the subject level contrast $\beta_{p, (c)}^{(s)}$. If the interval is entirely on the right side of zero, it indicates that ROI p is activated for subject s . In our analysis we calculated the $(1 - \alpha) \times 100\%$ credible intervals as the intervals between the $\alpha/2$ -th and $(1 - \alpha/2)$ -th quantiles obtained using the MCMC posterior samples.

Connectivity—For connectivity from ROI q to ROI p at the subject level we estimate the probability $\Pr(\xi_{pq}^{(s)} = \ell_j | Y)$, for $j = 0, 1, \dots, L$, using the posterior mean, and find the j that maximizes such probability. If $j = 0$, there is no connectivity from ROI q to p for subject s ; otherwise there is connectivity. Posterior medians and CIs of $\Phi_k^{(s)}(\ell)_{pq}$ are calculated to measure the connectivity strength at lag $\ell = 1, 2, \dots, j$ if $j > 0$. At the group level, we calculate the posterior medians of $\pi_{pq}^g(j)$ as the overall presence probability of the connectivity among the group. Further examination of the corresponding posterior CIs of $\mu_{\Phi_k(\ell), pq}^g$ for $\ell = 1, \dots, j$, provides inference on the direction and strength of the connectivity for each condition k .

HRFs—We derive posterior samples of the HRFs by calculating $Hd_p^{(s)}$ for each posterior sample of $d_p^{(s)}$; pointwise credible bands of the HRFs can hence be calculated. Group-level HRF estimates and credible bands can also be obtained similarly from $H\mu_{d_p}^g$. Neuroscientists are usually interested in the time to peak, width of the peak measured by full width at half maximum (FWHM), and time to post-stimulus undershoot. Thus, we can also calculate CIs of these quantities from the HRF posterior samples.

3.5 Computational Aspects

Given the large number of parameters involved in the model, MCMC sampling is computationally heavy. In order to increase the computational efficiency, we implemented the sampling algorithm for the subject-level parameters in C++ core using the R package *RcppArmadillo* (Eddelbuettel and Sanderson, 2014). The sampling for the remaining parameters was implemented in R. We also considered parallel computing so that the subject-level parameters were sampled simultaneously for all subjects.

For datasets with a large number of ROIs the current MCMC algorithm can be too expensive computationally. This is mainly due to multiplications and inversions of large matrices involved in the estimation of the connectivity parameters $\phi^{(s)}$, and the computation of the probabilities $\Pr(\xi_{pq}^{(s)} = \ell_j)$ given Y and the remaining parameters for $j = 0, \dots, L; p, q = 1, \dots,$

P , since they require likelihood evaluations sequentially for P^2 times. Computation time can be reduced by imposing special sparsity structures, e.g., rowwise sparsity, instead of element-wise sparsity for the spike and slab prior, at the cost of decreasing flexibility on the sparsity of the connectivity. Another way to speed up the estimation is to consider algorithms that approximate the joint posterior distribution such as variational Bayes (VB). A VB algorithm for approximate posterior inference is discussed in the next subsection.

3.6 Approximate Posterior Inference via Variational Bayes

Variational Bayes (VB) algorithms (Jordan et al., 1999) have been used in many applied scenarios, including fMRI data analysis (Woolrich et al., 2004; Luessi et al., 2014), as a computationally efficient tool for obtaining approximations to the joint posterior distribution. We begin by discussing a VB algorithm for a single subject s , and then discuss extensions to handle multiple subjects. For simplicity, the superscript for subject s will be dropped unless otherwise stated.

The posterior density for the model described in the previous section can be written as

$$p(\Theta|Y) \propto f(Y|\beta, \mathbf{d}, \phi, \xi, \Sigma) p(\beta) p(\mathbf{d}) p(\phi) \prod_{pq} p(\xi_{pq}) p(\Sigma),$$

with $f(Y|\cdot)$ being the likelihood function and $\Theta = (\beta, \mathbf{d}, \phi, \xi_{11}, \xi_{12}, \dots, \xi_{PP}, \Sigma)$. The approximation function we use has the form $Q(\Theta) = Q(\beta)Q(\mathbf{d})Q(\phi)\prod_{pq} Q(\xi_{pq})Q(\Sigma)$. Let θ denote one of the parameter subsets in Θ , i.e., θ could be one of $\beta, \mathbf{d}, \phi, \xi_{pq}$, for some p, q , or Σ ; and let θ^C denote $\Theta \setminus \theta$. In order to obtain $\log Q(\Theta)$, the approximate log posterior density for the parameter set θ that minimizes the Kullback-Leibler (KL) between the approximation and the posterior denoted as $KL(Q||p)$, we need to calculate $\mathbb{E}_{Q(\theta^C)} \log p(\Theta|Y)$ (see, e.g., Gelman et al., 2013). Then, we proceed as follows. Define $\|x\|_A^2 = x' A^{-1} x$. Given the conjugate structure in the model, when θ is equal to β, \mathbf{d} or ϕ we have

$$\log Q(\theta) = \text{const} - \frac{1}{2} \|\theta - \tilde{\mu}_\theta\|_{\tilde{\Sigma}_\theta}^2, \text{ implying that } Q(\theta) = \mathcal{N}(\tilde{\mu}_\theta, \tilde{\Sigma}_\theta). \text{ Similarly, when } \theta = \xi_{pq}, \text{ we have } \log Q(\theta) = \text{const} + \sum_{j=0}^L I(\theta = \ell_j) \log \tilde{\pi}_{j,\theta} \text{ (s. t. } \sum_{j=0}^L \tilde{\pi}_{j,\theta} = 1), \text{ and therefore } Q(\theta) = \text{Multinomial}(1, \{\tilde{\pi}_{j,\theta}\}_{j=0}^L). \text{ Finally, when } \theta = \Sigma, \text{ we have}$$

$$\log Q(\theta) = -\frac{1}{2} (\tilde{\nu}_\theta + P + 1) \log |\tilde{\Psi}_\theta| - \frac{1}{2} \text{tr}(\tilde{\Psi}_\theta \theta^{-1}) \text{ implying that } Q(\theta) = \text{InvWishart}(\tilde{\Psi}_\theta, \tilde{\nu}_\theta). \text{ The values of } \tilde{\mu}_\beta, \tilde{\Sigma}_\beta, \tilde{\mu}_\mathbf{d}, \tilde{\Sigma}_\mathbf{d}, \tilde{\mu}_\phi, \tilde{\Sigma}_\phi, \tilde{\pi}_{j,\xi_{pq}} \text{ for } j = 0, \dots, L, \text{ and } \tilde{\Psi}_\Sigma \text{ can be computed in an iterative manner until convergence, as briefly outlined below; and } \tilde{\nu}_\Sigma \text{ is a fixed constant. Details are given in the supplementary material.}$$

Define $\tilde{\pi}_{\xi,pq} = [\tilde{\pi}_{0,\xi,pq}, \dots, \tilde{\pi}_{L,\xi,pq}]$. Let $\tilde{\mu}_{\xi,pq}$ and $\tilde{\Sigma}_{\xi,pq}$ denote the mean and the covariance matrix for $\xi_{pq} = [\xi_{pq}(1), \dots, \xi_{pq}(L)]$ under $Q(\cdot)$; both quantities can be derived from $\tilde{\pi}_{\xi,pq}$ according to the Multinomial distributions. Let $\tilde{\mu}_\xi$ and $\tilde{\Sigma}_\xi$ denote $\{\tilde{\mu}_{\xi,pq} : p, q \in \{1, \dots, P\}\}$ and $\{\tilde{\Sigma}_{\xi,pq} : p, q \in \{1, \dots, P\}\}$, respectively; and let $\tilde{\mu}_{\xi,-pq}$ and $\tilde{\Sigma}_{\xi,-pq}$ denote $\{\tilde{\mu}_{\xi,p'q'} : p', q' \in \{1, \dots, P\}^2 \setminus \{(p, q)\}\}$ and $\{\tilde{\Sigma}_{\xi,p'q'} : p', q' \in \{1, \dots, P\}^2 \setminus \{(p, q)\}\}$. Define $\mu_{\Sigma^{-1}}$ to be $E_Q(\Sigma^{-1}) = \tilde{\Psi}_\theta^{-1} \tilde{\nu}_\theta$. Then, as shown in the supplementary material, it is possible to write, $\tilde{\mu}_\beta$

$= f_{\tilde{\mu}\beta}(\tilde{\mu}_d, \tilde{\mu}_\phi, \tilde{\mu}_\xi, \tilde{\mu}_{\Sigma^{-1}}), \tilde{\Sigma}\beta = f_{\tilde{\Sigma}\beta}(\tilde{\mu}_d, \tilde{\mu}_\phi, \tilde{\mu}_\xi, \tilde{\mu}_{\Sigma^{-1}}), \tilde{\mu}_d = f_{\tilde{\mu}_d}(\tilde{\mu}_\beta, \tilde{\mu}_\phi, \tilde{\mu}_\xi, \tilde{\mu}_{\Sigma^{-1}}), \tilde{\Sigma}_d = f_{\tilde{\Sigma}_d}(\tilde{\mu}_\beta, \tilde{\mu}_\phi, \tilde{\mu}_\xi, \tilde{\mu}_{\Sigma^{-1}}), \tilde{\mu}_\phi = f_{\tilde{\mu}_\phi}(\tilde{\mu}_\xi, \tilde{\Sigma}\xi, \tilde{\mu}_d, \tilde{\mu}_\beta, \tilde{\mu}_{\Sigma^{-1}}), \tilde{\Sigma}_\phi = f_{\tilde{\Sigma}_\phi}(\tilde{\mu}_d, \tilde{\mu}_\beta, \tilde{\mu}_\xi, \tilde{\Sigma}\xi, \tilde{\mu}_{\Sigma^{-1}}), \tilde{\pi}_{\xi,pq} = f_{\tilde{\pi}_{\xi,pq}}(\tilde{\mu}_\phi, \tilde{\Sigma}\phi, \tilde{\mu}_{\xi-pq}, \tilde{\Sigma}\xi_{-pq}, \tilde{\mu}_d, \tilde{\Sigma}_d, \tilde{\mu}_\beta, \tilde{\Sigma}\beta, \tilde{\mu}_{\Sigma^{-1}}), \tilde{\mu}_{\Sigma^{-1}} = f_{\tilde{\mu}_{\Sigma^{-1}}}(\tilde{\mu}_\beta, \tilde{\mu}_d, \tilde{\mu}_\phi, \tilde{\Sigma}\phi, \tilde{\mu}_\xi, \tilde{\Sigma}\xi_{pq})$ and $\tilde{\nu}_\Sigma = \nu_\Sigma + (T - L)R$. Note that, to simplify these calculations we approximate $E_Q(\theta\theta')$ by $E_Q(\theta)E_Q(\theta')$ when computing the moments of the approximate distributions of $\theta = \beta$, and $\theta = d$. However, we do not use such approximation when computing the moments of the approximate distributions of $\theta = \phi$, $\theta = \xi_{pq}$ and $\theta = \Sigma$, in order to get more accurate connectivity results. The VB algorithm works by first specifying the initial values for $\tilde{\mu}_\beta, \tilde{\Sigma}\beta, \tilde{\mu}_d, \tilde{\Sigma}_d, \tilde{\mu}_\phi, \tilde{\Sigma}\phi, \tilde{\pi}_{\xi}$ and $\tilde{\mu}_{\Sigma^{-1}}$, and then updating their values according to the updating functions repeatedly until convergence. Details of the VB algorithm are provided in the supplementary material. The joint posterior distribution can then be approximated by sampling $\beta \sim N(\tilde{\mu}_\beta, \tilde{\Sigma}\beta)$, $d \sim N(\tilde{\mu}_d, \tilde{\Sigma}_d)$, $\phi \sim N(\tilde{\mu}_\phi, \tilde{\Sigma}\phi)$, $\xi_{pq} \sim Multinomial(1, \tilde{\pi}_{\xi,pq})$, and $\Sigma \sim InvWishart(\tilde{\Psi}_\Sigma, \tilde{\nu}_\Sigma)$.

The VB algorithm described above can be easily extended to consider the full multi-subject model. In such case, the approximation becomes

$$Q(\Theta) = \prod_{s=1}^S Q(\Theta^{(s)}) \cdot \prod_{g=1}^G Q(\mu_\beta^{(g)}) Q(\mu_\phi^{(g)}) \left(\prod_{p,q} Q(\pi_{\xi,pq}^{(g)}) \right) Q(\mu_d^{(g)}) Q(\sum_\beta^{(g)}) Q(\sum_\phi^{(g)}) Q(\sum_d^{(g)}) Q(\sum_d^{(g)}),$$

where $Q(\Theta^{(s)})$ has the same form as that used for the single subject model and $\Theta^{(s)}$ is the set of subject-level parameters for subject s , i.e.,

$\Theta^{(s)} = (\beta^{(s)}, d^{(s)}, \phi^{(s)}, \xi_{11}^{(s)}, \xi_{12}^{(s)}, \dots, \xi_{pp}^{(s)}, \sum^{(s)})$ The calculations for $Q(\Theta^{(s)})$ are similar to those described above, except that the previously fixed group parameters are replaced with their expectations under $Q(\cdot)$. The calculations of $Q(\cdot)$ for the group parameters also use the conjugate structure in the model and so they are standard. In particular, we have that

$\mu_\beta^{(g)}, \mu_\phi^{(g)}$, and $\mu_d^{(g)}$ follow Gaussian distributions, $\pi_{\xi,pq}^{(g)}$ follow Dirichlet distributions,

$\sum_\beta^{(g)}(i, i)$, $\sum_\phi^{(g)}(i, i)$, and $\sum_d^{(g)}(i, i)$ follow Inverse-Gamma distributions, and $\Sigma^{(g)}$ follow

Inverse-Wishart distributions. Again, note that in the stroke study $\mu_d^{(g)}$ and $\sum_d^{(g)}$ are fixed as constants, and therefore we do not need the terms $Q(\mu_d^{(g)})$ and $Q(\sum_d^{(g)})$ in the equation above.

4. SIMULATION STUDY

In this section we show the performance of our proposed models in two simulation studies. Model performance based on MCMC posterior inference is assessed in a first simulation study. We then compare the MCMC and VB algorithms in a second simulation study. More specifically, Section 4.1 describes the simulation settings used to generate the datasets in the first simulation study; Section 4.2 discusses the prior distributions and provides some details about the MCMC sampler used to obtain posterior inference in this setting; Section 4.3 summarizes the results and discusses the performance of the proposed model; finally,

Section 4.4 compares the performance of the MCMC and VB algorithms in a second simulation study.

4.1 Simulation Settings

We simulated 30 datasets of a single group of 30 subjects, 2 conditions, 5 ROIs, 3 independent sessions with 48 time points for each session, and $TR = 2$ seconds. Each dataset was independently generated according to the model described in the previous section with $L = 1$. The BOLD amplitudes $\beta^{(s)}$, the connectivity parameters $\phi^{(s)}$ and the HRF basis coefficients $d^{(s)}$ for each subject were generated independently using multivariate normal distributions. ROI 2 was simulated to be not activated. The BOLD amplitudes for the two conditions in this ROI were the same, and therefore had a low signal to noise ratio due to a flat BOLD response shape. The connectivity network and the connectivity strength implied by the group-level connectivity parameter $\mu_\phi^{(g)}$ are summarized in Figure 5(a). The connectivities that are not present in Figure 5(a) are set to be zero for all subjects, i.e., the presence/absence of each connectivity is the same across all subjects. The mean HRFs for the five ROIs are displayed in Figure 5(b). ROIs 3 and 4 have the same HRF shapes. $\Sigma^{(s)}$ was constrained to be the same across the subjects. The exact values of the simulation parameters are listed in the supplementary material. Selected simulated time series for ROIs 1, 2, and 3 are shown in Figure 6.

4.2 Analysis

Here we discuss the prior distributions for the group parameters used in the analysis of the simulated data. We set $\nu_\Sigma = P+1$, the smallest value that leads to a proper Inverse-Wishart distribution for $\Sigma^{(s)}$. We also set $\mu_d^g = \mu_d^*$, $\sum_d^g = \sum_d^*$, with μ_d^* and \sum_d^* obtained from the constrained linear basis as described in Section 3.3. The remaining priors are as follows: $\mu_\beta^g \stackrel{ind}{\sim} N(\mathbf{0}_{3P}, 750 \cdot I_{3P})$, with the variance chosen to match the scale of the data, i.e., the variance was set to be a little larger than the maximum of all $(\mu_\beta^g)^2$; $\mu_\phi^g \stackrel{ind}{\sim} N(\mathbf{0}_{2LP^2}, I_{2LP^2})$, as we do not expect μ_ϕ^g to be much larger than 1; $\sum_\beta^g(i, i) \stackrel{ind}{\sim} InvGamma(1, 1)$ and $\sum_\phi^g(i, i) \stackrel{ind}{\sim} InvGamma(1, 1)$ chosen so that $\sum_\beta^g(i, i)$ and $\sum_\phi^g(i, i)$ are allowed to have large variability; $\pi_{pq}^g \stackrel{ind}{\sim} Dir(0.1 \cdot \mathbf{1}_{L+1})$, so that π_{pq}^g has a relatively large variability but does not favor in any specific number of lags $j \in \{0, 1, 2, \dots, L\}$. The maximum number of lags allowed in the model was set to $L = 1$.

4.3 Results

Here we summarize the posterior results based on the MCMC algorithm. The MCMC was run for 5,000 iterations after a burn-in period of 3,000 iterations. No convergence problems were detected.

Activation—All the ROIs were correctly identified as activated or not activated for all 30 datasets. The average bias in the group-level activation contrast $\mu_{\beta, p(c)}$, was 0.48; the average relative bias (i.e., the bias divided by the true value) was 2.9% for activated regions,

and specifically $[0.079, 0.041, 0.033, 0.044] \times 100\%$ for ROIs 1, 3, 4, and 5. The coverage of the 95% posterior CI for the contrast was 86.7% on average, and $[1.00, 1.00, 0.80, 0.67, 0.87] \times 100\%$ for ROIs 1–5. While ROI 4 had a relatively low coverage, it also had a relatively small bias compared to other ROIs, which suggests that the variability of the activation parameter at this ROI could be underestimated.

Connectivity—At the subject level, the full connectivity network was correctly inferred for 96.7% individuals across the datasets. Even for connectivity from ROI 5 to ROI 3, which is non-zero for the task condition and zero for the rest condition, the presence of connectivity was also correctly inferred. The overall sensitivity for the presence of connectivity was 99.5%, and the overall specificity was 100%. These results are based on a 0.5 threshold, i.e., we used $\Pr(\xi_{pq}^{(s)}=1|Y) > 0.5$ to determine the presence of connectivity. Other thresholds ranging from 0.05 to 0.95 were also applied, and the results were similar. At the group level, the bias in the connectivity presence probability $\pi_{pq}^g(1)$ was 0.0019, averaged across all $1 \leq p, q \leq P$. Specifically, this number was 0.0080 for connectivities that are present, and 2×10^{-5} for connectivities that are absent. For the connectivities that exist in the simulated network, the average coverage of the 95% posterior CI of the corresponding VAR parameters was 100%, with average bias around 4.69×10^{-5} , about 1.2% of the average of the elements of $\tilde{\mu}_\phi$, the group-level connectivity strength.

HRFs—The coverage of pointwise 95% posterior credible bands for subject and ROI specific HRFs was around 94.5% overall, and $[0.911, 1.000, 0.958, 0.935, 0.922] \times 100\%$ on average for each ROI. Pointwise 95% posterior CIs of average HRFs at ROIs 1, 2, and 5 for a specific dataset are shown in Figure 7. As expected, ROI 2, the non-activated ROI, shows a larger variability in HRF estimates than the other ROIs, due to low signal to noise ratio, however, the mean shape for this HRF is relatively well captured in spite of this.

Overall, our approach is able to correctly infer activation, connectivity and the HRF shapes. It also has relatively small biases and good coverages of the 95% posterior CIs in general, except that the variation of the activation parameters is a slightly underestimated for one of the ROIs, although the HRF for this ROI was appropriately estimated. Moreover, the connectivity results were found to be very robust to the choice of threshold for determining whether the connectivity is present or not at the subject level.

4.4 Comparison Between MCMC and VB

We ran a smaller and simpler simulation study to compare the performance of the Variational Bayes algorithm for approximate posterior inference, with that of the MCMC which is slower but allows us to obtain full posterior inference. More specifically, in this simulation we generated fMRI data only for 15 subjects using the same simulation settings described in Section 4.1, and we ignored the model structure at the group level, i.e., we conducted only single-subject analysis. The prior distributions used in this particular

simulation were $\beta^{(s)} \sim \mathcal{N}(\mathbf{0}, 50^2 I)$, $d^{(s)} \sim \mathcal{N}(\mu_d^*, \sum_d^*)$, $\phi^{(s)} \sim \mathcal{N}(\mathbf{0}, 0.5^2 I)$, $\pi_{pq}^{(s)} = [.5, .5]$, $\mathbf{v}_\Sigma = P + 1$, and $\Sigma^g = I$. We applied both the VB and MCMC algorithms to the simulated fMRI data

for each of the 15 subjects. For MCMC, we used a burn-in period of 5,000 iterations, and obtained a posterior sample of size 5,000. No convergence problems were detected.

The VB algorithm converged within 150 iterations for all subjects. The VB results are based on 5,000 samples from the approximate posterior. The MCMC took 104 seconds (for all the subjects) while the VB algorithm took 2 seconds, with the same hardware used in the previous simulation study, e.g., a Linux platform in a computer with 64 AMD64 processors and 256G memory.

The results in terms of bias, relative bias (“rbias”), 95% CI coverage (“cover”), true positive rate (“tp”) and true negative rate (“tn”) are summarized in Table 1. T_p , T_w and T_u denote time to peak, FWHM, and time to post-stimulus undershoot (in seconds), respectively. Overall, VB provides a reasonable approximation to the joint posterior distribution. VB led to larger biases than MCMC for most of the parameters, except for FWHM, T_w , and T_u which show slightly smaller biases. In terms of the 95% CIs, VB seems to yield intervals with smaller coverage than MCMC. Also, for inferring whether a connectivity is present, or inferring whether a region is activated, VB seems to report more false positives (lower β_{tn} and ϕ_{tn} compared to MCMC).

Overall, we found that VB approximations could greatly reduce the computation time at the cost of lower specificity and smaller coverage of the 95% CIs. In cases where false positives are of less concern and MCMC sampling is too slow for the problem at hand, such as fMRI datasets with a large number of ROIs and many subjects/treatments, the VB approximation provides a feasible and good alternative to the MCMC approach.

5. ANALYSIS FOR THE STROKE STUDY

We now consider the analysis of the real fMRI dataset described in Section 1.1. Full MCMC analysis is feasible in this setting given the dimension of the dataset, so we do not consider approximate VB inference in this case.

5.1 Analysis

We use $g = (\textit{healthy})$, (\textit{stroke}) to denote, respectively, the healthy group and the stroke group. The results below are based on a VAR structure with a maximum lag of $L = 1$. Diagnostic plots were generated to check the lag and other aspects of model fit, which will be discussed in the next subsection. Overall we found that a model with $L = 1$ was appropriate for most of the subjects in this study.

We used the same prior distributions used for the simulation study, with the exception of the prior for the activation strength. This prior was set to $\mu_{\beta}^g \sim N(0, 25 \cdot I)$ so that the standard deviation is about the largest $\mu_{\beta_p}^g$ obtained from the preliminary data analysis based on the method in Woolrich et al. (2004). In addition, in order to check the sensitivity of the results to the choice of prior distributions, the following alternative prior distributions were also considered. Sensitivity with respect to the choice of \mathbf{a}_{π} : we considered two alternative prior settings with $\mathbf{a}_{\pi} = [0.2, 0.1]$ and $\mathbf{a}_{\pi} = [0.1, 0.2]$. The first one assumes that the presence

probability of connectivity is most likely to be around $\frac{2}{3}$, the second assumes that such probability is most likely to be around $\frac{1}{3}$, while the original prior, $\mathbf{a}_\pi = [0.1, 0.1]$, assumes the probability is equally likely to be <0.5 or >0.5 . Sensitivity with respect to the priors on μ_β^g and μ_ϕ^g : we considered two alternative sets of priors, namely (1)

$$\mu_\beta^g \sim N(\mathbf{0}, 25 \times 25 I_{3P}), \mu_\phi^g \sim N(\mathbf{0}, 25 I_{2LP2}), \text{ and (2)}$$

$\mu_\beta^g \sim N(\mathbf{0}, \frac{1}{4} \times 25 I_{3P}), \mu_\phi^g \sim N(\mathbf{0}, \frac{1}{4} I_{2LP2})$. The first set has a larger variability compared to the original set of priors and the second set has a smaller variability. All of these priors led to very similar posterior results.

In addition to studying the posterior distributions of the activation and connectivity parameters within each group, we also obtained posterior inference on the difference in the presence probabilities of connectivity from ROI q to ROI p for stroke and healthy subjects, i.e., $\pi_{pq}^{(stroke)}(\ell) - \pi_{pq}^{(healthy)}(\ell)$, and also on the difference in connectivity strength,

$\mu_{\Phi_k(\ell), pq}^{(stroke)} - \mu_{\Phi_k(\ell), pq}^{(healthy)}$, for $\ell = 1$. The 95% posterior CIs of these quantities were computed. If the 95% posterior CI of the difference of a given quantity does not include 0, then we say that there is sufficient evidence to claim that such quantity is different for the stroke group and the healthy group.

A posterior sample of size 5,000 was obtained after a burn-in period of 3,000 iterations. The analysis was carried out on Linux platform in a computer with 64 AMD64 processors and 256G memory. The total computation time was 50 min.

5.2 Results

Activation—We found that all the ROIs were activated in response to the task condition for the stroke patients, and all the ROIs were activated for the healthy subjects except for RM1 (see Figures 8(a) and (b)). In both groups, LM1, which is primarily responsible for the motor function of the right side of the body, was found to be highly activated. A formal comparison of the activations between the groups is shown in Figure 8(c). Red indicates a positive difference. The plot shows there is sufficient evidence that the stroke group is more activated than the healthy group in RPMd, RM1 and SMA. This indicates a compensatory effect that requires M1 from the unaffected brain hemisphere and secondary motor regions to aid in executing the motor task.

Connectivity—Plots 9(a) and (b) list all the connectivities that are present in at least one subject for each group, and describe how likely is the presence of connectivity in each group in terms of the thickness of the arrow. Connectivity from LM1 to LPMd in healthy participants has the highest presence probability, around 1.00. Connectivity from RPMd to RM1 in stroke patients has the lowest presence probability, around 0.08, apart from the connectivities that are not listed. The solid arrows represent those connectivities that are predominant, i.e., those for which the presence probability exceeds 0.5, or $\pi_{pq}^g(1) | Y > 0.5$, while the dashed arrows correspond to connectivities that exist at least in one subject in the

group (i.e., $\Pr(\xi_{pq}^{(s)}(1)=1|Y) > 0.5$ for at least one s), but are not predominant in the group (i.e., $\pi_{pq}^g(1)|Y \leq 0.5$). All the listed connectivities have positive $\mu_{\Phi_k(1),pq}^g$ values, indicating a positive lagged association, i.e., the larger the current signal at ROI q is, the larger the future signal at ROI p is expected to be. In both groups, connectivity within LM1 is present in the majority of the subjects. This suggests that there may be positive feedback within M1 on the hemisphere corresponding to the moving hand. In the healthy group, the predominant connectivities are mostly interregional and all start from LM1. However, in the stroke group, the predominant connectivities are all intraregional.

Figure 9(c) lists all the differences in connectivity presence probability between the two groups ($\pi_{pq}^{(stroke)}(1) - \pi_{pq}^{(healthy)}(1)$) whose 95% posterior CI does not include 0. The red connectivities are more likely in the stroke group than the healthy group, while the blue connectivities are more likely in the healthy group than the stroke group. For the stroke group, there are larger connectivity values from secondary motor regions (LPMd, SMA) to themselves or to other regions; while for the healthy group, there are larger connectivity values from LM1, the region primarily responsible for the right hand movement, to other regions.

We also calculated the posterior samples of the difference in connectivity strength between the groups, $\mu_{\Phi_k(1),pq}^{(stroke)} - \mu_{\Phi_k(1),pq}^{(healthy)}$. However, all of the 95% posterior CIs of the difference in connectivity strength cover 0; i.e., there is not sufficient evidence for differences in the strength of connectivities across different groups.

HRFs—The estimated HRF (pointwise posterior median) for each region and subject is shown in Figure 10. There are no obvious differences between the two groups, except that in general there seems to be more interindividual variability in the subject-specific HRFs from the stroke group.

We checked the goodness of fit by examining the estimated mean, denoted by $\mathbf{y}^{*(s,r)}(t)$, and the temporally uncorrelated noise, $\mathbf{e}^{(s,r)}(t)$. For each subject s and session r , we obtained posterior samples of $\mathbf{y}^{*(s,r)}(t)$ by calculating $(\mathbf{y}^{*(s,r)}(t))^{(m)} = (\mathbf{X}^{(s)}(t))^{(m)} (\boldsymbol{\beta}^{(s)})^{(m)}$ for each iteration m in the MCMC algorithm; and then calculated a posterior estimate, $\hat{\mathbf{y}}^{*(s,r)}(t)$, by

taking the median of the posterior samples $\left\{ (\mathbf{y}^{*(s,r)}(t))^{(m)} \right\}_{m=1}^M$. Similarly, we obtained posterior samples of $\mathbf{e}^{(s,r)}(t)$ by calculating

$$(\mathbf{e}^{(s,r)}(t))^{(m)} = (\mathbf{u}^{(s,r)}(t))^{(m)} - (\boldsymbol{\Phi}^{*(s)}(1, t))^{(m)} (\mathbf{u}_p^{(s,r)}(t-1))^{(m)}, \text{ where}$$

$$(\boldsymbol{\Phi}^{*(s)}(1, t))^{(m)} = c_1(t-1)(\boldsymbol{\Phi}_1^{*(s)}(1))^{(m)} + c_2(t-1)(\boldsymbol{\Phi}_2^{*(s)}(1))^{(m)}, (\mathbf{u}^{(s)}(t))^{(m)} = \mathbf{y}^{(s,r)}(t) - (\mathbf{y}^{*(s,r)}(t))^{(m)}, \text{ and calculated } \hat{\mathbf{e}}^{(s,r)}(t) \text{ as the median. We inspected the time series plot of } \hat{\mathbf{e}}^{(s,r)}(t) \text{ for each ROI, and did not find any obvious temporal patterns in most of these time series plots. In addition, PACF plots also indicated that } L = 1 \text{ was sufficient for most subjects, sessions and regions, but some of these plots suggest that a model with } L = 2 \text{ may be required for a small number of subjects. Finally, plots of } \hat{\mathbf{y}}^{*(s,r)}(t) \text{ also show that the group model captures the mean trend of the data, especially for the ROIs that have relatively}$$

strong activations (e.g., LM1). Figure A.3 in the Appendix shows examples of time series plots of $\hat{y}^{(s,r)}(t)$ and $\hat{\varepsilon}^{(s,r)}(t)$, and the PACF plots of $\hat{\varepsilon}^{(s,r)}(t)$, for LM1 and two subjects, one healthy and the other one a stroke patient.

5.3 Discussion of the Analysis

Using our model, we obtained interesting results about the relative activation patterns in stroke patients compared to control subjects. First, as expected, we found that there was strong activation within LM1 for both the stroke and control groups. This is consistent with the fact that LM1 is primarily responsible for motor function of the right hand. All the ROIs were activated in response to the task condition (see Figure 8(a)) for the stroke patients; and all the ROIs were activated for healthy subjects except for RM1 (see Figure 8(b)), which is primarily responsible for the motor function of the left side of the body.

Activation of secondary motor areas in healthy controls initially seemed counterintuitive. However, the control subjects were age matched to the stroke group and, therefore, were of 50+ years of age. It is established that these regions are involved in the control of unilateral movement (Beaulé et al., 2012) and with aging we see increasing involvement of secondary brain regions to support simple motor movements (Ward, 2003).

Second, we found greater activation strength in the 3 motor regions, namely, SMA, RM1 and RPMd for stroke patients compared to healthy controls. The literature confirms this finding. After stroke, such secondary motor regions are routinely more activated to support execution of post-stroke movements, particularly in more impaired patients (Ward et al., 2003).

The connectivity modeling produced some unexpected results. After stroke, fMRI connectivity studies show that there are excitatory and inhibitory connections between primary motor cortex in the stroke-affected hemisphere and secondary motor regions within the affected and unaffected hemispheres (Rehme and Grefkes, 2013). Therefore, Figure 9(b), showing a lack of connectivity between LM1 and other regions like RM1 and RPMd in stroke patients is surprising.

Stroke is a highly heterogeneous disorder and perhaps in the stroke group there was sufficient noise and variation across the patients leaving the current statistical approach unable to detect predominant between-region connections. To our knowledge, there is a dearth of literature describing how activity within a region predicts subsequent activation in the region. A recent study from our group did identify that this intraregional activation prediction is present for LM1, LPMd, SMA, and RPMd (Gorrostieta et al., 2013), largely confirming what we observed in the current analyses. Together the data suggest possible positive feedforward connectivity in these regions. Furthermore, the greater likelihood of the connectivity being present in stroke patients versus healthy controls may lend support for these areas being more important to guiding and coordinating movement after stroke.

In healthy control subjects, our model identified probable connections from LM1 to the secondary motor regions. Although the exact neurobiological underpinnings of the connectivity cannot be determined from these analyses, considering the older age of the control group, the predominant connection from LM1 to LPMd is perhaps suggestive of

recruitment of a secondary motor area in the hemisphere responsible for movement of the right hand.

The activation and connectivity patterns specific to stroke patients that have been found in this analysis may be linked to compensatory effects that are related to post-stroke recovery. These findings may help neuroscientists in the design of future studies that can lead to a deeper understanding of the association between neurophysiology (activation and connectivity) and the degree of post-stroke recovery and therefore to a proper evaluation of the effect of a particular treatment.

6. CONCLUSIONS

Our approach presents several advantages with respect to currently available methods: (1) it is able to simultaneously infer activation, connectivity and HRFs; (2) it is able to provide ROI-specific and subject-specific HRFs, as well as condition-specific connectivity; (3) it is able to borrow information across subjects via hierarchical modeling to increase power for group comparison; this is particularly relevant for data with low signal to noise ratio; (4) it is able to easily incorporate relevant information derived from other studies but, unlike the dynamic causal modeling approach, it does not rely on biological assumptions that are difficult to verify from the data.

As a note of caution, the inference on connectivity patterns obtained using our approach is at the hemodynamic level instead of the neuronal level, i.e., connectivity from ROI q to ROI p here means there is an association between the current hemodynamic activity in ROI q to the future hemodynamic activity in ROI p . This type of connectivity is more difficult to interpret compared to neuronal-level connectivity mainly due to two limitations. One limitation is that the information transmission happening between the neurons is much faster than the fMRI sampling frequency, and spurious connectivity may be captured by VAR models due to downsampling. This type of behavior is illustrated in Fig. 7 in Valdes-Sosa et al. (2011). of further studies for measuring the degree One solution is to use vector autoregressive moving average models (VARMA) which are more robust to this issue, however, they are more computationally challenging than VAR models. The second limitation is that the sparsity of the neuronal-level connectivity may be distorted when transferred to the hemodynamic level through convolution (Valdes-Sosa et al., 2011; Smith et al., 2012). According to Penny et al. (2005), a solution to this limitation is to model the neuronal dynamics through a state space model. In a state space model, the fMRI signals are treated as the observed variables and the neuronal dynamics as the hidden states. Neuronal-level connectivity is thus modeled via the hidden states. Approximate posterior inference for this type of model can also be obtained efficiently by variational Bayes algorithms (Ryali et al., 2011; Luessi et al., 2014). However, we are concerned that the hidden states in the state space model are modeled at a much lower resolution than the actual neuronal activity, and thus do not accurately represent the underlying neuronal process. This leads us to suspect that such connectivity might still be difficult to interpret. Despite the limitations, our model seemed to appropriately capture the differential patterns between the stroke patients and the healthy subjects in this stroke study; we speculate that this is partially due to the fact that we modeled the region-specific HRFs.

Besides the limitations and possible solutions mentioned above, in future work we will make improvements to our statistical approach in the following ways. In the current approach we assume independence across groups and analyze the groups separately; in the future we will consider incorporating grouping factors into the model. This may further help increase the power of detecting the differences between groups since we can borrow information across more subjects. The current approach also assumes that the activation strength is the same across sessions; in future studies we will extend our model to take into account the variation in the activation strength among different sessions.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Beaulé V, Tremblay S, Théoret H. Interhemispheric Control of Unilateral Movement. *Neural Plasticity*. 2012; 2012:11. Article ID 627816.
- Bowman F, Caffo B, Bassett S, Kilts C. A Bayesian Hierarchical Frame-work for Spatial Modeling of fMRI Data. *NeuroImage*. 2008; 39:146–156. [PubMed: 17936016]
- Cribben I, Haraldsdottir R, Atlas L, Wager T, Lindquist M. Dynamic Connectivity Regression: Determining State-Related Changes in Brain Connectivity. *NeuroImage*. 2012; 61:907–920. [PubMed: 22484408]
- David O, Guillemain I, Sallet S, Reyt S, Deransart C, Segebarth C, Depaulis A. Identifying Neural Drivers with Functional MRI: An Electrophysiological Validation. *PLOS Biology*. 2008; 6:e315.
- Degras D, Lindquist M. A Hierarchical Model for Simultaneous Detection and Estimation in Multi-Subject fMRI Studies. *NeuroImage*. 2014; 98:61–72. [PubMed: 24793829]
- Eddelbuettel D, Sanderson C. RcppArmadillo: Accelerating R with high-performance C++ linear algebra. *Computational Statistics and Data Analysis*. 2014; 71:1054–1063.
- Friston K, Harrison L, Penny W. Dynamic Causal Modelling. *NeuroImage*. 2003; 19:1273–1302. [PubMed: 12948688]
- Friston K, Jezzard P, Turner R. Analysis of Functional MRI Time-Series. *Human Brain Mapping*. 1994; 1:153–171.
- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., Rubin, D. *Bayesian Data Analysis*. 3. Boca Raton, FL: Chapman and Hall/CRC; 2013.
- George EI, McCulloch RE. Variable Selection via Gibbs Sampling. *Journal of the American Statistical Association*. 1993; 88:881–889.
- Gorrostieta C, Fiecas M, Ombao H, Burke E, Cramer S. Hierarchical Vector Auto-Regressive Models and Their Applications to Multi-Subject Effective Connectivity. *Front Comput Neurosci*. 2013; 7
- Guo Y, Pagnoni G. A Unified Framework for Group Independent Component Analysis for Multi-subject fMRI Data. *NeuroImage*. 2008; 42:1078–1093. [PubMed: 18650105]
- Harrison L, Penny WD, Friston K. Multivariate Autoregressive Modeling of fMRI Time Series. *NeuroImage*. 2003; 19:1477–1491. [PubMed: 12948704]
- Ishwaran H, Rao S. Spike and Slab Variable Selection: Frequentist and Bayesian Strategies. *The Annals of Statistics*. 2005; 33:730–773.
- Jordan MI, Ghahramani Z, Jaakkola TS, Saul LK. An Introduction to Variational Methods for Graphical Models. *Machine Learning*. 1999; 37(2):183–233. <http://link.springer.com/article/10.1023/A%3A1007665907178>.
- Kang H, Ombao H, Linkletter C, Long N, Badre D. Spatio-Spectral Mixed-Effects Model for Functional Magnetic Resonance Imaging Data. *Journal of the American Statistical Association*. 2012; 107:568–577. [PubMed: 25400305]

- Koop, G., Korobilis, D. Manual to Accompany MATLAB Package for Bayesian VAR Models. Glasgow; United Kingdom: 2009. <http://personal.strath.ac.uk/gary.koop/KoKoManual.pdf>
- Lindquist M. The Statistical Analysis of fMRI Data. *Statistical Science*. 2008; 23:439–464.
- Lindquist M, Loh J, Atlas L, Wager T. Modeling the Hemodynamic Response Function in fMRI: Efficiency, Bias and Mis-Modeling. *NeuroImage*. 2008; 45:187–98.
- Luessi M, Babacan SD, Molina R, Booth JR, Katsaggelos AK. Variational Bayesian causal connectivity analysis for fMRI. *Frontiers in Neuroinformatics*. 2014; 8:45. [PubMed: 24847244]
- Luo, X. A Hierarchical Graphical Model for Big Inverse Covariance Estimation with an Application to fMRI. Brown University, Dept. of Biostatistics; 2014. arXiv preprint <http://arxiv.org/pdf/1403.4698v2.pdf>
- Mitchell TJ, Beauchamp JJ. Bayesian Variable Selection in Linear Regression. *Journal of the American Statistical Association*. 1988; 83:1023–1032.
- Ollinger J, Shulman G, Corbetta M. Separating Processes Within a Trial in Event-Related Functional MRI. *NeuroImage*. 2001; 13:210–217. [PubMed: 11133323]
- Penny W, Ghahramani Z, Friston K. Bilinear dynamical systems. *Philosophical transactions of the Royal Society of London Series B, Biological Sciences*. 2005; 360:983–993. [PubMed: 16087442]
- Rehme AK, Grefkes C. Cerebral Network Disorders After Stroke: Evidence From imaging-Based Connectivity Analyses of Active and Resting Brain States in Humans. *The Journal of Physiology*. 2013; 591:17–31. [PubMed: 23090951]
- Ryali S, Supekar K, Chen T, Menon V. Multivariate dynamical systems models for estimating causal interactions in fMRI. *NeuroImage*. 2011; 54:807–823. [PubMed: 20884354]
- Smith M, Fahrmeir L. Spatial Bayesian Variable Selection With Application to Functional Magnetic Resonance Imaging. *Journal of the American Statistical Association*. 2007; 102:417–431.
- Smith SM, Bandettini PA, Miller KL, Behrens TEJ, Friston KJ, David O, Liu T, Woolrich MW, Nichols TE. The danger of systematic bias in group-level FMRI-lag-based causality estimation. *NeuroImage*. 2012; 59(2):1228–9. [PubMed: 21867760]
- Valdes-Sosa PA, Roebroeck A, Daunizeau J, Friston K. Effective connectivity: Influence, causality and biophysical modeling. *NeuroImage*. 2011; 58:339–361. [PubMed: 21477655]
- Ward N. Age-Related Changes in the Neural Correlates of Motor Performance. *Brain*. 2003; 126:873–888. [PubMed: 12615645]
- Ward N, Brown M, Thompson A, Frackowiak R. Neural Correlates of Outcome After Stroke: A Cross-Sectional fMRI Study. *Brain*. 2003; 126:1430–1448. [PubMed: 12764063]
- Wellcome Trust Center for Neuroimaging, UCL. SPM8 Software; 2009. www.fil.ion.ucl.ac.uk/spm
- Woolrich M, Behrens T, Smith S. Constrained Linear Basis Sets For HRF Modelling Using Variational Bayes. *NeuroImage*. 2004; 21:1748–1761. [PubMed: 15050595]
- Worsley K, Friston K. Analysis of fMRI Time-Series Revisited-Again. *NeuroImage*. 1995; 2:173–181. [PubMed: 9343600]
- Xu L, Johnson T, Nichols T, Nee D. Modeling Inter-Subject Variability in fMRI Activation Location: A Bayesian Hierarchical Spatial Model. *Biometrics*. 2009; 65:1041–1051. [PubMed: 19210732]
- Yu, Z., Ombao, H., Prado, R., Quinlan, EB., Cramer, SC. A Bayesian Model for Activation and Connectivity in Task-related fMRI Data. *Bayesian Inference in Science and Engineering*; 2015.
- Zarahn E. Using Larger Dimensional Signal Subspaces to Increase Sensitivity in fMRI Time Series Analyses. *Human Brain Mapping*. 2002; 17:13–16. [PubMed: 12203684]
- Zhang L, Guindani M, Vannucci M. Bayesian models for functional magnetic resonance imaging data analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*. 2015; 7:21–41. <http://dx.doi.org/10.1002/wics.1339>. [PubMed: 25750690]
- Zhang T, Li F, Beckes L, Coan J. A Semi-Parametric Model of the Hemodynamic Response for Multi-Subject fMRI Data. *NeuroImage*. 2001; 75:136–145.

APPENDIX

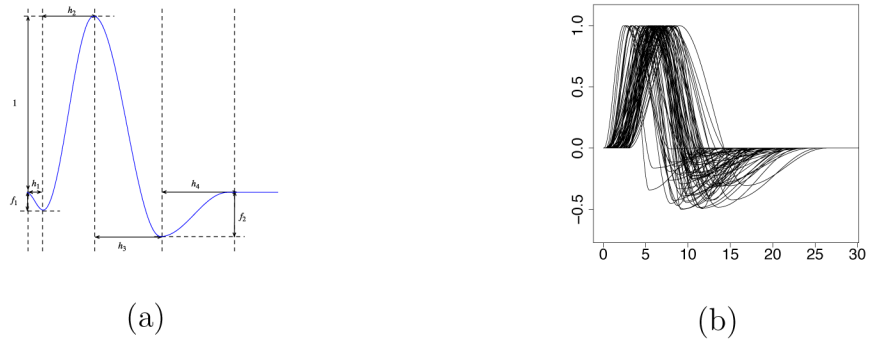


Figure A.1. (a) Half-cosine parameterization of the HRF. (b) Simulated HRFs.

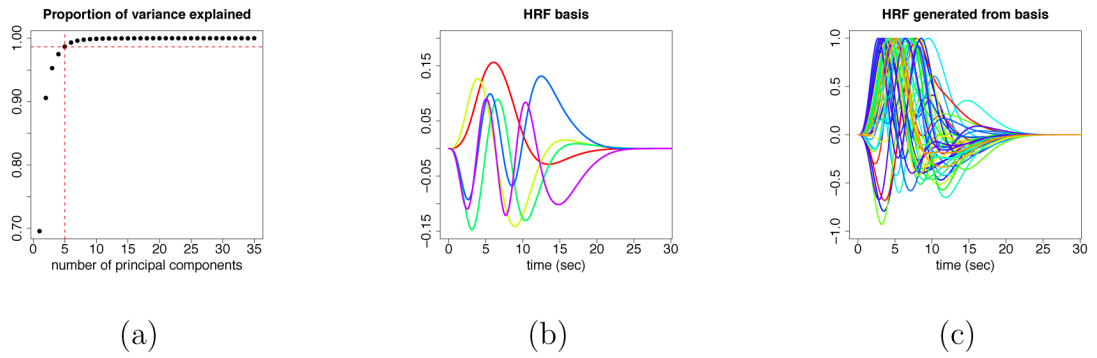


Figure A.2. (a) Variability explained by the basis vectors; (b) HRF basis set; (c) HRFs generated by the constrained linear basis.

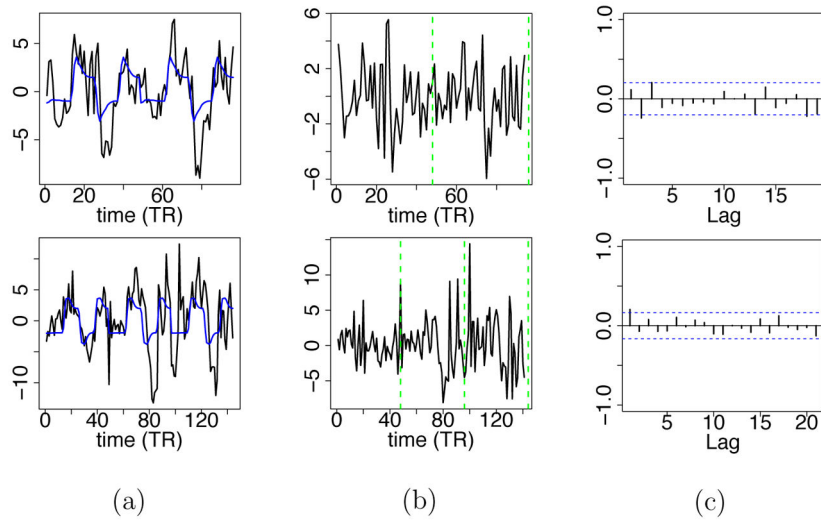


Figure A.3.

(a) $y_p^{(s,r)}$ (black) and $\hat{y}_p^{*(s,r)}$ (blue). (b) $\hat{\varepsilon}_p^{(s,r)}(t)$. (c) PACFs of $\hat{\varepsilon}_p^{(s,r)}(t)$. The first row is for LM1 of a healthy subject, and the second row is for LM1 of a stroke patient.

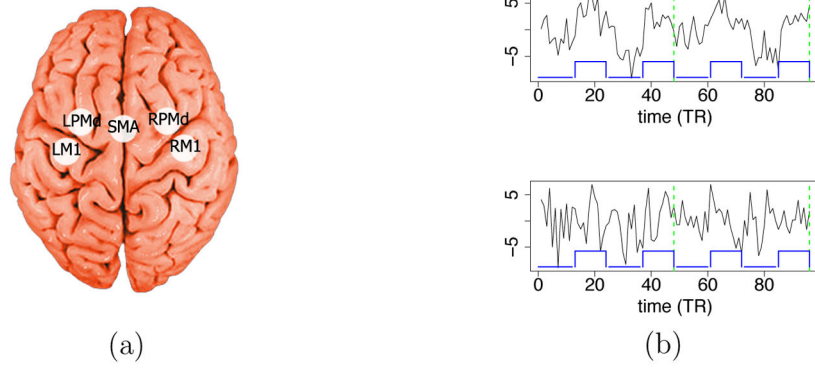


Figure 1.

(a) Locations of the five regions of interest: LM1 and RM1 (left and right primary motor cortex), LPMd and RPMd (left and right dorsal premotor cortex), and SMA (supplementary motor area). (b) FMRI time series from a healthy subject at LM1 (top) and RM1 (bottom); the lines at the bottom of each plot indicate the task condition (high value) and the rest condition (low value). The time series of the two independent sessions (each with 48 scans) are concatenated on the plot. The dashed lines separate the two sessions.

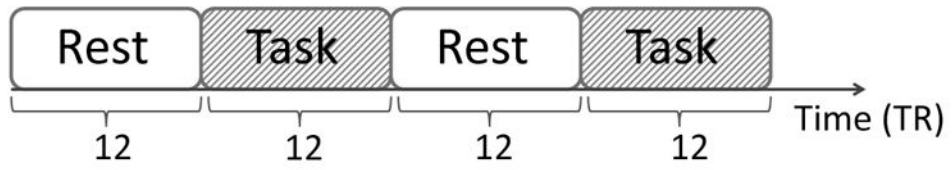


Figure 2.
Block design of the experiment for each session.

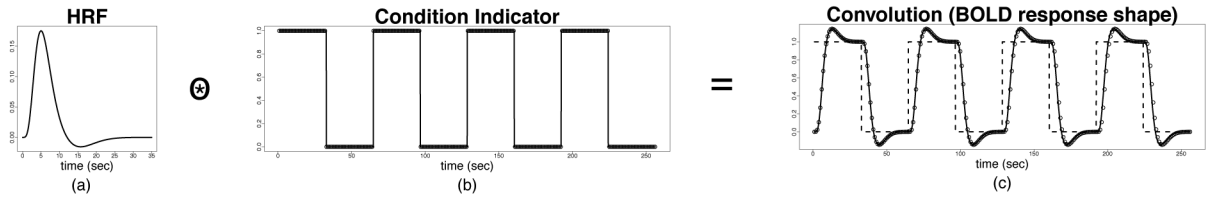


Figure 3.
(a) Canonical HRF. (b) Condition indicator function. (c) BOLD response.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

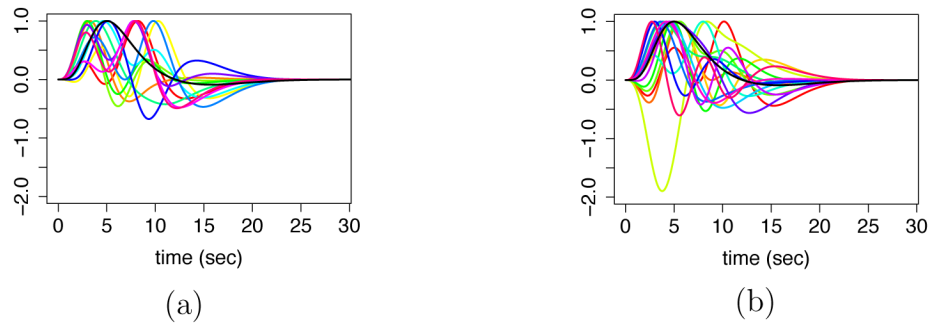


Figure 4. A preliminary analysis of HRFs. (a) HRFs of all healthy subjects at a given ROI. (b) HRFs of all stroke subjects at the same ROI. The black curve is canonical HRF.

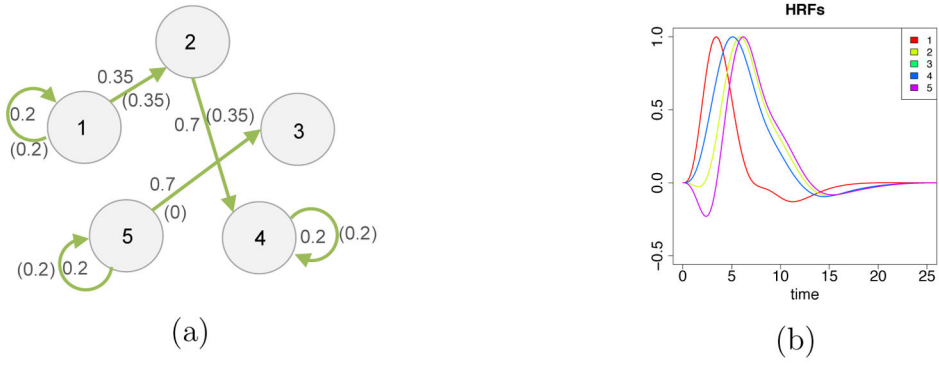


Figure 5. (a) Simulated connectivity network; connectivities not shown in the graph are not present in any of the subjects. The numbers without parentheses represent the group-level connectivity strength under task condition; the numbers in parentheses represent the group-level connectivity strength under rest condition. (b) Simulated mean HRFs.

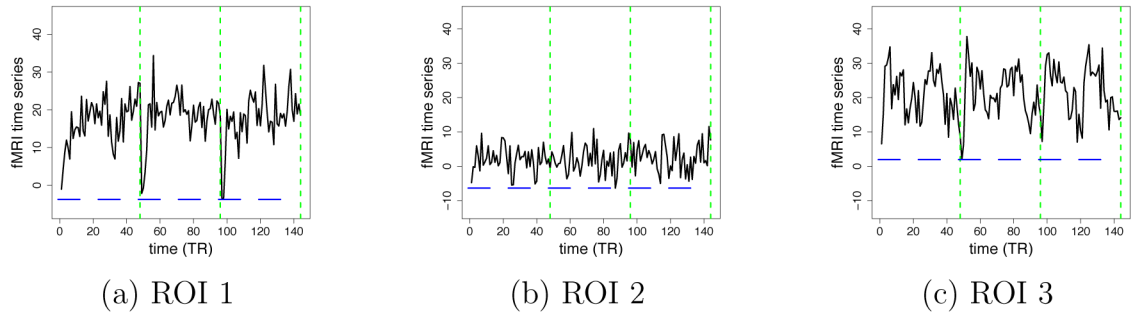


Figure 6. Examples of simulated fMRI time series at different ROIs from a single subject. Green lines separate different sessions. Blue lines indicate when the task condition is on.

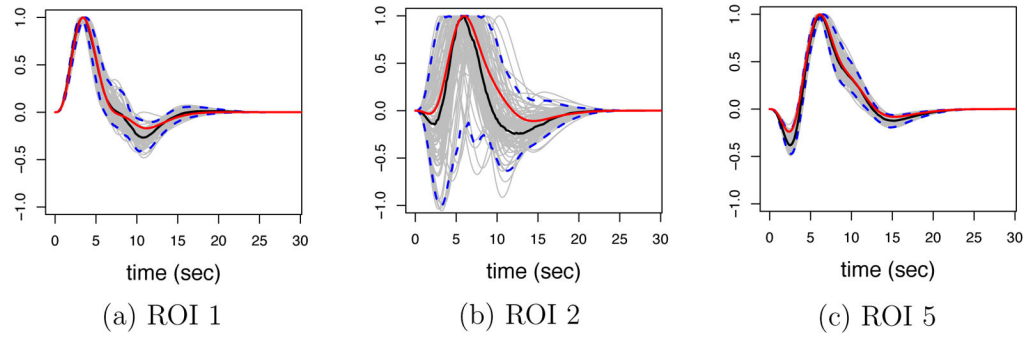


Figure 7.

Posterior HRF results for ROIs 1, 2, and 5 in a given subject. Gray curves correspond to posterior samples of the HRFs, blue curves are the bounds of the 95% pointwise posterior credible bands, the red curves are the true HRFs, the black curves are the posterior median HRFs.

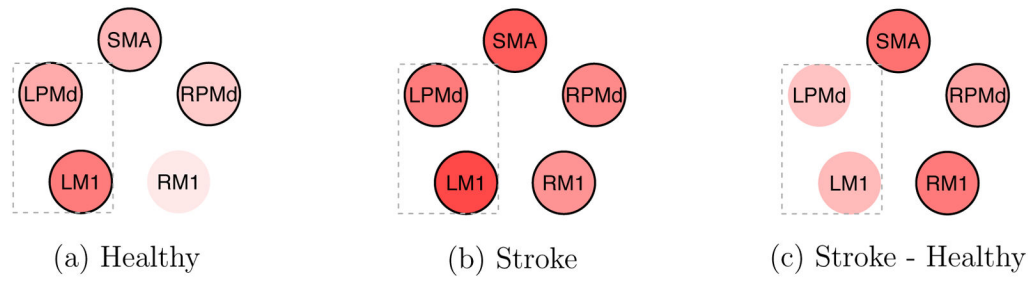


Figure 8.

Activation strength for the two groups and activation difference between the groups. Darker color indicates higher activation (or activation difference). Red indicates a positive value, and white represents zero. The black line outside the circle indicates that the corresponding 95% posterior CI does not include zero. The rectangular frame indicates the ROIs in the hemisphere corresponding to the moving hand.

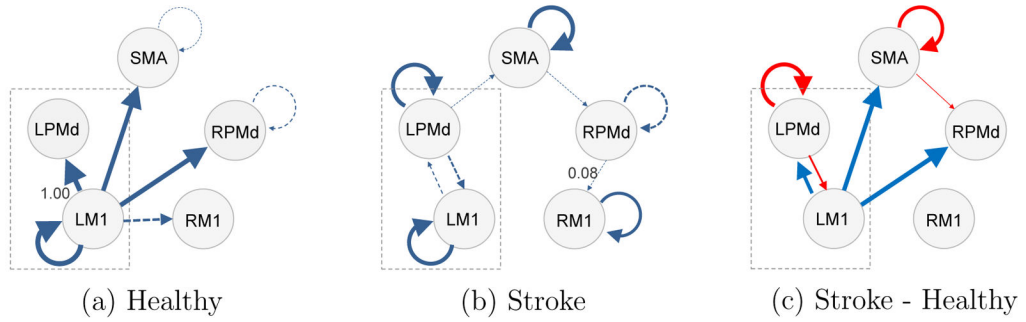


Figure 9.

The probabilities of the presence of connectivity for the two groups and the probability differences between the groups. A thicker arrow indicates a higher chance that connectivity is present in the group (plots (a) and (b)), or greater difference in the probability of connectivity between the two groups (plot (c)). In (a) and (b), those connectivities that are absent in all subjects in each group are not shown; a dashed arrow represents a connectivity whose presence probability does not exceed 0.5; the probabilities of the most likely and the least likely connectivities among the listed connectivities are also provided. In (c), red indicates that the connectivity is more likely to be present in the stroke group than the healthy group, and blue indicates that the connectivity is less likely to be present in the stroke group; only those differences whose 95% posterior CI does not contain 0 are included.

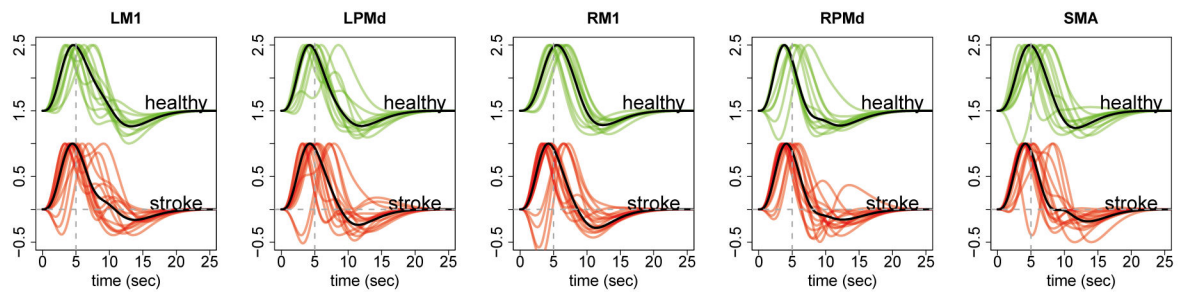


Figure 10.

Estimated HRF (pointwise posterior median) for each ROI and subject for the two groups.

The vertical dashed line is at 5 sec and the horizontal dashed line is at the baseline value.

The black curve is the median HRF of the corresponding group.

Comparison of results of VB and MCMC.

Table 1

	β_{bias}	β_{rbias}	β_{tp}	β_{tn}	$\beta_{\text{cover}\Sigma}$	$^{-1}\text{bias}\Sigma$	$^{-1}\text{rbias}\Sigma$	$^{-1}\text{cover}$	hrf.cover	$T_p\text{bias}$	$T_p\text{cover}$	$T_w\text{bias}$	$T_w\text{cover}$	$T_u\text{bias}$	$T_u\text{cover}$	ϕ_{bias}	ϕ_{rbias}	ϕ_{tp}	ϕ_{tn}	ϕ_{cover}
VB	-0.011	-0.0344	0.856	0.828	0.972	0.106	0.884	0.876	0.876	-0.156	0.907	-0.239	0.893	-1.095	0.893	0.191	0.0156	0.983	0.867	0.773
MCMC	-0.006	-0.0183	0.889	0.898	0.961	0.062	0.929	0.967	0.967	-0.053	0.987	-0.403	0.947	-1.165	0.973	-0.010	-0.0012	0.983	1.000	0.960