Modeling the spatial and temporal dependence in fMRI data

Gordana Derado[1,*], **F. DuBois Bowman**[1], and **Clinton D. Kilts**[2]

[1] Department of Biostatistics and Bioinformatics, The Rollins School of Public Health, Emory University 1518 Clifton Road, Atlanta, GA 30322, U.S.A

[2] Department of Psychiatry and Behavioral Sciences, Emory University School of Medicine, 1639 Pierce Drive, Atlanta, GA 30322, U.S.A

Summary

Functional magnetic resonance imaging (fMRI) data sets are large and characterized by complex dependence structures driven by highly sophisticated neurophysiology and aspects of the experimental designs. Typical analyses investigating task-related changes in measured brain activity use a two-stage procedure in which the first stage involves subject-specific models and the second-stage specifies group (or population) level parameters. Customarily, the first-level accounts for temporal correlations between the serial scans acquired during one scanning session. Despite accounting for these correlations, fMRI studies often include multiple sessions and temporal dependencies may persist between the corresponding estimates of mean neural activity. Further, spatial correlations between brain activity measurements in different locations are often unaccounted for in statistical modeling and estimation. We propose a two-stage, spatio-temporal, autoregressive model which simultaneously accounts for spatial dependencies between voxels within the same anatomical region and for temporal dependencies between a subject's estimates from multiple sessions. We develop an algorithm that leverages the special structure of our covariance model, enabling relatively fast and efficient estimation. Using our proposed method, we analyze fMRI data from a study of inhibitory control in cocaine addicts.

Keywords

Cocaine-addiction; covariance modeling; neuroimaging; simultaneous autoregressive model; spatio-temporal model

1. Introduction

Functional magnetic resonance imaging (fMRI) studies yield large data sets that contain temporal correlations from repeated scanning (within and between scanning sessions) and complex spatial correlations. Ignoring correlations affects the precision of estimates of model parameters and consequently may lead to inaccurate statistical tests. Spatio-temporal modeling may mitigate these shortcomings by incorporating more physiologically plausible assumptions and by borrowing strength across related measures of neural activity. One session of a typical fMRI neuroactivation study acquires 3-D scans every 2–3 seconds while the subject performs different experimental tasks. fMRI studies may also involve multiple sessions (e.g corresponding to pre- and post-treatment periods). Typical analyses

investigating task-related changes in measured brain activity use a two-stage procedure in which the first stage involves subject-specific and voxel-specific models relating neural processing to experimental tasks. The second stage specifies voxel-specific models for group (or population) level parameters. The stage I analysis accounts for temporal correlations between serial scans within one session using variants of autoregressive models (Bullmore et al., 1996; Purdon et al., 2001; Friston et al., 2002). Despite addressing scan-to-scan correlations within a session, the stage I estimates for multiple sessions may also exhibit correlations, typically not accounted for at the second stage. Our focus is on modeling these latter correlations.

Some attempts have been made at modeling the repeated measures correlations and spatial correlations in fMRI data, though typically not in the same statistical model. Worsley et al. (2002) introduce a random effects analysis for combining sessions and the method removes trends in effects over time. Bowman and Kilts (2003) fit repeated measures covariance structures within a linear model to address between-session correlations in positron emission tomography (PET) neuroimaging data. Neither of these two approaches addresses the spatial correlations between localized brain activity measurements, and the method by Bowman and Kilts (2003) involves long computations for PET data, making it less suitable for higher dimensional fMRI data. Bowman (2005) presents a two-stage model in which the second stage accounts for spatial dependencies within brain regions or networks, and Bowman et al. (2008) give a more flexible Bayesian model to capture correlations both within and between brain regions. Neither of these approaches, however, accounts for temporal or repeated measures associations between the multiple experimental effects for a subject. There have been a few attempts to model both temporal and spatial correlations simultaneously, but these approaches are generally hampered by long computations. The computations generally involve inverting large covariance matrices, and separable models (over space and time) are often considered (Benali et al., 1997; Hartvig, 2002). Gössl et al. (2001) use a Bayesian approach to fit a semi-parametric spatial and temporal model for fMRI data. Their approach requires extensive computations, limiting its applicability in practice. Woolrich et al. (2004) propose a Bayesian framework to model the noise via a non-separable space-time, simultaneous autoregressive model. This approach is very time consuming for the fMRI datasets, taking roughly 6 hours for processing a single slice of 3-D fMRI data.

We propose a two-stage model that accounts for both spatial and temporal correlations in fMRI data, and our model leads to fast parameter estimation. In the second stage, we construct a simultaneous autoregressive model to capture spatio-temporal correlations between the multiple (session) effects at a given location and between pairs of voxels within defined anatomical regions. We use maximum likelihood (ML) methods to estimate parameters from our spatio-temporal model. We overcome computational challenges involved with estimation by deriving an algorithm that simplifies the calculations of inverses and determinants of large matrices, leading to fast estimation of the model parameters. Our spatio-temporal model provides a unified framework for both voxel-level and region-level inferences. Using our proposed model, we analyze fMRI data from a study of inhibitory control in cocaine addicts to evaluate the effects of behavioral therapy on neural processing related to inhibitory control. To further delineate benefits of our spatio-temporal model, we conduct two simulation studies: one to validate the accuracy of our estimation methods and the other to evaluate the relative efficiency of our proposed spatio-temporal model compared to the general linear model (GLM (not to be confused with *generalized linear model*)).

2. Experimental Data

To illustrate the use of our proposed model and to give motivation for its development, we apply it to data from an fMRI study evaluating the impact of cocaine addiction and

treatment-related abstinence on neural responses to motor inhibition tasks. The sample included 12 cocaine addicts enrolled in an intensive outpatient behavioral treatment program and 15 healthy controls, matched by age, sex (all males), race, handedness, education and early life adversity. All the subjects were scanned while performing an inhibitory control task used to model a characteristic hallmark of drug addiction: an impairment in the ability to inhibit behaviors (e.g. drug-seeking behaviors). The cocaine addicts were scanned in two separate sessions, before and after treatment, and similarly control subjects had baseline and follow-up scans.

Experimental conditions

The inhibitory control task, referred to as the *STOP-signal* task, was designed to evaluate the ability to cancel a prepotent motor response (Aron and Poldrack, 2006). In this task, subjects are presented with visual *GO-stimuli*, consisting of uppercase alphabetical letters appearing on the screen for 0.5 seconds with an inter-stimulus interval of 2.3 seconds. They are instructed to respond to this signal by pressing a button as quickly as possible. An auditory STOP-signal (also lasting 0.5 seconds) appears randomly in 16% of the trials. The occurrence of the STOP-signal following a GO-stimulus is an indicator to refrain from pushing the button. Hence, a successful performance requires the inhibition of a prepotent behavior. The aim of our analysis is to identify brain locations where there are inhibitory control-related increases in brain activity following treatment for cocaine addiction.

3. Methods

We represent a single scan from a subject as a $53\times63\times46$ 3-D rectangular lattice consisting of volume elements (voxels) indexed by a mapping of coordinates (x, y, z) to $v = 1, …, V$ ($V > 150, 000$ in our study). From one scanning session, each voxel contains serial measures of localized brain activity called Blood-Oxygen-Level Dependent (BOLD) fMRI responses. The number of intracranial voxels is too large to estimate a global spatial correlation matrix including all voxel pairs. While retaining the ability to address voxel-level inference, we partition the voxels into mutually exclusive neuroanatomical regions known as Brodmann areas (BA's) (Brodmann, 1909), additionally separating BA's in the left and right hemispheres. For broader coverage, we add several subcortical regions to the BA map, obtained from the Automated Anatomical Labeling (AAL) system (Tzourio-Mazoyer et al., 2002). Building our model based on well-established anatomical parcellations, rather than data-driven functional clusters aids interpretability.

3.1 Statistical Model

Following the conventional two-stage modeling approach for fMRI data, we fit a GLM at the first stage for each individual's vector of serial BOLD responses, separately for each voxel. The GLM regresses each voxel's BOLD responses for subject i ($i = 1, …, K$) on within-subject design variables (e.g. session or stimulus indicators) and on covariates that are not of intrinsic interest, such as high-pass filtering variables. The model accounts for short-range scan-to-scan correlations within a session using a first-order autoregressive process with white noise (Purdon et al., 2001; Friston et al., 2002). The regression

coefficient, $B_{igs}^{(p)}(v)$, for voxel v represents a summary measure of an individual's (mean) neural activity associated with session or stimulus p. We add a subscript g to denote the neuroanatomic region to which voxel v belongs ($g = 1, …, G$), with the region consisting of V_g voxels in total.

At the second stage, we propose a spatio-temporal (ST) autoregressive model to capture temporal correlations between the multiple sessions and spatial correlations between pairs of intra-regional voxels. We express the spatio-temporal model as follows:

$$B_{igs}^{(p)}(v) = \mathbf{x}_i'\beta_{gs}^{(p)}(v) + \rho_{gs}\frac{1}{V_g - 1}\sum_{v^* \in \mathcal{N}_v}\left(B_{igs}^{(p)}(v^*) - \mathbf{x}_i'\beta_{gs}^{(p)}(v^*)\right) + \xi_{gs}\frac{1}{q-1}\sum_{p^* \neq p}\left(B_{igs}^{(p^*)}(v) - \mathbf{x}_i'\beta_{gs}^{(p^*)}(v)\right) + e_{igs}(v),$$

(1)

where $\mathbf{x}_i$ is a vector of between-subject design variables, $\beta_{gs}^{(p)}(v)$ are the group-level parameters, $e_{igs}(v) \sim N(0, \varphi_{gs}^2)$ (φ_{gs}^2 is the between-subject variation), ρ_{gs} is a spatial dependence parameter, ξ_{gs} reflects the temporal dependence between sessions, $\mathcal{N}_v$ is the set of all voxels within the same neuroanatomic region as v, and s denotes a subject group (e.g., patients, controls). When $\rho_{gs} = \xi_{gs} = 0$, model (1) resembles a typical stage II GLM from a random effects analysis of brain imaging data, with constant variances within each region.

3.2 Parametric Covariance Model

Our model involves three parameters defining the variance-covariance structure: the spatial (ρ) and temporal (ξ) dependence parameters and the variance parameter (φ^2). In this section, we present details about the assumed covariance/correlation model. Let

$\mathbf{B}_{igs} = \left(\mathbf{B}_{igs}^{(1)'}, \ldots, \mathbf{B}_{igs}^{(q)'}\right)'$, where each $\mathbf{B}_{igs}^{(p)}$ contains individualized parameters for all voxels in

region g, and similarly $\beta_{gs} = \left(\beta_{gs}^{(1)'}, \ldots, \beta_{gs}^{(q)'}\right)'$. Model (1) implies that $\mathbf{B}_{igs}$ has the following multivariate normal distribution $\mathbf{B}_{igs} \sim \text{MVN}(\beta_{gs}, \Omega_{gs})$ (see Appendix), where

$\Omega_{gs} = \varphi_{gs}^2 \Psi_{gs}\Psi_{gs}$, and Ψ_{gs} is a $qV_g \times qV_g$ matrix with

$$\Psi_{gs}^{-1} = \begin{bmatrix} \mathbf{I}_{V_g} - \rho_{gs}\mathbf{W}_g & -\frac{\xi_{gs}}{q-1}\mathbf{I}_{V_g} & \cdots & -\frac{\xi_{gs}}{q-1}\mathbf{I}_{V_g} \\ -\frac{\xi_{gs}}{q-1}\mathbf{I}_{V_g} & \mathbf{I}_{V_g} - \rho_{gs}\mathbf{W}_g & \cdots & -\frac{\xi_{gs}}{q-1}\mathbf{I}_{V_g} \\ \vdots & \vdots & \ddots & \vdots \\ -\frac{\xi_{gs}}{q-1}\mathbf{I}_{V_g} & -\frac{\xi_{gs}}{q-1}\mathbf{I}_{V_g} & \cdots & \mathbf{I}_{V_g} - \rho_{gs}\mathbf{W}_g \end{bmatrix}_{qV_g \times qV_g}$$

(2)

($\mathbf{I}_{V_g}$ and $\mathbf{J}_{V_g}$ denote identity and unity matrices of size V_g). In (2), $\mathbf{W}_g = [1/(V_g - 1)](\mathbf{J}_{V_g} - \mathbf{I}_{V_g})$ specifies the neighborhood structure consisting of all voxels that fall in the same anatomical region. The blocks along the main diagonal of Ψ_{gs}^{-1} allow for spatial correlations between measures of task-related brain activity for voxels in region g, i.e. between elements of $\mathbf{B}_{igs}^{(p)}$. The off-diagonal blocks of Ψ_{gs}^{-1} capture correlations between the summary measures of brain activity associated with various scanning sessions. Our model assumes an exchangeable covariance structure between voxels in the same neuroanatomic region. Although the complexity of human brain function is likely to render departures from this assumption, the exchangeable structure provides an improvement over the often used independence assumption and seems reasonable for statistical modeling purposes based on descriptive empirical results (not shown here; see Web Appendix A).

Calculating the matrix Ψ_{gs}—Calculating Ψ_{gs} involves inverting a matrix of size $qV_g \times qV_g$, which can become unwieldy for large regions, and this calculation is performed iteratively during estimation of the model. We derive an algorithm to facilitate calculations of inverses and determinants of our large highly structured covariance matrices, thereby enabling estimation of our model for fMRI applications. By recursive calculation on q, the number of blocks in Ψ_{gs}^{-1}, we represent Ψ_{gs}^{-1} as follows

$$\boldsymbol{\Psi}_{gs} = \begin{bmatrix} d\mathbf{I}_{V_g} + f\mathbf{J}_{V_g} & u\mathbf{I}_{V_g} + z\mathbf{J}_{V_g} & \cdots & u\mathbf{I}_{V_g} + z\mathbf{J}_{V_g} \\ u\mathbf{I}_{V_g} + z\mathbf{J}_{V_g} & d\mathbf{I}_{V_g} + f\mathbf{J}_{V_g} & \cdots & u\mathbf{I}_{V_g} + z\mathbf{J}_{V_g} \\ \vdots & \vdots & \vdots & \ddots \\ u\mathbf{I}_{V_g} + z\mathbf{J}_{V_g} & u\mathbf{I}_{V_g} + z\mathbf{J}_{V_g} & \cdots & d\mathbf{I}_{V_g} + f\mathbf{J}_{V_g} \end{bmatrix}_{qV_g \times qV_g} \tag{3}$$

Hence, $\boldsymbol{\Psi}_{gs}$ is determined by four functions: $d(\rho_g, \xi_g, V_g, q), f(\rho_{gs}, \xi_{gs}, V_g, q), u(\rho_{gs}, \xi_{gs}, V_g, q)$, and $z(\rho_{gs}, \xi_{gs}, V_g, q)$. The explicit formulas for d, f, u and z are quite lengthy and are given in the Web Appendix B. We employ (3) to circumvent issues of long computing times and excessive memory consumption.

Estimated spatial and temporal parameters are not directly interpretable as spatial and temporal correlations. The spatial and temporal correlations, denoted η_ρ and η_ξ, respectively, require calculation of the covariance matrix $\boldsymbol{\Omega}_{gs}$, i.e. of $\boldsymbol{\Psi}_{gs}\boldsymbol{\Psi}_{gs}$. Leveraging the special structure of $\boldsymbol{\Psi}_{gs}$, we obtain $\boldsymbol{\Psi}_{gs}\boldsymbol{\Psi}_{gs}$ as a partitioned matrix, with equal diagonal blocks $A\mathbf{I}_{V_g} + B\mathbf{J}_{V_g}$, and off-diagonal blocks $C\mathbf{I}_{V_g} + D\mathbf{J}_{V_g}$, where

$$\begin{aligned} A &= d^2 + (q-1)u^2 & B &= 2df + V_g f^2 + (q-1)(2uz + V_g z^2) \\ C &= 2du + (q-2)u^2 & D &= 2(fu + dz + V_g fz) + (q-2)(2uz + V_g z^2). \end{aligned} \tag{4}$$

The expressions for the spatial correlation, η_ρ, and the temporal correlation, η_ξ, are then given by $\eta_\rho = B/(A + B)$ and $\eta_\xi = (C + D)/(A + B)$.

3.3 Estimation

We perform estimation of the second stage spatio-temporal model using ML methods. The log-likelihood function for model (1) is given by

$$\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\rho}, \boldsymbol{\xi}, \boldsymbol{\beta}|\mathbf{B}) \propto \sum_{s=1}^{2} \sum_{i \in K_s} \sum_{g=1}^{G} \left\{ -\frac{qV_g}{2} \ln(\varphi_{gs}^2) + \ln(\mathrm{abs}|\boldsymbol{\Psi}_{gs}^{-1}|) - \frac{1}{2\varphi_{gs}^2} (\mathbf{B}_{igs} - \mathbf{X}_i \beta_{gs})' \boldsymbol{\Psi}_{gs}^{-1} \boldsymbol{\Psi}_{gs}^{-1} (\mathbf{B}_{igs} - \mathbf{X}_i \beta_{gs}) \right\}, \tag{5}$$

where K_s denotes the number of subjects in subgroup s (cocaine addicts and controls). The

ML estimator of the mean parameter vector $\boldsymbol{\beta}_{gs}$ is given by $\widehat{\beta}_{gs} = \frac{1}{K_s} \sum_{i=1}^{K_s} \mathbf{B}_{igs}$, $g = 1, \ldots, G$. $\hat{\boldsymbol{\beta}}_{gs}$ is unbiased (as is $\hat{\boldsymbol{\beta}}_{gs}$ from a GLM analysis) and does not depend on the covariance parameters. Estimation of the covariance parameters, therefore, proceeds using the partially maximized likelihood function $\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\rho}, \boldsymbol{\xi}, \hat{\boldsymbol{\beta}}|\mathbf{B})$. We use the Fisher scoring algorithm to estimate the covariance parameters $\boldsymbol{\rho}_s = (\rho_{1s}, \ldots, \rho_{Gs})'$, $\boldsymbol{\xi}_s = (\xi_{1s}, \ldots, \xi_{Gs})'$, and $\boldsymbol{\varphi}_g = (\varphi_{1s}, \ldots, \varphi_{Gs})'$.

Note on the parameter space—To apply our model, the matrix $\boldsymbol{\Psi}_{gs}^{-1}$ must be nonsingular, or equivalently $\boldsymbol{\Omega}_{gs}$ must be positive definite. In the context of our fMRI data, with $V_g > 20$, for all g, and $q = 2$, the resulting explicit parameter constraints are as follows:

$$\rho_{gs}+(V_g-1)(1+\xi_{gs})\neq 0 \quad \rho_{gs}+(V_g-1)(1-\xi_{gs})\neq 0$$
$$\rho_{gs}-\xi_{gs}\neq 1 \qquad\qquad \rho_{gs}+\xi_{gs}\neq 1.$$

We monitor these boundary constraints during our iterative estimation procedure.

Estimation involves iterative calculations of the determinant and the inverse of a $qV_g \times qV_g$ matrix (up to roughly 5,800×5,800 in our data application). Employing the simplifications discussed in Section 3.2 enables fast and efficient computations. Web Appendix C contains details regarding the score functions necessary for estimation.

3.4 Inferences

Inferences that are commonly sought in functional neuroimaging studies target the identification of differences in neural processing between experimental tasks, sessions (e.g. treatment periods), or groups. One may pursue these inferences at an extremely localized (voxel) level or at a more spatially coarse regional level. In either case, estimation and inferences in this setting produce maps of distributed brain activity and the corresponding thresholded maps showing statistically significant (or highly probable, in a Bayesian context) differences. We seek inferences about linear functions of $\beta_{gs}^{(p)}$, where for the inhibitory control study in cocaine addicts, $p = 1$ (baseline or pre-treatment) or $p = 2$ (follow-up or post-treatment) and $s = 1$ (cocaine addicts) or $s = 2$ (controls).

Voxel-level inferences—Following estimation using our spatio-temporal model, we obtain t-statistic images from the voxel-specific contrast estimates, then threshold the t-statistics to determine voxels exhibiting statistically significant changes (or differences) in brain activity. To adjust for typical changes observed with repeated scanning sessions in the inhibitory control study, we identify voxels for which the changes in brain activity following treatment are larger in cocaine addicts than the corresponding changes between follow-up and baseline activity in control subjects. Specifically, we estimate and test hypotheses about

$\theta_g^*=\mathbf{C}_g^*\beta_g$, where $\mathbf{C}_g^*=[-\mathbf{I}_{V_g}\ \mathbf{I}_{V_g}\ \mathbf{I}_{V_g}\ -\mathbf{I}_{V_g}]_{(V_g\times 4V_g)}$ and $\beta_g=[\beta_{g1}^{(1)'},\beta_{g1}^{(2)'},\beta_{g2}^{(1)'},\beta_{g2}^{(2)'}]'_{(4V_g\times 1)}$. Each element of θ_g^* represents a voxel-specific parameter. We construct Wald-type statistics for hypothesis testing, with $\mathrm{Var}(\widehat{\theta_g^*})=\mathbf{C}_g^*\mathrm{Var}(\widehat{\beta_g})\mathbf{C}_g^{*'}$ and $df = 2(n_p + n_c - 2)$, where n_p is the number of addicts, and n_c is the number of controls. We apply the stringent threshold of $\alpha = 0.005$, which is one approach that has a strong precedent in the neuroimaging literature, but other thresholding approaches, such as false discovery rate (Benjamini and Hochberg, 1995) and random field theory (Friston et al., 1995) are also available.

Region-level inferences—Our spatio-temporal model also enables analyses targeting an entire anatomical brain region corresponding to the underlying anatomical parcellation. For inferences in region g, define $\theta_g = \mathbf{C}_g\boldsymbol{\beta}_g$, with $\mathbf{C}_g=\frac{1}{V_g}[-\mathbf{1}_{V_g}'\ \mathbf{1}_{V_g}'\ \mathbf{1}_{V_g}'\ -\mathbf{1}_{V_g}']_{(1\times 4V_g)}$, and construct the t-statistic using $\mathrm{Var}(\widehat{\theta_g})=\mathbf{C}_g\mathrm{Var}(\widehat{\beta_g})\mathbf{C}_g'$ and degrees of freedom $df_r = 2V_g(n_p + n_c - 2)$.

Combining both sessions for subgroups s, i.e., $\beta_{gs}=[\beta_{gs}^{(1)'},\beta_{gs}^{(2)'}]'$, the variance of $\hat{\boldsymbol{\beta}}_{gs}$ is given by $\mathrm{Var}(\widehat{\beta_{gs}})=\frac{1}{n_p+n_c}\varphi_{gs}^2\mathbf{\Psi}_{gs}\mathbf{\Psi}_{gs}$. A notable advantage of modeling spatial correlations using our model is that the resulting regional level analyses account for spatial dependencies between intra-regional voxel pairs, rather than implicitly assuming independence.

Statistical significance of spatial and temporal correlations—It is often informative to examine the magnitudes of spatial and temporal correlations as discussed in Section 3.2. Given the complexity of the estimators of these quantities, it is difficult to derive analytical expressions for the variances of these estimators. If one has inherent interest in testing hypotheses about these parameters, then standard bootstrapping methods are applicable. However, this approach may be computationally demanding. Given that testing hypotheses about correlations was not of chief interest in the inhibitory control study, we present estimates of correlations mainly for descriptive purposes. We estimate standard errors of the correlations using 30 bootstrap samples drawn (with replacement) from the subjects in our data, separately for controls and patients. We apply our model and estimate the spatial and temporal correlations for each sample, and then calculate the bootstrap standard errors as the standard deviations of the 30 estimated values for each parameter, which should provide reasonable estimates of variability (Efron and Tibshirani, 1998).

4. Application to Study of Inhibitory Control in Cocaine Addicts

We used our spatio-temporal model to analyze fMRI data from the inhibitory control study in cocaine addicts. We estimate separate spatial and temporal parameters for cocaine addicts and controls. We carry out image pre-processing (slice timing, realigning, normalizing, smoothing) and first level, single subject analyses in SPM5 (SPM5 is a MATLAB software package implementing Statistical Parametric Mapping for neuroimaging data available for download from the Wellcome Trust Centre for Neuroimaging web page http://www.fil.ion.ucl.ac.uk/spm/software/spm5/). To account for inter-subject neuroanatomic differences that may persist after spatial normalization, we apply very focal spatial smoothing (6 mm [2 voxels] full width at half maximum (FWHM) Gaussian kernel) rather than the more spatially expansive smoothing often applied in practice. Our smoothing should only have a small effect on subsequent spatial modeling, since most of the Brodmann areas (within which we model spatial correlations) are quite large, relative to the small size of the smoothing kernel. Our analysis includes 38 BA's in each hemisphere, after excluding those with fewer than 20 voxels, and we supplement the BA's with six subcortical regions, resulting in a parcellation consisting of 82 brain regions. We perform exploratory analyses, computing crude estimates of spatial and temporal correlations (Web Appendix A), which provide support for the existence of correlations from both sources as well as for the exchangeability assumption in our covariance model.

4.1 Voxel-level inferences

Figure 1(a) shows voxels that reveal statistically significant increases in inhibitory control-related brain activity following treatment at $\alpha = 0.005$. The axial slices correspond to 5 mm, 14 mm, 44 mm, and 65 mm above the anterior-posterior commissural plane. The significant areas include the right frontopolar cortex (BA 10), left middle temporal gyrus (BA 21), and retrosubicular area (BA 48) in the 5 mm slice; left and right thalamus and the right inferior frontal cortex (BA 45) at 14 mm; left visual association cortex V3 (BA 19) and somatosensory association cortex (BA 7) and the left angular gyrus (BA 39) at 44 mm; and the right pre-supplementary motor area (pre-SMA, BA 6) at 65 mm. Other regions (not shown) exhibiting significant treatment-related increases in brain activity include the right inferior prefrontal gyrus (BA 47), left and right fusiforum gyrus (BA 37), and left primary auditory cortex (BA 41).

Figure 1(b) shows results from the corresponding GLM–based analysis commonly used in the neuroimaging literature. Many of the areas showing statistical significance agree with the results from our approach (e.g. Brodmann areas 6, 7, 10, and 19 emerged in both analyses), but there are some differences. For example, our method identifies increased

treatment-related brain activity in BA 48 and BA's 21 and 22 (5 mm), which do not emerge from the GLM analysis. Also, both methods identify areas in the thalamus, but our method detects a large right thalamic activation not revealed by the GLM. Similarly, both methods identify voxels in BA's 6 (pre-SMA) and 39, but our model yields more spatially extensive activations. Overall, our ST model produces more statistically significant voxels than the GLM. However, neither method produces significant voxels when applying more conservative family-wise error or false discovery rate multiple testing procedures. The differences between the two models become more apparent at larger significance levels (e.g. $\alpha = 0.01$, not shown), suggesting that with a larger sample size and stronger activation signal, greater differences are likely to emerge. We provide interpretive remarks about our results in Section 4.4.

4.2 Region-level inferences

Our region-level analysis of treatment-emergent changes in neural processing, using methods described in Section 3.4, revealed only one region (left BA 39, angular gyrus) achieving statistical significance at $\alpha = 0.05$, uncorrected (see Figure 1(c)). The comparative region-level GLM analysis (that implicitly assumes spatial and temporal independence) did not yield any statistically significant differences. Neither method reveals significant differences at more stringent thresholds such as $\alpha = 0.005$ (uncorrected).

4.3 Spatial and temporal correlations

Figure 2 shows the model-based estimates of spatial and temporal correlations for both the cocaine addicts and control subjects in one axial slice of the brain. The color scale indicates the strength of correlations, with brighter shades indicating correlations of larger magnitude. For temporal correlations, positive values are shown in one color (top colorbar) and negative values are shown in a separate color (bottom colorbar). For clearer distinction, we add stripes to regions with negative correlations. All of the correlations in Figure 2 would have the value zero in a GLM analysis, signaling the need to account for them in our model.

Table 1 gives the model-based estimates, along with bootstrap standard errors (in parentheses), of both temporal and spatial correlations for select regions in the inhibitory control study. The associated spatial (ρ) and temporal (ξ) dependence parameters achieve statistical significance for all regions at $\alpha = 0.01$, except for BA 29 (L), which is significant at $\alpha = 0.05$. The selected regions in Table 1 are those that have correlations with the largest magnitude. The temporal correlations for cocaine addicts range from -0.24 to 0.45 and tend to be larger in control subjects for whom they range from -0.28 to 0.64. The regional spatial correlations are not consistently higher for either group, ranging from 0.08 to 0.54 among cocaine addicts and from 0.04 to 0.65 for healthy controls.

The brain regions with highest temporal correlations differ between the two groups. For the controls, the highest temporal correlations (~ 0.6) are in right BA 42, partially covering the auditory association cortex, and in left BA 21, left BA 22, and right BA 22, regions in the temporal lobe which are also revealed by our voxel-level analysis. For the cocaine addicts, the highest temporal correlations are somewhat smaller (~ 0.4) and appear in left BA 43 and in right BA 22. Spatial correlations for the control group are highest (≥ 0.6) in left and right BA 26, a retrosplenial region in the cingulate cortex. For the cocaine addicts, the highest estimated spatial correlations (~ 0.5) are in left BA 29, also in the retrosplenial region, and in left BA 43. For both groups, the region with the lowest estimated spatial correlation (≤ 0.08) is the right BA 48, which is the second largest region (2,875 voxels). This suggests that it is the least spatially homogeneous region and that subdividing large BA's may be warranted.

4.4 Implications of Results

The study on inhibitory control in cocaine addicts that is the subject of this novel image analysis approach is, to the best of our knowledge, the first to examine the changes in brain activity associated with an intensive addiction behavioral therapy. Inhibiting prepotent or automatic responses is critical to the organization of successful goal-directed behaviors. Aron et al. (2007) demonstrated that the ability to stop motor responses depends critically on a network of structures including the right inferior frontal cortex and the subthalamic nucleus, both of which are connected to the pre-SMA.

Our results agree, in several aspects, with the findings from these and other previous neuroimaging studies. Greater activation of the right orbitofrontal cortex is necessary for behavioral inhibition in impulsive individuals (Horn et al., 2003). Additionally, the right ventral prefrontal cortex has been selectively implicated in the neuropathophysiology of drug addiction (Volkow et al., 2005). Voxel-level analysis of our data reveals that an important effect of treatment is an enabling of the right inferior prefrontal gyrus (BA 47) and right orbitofrontal cortex (BA 11) responses to a demand for inhibition of habitual motor responses. Our findings are promising since they suggest a positive effect of the behavioral therapy on the neural processing deficits associated with cocaine-addiction, though we do not strictly regard these effects to be causal.

The pre-SMA is identified as a significant treatment-related area in our voxel-level analysis. The caudal pre-SMA is critical to controlled action selection (Nachev et al., 2007) and is functionally compromised in cocaine addicts (Kaufman et al., 2003). Enhancing the neural response of the pre-SMA to a demand for a shift from habitual responses to controlled response inhibition represents a plausible neural correlate of drug refusal skills acquired in behavioral therapies targeting relapse prevention. Our region level analysis identifies the angular gyrus (BA 39), which is important in visuospatial attention (Cattaneo et al., 2009).

In areas exhibiting significant voxel-level post-treatment increases in task-related brain activity, we observe distinct patterns of intra-regional correlations between patients and controls. Based on the estimates of the spatial correlations, the controls exhibit higher spatial coherence in, for example, right BA's 11 (controls = 0.24(0.11), addicts = 0.14(0.04)) and 47 (controls = 0.28(0.13), addicts = 0.18(0.10)), which are frontal areas involved in executive functioning, planning and decision making, as well as in both left and right thalamus (e.g. in right thalamus, controls = 0.32(0.09), addicts = 0.17(0.07)), which play a central role in the flow of information to the cortex (Sherman and Guillery, 2002). Cocaine addicts showed more coherent functioning in left BA's 19 (addicts = 0.16(0.02), controls = 0.09(0.07)) and 39 (addicts = 0.31(0.10), controls = 0.16(0.10)), which are involved in visual attention. The latter finding is particularly interesting since BA 39 was identified in our region-level analysis.

5. Simulations

We report results of two simulation studies but defer many of the details to Web Appendix D. The goal of the first simulation is to evaluate the accuracy of our estimation procedure. In the second simulation study, we evaluate the relative efficiency of our spatio-temporal model compared to the GLM, both for voxel-level and region-level estimators of secondary parameters of interest. Below we briefly summarize the main findings.

In the first simulation study, we find that estimates for both ρ_{gs} and ξ_{gs} (spatial and temporal parameters) are quite accurate, while the between subject variability φ_{gs}^2 is slightly underestimated (roughly 6%), which likely stems from the well-known downward bias of ML estimates of variance components (Laird and Ware, 1998). Restricted ML estimation

can be used for our model and would presumably mitigate the observed bias. We favor ML estimation in our context, however, because it substantially facilitates computations. The second simulation reveals that the GLM-based voxel-level variance estimates are on average higher than the corresponding estimates obtained from our spatio-temporal model. At the region level, our model also gives more precise estimates compared to GLM. The increased precision of our model estimates stems from the fact that we borrow strength across within region measures of neural activity, which will tend to yield gains in statistical power.

6. Discussion

We propose a novel spatio-temporal modeling framework for functional neuroimaging data, overcoming unsupported independence assumptions of standard GLM-based methods. We use Brodmann area templates to define neuroanatomic regions, but other maps are applicable. Our approach provides a unified framework for voxel-level and region-level inferences. An important contribution of our work is that we derive efficient computational solutions to facilitate implementation.

One limitation of our model is that it does not account for correlations between regions. Doing so within our current framework would lead to substantially increased, perhaps prohibitive, computations. Bowman et al. (2008) establish a Bayesian framework that models between-region correlations, in addition to within-region correlations, but the number of regions included in the analysis is consequently constrained by the sample size. An advantage of our proposed model is that the number of regions included is unconstrained. Our current ML estimation procedure requires close monitoring to ensure convergence and confinement within the parameter space.

In summary, our proposed spatio-temporal model provides an appealing, computationally efficient alternative to standard GLM-based methods for analyzing fMRI data. Our model is based on assumptions that are more neurophysiologically plausible, capturing correlations between different brain locations and between estimates of neural activity at different scanning sessions. These correlations lead to interpretive advantages over the GLM, e.g. revealing information about the degree of coherence in brain activity within defined neuroanatomic regions. Our simulation studies demonstrate that our model estimates are quite accurate and that the standard errors associated with estimates of our mean model parameters are on average smaller than those from a GLM. This increase in efficiency will often lead to more powerful statistical tests and the detection of more statistically significant voxels.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

References

Aron AR, Behrens TE, Smith S, Frank MJ, Poldrack RA. Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. The Journal of Neuroscience. 2007; 14:3743–3752. [PubMed: 17409238]

Aron AR, Poldrack RA. Cortical and subcortical contributions to stop signal response inhibition: role of the subthalamic nucleus. Journal of Neuroscience. 2006; 26:2424–2433. [PubMed: 16510720]

Benali, H.; Buvat, I.; Anton, JL.; Pélégrini, M.; Di Paola, M.; Bittoun, J.; Burnod, Y.; Di Paola, R. Space-time statistical model for functional MRI image sequences. In: Duncan, J.; Gindi, G., editors. Information Processing in Medical Imaging. 1997. p. 285-298.

Benjamini Y, Hochberg Y. Controlling the false discovery rate: A practical and powerful approach to multiple testing. Journal of the Royal Statistical Society Series B (Methodological). 1995; 57:289–300.

Bowman F. Spatio-temporal modeling of localized brain activity. Biostatistics. 2005; 6:558–575. [PubMed: 15843592]

Bowman F, Caffo B, Bassett S, Kilts C. A Bayesian hierarchical framework for spatial modeling of fMRI data. Neuro Image. 2008; 39:146–156. [PubMed: 17936016]

Bowman F, Kilts C. Modeling intra-subject correlation among repeated scans in positron emission tomography (PET) neuroimaging data. Human Brain Mapping. 2003; 20:59–70. [PubMed: 14505332]

Brodmann, K. Kapitel in Vergleichende Lokalisationslehre der Grosshirnrinde. Vol. IV. Leipzig: Verlag von Johann Ambrosias Barth; 1909. Beschreibung der einzelnen Hirnkarten.

Bullmore E, Brammer M, Williams S, Rabe-Hesketh S, Janot N, David A, Mellers J, Howard R, Sham P. Statistical methods of estimation and inference for functional MR image analysis. Magnetic Resonance in Medicine. 1996; 35:261–277. [PubMed: 8622592]

Cattaneo Z, Silvanto J, Pascual-Leone A, Battelli L. The role of the angular gyrus in the modulation of visuospatial attention by the mental number line. Neuro Image. 2009; 44:563–568. [PubMed: 18848630]

Efron, B.; Tibshirani, RJ. An Introduction to the Bootstrap. Chapman & Hall; 1998.

Friston K, Glaser DE, Henson RN, Kiebel S, Phillips C, Ashburner J. Classical and Bayesian inference in neuroimaging: applications. Neuroimage. 2002; 16:484–512. [PubMed: 12030833]

Friston KJ, Holmes AP, Worsley KJ, Poline JP, Frith CD, Frackowiak RSJ. Statistical parametric maps in functional imaging: A general linear approach. Human Brain Mapping. 1995; 2:189–210.

Gössl C, Auer D, Fahrmeir L. Bayesian spatiotemporal inference in functional magnetic resonance imaging. Biometrics. 2001; 57:554–562. [PubMed: 11414583]

Hartvig NV. A Stochastic Geometry Model for Functional Magnetic Resonance Images. Scandinavian Journal of Statistics. 2002; 29:333–353.

Horn NR, MD, Elliot R, Deakin JFW, Woodruff PWR. Response inhibition and impulsivity: an fMRI study. Neuropsychologia. 2003; 41:1959–1966. [PubMed: 14572528]

Kaufman JN, Ross TJ, Stein EA, Garavan H. Cingulate hypoactivity in cocaine users during a GO-NOGO task as revealed by event-related functional magnetic resonance imaging. The Journal of Neuroscience. 2003; 23:7839–7843. [PubMed: 12944513]

Laird NM, Ware JH. Random-effects models for longitudinal data. Biometrics. 1998; 38:963–974. [PubMed: 7168798]

Nachev P, Wydell H, O'Neil K, Husain M, Kennard C. The role of the pre-supplementary motor area in the control of action. Neuro Image. 2007; 36:T155–T163. [PubMed: 17499162]

Purdon P, Solo V, Weisskoff R, Brown E. Locally regularized spatiotemporal modeling and model comparison for functional MRI. Neuro Image. 2001; 14:912–923. [PubMed: 11554810]

Sherman SM, Guillery RW. The role of the thalamus in the flow of information to the cortex. Phil Trans R Soc Lond B. 2002; 357:1695–1708. [PubMed: 12626004]

Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M. Automated anatomical labelling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single subject brain. Neuro Image. 2002; 15:273–289. [PubMed: 11771995]

Volkow ND, Wang GJ, Fowler JS, Wong C, Ding YS, Hitzemann R, Swanson JM, Kalivas P. Activation of orbital and medial prefrontal cortex by methylphenidate in cocaine-addicted subjects but not in controls: relevance to addiction. The Journal of Neuroscience. 2005; 25:3932–3939. [PubMed: 15829645]

Woolrich MW, Jenkinson M, Brady JM, Smith SM. Fully Bayesian spatio-temporal modeling of fMRI data. IEEE Transactions on Medical Imaging. 2004; 23:213–231. [PubMed: 14964566]

Worsley K, Liao CH, Aston J, Petre V, Duncan GH, Morales F, Evans AC. A general statistical analysis for fMRI data. Neuro Image. 2002; 15:1–15. [PubMed: 11771969]

7. Appendix

Our stage-two model (1) can be written in matrix form as follows:

$$\mathbf{B}_{igs} = \beta_{gs} + \rho_{gs}\mathbf{W}^*(\mathbf{B}_{igs} - \beta_{gs}) + \xi_{gs}\mathbf{M}^*(\mathbf{B}_{igs} - \beta_{gs}) + e_{igs},$$

where $e_{igs} \sim \mathrm{MVN}(\mathbf{0}, \varphi_{gs}^2\mathbf{I})$ and where $\mathbf{W}^* = \mathbf{I}_q \otimes \mathbf{W}_g$, $\mathbf{M}^* = 1/(q-1)(\mathbf{J}_q \otimes \mathbf{I}_{V_g} - \mathbf{I}q \otimes \mathbf{I}_{V_g})$ and $\mathbf{W}_g = 1/(V_g - 1)(\mathbf{J}_{V_g} - \mathbf{I}_{V_g})$. We can rewrite the above equation as

$$\mathbf{B}_{igs} = \beta_{gs} + (\mathbf{I} - \rho_{gs}\mathbf{W}^* - \xi_{gs}\mathbf{M}^*)^{-1} e_{igs} = \beta_{gs} + \boldsymbol{\Psi}_{gs} e_{igs} \tag{6}$$

Therefore, $\mathbf{B}_{igs} \sim \mathrm{MVN}(\beta_{gs}, \varphi_{gs}^2\boldsymbol{\Psi}_{gs}\boldsymbol{\Psi}_{gs})$.

Figure 1.
Voxel-level and region-level inferences for two models: the proposed spatio-temporal (ST) model (top) and the GLM (bottom). (a): voxels that achieve statistical significance in the ST model analysis at $\alpha = 0.005$ in favor of a one-sided alternative hypothesis of increased activity following treatment for cocaine addiction. Significant increases occur in the right frontopolar area (BA 10), left middle and superior temporal gyri (BA's 21/22), right retrosubicular area (BA 48), right inferior frontal cortex (BA 45), left and right thalamus (Thal), left somatosensory and visual association cortices (BA's 7 and 19), left angular gyrus (BA 39), and right BA 6 (pre-SMA); (b): voxel-level results from a GLM analysis; (c): region, left BA 39 (angular gyrus), that achieves significance in our ST model analysis at $\alpha = 0.05$; (d): corresponding region-level results from the GLM. Slice labels denote the distance, in mm, above the anterior-posterior commissural plane.
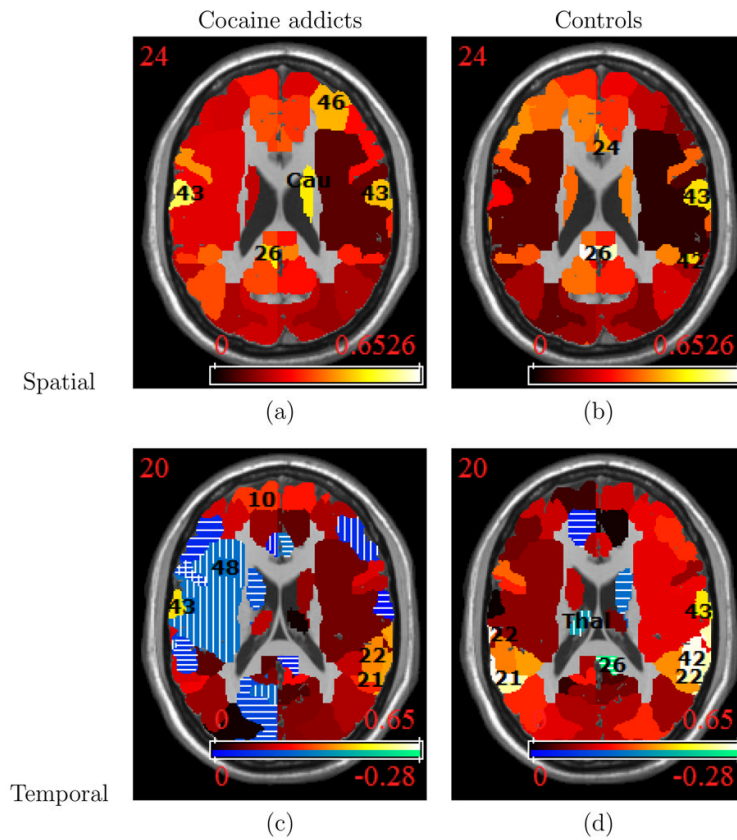
Figure 2.
Model-based estimates of the spatial and temporal correlations for the inhibitory control study (a) addicts: spatial correlations, (b) controls: spatial correlations, (c) addicts: temporal correlations, (d) controls: temporal correlations. For (a) and (b), the darker to lighter intensities in the colorbar represent lower to higher correlations. For (c) and (d), the top colorbar represents positive correlations, again with lighter intensities reflecting higher correlations, and the bottom colorbar represents negative correlations (stripes added to maps) with brighter intensities indicating stronger negative correlations. The temporal correlations for control subjects range from −0.28 to 0.64 and are in general larger than those for the cocaine addicts, which range from −0.24 to 0.45. The regional spatial correlations are not consistently higher for either group, ranging from 0.08 to 0.54 among cocaine addicts and from 0.04 to 0.65 for healthy controls.

Table 1

Model-based estimates of spatial and temporal correlations and estimated bootstrap standard errors (in parentheses).

Group	Region	Temporal correlations (std. err.)	Region	Spatial correlations (std. err.)
Addicts				
	BA 43 (L)	0.45 (0.19)	BA 43 (L)	0.54 (0.10)
	BA 22 (R)	0.39 (0.15)	BA 29 (L)	0.50 (0.12)
	BA 42 (R)	0.37 (0.18)	BA 3 (L)	0.47 (0.10)
	BA 3 (L)	0.37 (0.17)	caudate (R)	0.47 (0.10)
	BA 9 (R)	0.34 (0.28)	BA 43 (R)	0.43 (0.07)
	BA 35 (R)	−0.24 (0.06)	BA 48 (R)	0.08 (0.07)
Controls				
	BA 22 (L)	0.64 (0.07)	BA 26 (L)	0.65 (0.05)
	BA 42 (R)	0.63 (0.09)	BA 26 (R)	0.60 (0.08)
	BA 22 (R)	0.60 (0.09)	BA 5 (R)	0.50 (0.12)
	BA 21 (L)	0.58 (0.10)	putamen (L,R)	0.50 (0.08, 0.06)
	BA 43 (R)	0.46 (0.10)	BA 38 (L)	0.47 (0.10)
	BA 26 (R)	−0.28 (0.11)	BA 48 (R)	0.04 (0.02)