



University of Warwick institutional repository: <http://go.warwick.ac.uk/wrap>

This paper is made available online in accordance with publisher policies. Please scroll down to view the document itself. Please refer to the repository record for this item and our policy information available from the repository home page for further information.

To see the final version of this paper please visit the publisher's website. Access to the published version may require a subscription.

Author(s): Robert Leech and Dennis Leech

Article Title: Testing for spatial heterogeneity in functional MRI using the multivariate general linear model

Year of publication: 2010

Link to published article:

http://www2.warwick.ac.uk/fac/soc/economics/research/workingpapers/2010/twerp_938.pdf

Publisher statement: None

Testing for spatial heterogeneity in functional MRI using the
multivariate general linear model

Robert Leech and Dennis Leech

No 938

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Testing for spatial heterogeneity in functional
MRI using the multivariate general linear model

Robert Leech¹ and Dennis Leech²

¹The Computational, Cognitive and Clinical Neuroimaging
Laboratory, The Division of Experimental Medicine, Imperial
College London, Hammersmith Hospital Campus, Du Cane
Road, London, W12 0NN, UK

²Department of Economics, University of Warwick, Coventry,
CV4 7AL

1 Abstract

Much current research in functional MRI employs multivariate machine learning approaches (e.g., support vector machines) to detect fine-scale spatial patterns from the temporal fluctuations of the neural signal. The aim of many studies is not classification, however, but investigation of multivariate spatial patterns, which pattern classifiers detect only indirectly. Here we propose a direct statistical measure for the existence of fine-scale spatial patterns (or spatial heterogeneity) applicable for fMRI datasets. We extend the univariate general linear model (typically used in fMRI analysis) to a multivariate case. We demonstrate that contrasting maximum likelihood estimations of different restrictions on this multivariate model can be used to estimate the extent of spatial heterogeneity in fMRI data. Under asymptotic assumptions inference can be made with reference to the χ^2 distribution. The test statistic is then assessed using simulated timecourses derived from real fMRI data. This demonstrates the utility of the proposed measure of heterogeneity as well as considerations in its application. Measuring spatial heterogeneity in fMRI has important theoretical implications in its own right and has potential uses for better characterising neurological conditions such as stroke and Alzheimer's disease.

Key wordsNeuroimaging, Multivariate pattern analysis, Maximum likelihood estimation, Seemingly unrelated regression.

2 Introduction

The traditional approach to quantifying changes in brain activation with functional magnetic resonance imaging (fMRI) is massively univariate; a separate general linear model is fitted to the time series of each of the many tens of thousands of spatially distinct measurement locations (voxels). Given the relatively poor signal-to-noise of the measurable signal, Gaussian spatial smoothing across neighboring voxels is normally applied. However, this smoothing assumes that the detectable neural signal is spatially homogeneous (i.e., approximately the same general linear model should fit adjacent voxels). This assumption has been the subject of intense recent debate (de Beeck and Hans 2009, Kamitani and Sawahata 2009, de Beeck and Hans 2010). Some multi-voxel analysis techniques indicate that brain regions may be highly heterogeneous (Haxby et al. 2001, Kamitani and Tong 2005, Haynes et al. 2007) and demonstrate huge improvements in detecting signal over the traditional uni-

variate approach. Furthermore, whether areas of the brain are highly heterogeneous (with fine-scale spatial pattern information) is important theoretically for understanding the roles of regions of cortex, with potential implications for better understanding neurological conditions such as Alzheimer’s disease or stroke. However, the multivariate pattern classification techniques used to date (e.g., support vector machines and linear discriminant analysis) are indirect ways to measure the heterogeneity of the fMRI signal across brain regions.

In this paper, we propose a framework and test statistic to directly measure the extent to which there is a heterogeneous pattern of activation across neighboring voxels. The proposal takes as its starting point the information-based approach of Kriegeskorte and colleagues (Kriegeskorte et al. 2006a). Instead of using multivoxel pattern classifiers to decode which state the brain is in as a proxy for the amount of information available in a region, Kriegeskorte and colleagues proposed using a variant on the Mahalanobis Distance to provide a measure of the amount of information available across voxels in a patch of cortex. The Mahalanobis distance statistic was shown to be substantially more sensitive at detecting signals than standard univariate t-statistics or Euclidean distance measures. Variants of this statistic have been used suc-

cessfully in several fMRI experiments within the visual domain (Kriegeskorte et al. 2006a, Serences and Boynton 2007a;b, Stokes et al. 2009).

The Mahalanobis distance (Mahalanobis 1936) is a multivoxel similarity measurement that unlike Euclidean distance is scale invariant and controls for covariance across datapoints. In fMRI datasets, the Mahalanobis distance can calculate the similarity between different experiment conditions, e.g., how dissimilar are visual and auditory processing for a given set of voxels. This statistic controls for error covariance across voxels; spatially correlated error is expected in fMRI datasets given spatial patterns of sources of noise affecting the BOLD signal, such as vasculature, movement artifacts etc. To calculate the Mahalanobis distance, Kriegeskorte et al. (2006b) fitted a separate GLM analysis to the unsmoothed functional data for each voxel. The Mahalanobis distance was then calculated using the resulting vectors of beta weights and the residual error covariance matrix across voxels

$$\Delta = (\hat{\beta}^2 - \hat{\beta}^1)' \Sigma^{-1} (\hat{\beta}^2 - \hat{\beta}^1) \quad (1)$$

where $\hat{\beta}^1$ and $\hat{\beta}^2$ are vectors of estimated regression coefficients associated

with two types of stimuli and $\hat{\Sigma}$ is the error covariance matrix.

In this paper we start by demonstrating how the Mahalanobis distance relates to the multivariate general linear model. We then set up different restrictions on the GLM to compare models under an assumption of spatial homogeneity and spatial heterogeneity resulting in a chi-square statistic measuring the heterogeneity in a patch of cortex. This statistic explains the amount of fine-scale information available across a number of voxels. The statistic is then assessed with synthetic data derived from fMRI datasets to demonstrate its utility and to explore limitations in its application.

3 Theory

3.1 The statistical model

We start with the GLM as applied to fMRI datasets (Friston et al. 1994).

We assume the response by each voxel can be modelled by a linear regression model written as follows.

$$y_{it} = \sum_{h=1}^k \beta_{hi} x_{ht} + u_{it} \quad (2)$$

where i is the voxel subscript, $i = 1, \dots, n$, and t is the time subscript, $t = 1, \dots, T$. β_{hi} represents the response of voxel i to the stimulus measured by regressor h .

This can be written using matrix algebra as,

$$\mathbf{y}_i = \mathbf{X}\boldsymbol{\beta}_i + \mathbf{u}_i \quad (3)$$

where \mathbf{y}_i is the T element observation vector for voxel i , \mathbf{X} the $T \times k$ input matrix, \mathbf{u}_i the T element error vector, and $\boldsymbol{\beta}_i$ the k element coefficient vector. There is a different regression model for each voxel but all models have a common regressor matrix.

The error for each voxel equation is assumed to have zero mean, be serially independent and homoscedastic, but correlated across equations. That is, $E(\mathbf{u}_i) = \mathbf{0}$ for all i and $E(\mathbf{u}_i \mathbf{u}_j') = \sigma_{ij} \mathbf{I}_T$ for all i and j .

This system of equations can be written more compactly as:

$$\mathbf{y} = \mathbf{Z}\boldsymbol{\beta} + \mathbf{u}, \quad E(\mathbf{u}) = \mathbf{0}, \quad E(\mathbf{u}\mathbf{u}') = \boldsymbol{\Omega}, \quad (4)$$

where

$$\mathbf{y} = \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \vdots \\ \mathbf{y}_n \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_n \end{pmatrix}, \quad \boldsymbol{\beta} = \begin{pmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \\ \vdots \\ \boldsymbol{\beta}_n \end{pmatrix}, \quad \mathbf{Z} = \begin{pmatrix} \mathbf{X} & 0 & \dots & 0 \\ 0 & \mathbf{X} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathbf{X} \end{pmatrix} = \mathbf{I}_n \otimes \mathbf{X},$$

$$\boldsymbol{\Omega} = \begin{pmatrix} \sigma_{11}\mathbf{I}_T & \sigma_{12}\mathbf{I}_T & \dots & \sigma_{1n}\mathbf{I}_T \\ \sigma_{12}\mathbf{I}_T & \sigma_{22}\mathbf{I}_T & \dots & \sigma_{2n}\mathbf{I}_T \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{1n}\mathbf{I}_T & \sigma_{2n}\mathbf{I}_T & \dots & \sigma_{nn}\mathbf{I}_T \end{pmatrix} = \boldsymbol{\Sigma} \otimes \mathbf{I}_T, \quad \text{and } \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \dots & \sigma_{1n} \\ \sigma_{12} & \sigma_{22} & \dots & \sigma_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{1n} & \sigma_{2n} & \dots & \sigma_{nn} \end{pmatrix}.$$

It is well known (for example (Zellner 1962, Greene 2003)) that, when there is a common set of regressors, the correlation of the errors in different equations is irrelevant to coefficient estimation. The efficient estimator, equivalent to generalized least squares, is ordinary least squares applied to each equation separately. That is, the efficient estimator of the complete system is written:

$$\hat{\boldsymbol{\beta}}_i = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}_i, \text{ for all } i, \text{ and therefore } \hat{\boldsymbol{\beta}} = (\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'\mathbf{y}. \quad (5)$$

The covariance matrix of $\hat{\boldsymbol{\beta}}$ can be shown to be

$$\mathbf{V}(\hat{\boldsymbol{\beta}}) = (\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'\boldsymbol{\Omega}\mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}. \quad (6)$$

Noting that

$$\mathbf{Z}'\mathbf{Z} = [\mathbf{I}_n \otimes \mathbf{X}'] [\mathbf{I}_n \otimes \mathbf{X}] = \mathbf{I}_n \otimes \mathbf{X}'\mathbf{X}, \text{ and hence that } (\mathbf{Z}'\mathbf{Z})^{-1} = \mathbf{I}_n \otimes (\mathbf{X}'\mathbf{X})^{-1},$$

and that

$$\mathbf{Z}'\boldsymbol{\Omega}\mathbf{Z} = [\mathbf{I}_n \otimes \mathbf{X}'] [\boldsymbol{\Sigma} \otimes \mathbf{I}_T] [\mathbf{I}_n \otimes \mathbf{X}] = \boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X}),$$

we can rewrite (6) as

$$\mathbf{V}(\hat{\boldsymbol{\beta}}) = [\mathbf{I}_n \otimes (\mathbf{X}'\mathbf{X})^{-1}] [\boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X})] [\mathbf{I}_n \otimes (\mathbf{X}'\mathbf{X})^{-1}] = \boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X})^{-1}. \quad (7)$$

This gives a simple expression for the covariance between any pair of estimated coefficients. Letting $\mathbf{W} = (\mathbf{X}'\mathbf{X})^{-1}$, we can write

$$\text{Cov}(\hat{\beta}_{gi}, \hat{\beta}_{hj}) = \sigma_{ij} w_{gh} \quad (8)$$

for all pairs of voxels $i, j = 1, \dots, n$, and all pairs of regression coefficients $g, h = 1, \dots, k$.

All this assumes that the covariance matrix of equation errors, $\boldsymbol{\Sigma}$, is known.

In practice it must be estimated consistently from the residuals, using,

$$\hat{\sigma}_{ij} = (\mathbf{y}_i - \mathbf{X}\hat{\boldsymbol{\beta}}_i)'(\mathbf{y}_j - \mathbf{X}\hat{\boldsymbol{\beta}}_j)/T, \quad (9)$$

(or, $\hat{\sigma}_{ij} = (\mathbf{y}_i - \mathbf{X}\hat{\boldsymbol{\beta}}_i)'(\mathbf{y}_j - \mathbf{X}\hat{\boldsymbol{\beta}}_j)/(T - k)$) and $\boldsymbol{\Sigma}$ replaced with $\hat{\boldsymbol{\Sigma}}$ throughout in the manner of FGLS.

3.2 Inference

Our focus is on testing a set of linear restrictions on coefficients across the separate voxel equations. We assume there are g restrictions written in the form $\mathbf{R}\boldsymbol{\beta} = \mathbf{0}$ where \mathbf{R} is a $g \times nk$ matrix of given constants of rank g . (We do not need to consider inhomogeneous restrictions here because units of measurement are arbitrary in the context.)

This framework allows us to devise tests for heterogeneity versus homogeneity (that is, constancy) of the effects of different stimulæ across voxels, as well as to test if stimulæ are significantly different from zero.

If the equation errors u_{it} are Gaussian then the efficient estimator of $\boldsymbol{\beta}$ given in equation (5) is normally distributed and the well known (see for example Rao 1973) inferential procedures for generalized linear regression models are available based on the covariance matrix in equation (6), (7) or (8). These hold asymptotically if $\boldsymbol{\Sigma}$ is estimated consistently.

We characterise the inferential problem as testing a null hypothesis of the

general type

$$H_0 : \mathbf{R}\boldsymbol{\beta} = \mathbf{0}. \quad (10)$$

which can be approached in various ways, for example by a Wald or LR principle.

Letting $\mathbf{R}\boldsymbol{\beta} = \boldsymbol{\delta}$ the null hypothesis of interest becomes $H_0 : \boldsymbol{\delta} = \mathbf{0}$. Now define $\hat{\boldsymbol{\delta}} = \mathbf{R}\hat{\boldsymbol{\beta}}$. Then $\hat{\boldsymbol{\delta}}$ is asymptotically normally distributed with expectation $E(\hat{\boldsymbol{\delta}}) = \mathbf{0}$ and covariance matrix $V(\hat{\boldsymbol{\delta}}) = \mathbf{R}V(\hat{\boldsymbol{\beta}})\mathbf{R}'$. We call the test statistic Δ , given by

$$\Delta = \hat{\boldsymbol{\delta}}'[\mathbf{R}V(\hat{\boldsymbol{\beta}})\mathbf{R}']^{-1}\hat{\boldsymbol{\delta}} \quad (11)$$

which has an asymptotic χ^2 distribution with g degrees of freedom if H_0 is true.

Alternatively we employ the likelihood principle and LR tests based on the loglikelihood function

$$\ln L = -\frac{Tn}{2}\ln(2\pi) - \frac{T}{2}\ln|\boldsymbol{\Omega}| - \frac{1}{2}Q \quad (12)$$

where Q is the appropriate generalized sum of squared residuals,

$$Q = \mathbf{u}'\boldsymbol{\Omega}^{-1}\mathbf{u}. \quad (13)$$

We can specify a LR test of a null hypothesis in the form of (10) as twice the difference between the maximized value of (12) under the null and alternate hypotheses, $LR = -2(\ln L_0 - \ln L_1)$. The test statistic is then

$$LR = Q_0 - Q_1, \tag{14}$$

where Q_0 is the minimum of (13) under H_0 and Q_1 the unconstrained minimum (or, the minimum under the alternate H_1). This assumes Σ to be known or the same estimate of it used under both the null and alternate. Standard theory gives the statistic (14) an asymptotic chi-squared distribution with g degrees of freedom. (See (Greene 2003).)

3.3 Testing heterogeneity across voxels

We now describe tests of spatial characteristics of the multivariate signal across voxels. In particular, we are interested in investigating heterogeneity of response coefficients. We describe a procedure for comparing the unrestricted model, allowing heterogeneity between voxels, with a restricted model, which we say has homogeneity between voxels.

We characterize the null hypothesis of homogeneity in terms of a linear function of the coefficients. For the i th voxel the function of interest is written

$\mathbf{r}'\boldsymbol{\beta}_i$ for suitable constants $\mathbf{r}' = (r_1, r_2, \dots, r_k)$. The null hypothesis states that $\mathbf{r}'\boldsymbol{\beta}_i$ is constant for all i . That is, $\mathbf{r}'\boldsymbol{\beta}_i = \mathbf{r}'\boldsymbol{\beta}_j$ for all $i \neq j$. Therefore, we can write the restrictions to be tested in the form: $H_0 : \mathbf{R}\boldsymbol{\beta} = \mathbf{0}$

with

$$\mathbf{R} = \begin{pmatrix} \mathbf{r}' & -\mathbf{r}' & \mathbf{0}' & \cdots & \mathbf{0}' \\ \mathbf{r}' & \mathbf{0}' & -\mathbf{r}' & \cdots & \mathbf{0}' \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{r}' & \mathbf{0}' & \mathbf{0}' & \cdots & -\mathbf{r}' \end{pmatrix}.$$

An important case is where interest centers on the *difference* between two effects. If the two effects are assumed to be the first two coefficients, then $\mathbf{r}'\boldsymbol{\beta}_i = \beta_{1i} - \beta_{2i}$ and $\mathbf{r}' = (1, -1, 0, \dots, 0)$.

Let the common difference be θ . Then the restricted model can be written

$$H_0 : \beta_{1i} - \beta_{2i} = \theta, \text{ constant for all } i.$$

The number of coefficients under the unrestricted model, H_1 , is nk (the

β s). Under the restrictions in H_0 there are $nk - n + 1$ coefficients (the nk β 's less the n β_2 's plus θ). Therefore the number of parameter restrictions under H_0 is $n - 1$.

The LR test requires estimating the model, and evaluating the log likelihood, separately under H_1 and H_0 . We proceed as if Σ is known; it is estimated from the residuals of the unrestricted model. The unrestricted model is efficiently estimated as a set of ordinary least squares regressions. The restricted model can be estimated using a simple reparameterization as follows.

Write $\beta_{1i} = \beta_{2i} + (\beta_{1i} - \beta_{2i})$. Under H_0 , this becomes $\beta_{1i} = \beta_{2i} + \theta$.

Substituting into (2) and collecting terms in β_{2i} gives the restricted set of equations,

$$y_{it} = \theta x_{1t} + \beta_{2i}(x_{1t} + x_{2t}) + \sum_{h=3}^k \beta_{hi} x_{ht} + u_{it} \quad (15)$$

The system of equations (15) is a set of seemingly unrelated regressions with a common set of regressors (a multivariate regression) with cross-equation restrictions. It can be written, in matrix notation, as

$$\mathbf{y} = \theta \tilde{\mathbf{x}} + \tilde{\mathbf{Z}}\boldsymbol{\gamma} + \mathbf{u}, \quad (16)$$

where we define

$$\boldsymbol{\gamma} = \begin{pmatrix} \gamma_1 \\ \gamma_2 \\ \vdots \\ \gamma_n \end{pmatrix}, \boldsymbol{\gamma}_i = \begin{pmatrix} \beta_{2i} \\ \beta_{3i} \\ \vdots \\ \beta_{ki} \end{pmatrix}, i = 1, \dots, n, \text{ and } \tilde{\mathbf{Z}} = \begin{pmatrix} \tilde{\mathbf{X}} & 0 & \dots & 0 \\ 0 & \tilde{\mathbf{X}} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \tilde{\mathbf{X}} \end{pmatrix} = \mathbf{I}_n \otimes \tilde{\mathbf{X}},$$

$\tilde{\mathbf{X}}$ is the $T(k-1)$ matrix, $\tilde{\mathbf{X}} = [\mathbf{x}_1 + \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_k]$, where \mathbf{x}_i is the i th column of the design matrix \mathbf{X} ,

$$\text{and } \tilde{\mathbf{x}} = \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_1 \end{pmatrix} = \boldsymbol{\iota}_n \otimes \mathbf{x}_1, \text{ where } \boldsymbol{\iota}_n = \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix}.$$

We can rewrite the system

$$\mathbf{y} - \theta \tilde{\mathbf{x}} = \tilde{\mathbf{Z}}\boldsymbol{\gamma} + \mathbf{u}, \quad (17)$$

which can also be written,

$$\mathbf{y}_i - \theta \mathbf{x}_1 = \tilde{\mathbf{X}}\boldsymbol{\gamma}_i + \mathbf{u}_i, i = 1, \dots, n,$$

a system with common regressor set for given θ , and therefore the γ_i s are estimated efficiently by OLS applied separately to each equation.

This is equivalent to estimating (17) by ordinary least squares giving,

$$\hat{\gamma} = (\tilde{\mathbf{Z}}' \tilde{\mathbf{Z}})^{-1} \tilde{\mathbf{Z}}' (\mathbf{y} - \theta \tilde{\mathbf{x}}). \quad (18)$$

which is the maximum likelihood estimator of γ conditional on θ .

The likelihood function (12) is maximized when the generalized least squares criterion (13) is minimized. Write,

$$Q(\theta, \gamma) = \mathbf{u}' \boldsymbol{\Omega}^{-1} \mathbf{u} = (\mathbf{y} - \theta \tilde{\mathbf{x}} - \tilde{\mathbf{Z}} \gamma)' \boldsymbol{\Omega}^{-1} (\mathbf{y} - \theta \tilde{\mathbf{x}} - \tilde{\mathbf{Z}} \gamma).$$

Substituting (18) for γ in this, we concentrate the generalized sum of squared residuals function (13) with respect to γ and obtain a function of θ only, say $Q^*(\theta)$:

$$Q^*(\theta) = [\mathbf{y} - \theta \tilde{\mathbf{x}} - \tilde{\mathbf{Z}} (\tilde{\mathbf{Z}}' \tilde{\mathbf{Z}})^{-1} \tilde{\mathbf{Z}}' \mathbf{y} - \theta \tilde{\mathbf{Z}} (\tilde{\mathbf{Z}}' \tilde{\mathbf{Z}})^{-1} \tilde{\mathbf{Z}}' \tilde{\mathbf{x}}]' \boldsymbol{\Omega}^{-1} [\dots].$$

If we write $\mathbf{M} = \mathbf{I} - \tilde{\mathbf{Z}} (\tilde{\mathbf{Z}}' \tilde{\mathbf{Z}})^{-1} \tilde{\mathbf{Z}}'$, we can define $\mathbf{y}^r = \mathbf{M} \mathbf{y}$, and $\tilde{\mathbf{x}}^r = \mathbf{M} \tilde{\mathbf{x}}$ as residuals from least squares regressions on $\tilde{\mathbf{Z}}$.

Hence,

$$Q^*(\theta) = [\mathbf{y}^r - \theta \tilde{\mathbf{x}}^r]' \boldsymbol{\Omega}^{-1} [\mathbf{y}^r - \theta \tilde{\mathbf{x}}^r], \quad (19)$$

which is minimized by the maximum likelihood estimator.

The estimator of θ , which minimizes (19), and is also the ML estimator, is easily shown to be found by a bivariate generalized least squares regression of \mathbf{y}^r on $\tilde{\mathbf{x}}^r$, giving:

$$\hat{\theta} = \frac{\tilde{\mathbf{x}}^{r'} \boldsymbol{\Omega}^{-1} \mathbf{y}^r}{\tilde{\mathbf{x}}^{r'} \boldsymbol{\Omega}^{-1} \tilde{\mathbf{x}}^r}. \quad (20)$$

The LR criterion (14) is then simply evaluated.

4 Simulations

To test the proposed log-likelihood-ratio measure of heterogeneity, we applied it to simulated data generated from an fMRI dataset. Using synthetic data allowed us to: (1) systematically vary the spatial characteristics of the signal; (2) test the validity of the asymptotic chi-square distribution of the test statistic under different conditions (i.e., with different numbers of voxels and timepoints); and (3) investigate violations of the assumptions of the GLM, i.e., autocorrelation of error.

FMRI data was taken from a 32-year-old neurologically healthy male participant who provided informed consent according to local ethics procedures. Two-hundred and seventy-six echoplanar images of resting data (i.e., the participant lay in the scanner doing nothing) were acquired (i.e., 276 timepoints at each voxel) with a resolution of $2.19 \times 2.19 \times 4.0$ mm. The covariance matrix across neighboring voxels (chosen from within superior temporal cortex) was calculated. This covariance matrix was used to generate random noise sampled from a multivariate normal distribution.

4.1 Simulating spatial heterogeneity

In the first set of simulations, we demonstrate that the test statistic has power with respect to multivariate heterogeneity. A design matrix was constructed with three columns ($k=3$), consisting of two dichotomous variables x_{1t} and x_{2t} and a constant x_{3t} . At each timepoint either $x_{1t} = 1$ (and $x_{2t} = 0$) or $x_{2t} = 1$ (and $x_{1t} = 0$). For a third of timepoints, both $x_{1t} = 0$ and $x_{2t} = 0$. In the first case, seven timeseries (this corresponds to a central voxel and adjacent voxels) were generated by drawing 276 timepoints from a multivariate normal distribution with a real covariance matrix. The design matrix multiplied by

β values was added to the random noise. In the homogeneous case, all seven timeseries had the same signal (i.e., the betas associated with experimental condition one were either positive in all voxels or negative in all voxels). In the heterogenous case, half of the timeseries had positive β values for condition one and half had negative beta values. Note the overall signal was the same in all situations (i.e., the absolute difference between condition 1 and condition 2 was constant).

Figure 1 shows the results of 500 simulations with different covariance matrices and randomly generated timeseries. The x-axis represents the number of timeseries with homogeneous signals. As the number of timeseries with β_1 positive and β_2 negative increases the heterogeneity across the timeseries increases. This is reflected in the increase in the estimated rejection probability. As such, the test statistic reflects the degree of heterogeneity in the β coefficients across the timeseries.

4.2 Asymptotic χ^2 assumption

The χ^2 properties of the test statistic only hold asymptotically, and previous work on seemingly unrelated regression (Dufour and Khalaf 2000) suggest that

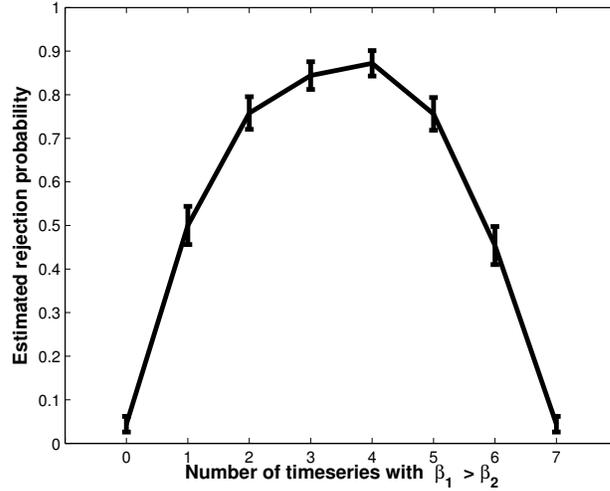


Figure 1: Varying heterogeneity across the seven timeseries. At the extreme values of the x-axis all timeseries (voxels) have the same β coefficients: the homogeneous cases. In the center of the x-axis, half the voxels have $\beta_1 > \beta_2$ and half the reverse. The y-axis is the log-likelihood ratio comparing the unrestricted and restricted models. The dotted line marks the χ^2 value of $p < 0.05$. Mean and standard error of the estimated rejection probability across 500 simulations are shown.

as the covariance matrix increases in size, this assumption becomes untenable except with very large numbers of timepoints. With smaller numbers of timepoints, overrejection of the null hypothesis becomes an issue.

We investigated the validity of the χ^2 assumption in two situations: 1) one with 7 timeseries (equivalent to a sphere of 1 voxel radius around a central voxel); and 2) one with 33 timeseries (equivalent to a sphere of 2 voxel radius). In both situations, multivariate random noise was generated with a real covariance structure. Since we are interested in the assessing the overrejection rate of estimated χ^2 values, no signal was inserted into the background noise. The number of timepoints was varied in from 50 to 1500, and a hundred simulations and models were fitted with each number of timepoints.

Figure 2 shows the proportion of models that are judged to be significant at a $p < 0.05$ significance level. Given the absence of signal, this value should be approximately 0.05. We see that for the seven timeseries case, there is not substantial overrejection of the null hypothesis, even with only 50 timepoints. However, for the larger 33 timeseries scenario, incorrect overrejection of the null hypothesis occurs frequently; with only 50 timepoints, rejection of the

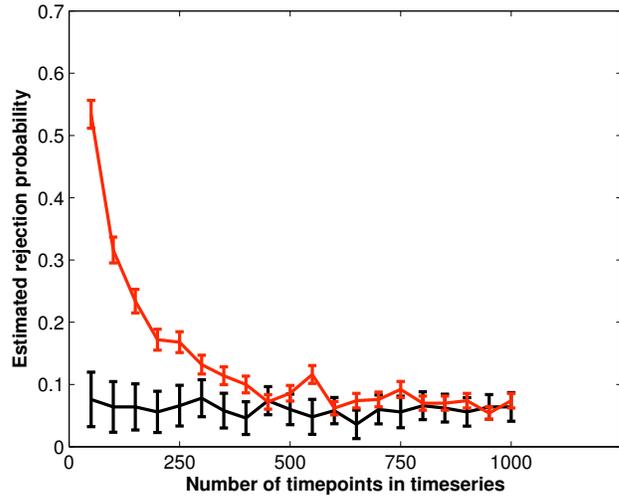


Figure 2: The proportion of models that incorrectly reject the null hypothesis assuming the difference in log likelihoods is χ^2 distributed for different numbers of timepoints. The red line is the case with 33 timeseries and the blue line with 7 timeseries. Values are the mean and standard error of 500 simulations. null occurs half the time. Only after approximately 500 timepoints does the rejection rate approximate that expected given the null hypothesis. This finding highlights the problem of accurately estimating large covariance matrices without large numbers of datapoints. In these situations, the log-likelihood ratio does not adequately approximate the χ^2 distribution.

4.3 Autocorrelation of the residuals

One common problem in applying the GLM to fMRI timeseries is the presence of autocorrelation in the residuals. In general, this can lead to underestimation of the error variance and inefficient model estimation. We conducted simulations to investigate the effects of error autocorrelation on the proposed log-likelihood ratio statistic. Fifteen-order autocorrelation estimates were calculated using standard analysis software `fsl` (Woolrich et al. 2001). Multivariate normally-distributed random data was then filtered using these autocorrelation estimates to simulate realistic autocorrelation in an fMRI dataset. Two situations were considered: one with a rapidly varying design matrix (with timepoints assigned to either condition one or condition two intermixed at random; equivalent to an event-related fMRI experimental design). The second was equivalent to a blocked fMRI design, with a continuous sequence of timepoints (one third of the total) corresponding to one condition, followed by another third corresponding to the other condition.

Figure 3 shows the performance of the blocked and the event-related experimental design in the presence of autocorrelation of the residuals. Inflation of type-II errors occurred for the blocked design; whereas, the event-related

design was similar (slightly more conservative) than a model with no autocorrelation. This result highlights that correction of autocorrelated noise is an important step for the proposed statistic (as for many fMRI analyses), especially for blocked experimental designs. Standard techniques for prewhitening (Woolrich et al. 2001) non-white noise can be used in conjunction with the statistics proposed here.

5 Discussion

In this paper, we have presented a measure of multivariate heterogeneity across multiple timeseries, that is designed to be able to assess the amount of multivoxel information available in functional MRI datasets. We extended the univariate GLM to a multivariate case, an example of seemingly unrelated regression, equivalent to a Mahalanobis distance measure previously used with fMRI. An unrestricted version of the multivariate GLM provides the maximum likelihood model across timeseries (allowing β values to be different for each timeseries). A restricted model assessed the homogeneous case that all timeseries had the same β coefficients (i.e., the maximum likelihood β coefficients that are constrained to be the same in all voxels). Comparing the restricted

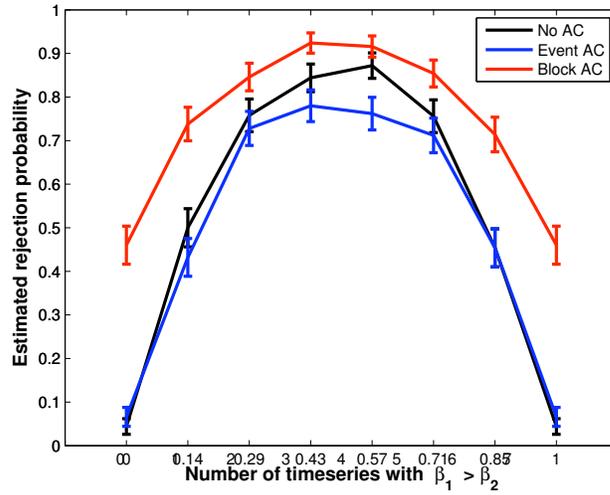


Figure 3: The effects of autocorrelated residuals on the log-likelihood ratio statistics. The three lines correspond to either the situation with no autocorrelation of the residuals, or else realistic error autocorrelation, but with two different design matrices corresponding to standard experimental paradigms. Values are the mean estimated rejection probability and standard error of 500 simulations.

and unrestricted cases measures the amount of heterogeneity across timeseries and is asymptotically χ^2 distributed.

Many recent and high-profile fMRI studies (Kay et al. 2008, Mitchell et al. 2008, Kriegeskorte et al. 2008) have taken advantage of multivariate pattern analysis techniques, however, until now, no specific statistical measure has been formulated to assess whether such approaches are picking up fine-scale (or distributed) spatial patterns across voxels. The measure of heterogeneity proposed here fulfils this role, providing a statistically well-grounded quantification of multivoxel spatial heterogeneity. This measure is intended to be applied in a searchlight approach, whereby patches of cortex (e.g., spheres of voxels) are considered in turn, and the spatial heterogeneity of the neural pattern mapped out. The statistic could also be used to consider heterogeneity across multiple peak voxels within a single cluster of activation, to assess whether it is truly a single cluster.

The spatial distribution of patterns of neural activation has important theoretical implications in its own right. Spatially-distributed patterns of neural activation suggest complex underlying neural processing of a given task.

Highly spatially homogeneous patterns of activation suggest underlying processing dedicated to a specific task or modality; this feeds into long-standing debates in cognitive neuroscience about regional neural specialization for a given task. The proposed heterogeneity statistic is also relevant for better understanding various neurological conditions. Changing spatial heterogeneity may be a hallmark of changing neural processing following neurological insult (e.g., stroke) or in Alzheimer’s disease and could be useful in diagnosis and monitoring of disease and recovery from disease. Spatial heterogeneity may provide a measure of redundancy in a given cortical region, which may in turn predict resilience from neurological insult. The proposed statistic is useful in conjunction with other multivariate analysis techniques, by providing evidence for useful heterogeneity across voxels that can be taken advantage of with, for example, subsequent pattern classification. Finally, the proposed technique, although designed for functional neuroimaging, may have applications in other domains where heterogeneity across timeseries (or equivalent) is important. For example, in economics, it may be important to assess whether different countries GDP responds to external events in a homogeneous way or not.

The proposed approach is computationally efficient and potentially powerful; however, as demonstrated in the simulations, there are several potential limitations to asymptotic use of the log-likelihood ratio measure of heterogeneity for inference. First, the presence of autocorrelation of the errors may result in overrejection of the null hypothesis and/or lack of sensitivity. However, well-known techniques (e.g., prewhitening) can be used to address this problem. The increase in the over-rejection rate when estimating larger covariance matrices is a more serious problem. For the 33 voxel case, the number of timepoints necessary for the log-likelihood ratio statistic to approximate the chi-square distribution is unrealistically large for many fMRI designs (that typically have fewer than 500 timepoints). This problem is known in the literature on multivariate regression with finite sample size and has been addressed using parametric bootstrapping to assess significance (Dufour and Khalaf 2000). An alternative approach is to use the log-likelihood ratio measure of heterogeneity as a surrogate descriptive statistic to make comparisons across multiple participants, rather than to make formal inference. For example, data from two groups (patients and controls) could be transformed into a standard space, and the log-likelihood ratio at a given location could be calculated for each participant and compared across groups using parametric or non-parametric

statistics. This approach is feasible for many fMRI designs and has been used previously for statistics whose distribution is not known (Filippini et al. 2009).

References

de Beeck, O. and Hans, P. (2009). Probing the mysterious underpinnings of multi-voxel fMRI analyses. *NeuroImage*.

de Beeck, O. and Hans, P. (2010). Against hyperacuity in brain reading: Spatial smoothing does not hurt multivariate fMRI analyses? *NeuroImage*, 49(3):1943–1948.

Dufour, J. and Khalaf, L. (2000). *Simulation-based finite and large sample tests in multivariate regressions*. Université de Montréal, Centre de recherche et développement en économique.

Filippini, N., MacIntosh, B., Hough, M., Goodwin, G., Frisoni, G., Smith, S., Matthews, P., Beckmann, C., and Mackay, C. (2009). Distinct patterns of brain activity in young carriers of the APOE- ϵ 4 allele. *Proceedings of the National Academy of Sciences*, 106(17):7209.

Friston, K., Holmes, A., Worsley, K., Poline, J., Frith, C., Frackowiak, R.,

- et al. (1994). Statistical parametric maps in functional imaging: a general linear approach. *Human Brain Mapping*, 2(4):189–210.
- Greene, W. (2003). *Econometric analysis*. prentice Hall Upper Saddle River, NJ.
- Haxby, J., Gobbini, M., Furey, M., Ishai, A., Schouten, J., and Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293(5539):2425–2430.
- Haynes, J., Sakai, K., Rees, G., Gilbert, S., Frith, C., and Passingham, R. (2007). Reading hidden intentions in the human brain. *Current Biology*, 17(4):323–328.
- Kamitani, Y. and Sawahata, Y. (2009). Spatial smoothing hurts localization but not information: Pitfalls for brain mappers. *NeuroImage*.
- Kamitani, Y. and Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, 8(5):679–685.
- Kay, K., Naselaris, T., Prenger, R., and Gallant, J. (2008). Identifying natural images from human brain activity. *Nature*, 452(7185):352–355.
- Kriegeskorte, N., Goebel, R., and Bandettini, P. (2006a). Information-based

- functional brain mapping. *Proceedings of the National Academy of Sciences*, 103(10):3863.
- Kriegeskorte, N., Goebel, R., and Bandettini, P. (2006b). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences*, 103(10):3863.
- Kriegeskorte, N., Mur, M., Ruff, D., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., and Bandettini, P. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, 60(6):1126–1141.
- Mahalanobis, P. (1936). On the generalized distance in statistics. 12(1):49–55.
- Mitchell, T., Shinkareva, S., Carlson, A., Chang, K., Malave, V., Mason, R., and Just, M. (2008). Predicting human brain activity associated with the meanings of nouns. *science*, 320(5880):1191.
- Rao, C. (1973). *Linear statistical inference and its applications*.
- Serences, J. and Boynton, G. (2007a). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron*, 55(2):301–312.

- Serences, J. and Boynton, G. (2007b). The representation of behavioral choice for motion in human visual cortex. *Journal of Neuroscience*, 27(47):12893.
- Stokes, M., Thompson, R., Cusack, R., and Duncan, J. (2009). Top-Down Activation of Shape-Specific Population Codes in Visual Cortex during Mental Imagery. *Journal of Neuroscience*, 29(5):1565.
- Woolrich, M., Ripley, B., Brady, M., and Smith, S. (2001). Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage*, 14(6):1370–1386.
- Zellner, A. (1962). An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association*, 57(298):348–368.