

The analysis of functional MRI time-series

Friston KJ¹, Jezzard P² & Turner R²

1. The Neurosciences Institute, La Jolla, CA 92037, USA

2, Laboratory of Cardiac Energetics, NHLBI, NIH, Bethesda MD 20892, USA

Address for correspondence

MRC Cyclotron Unit,

Hammersmith Hospital,

DuCane Road,

LONDON W12 0HS

Tel (011 44 81) 740 3162

Fax (011 44 81) 743 3987

Key words Functional MRI, Time-series, Statistical parametric mapping, significance, visual, cross-correlations, autocorrelations

Abstract

We present a method for detecting significant and regionally specific correlations between sensory input and the brain's physiological response, as measured with functional MRI. The method involves testing for correlations, between sensory input and the hemodynamic response, after convolving the sensory input with an estimate of the hemodynamic response function. This estimate is obtained without reference to any assumed input.

To lend the approach statistical validity, it is brought into the framework of statistical parametric mapping by using a measure of cross-correlations, between sensory input and hemodynamic response, that is valid in the presence of intrinsic autocorrelations. These autocorrelations are necessarily present, due to the hemodynamic response function or temporal point spread function.

Introduction

This article is about the analysis of functional MRI data obtained during activation studies of the human brain. The problem addressed is how to identify regionally specific and significant correlations between a time-dependent sensorimotor or cognitive parameter and measured changes in neurophysiology (Bandettini et al 1993a, Friston et al 1993a). This problem has a number of well established solutions in PET functional imaging, where activation studies are almost universally analyzed using some form of statistical parametric mapping (Friston et al 1990). However there are some special issues to be considered when dealing with MRI data.

Statistical parametric maps (SPMs) are images whose voxel values are distributed, under the null hypothesis, according to some known probability density function. The null hypothesis is usually that there are no significant activations, or correlations between sensorimotor parameters and central physiology. SPMs are treated as *smooth* multidimensional statistical processes and thresholded such that the probability of getting an activated region by chance (over the entire SPM) is suitably small (e.g. 0.05). The methods for determining the appropriate threshold have only recently been developed (Friston et al 1991 and Worsley et al 1992).

Two fundamental aspects of MRI data that bear directly on the detection of significant correlations are: (i) The hemodynamic response to sensory input (evoked changes in neuronal activity) is transient, delayed and dispersed in time. (ii) Because the sampling interval of some MRI techniques is typically much shorter than the time-constants of hemodynamic changes, the resulting time-series can show substantial autocorrelation.

Transient aspects of the hemodynamic response

The transient nature of the hemodynamic response is usually attributed to a physiological uncoupling of regional cerebral perfusion and oxygen metabolism (Fox et al 1988). The result is a

time-dependent (but uncharacterised) change in the relative amounts of venous oxy- and deoxy-hemoglobin. Due to the differential magnetic susceptibility of oxy- and deoxy-hemoglobin, a transient change in intra-voxel dephasing is observed. This change subtends the measured signal (Kwong et al 1992, Ogawa et al 1992, Bandettini et al 1992). The physiological mechanisms which mediate between neuronal activity and physiology, at the level of perfusion and cerebral metabolism, are known to have time-constants in the millisecond (ms) to seconds range. *In-vivo* optical imaging of microcirculatory events, in the visual cortex of monkeys, suggest the following sequence of activity-dependent changes: 200-400 ms after the onset of neuronal activity, highly localized oxygen delivery occurs, followed 300-400 ms later by an increase in blood volume. After 1000 ms a substantial rise in oxyhemoglobin is seen (Frostig et al 1990). MRI data concur with this latter observation, suggesting a rise in relative oxyhemoglobin, that is maximal after 4 - 10 seconds (Bandettini 1993b). The delay and dispersion associated with the hemodynamic response, to the onset of neuronal activity, is therefore substantial when compared to the sampling interval typical of fast (e.g. EPI) MRI times series (100 ms - 5 s). This is important because the correlations between evoked changes in neuronal activity and measured hemodynamics will be displaced and dispersed (smeared) in time. Simply correlating a sensory parameter (reflecting input at a neuronal level) with the hemodynamic response, will miss significant cross-correlations that are distributed in time according to the delay and dispersion of the *hemodynamic response function*. For example if the delay rendered a sinusoidal sensory input and the hemodynamic response out of phase by $\pi/2$, the correlation would be zero. The hemodynamic response function can be thought of as a *temporal point spread function* which not only smooths sensory input, but also applies a shift in time. In other words these functions describe the physiological response to a point (delta function or impulse) input, if one were able to present such a stimulus. Clearly to assess the true 'correlation', between a sensory parameter and hemodynamic response, the sensory parameter must first be subject to the same delay and dispersion as that mediating between neuronal

activity and hemodynamics. More formally, the correlation of interest is between the MRI time-series and the sensory input convolved with the hemodynamic response function. This is the first observation on which the proposed approach is predicated.

Stationary aspects of the hemodynamic response

Due to the dispersive nature of the response function, the hemodynamics at any voxel will be inherently smooth or auto-correlated in time. This autocorrelation will be evident even in the absence of evoked changes in neuronal activity. One can think of hemodynamics as the result of convolving a neuronal process with an effective hemodynamic response function, where the neuronal process comprises *evoked transients* and *intrinsic activity*. In mathematical terms this can be modeled as:

$$x(t) = \Lambda\{ \varepsilon(v) + z(t) \} \quad 1$$

Where $x(t)$ is the observed signal at a particular time (t), $\Lambda\{.\}$ is a convolution operator modeling the effect of the response function, $\varepsilon(v)$ is the evoked neuronal transient as a function of time from stimulus onset (v) and $z(t)$ is an *uncorrelated* neuronal process representing fast intrinsic neuronal dynamics. The reason for characterizing the intrinsic component in this way comes from the observation that autocorrelations in neuronal dynamics are most substantial in the 10 - 100 ms range (Aertsen and Preissl 1991, Nelson et al 1992) and are therefore very fast compared to hemodynamics. Clearly if neural dynamics showed slower autocorrelations, the effective hemodynamic response would embody some elements attributable to these slow neuronal changes. This model highlights the fact that observed autocorrelations have at least two components; (i) an evoked component due to the convolved neuronal transient that is phase-locked to the stimulus or task onset and (ii) intrinsic autocorrelations that result from intrinsic neuronal activity.

To the extent that the above model is true, the autocorrelations and response function are directly related. In particular when there are no evoked changes in neuronal activity (or any components phase-locked to stimulus onset have been removed) the intrinsic autocorrelations can be used to determine the response function. The relationship (between intrinsic autocorrelations and the effective response function) furnishes a way of estimating the response function directly from the physiological data that eschews any assumptions about the input or evoked neuronal activity. This is the second tenet upon which the proposed approach is based and is important because it ensures that there is no bias when testing these assumptions *post hoc*.

Finally temporal smoothness or autocorrelations are important from the point of view of detecting significant correlations between the input and observed response. Even if no neuronal response is evoked, the presence of intrinsic autocorrelations will increase the probability of a spurious and high correlation coefficient. This follows from the fact that the *effective degrees of freedom* associated with a measured correlation will be smaller than the number of time points or scans. In simple terms, for two time-series of fixed length, the probability of obtaining a high correlation coefficient, *by chance*, increases with smoothness. When assessing the significance of cross-correlations between two time-series it is therefore necessary to explicitly account for the autocorrelations expected under the null hypothesis (in this case the intrinsic autocorrelations). This is the third tenet of our approach.

The stationariness of the intrinsic autocorrelations (and implicitly the response function) referred to above are in time. Stationariness in time does not imply stationariness in space. In other words it is possible for the response function and the autocorrelative behaviour of hemodynamics to vary from region to region, or voxel to voxel.

The relationship between intrinsic autocorrelations (both spatial and temporal) and the effective degrees of freedom is a very general one and affects the analysis of all functional imaging data,

using any form of statistical parametric mapping. In what follows this theme occurs twice (i) in deriving a statistical quotient which tests for significant temporal cross-correlations in the presence of intrinsic autocorrelations and (ii) thresholding the resulting statistical parametric maps, which are spatially auto-correlated.

We present below (i) the theoretical aspects of assessing the significance of correlations between sensory (cognitive or motor) parameters and hemodynamic responses measured with MRI, (ii) an application to real data and (iii) an evaluation of the hemodynamic response function used in the preceding two sections.

Theory

The objective, of the described approach, is to produce a statistical parametric map which can be thresholded to identify significant correlations between a given input parameter and regional hemodynamics. This requires; (i) a statistical parameter that tests the significance of the above correlation and (ii) a technique for thresholding the SPM which renders the probability of identifying a significant region by chance suitably small (e.g. 0.05). The latter techniques have already been established in the context of functional imaging, using the theory of stationary Gaussian processes with known or measurable autocorrelation (Friston et al 1991). The problem therefore reduces to defining a statistic that tests for the presence of significant cross-correlations in the presence of intrinsic autocorrelations.

The following relies on the theory of stationary stochastic processes and, in particular, their representation in the frequency domain. The interested reader will find an excellent review of the important standard results (a number of which are used below) in Cox and Miller (1980, pages 309 - 337). Covariance refers to the average or expectation of the product of two processes x and y ($E\{x(t).y(t)\}$), after normalization to zero mean. Correlations refer to the same thing but when the processes are normalized to unit variance. Correlations can be thought of as normalized covariances. Cross-correlation $\rho_{xy}(\tau)$ and cross-covariance functions $\gamma_{xy}(\tau)$ are functions which represent the magnitude of correlation or covariance, as a function of the temporal displacement τ (lag) between the processes $x(t)$ and $y(t)$ ($E\{x(t).y(t + \tau)\}$). Autocorrelation and auto-covariance functions, $\rho_{xx}(\tau)$ and $\gamma_{xx}(\tau)$ are simply the correlations or covariances of a process with itself at some time τ later ($E\{x(t).x(t + \tau)\}$). Clearly $\gamma_{xx}(0)$ is the variance of $x(t)$.

The hemodynamic response function

In what follows the expressions are in continuous time and assume that real fMRI time-series are reasonably approximated by continuous space-time processes. Let the observed MRI time-

series, at any point in the brain, be modeled by a function of time $x(t)$ and let a time-dependent sensorimotor or cognitive parameter of interest be similarly modeled by $c(t)$. We will refer to this parameter as the *contrast*. In analyses of (co-)variance the 'contrast' is used to test for a specific profile of time-dependent changes. Both here and in general the contrast has zero mean and unit variance. The effects of delay and dispersion of the hemodynamic response are modeled by a convolution operator $\Lambda\{.\}$. e.g:

$$\Lambda\{c(t)\} = \int c(t - \tau)h\{\tau\} d\tau \quad 2$$

The function $h\{\tau\}$ is the (impulse) response function and has an equivalent representation in frequency space ($\omega = 2\pi f$ where f is frequency) called the transfer function $H(\omega)$: The transfer function is simply the Fourier transform of the response function $h\{\tau\}$. The particular form of $h\{\tau\}$ reflects our assumptions (or knowledge) about the effective hemodynamic response function. It is a function of delay τ . The average or expected delay is simply the expectation associated with $h\{\tau\}$. Similarly the dispersion is the variance of $h\{\tau\}$. In what follows we assume a Poisson form for $h\{\tau\}$:

$$h\{\tau\} = \frac{\lambda^\tau e^{-\lambda}}{\tau!} \quad 3$$

The Poisson distribution is a parsimonious choice because its mean and variance are equal and are defined by a single parameter λ . It should be born in mind that the values of τ in eqn(3) are real positive integers. In what follows τ has units of seconds (if the repeat time is not an integer valued number of seconds than eqn(3) could be replaced with a gamma distribution with appropriate moments). It is of course possible to use forms of $h\{\tau\}$ that do not assume a

relationship between delay and dispersion (see discussion). We validate the choice of a Poisson form in a subsequent section.

To account for the effects of delay and dispersion on the cross-correlations between $x(t)$ and $c(t)$ we apply $\Lambda\{\cdot\}$ to $c(t)$ and examine the correlation between $x(t)$ and $\Lambda\{c(t)\}$. The contrast $c(t)$ convolved with the hemodynamic response function $h\{\tau\}$ represents a new contrast, that has been corrected for delay and dispersive effects, and will be referred to as $y(t)$. It should be noted that the (power) spectral density of $y(t)$ [$g_y(\omega)$] is simply related to the spectral density of $c(t)$ [$g_c(\omega)$] and the transfer function $H(\omega)$ (Cox and Miller 1980):

$$g_y(\omega) = |H(\omega)|^2 g_c(\omega) \quad 4$$

Testing for significant cross-correlations in the presence of intrinsic autocorrelations

In this section we observe that $\gamma_{xy}(0)$, the covariance of interest has, approximately, a Gaussian distribution. The variance of this distribution (over many realizations at one voxel or many voxels in one realization) is the same as the variance of $\gamma_{xy}(\tau)$ over time τ . The variance of $\gamma_{xy}(\tau)$ is its spectral density, integrated over frequencies. Because $\gamma_{xy}(\tau)$ is effectively the contrast, convolved with the response function, convolved with the MRI time-series, its spectral density, *under the null hypothesis*, is simply the three way product of the spectral densities of $c(t)$, $h(\tau)$ and $x(t)$. The spectral density of $x(t)$, under the null hypothesis of no evoked neuronal activity, is specified by the intrinsic autocorrelations. Given that these spectral densities are either known or can be estimated the distribution of $\gamma_{xy}(\tau)$, under the null hypothesis, can be determined. These arguments are now presented more formally:

The cross-covariance between the MRI time-series $x(t)$ and the convolved contrast $y\{t\}$ will itself be a stationary stochastic process with zero mean. Furthermore because $\gamma_{xy}(\tau)$ is obtained by

convolving $x(t)$ with $y(t)$, $\gamma_{xy}(\tau)$ represents a weighted sum of $x(t)$ and, by central limit theorem, will tend to a Gaussian distribution [this assumption is only strictly true when $c(t)$ has non-zero values over an interval which is large compared to the width of $\gamma_{xx}(\tau)$]. The variance of $\gamma_{xy}(\tau)$ is:

$$\text{var}\{\gamma_{xy}(\tau)\} = \int_{-\infty}^{\infty} g_{xy}(\omega) d\omega \quad 5$$

where $g_{xy}(\omega)$ is the cross-spectral density of processes $x(t)$ and $y(t)$ (Cox and Miller 1980). $g_{xy}(\omega)$ can be interpreted as the second order probabilistic characteristics (variance) of the orthogonal process $s_{xy}(\omega)$, in the frequency domain, that is equivalent to $\gamma_{xy}(\tau)$ in the time domain, where (c.f. Parseval's Theorem):

$$\begin{aligned} g_{xy}(\omega) &= E\{s_{xy}(\omega)s_{xy}(\omega)^*\} \\ \text{and } s_{xy}(\omega) &= s_x(\omega)s_y(\omega) \end{aligned} \quad 6$$

* denotes complex conjugate and $s_x(\omega)$ and $s_y(\omega)$ are the spectral representations of processes $x(t)$ and $y(t)$. Let $[u] = u.u^*$. Now if $s_x(\omega)$ and $s_y(\omega)$ are independent (equivalently if $x(t)$ and $y(t)$ are independent):

$$g_{xy}(\omega) = E\{[s_{xy}(\omega)]\} = E\{[s_x(\omega)]\}.E\{[s_y(\omega)]\} = g_x(\omega).g_y(\omega) \quad 7$$

(c.f. Theorem 6F in Grimmett & Welsh 1986). Eqn(7) expresses the fact that the cross-spectral density is the product of the densities of the two underlying processes, given that they are independent (i.e. there is no systematic phase relationship). Combining eqn(5) and eqn(7):

$$\text{var}\{\gamma_{xy}(\tau)\} = \int_{-\infty}^{\infty} g_x(\omega).g_y(\omega).d\omega \quad 8$$

Eqn(8) expresses the variance of the cross-covariance function as an integral (sum) over all frequencies. The variance at each frequency, under the null hypothesis, is the product of the variances (spectral densities) of the underlying processes.

The stationariness of $\gamma_{xy}(\tau)$ implies that the variance of $\gamma_{xy}(\tau)$ is not a function of time and therefore $V\{\gamma_{xy}(0)\}$ (over multiple realizations) = $V\{\gamma_{xy}(\tau)\}$ (over time). So the quotient $\zeta(0)$

$$\begin{aligned}\zeta(0) &= \gamma_{xy}(0) \cdot \text{var}\{\gamma_{xy}(\tau)\}^{-1/2} \\ &= \rho_{xy}(0) \cdot \left(\frac{\int g_x(\omega) d\omega \cdot \int g_y(\omega) d\omega}{\int g_x(\omega) g_y(\omega) d\omega} \right)^{1/2} \\ &= \rho_{xy}(0) \cdot v^{1/2}\end{aligned}\tag{9}$$

will, under the null hypothesis, have (roughly) a Gaussian distribution of unit variance and zero mean. This is also called the z-score and a statistical parametric map of $\zeta(0)$ is referred to as a SPM{z}. $\zeta(0)$ can be thought of as the *correlation* between x(t) and y(t) times the square root of the effective degrees of freedom (v).

In order to compute v one needs to estimate $g_y(\omega)$ and $g_x(\omega)$. The estimation of $g_y(\omega)$ is straightforward, however the estimation of $g_x(\omega)$ due to intrinsic autocorrelations ($g_x(\omega)$ under the null hypothesis) requires a little more thought:

Estimating $g_y(\omega)$, $g_x(\omega)$ and the hemodynamic response function $h(\tau)$

Clearly if c(t) and h(τ) were sufficiently well behaved $g_y(\omega)$ could be determined analytically as in eqn(4) [using $g_c(\omega)$ and H(ω)]. However in practice it is considerably easier to simply compute the spectral density of y(t) [after convolving c(t) with h(τ)] using fast Fourier transforms and thereby estimate $g_y(\omega)$ directly.

The estimation of $g_x(\omega)$, the spectral density of the MRI time-series, under the null hypothesis, is more problematic, given (i) the short and noisy time-series that are typically available and (ii) the null hypothesis is (hopefully) seldom true and the autocorrelations (and spectral densities) of $x(t)$ will contain both evoked and intrinsic components. To obtain an estimate $g_x(\omega)$ *under the null hypothesis* it is necessary to discount evoked contributions. This is simply achieved by removing components that are phase-locked to the stimulus onset. In our work this is done by analyzing the time-series from one stimulus-cycle minus the time-series from another, namely the *phase-shifted differences*. A related device, called the 'shift predictor', is found in the analysis of separable neuronal spike-trains recorded with multi-unit electrodes. The shift predictor is used to remove the confounding effect of stimulus-locked firing, when assessing the *cross-correlations* between two units. In the present case we wish to remove confounding stimulus-locked effects on estimated *autocorrelations*. In what follows $x_p(t)$ will refer to the difference in (p) phase-shifted MRI time-series. From eqn(1) and eqn(2):

$$x_p(t) = x(t) - x(t + q.p) = \Lambda\{z(t) - z(t + q.p)\} \quad 10$$

where q is the stimulus or task period (or cycle length) and p is an integer greater than 0. Because $z(t)$ is modeled as an uncorrelated stochastic process (innovation) the spectral and auto-correlative properties of $x_p(t)$ correspond to those of $x(t)$ in the absence of evoked transients. Equivalently, under the null hypothesis $\rho_{xx}(h) = \rho_{x_px_p}(h)$ and $g_x(\omega)/\text{var}\{x(t)\} = g_{x_p}(\omega)/\text{var}\{x_p(t)\}$.

Although it is possible to estimate $g_{x_p}(\omega)$ by subjecting $x_p(t)$ to Fourier transformation, we use the following expedient method. If one assumes that the intrinsic autocorrelations [$\rho_{x_px_p}(h)$] can be modeled by a Gaussian autocorrelation function of the form:

$$\rho_{xx}(\tau) = \rho_{x_px_p}(\tau) = e^{-\tau^2/4\sigma^2}$$

$$\text{then: } g_x(\omega) \propto g_{x_p}(\omega) \propto e^{-\omega^2 \sigma_t^2} \quad 11$$

(Cox and Miller 1980). The constants of proportionality implicit in eqn(11), cancel in eqn(9) and do not affect v . The parameter σ_t reflects the intrinsic smoothness of $x(t)$ and is estimated in a computationally efficient way using the equality:

$$\sigma_t \approx \text{var}\{x_p(t)\} \cdot [2 \text{var}\{dx_p/dt\}]^{-1/2} \quad 12$$

This relationship follows from eqn(11) and the standard result: $V\{dx/dt\} = -\gamma_{xx}''(0)$ (Cox and Miller 1980). See Friston et al (1991) for a fuller discussion. It should be noted that whatever phase-shifting strategy is employed $x_p(t)$ will always be shorter than $x(t)$. From a practical point of view this means that the low frequency components of $g_x(\omega)$ cannot be estimated directly. Modeling $g_x(\omega)$ with a Gaussian (or any other) distribution provides for an estimate of these low components.

A process with auto-covariance given by eqn(11) is obtained by convolving a random uncorrelated process with a Gaussian point spread function of parameter σ_t . If one were to assume [see eqn(10)] that the MRI time-series was a result of convolving an unknown but uncorrelated underlying neuronal process with the temporal point spread function $h(\tau)$, then the estimate of σ_t also specifies λ (the delay or dispersion). This follows from the fact that the Gaussian and Poisson functions are very similar for large values of λ (they are identical in the limit) where:

$$\lambda \approx \sigma_t^2 \quad 13$$

The empirically derived value of σ_t , based on the data set described below, was 0.92 scans. This corresponds to a delay and dispersion (λ) of 7.69 seconds, which falls comfortably in the reported ranges (Bandettini 1993b). It should be noted that there is no direct information about delay (as opposed to dispersion) in σ_t , the reason that our hemodynamic response function can be entirely specified by σ_t depends on the Poisson assumption implicit in eqn(3).

To summarise this section; a single parameter σ_t can be estimated from the ratio of the variances of the time-series and its derivative. This parameter (temporal smoothness) not only provides a very efficient and complete estimate of the spectral density of the time-series (assuming the density is Gaussian) but it can also be used to fix the parameter of the response function (assuming the function is Poisson). This single parameter is all that one needs to estimate, in calculating the effective degrees of freedom (ν) for the correlation of interest.

Thresholding the SPM{z}

The final stage of analysis involves identifying a threshold which protects against experiment-wise false positives. The key thing to note here is that due to the spatial autocorrelations in the SPM{z} the number of voxels that are expected to be above a threshold (u) is far more than the number of regions, (or maxima) above threshold. As we are concerned with activated regions (connected subsets of the total excursion set $\equiv \zeta(0) \geq u$) and not voxels, a Bonferroni correction is not appropriate (a Bonferroni correction is based on the expected number of voxels). Current approaches (Friston et al 1991; Worsley et al 1992) use the fact that the relationship between the *probability of getting at least one region above u* (P) and the expected number of regions (or maxima) (μ), tends to equality as u gets very big. The expected number of maxima depends on the threshold (u), smoothness (σ_s) and size of the SPM (S). For an SPM{z} of D dimensions:

$$P \leq \mu \approx S (2)^{-(D+1)/2} (2\sigma_s^2)^{-D/2} u^{D-1} e^{-u^2/2} \quad 14$$

(Hasofer 1978, Adler 1981). σ_s is the spatial smoothness of the SPM and is estimated according to eqn(12) but using partial derivatives *in space* over rows and columns. The appropriate threshold is chosen such that $P \approx 0.05$. See Friston et al (1991) for a derivation of an equivalent equation (directly in terms of regions) in the context of functional imaging. A related approach based on the Euler characteristic has also been devised (Worsley et 1992). It is of course acceptable to present SPMs thresholded without this correction for multiple non-independent comparisons. In this case the SPM is referred to as 'descriptive' and an uncorrected one-tailed threshold of $p=0.001$ is usually recommended (i.e. 3.09 for $SPM\{z\}$). It should be noted that the thresholding and interpretation of SPMs is a branch of spatial statistics that is receiving a lot of attention. The thresholding described in this paper represents a simple and established approach, but is not necessarily the most powerful. See Friston et al (submitted) for further a more current analysis.

The theories presented above used processes in continuous time and relied on a number of approximations (usually exact in the limit). In order to assess the robustness of the expressions presented here we applied the results to a simulated MRI data set.

Validation using simulated data

The parameters of the simulated data were chosen to correspond to the real MRI data analyzed in the next section. 60 images of 32×32 voxels were created by convolving an uncorrelated random number field with a 3-dimensional Gaussian kernel of parameter 1.4 voxels, within the images, and 0.9 over time. The data were phase-shifted and subtracted, assuming a stimulus-cycle period of 20 scans. In the example presented here, and in those below, the differences between all possible pairs of stimulus-cycles were analyzed and the results averaged over all pairs. According to eqn(12) the estimate of temporal smoothness was 0.905 scans and compares well with the

expected value of 0.9. Assuming a repeat time of 3 seconds, these simulated data would have a temporal smoothness (σ_t) of 2.71 seconds. This degree of smoothness would be produced by a temporal point spread function of Poisson parameter 7.37 seconds. $\zeta(\tau)$ was computed, according to eqn(9), for all voxels, using a square wave contrast $c(t)$ of period 20 scans (10 scans 'on' and 10 scans 'off') and a Poisson response function with $\lambda = 7.37$ seconds. The resulting distribution of $\zeta(\tau)$ was determined by pooling across voxels and lag (τ). This simulated distribution was compared to the theoretically predicted Gaussian distribution.

The results of this simulation are shown in figure 1. The response function is shown (top left) with the resulting contrast before $[c(t)]$ and after $[y(t)]$ convolution with $h(\tau)$ (top right). The agreement between the empirical and theoretical distributions of $\zeta(\tau)$ is clearly evident (lower left). The corresponding SPM $\{z\}$ of $\zeta(0)$ is shown on the lower right. The highest value of $\zeta(0)$ obtained in this realization was 3.57. The time-series from the corresponding voxel is displayed with the contrasts (dotted line - top right). It is important to note that these high values (and apparent 'following' of the contrast) can easily occur by chance and should not be over-interpreted in real data. No value exceeded the corrected threshold (broken vertical line - bottom left), which was 3.79 according to eqn(14) and the *post hoc* estimate of spatial smoothness [eqn(12)]. This estimated spatial smoothness (averaged over both dimensions) was 1.45 voxels and compares favourably with the expected value of 1.4 voxels.

Analysis of MRI time-series

Data Acquisition

The data used to illustrate the method were a times series of 64 gradient-echo EPI single coronal slices (5mm thick, 64x64 voxels) through the calcarine sulcus and extrastriate areas. Images were obtained every 3 seconds from a single subject using a 4.0T whole body system, fitted with a small (27cm diameter) z-gradient coil (TE 25ms, acquisition time 41ms). Photic stimulation (at 16 Hz) was provided by goggles fitted with light emitting diodes. The stimulation was off for the first 10 scans (30 seconds), on for the second 10, off for the third, and so on. Images were reconstructed without phase correction. The data were interpolated from 64x64 voxels to 128x128 voxels. Each voxel thus represented 1.25 x 1.25 x 5 mm of cerebral tissue.

Data Pre-processing

Image manipulations and data analysis was performed in Matlab (MathWorks Inc. Sherborn MD, USA). The first 4 scans were removed to eschew magnetic saturation effects. The scans were corrected for (slight) subject movement using nonlinear minimisation and a computationally efficient cubic interpolation algorithm (Keys 1981). Images were translated and rotated to minimize the sum of squares between each of the 64 images and their average (both scaled to the same mean intensity) using the Levenberg-Marquardt method (More 1977). Only the 36 x 60 voxel subpartitions (of the original images) containing the brain were subject to further analysis. All the time-series, for each voxel, were normalized to a mean of zero and convolved, in the time domain, with a Gaussian filter (FWHM = 1.5 voxels) to suppress thermal noise.

Computing $\zeta(0)$

The real data were treated in an identical fashion to the simulated data. The temporal smoothness (σ_τ) was estimated using eqn(12) to be 0.92 as mentioned above. This estimate was the mean over all intracerebral voxels and corresponds to a value of 7.69 seconds for λ , the parameter of our assumed Poisson response function $h(\tau)$. In order to demonstrate that modeling intrinsic autocorrelations with a Gaussian form is reasonable, the empirical intrinsic autocorrelation functions and corresponding spectral density [$g_x(\omega)$] are shown in Figure 2. These functions were obtained by averaging estimates over the differences between all pairs of 20 scan stimulation periods (phase-shifted differences). The data were windowed with a Hanning function before being subject to Fourier transform. The solid lines in Figure 2 correspond to the autocorrelations (and spectral density) that would be obtained by convolving an uncorrelated innovation with a Poisson function of parameter 7.69 seconds. It should be noted that the determination of the appropriate Poisson parameter does not require computation of either the autocorrelation function or spectral density, but uses the variance of the first derivatives as in eqn(12). These data are presented only to validate the Gaussian approximation for $g_x(\omega)$.

The contrast $c(t)$ used in this example was, again, a square wave which was positive during photic stimulation and negative otherwise. $c(t)$ had zero mean and unit variance. The measured correlation between $y(t)$ and the MRI time-series $x(t)$ was computed for all voxels and scaled by \sqrt{v} according to eqn(9) to give $\zeta(0)$. \sqrt{v} was 4.58, corresponding to 21 effective degrees of freedom. In other words one accrues a 'degree of freedom' every 2.85 scans. The reason v is the same for all voxels is that we assumed the response function and intrinsic autocorrelations (and spectral density) were the same everywhere. Clearly one does not have to assume this (see discussion).

The results of this analysis are seen in figure 3. The hemodynamic response function, based on the empirically determined 'smoothness' of the time-series is displayed to the left of the contrast before and after convolution. The dotted line in the top right graph is the response of the voxel showing the greatest correlation. The distribution of $\zeta(0)$ (bottom left) is dramatically skewed,

with almost all voxels showing a marked positive correlation with the convolved contrast. The highest value of $\zeta(0)$ was 4.44. The corresponding SPM{z} is displayed on the bottom right.

Thresholding the SPM

The spatial smoothness (σ_s) was computed according to eqn(12) over all the rows and columns of the SPM{z}. The means of these estimates was 1.46 voxels or 1.81 mm. With a search volume (S) of 2160 voxels, and a P value of 0.05, the threshold according to eqn(14) was 3.97. Figure 4 shows the results of this thresholding. The picture on the left makes the point that the number of voxels above threshold far exceeds the number of regions. The reason for this, is of course, spatial smoothness in the underlying process. This discrepancy between the expected number of supra-threshold voxels and the expected number of regions, in the excursion set, means that a Bonferroni correction based on *voxel* expectation is not appropriate. This follows from the fact that one requires the expected number of *regions* to be 0.05. The striate and extra-striate regions constituting the excursion set can all be considered significant in a strict statistical sense. They include V1 bilaterally and (on the right) V2 dorsal to the calcarine fissure. Extrastriate regions in the fusiform, lingual and dorsolateral cortices are also evident.

A null analysis

In order to demonstrate the specificity of the approach we repeated an identical analysis but changed the period of the contrast $c(t)$ from 20 scans to 16 scans. This simple manipulation eliminated all the significant voxels. In fact the highest value of $\zeta(0)$ was only 2.719. The results of this analysis are shown in figure 5 using the same format as figures 1 and 3.

Validation of the hemodynamic response function

Validation in the time Domain

Thus far it has been assumed that the hemodynamic response function $h(\tau)$ is reasonably 'correct' and we have proceeded to assess the ability of the contrast $c(t)$ to explain the observed changes. One can take the completely opposite approach and estimate $h(\tau)$ by assuming $c(t)$ is correct. In previous sections we have assumed a Poisson form for $h(\tau)$ and have used the empirically determined temporal smoothness to specify an appropriate parameter. In this section we present a *post hoc* validation of $h(\tau)$ by computing it directly; using the contrast $c(t)$ as the input function and the brain's response as the output function. Although this could be done at each voxel, a more robust estimate is obtained if some global response is used. This global or distributed output function was the response of the first spatial mode (principal component or eigenimage) based on the *between voxel* covariances.

The first spatial mode or eigenimage is defined by the first eigenvector of the covariance matrix of the MRI time-series. If \mathbf{x} corresponds to a matrix with 60 rows (one for each scan) and 2160 columns (one for each voxel) of the time-series for each voxel then the eigenvector solution \mathbf{e} is simply computed using singular value decomposition (Golub and Van Loan 1991) such that:

$$\mathbf{x}^T \mathbf{x} \cdot \mathbf{e} = \mathbf{e} \cdot \mathbf{\Omega} \quad 15$$

where τ denotes transposition and $\mathbf{\Omega}$ is a diagonal matrix of eigenvalues. The column of \mathbf{e} corresponding to the largest eigenvalue is the first spatial mode or eigenimage \mathbf{e}_1 . This mode constitutes a distributed brain system with high intra-covariance or functional connectivity. See Friston et al (1993b) for a full discussion. The response of \mathbf{e}_1 is given by:

$$\mathbf{r} = \mathbf{x} \cdot \mathbf{e}_1 \quad 16$$

where \mathbf{r} is a column vector of length 60. Using least squares deconvolution one can directly estimate the response function that best transforms the input $c(t)$ to output \mathbf{r} . The results of this deconvolution are shown in figure 6, where the 'memory' of the convolution has been restricted to 10 scans. The first eigenimage or mode (top left) reveals high loadings in the same areas identified by the more conventional (hypothesis led) analysis of the previous sections. The corresponding response is shown with the contrast (top right). The results of the least squares deconvolution (dots) agree remarkably with the hemodynamic response function based only on the temporal autocorrelations (broken line). The predicted and actual responses of the first mode are compared on the lower right. The predicted time courses are simply the contrast convolved with $h(\tau)$ (broken line) and the empirically determined response function (solid line). It is apparent that the empirical response function has a slightly more protracted tail with some hint of biphasic structure. One might conjecture that this feature may reflect macrovascular effects as remote regions are drained.

Validation in the frequency domain

The hemodynamic response function 'behaves essentially as a low pass filter in the detection of neuronal firing' (Bandettini 1993b). In other words higher frequency inputs are more severely attenuated than low frequency inputs. The dependency of output variance on input frequency is captured by the (squared) transfer function $|H(\omega)|^2$ seen in Figure 7 (left). This example is for a Poisson parameter of 5 seconds and demonstrates the downwards modulation of output variance with increasing frequency of a periodic input. For example the variance of a periodic input at 0.1 cycles per second (a stimulus cycle length of 10 seconds) is attenuated to about 20%. The hemodynamic response function can provide quite a powerful prediction of the brain's capacity to follow a periodic stimulus when presented at increasing frequencies. Bandettini et al have

examined this phenomenon using fast self-paced sequential finger-thumb opposition at different on-off switching frequencies, ranging from period lengths of 50 (e.g. 25 seconds of finger opposition and 25 at rest) to 2 seconds. The data reported in Bandettini (1993b) are in terms of relative activation or amplitude of response. These data are in excellent agreement with theoretical predictions based on a Poisson hemodynamic response function (Figure 7 - left). The theoretical values were derived by simply convolving a square wave contrast, of appropriate frequency, with a Poisson response function (parameter 5 seconds) and taking the maximum amplitude difference. The Poisson parameter used here (5 seconds) is smaller than that calculated on the basis of our photic stimulation data and may reflect a number of differences, including: (i) differences in the neuronal response to varying task conditions (ii) differences due to relative microvascular and macrovascular contributions to the hemodynamic response (secondary to different acquisition parameters) or (ii) differences in motor, versus visual cortex, hemodynamic response.

Validation in terms of phase shifts

A simple and powerful way to estimate the delay associated with the hemodynamic response is to compute the phase shifts of the largest Fourier component of the input sequence at each voxel. The results of this analysis are shown in Figure 8, as a distribution of phase differences between the contrast $[c(t)]$ and the signal $[x(t)]$, over all intracranial voxels. It is seen that distribution is broad, but centered on the delay estimated on the basis of temporal smoothness alone. Anecdotally, it appears that the longer delays (greater phase-differences) are characteristic of regions that would be more subject to macrovascular effects (see next section).

The effects of changing λ

In this final section we present a brief analysis of the effects of changing λ . To assess the sensitivity, of detecting significant correlations, to changes in the parameter λ , we repeated the

analysis based on $\zeta(0)$ while systematically increasing λ from 0 to 16 seconds in 0.1 second steps. We used the 'exceedence proportion' as an index of the overall or 'omnibus' sensitivity to significant correlations. The exceedence proportion is the proportion of voxels in the SPM $\{z\}$ that are above threshold. This exceedence proportion has been used for some years as a rough guide to the omnibus significance of SPMs (Friston et al 1990) and has recently been the subject of some rigorous statistical analysis (Worsley & Vandal - submitted).

The results of this analysis are presented in figure 9 (top). The first thing to note is that the greatest exceedence proportion is obtained for values of λ in the range 6 - 8 seconds. The value of used in previous section (7.69) is marginally suboptimal (the best value was 6.6 seconds). One explanation for this was found in a closer analysis of these results:

If the optimal value of λ is computed for each voxel, one observes a systematic and spatially ordered change as one moves from striate cortex to extrastriate cortex, with the extrastriate cortices preferring a *shorter delay and dispersion*. In other words some regions of occipital cortex respond in a more acute and phasic way, to the same sensory input, than do others. Figure 8 (lower right) shows these voxel-specific optimal values of λ , for all voxels that showed a reasonable correlation with the convolved contrast (non corrected threshold of $p = 0.001$ - lower left). This result is not surprising given the reported evidence for phase shifts in response profiles (of several seconds) in MRI data (Bandettini 1993b). What is surprising here, is that the extrastriate regions appear to respond *before* primary visual cortex does. There are obviously many potential neuronal and hemodynamic explanations for this (e.g. slice reentrant large draining vessels), however one interesting possibility is that V1 responses are modulated by reentrant connections (Edelman 1978) from extrastriate regions. This particular aspect of inter-regional dynamics will be addressed in a subsequent paper.

In conclusion, based on the exceedence proportion, the approach appears robust as long as λ is within a second or so of the real delay and dispersion.

Discussion

We have described what we consider to be a reasonable characterization of MRI time-series in terms of regionally specific correlations between a sensorimotor or cognitive input and the hemodynamic response. The approach is based on the correlations between the MRI signal and the input sequence convolved with an estimate of the hemodynamic response function. The hemodynamic response function (in this paper) was assumed to be Poisson and was determined using the observed autocorrelations in the physiological data (having removed stimulus locked components). In order to give this approach statistical validity it was brought into the framework of statistical parametric mapping by (i) using a statistical parameter that tests for significant cross-correlations (between sensory input and physiological response) in the presence of intrinsic autocorrelations and (ii) using established thresholding techniques to render the probability of finding a significant region, by chance, suitably small (0.05). In particular the proposed statistic $\zeta(0)$ has a Gaussian distribution under the null hypothesis and the resulting map of $\zeta(0)$ constitutes a SPM $\{z\}$. $\zeta(0)$ is simply the correlation between convolved sensory input and hemodynamic response, scaled by a factor which can be thought of as the square root of the effective degrees of freedom v . v is a function of the distribution and overlap of the two processes in frequency space.

One important aspect of this work is that the hemodynamic response function, or temporal point spread function, is estimated without reference to any assumed input (it is estimated using intrinsic autocorrelations in the physiological data only). Any estimation based on the sensory input would clearly result in a rather circular exercise, in that the technique would be biased towards the input function used.

There are a number of alternative approaches to the problem of detecting significant cross-correlations, some of which we have evaluated. One obvious approach is to use the maximum cross-correlation between $c(t)$ and $x(t)$ over some reasonable interval $\gamma_{cx}(\tau)_{\max}$. This is simply

achieved by computing $\gamma_{cx}(\tau)$ and thresholding in space and time according to eqn(14). For slice data the $SPM\{\gamma_{cx}(\tau)\}$ represents a three dimensional process and for volume MRI data the corresponding SPM would be four dimensional. We have tried this approach and found it to be less sensitive and less robust than the one we have presented. Another test for significant cross-correlations would be to integrate $\gamma_{cx}(\tau)$ over some reasonable interval, or more generally take a weighted sum over an interval. In fact the method described in this paper is exactly equivalent to taking a weighted sum of $\gamma_{cx}(\tau)$, where the weighting is defined by the hemodynamic response function $h(\tau)$. This can be seen by noting that the weighted sum of the cross-covariance function is obtained by the three-way convolution $[c(t) \otimes x(t)] \otimes h(\tau)$. By the commutivity of the convolution operator (\otimes) this is equal to $[c(t) \otimes h(\tau)] \otimes x(t)$. The latter expression represents the cross-correlation between the convolved contrast and $x(t)$, which is precisely what has been used.

The choice of a Poisson form for the hemodynamic response function means that delay and dispersion are totally confounded. We have examined the effects of varying the assumed delay and dispersion separately, using the exceedence proportion as an measure of omnibus sensitivity. In these analyzes we used the gamma distribution, which has two parameters and allows the independent manipulation of mean and variance. We do not report the results of this analysis because they did not contribute any significant insight. In general the exceedence proportion showed exactly the same dependence on delay as for the Poisson form. The dependence on dispersion was much less marked. We could have used a Gaussian response function and set the delay equal to dispersion (for large values of λ this and the Poisson function are the same). The aesthetic advantage of the Poisson function is that it only exists for positive lag (τ).

There are obviously many variants on the details we have presented here. The central tenets of the approach are: (i) The assessment of regional specific associations between neuronal activity and

a stimulus or task parameter in terms of correlations. (ii) The appropriate correlation is between the observed hemodynamic response and the hypothesized input which has been convolved with a response function. (ii) The response function is specified without reference to an assumed input. (iv) Care is taken to ensure that, in the absence of any response, intrinsic temporal and spatial autocorrelations do not subtend spurious correlations (i.e. the experiment-wise false positive rate is maintained at a suitably low level). Within these constraints there are a number of variants which deserve mention. First, in the analysis above, it was assumed that the response function and intrinsic autocorrelations were the same for all voxels. An identical analysis performed, in parallel for each voxel, would allow for regional variation in the response function. Our analyses (see last section) suggest these regional variations can be substantial. This extension would certainly increase sensitivity but would require long time-series, to ensure robust estimates of the intrinsic spectral density $g_x(\omega)$. A second variant would involve estimating the response function on the basis of independent data; for example a second study or by splitting the same study into two halves and using the SVD or PCA approach and least squares deconvolution - Jim Haxby (personal communication).

Although we used temporal smoothness to estimate both the intrinsic spectral density of the MRI signal and the response function, it is important to note that the first application is mandatory, the second is not: Any response function will, under the null hypothesis, render the distribution of $\zeta(0)$ Gaussian. Despite the fact that sensitivity to real change will be optimized when the dispersion or variance of the response function 'matches' the intrinsic autocorrelations, there is no statistical requirement for this matching to be implemented. In other words one is free to choose whatever response function one likes. Complicated biphasic response functions are attractive in the sense they are capable of modeling the 'undershoot' effect. However the data did not suggest that this extension was justified in the examples above.

One particular constraint imposed by a Poisson form is that delay and dispersion are yoked. This may be inappropriate for some data and this constraint should not be considered a cornerstone of the proposed method. One particular fallacy of linking the delay and dispersion is revealed when considering the effect of temporal smoothing of the data to increase signal to noise. Here the underlying response function will be subject to this smoothing and its effective dispersion will increase. This would be correctly detected and modeled by *post hoc* estimates of smoothness. However the 'correct' delay will not have changed. In some instances it may be best to fix the response function to some appropriate delay and allow its variance to be dictated by σ_{t^2} (e.g. using the gamma distribution).

In the method presented the impact of thermal noise was suppressed by temporal smoothing at a data pre-processing stage. With longer time-series it may be possible to adopt a more formal approach, by segregating the data into an uncorrelated white (or Raleigh) noise and a physiological autocorrelated process. Many linear devices exist for this sort of analysis (e.g. auto-regressive and moving average models), however they are usually not considered in the context of very short time-series.

Correlations or categories ?

The principle of correlating sensorial or behavioural parameters with cerebral excitation has a long history in neuroscience, from the days of Goltz and Ferrier (e.g. Ferrier 1875, Goltz 1881, Phillips 1984) to modern day approaches which use exogenous stimulation; e.g. magneto-stimulation or endogenous activity (PET and MRI). Correlations in micro-electrode recording and functional imaging data have received special attention recently, as they are a possible index of functional and effective connectivity (Gerstein et al 1969, Friston 1993b).

Although there is a fundamental equivalence between testing for specific time-dependent changes with contrasts in analysis of (co-)variance of PET data and correlating a time-dependent parameter

with MRI data (see Friston et al 1993a for a concrete illustration of this point), it is important to acknowledge the differences: MRI data represent time-series that lend themselves directly to signal processing. For example in this paper, using a few standard results from the theory of stochastic processes, one can solve the problem of detecting cross-correlations in the presence of autocorrelations and avoid what could have been a discursive and dense statistical analysis.

This paper has focused exclusively on correlations as a means of detecting significant activations. An alternative, found in the recent functional MRI literature, is to categorize the MRI time-series into activation and baseline scans and express the mean difference in some form. There have been profound reservations about this approach related to the observation that the residual MRI signal, around the mean for each stimulus-cycle, is not normally distributed (John Baker - personal communication). The analysis presented in this paper concurs with these reservations: While the residual variability about the convolved contrast $y(t)$ is normally distributed, the probability that it will be similarly distributed about a horizontal line running from stimulus onset to offset is low. The MRI time-series reflects a true dynamic response to an input, characterized by its response function. To pretend that this response, following stimulus onset, constitutes a single category is not really tenable. The proper way to conceptually categorize these time-series, is in terms of offset from stimulus onset. For example the set of first scans after stimulus onset is a different category from the set of second scans following stimulus onset. The residuals about the means of these categories will be normally distributed. Note that the convolved contrast implicitly assumes the latter form of categorisation because only these scans have the same values of $y(t)$.

Implicit in phase-shifting (as a tool to estimate intrinsic autocorrelations) is the requirement that stimulus cycles are repeated over and over again. There is a fundamental point to be made here: If one wants to implement intra-subject signal averaging in its broadest sense, then multiple presentations of the same stimulus-cycle are obligatory (the more the better). This places significant constraints on experimental design. The key theme is, again, that *the object which is*

averaged is a complete stimulus-cycle not repeated scans within a cycle. A failure to recognize this will sacrifice all the rich nonlinear information, about an evoked neuronal and hemodynamic transient, for the apparent ease of simple categorical comparisons. Ironically this nonlinearity invalidates a parametric approach to averaging scans from the same stimulus-cycle (see above).

In the context of MRI data, correlational approaches offer a number advantages over alternative approaches. One key issue is that the correlation between sensory input and hemodynamics will not be sensitive to any artifact that is orthogonal (independent) to the input. These artifacts include movement artifacts and low frequency aliasing artifacts due to an interaction between the repeat time and cardiac, respiratory and heart rate variability (HRV) cycles. The proposed approach has one interesting feature with regard to movement artifacts; because the effect of movement on signal is not subject to delay or dispersion (whereas the effect of evoked neuronal activity is), convolving the contrast with a hemodynamic response function will render the correlations insensitive to *movement associated with changing stimuli or tasks.*

In conclusion we hope to have presented some reasonable solutions to a relatively simple problem, namely how to identify regionally specific correlations between sensory, cognitive or behavioural parameters and cerebral physiology measured with functional MRI.

Acknowledgments

We wish to thank Jack Belliveau, John Baker, Jim Haxby and Nick Lange for helpful discussions about these issues. The work by KJF was carried out as part of the Institute Fellows in Theoretical Neurobiology research program at The Neurosciences Institute, which is supported by the Neurosciences Research Foundation. The Foundation receives major support for this research from the J.D. and C.T. MacArthur Foundation, the Lucille P. Markey Charitable Trust. KJF was a W.M. Keck Foundation Fellow.

References

- Adler RJ & Hasofer AM (1981) The geometry of random fields. Wiley, New York.
- Aertsen A & Preissl H (1991) Dynamics of activity and connectivity in physiological neuronal Networks in Non Linear Dynamics and Neuronal Networks Ed Schuster HG VCH publishers Inc., New York NY USA p281-302
- Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, Hyde JS (1992) Time course EPI of human brain function during task activation. *Mag. Res. Med.* 25:390-397
- Bandettini PA (1993a) *Mag. Res. Med.* 30:161-173
- Bandettini PA (1993b) MRI studies of brain activation: Temporal characteristics. In *Functional MRI of the Brain*, Society of Magnetic Resonance in Medicine, Berkely CA p143-151
- Cox DR & Miller HD (1980) The theory of stochastic processes, Chapman and Hall, London and New York.
- Edelman GM (1978) Group selection and phasic reentrant signaling: a theory of higher brain function In: *The Mindful Brain*. Edelman GM & Mountcastle VB (eds) p55-100. MIT Press.
- Ferrier D, (1875) Experiments on the brain of monkeys. *Proc. Royal Soc. Lond.* 23: 409-430
- Fox PT Raichle ME Mintun MA & Dence C (1988) *Science* 241:462
- Friston KJ Frith CD Liddle PF Dolan RJ Lammertsma AA & Frackowiak RSJ (1990) The relationship between global and local changes in PET scans *J. Cereb. Blood Flow Metab.* 10:458-466
- Friston, K.J Frith, C.D Liddle, P.F and Frackowiak, R.S.J. (1991) Comparing functional (PET) images: The assessment of significant change. *J. Cereb. Blood Flow Metab.* 11:690-699.
- Friston KJ Jezzard P Frackowiak RSJ & Turner R (1993a) Characterizing focal and distributed physiological changes with MRI and PET. In *Functional MRI of the Brain - Syllabus: Functional MRI of the Brain*, Society of Magnetic Resonance in Medicine, Berkely CA p207-216

- Friston KJ Frith CD Liddle PF & Frackowiak RSJ (1993b) Functional connectivity: The principal component analysis of large (PET) data sets. *J. Cereb. Blood Flow Metab.* 13:5-14
- Friston KJ, Worsley KJ, Frackowiak RSJ, Mazziotta JC & Evans, AC (1994) Assessing the significance of focal activations using their spatial extent. submitted.
- Frostig RD Lieke RR Ts'o DY & Grindvald A (1990) Cortical functional architecture and local coupling between neuronal activity and the microcirculation revealed by in vivo high resolution optical imaging of intrinsic signals *Proc. Natl. Acad. Sci.* 87: 6082-6086
- Gerstein GL & Perkel DH (1969) Simultaneously recorded trains of action potentials: Analysis and functional interpretation *Science* 164:828-830
- Goltz F (1881) In: Transactions of the 7th international medical congress. Volume 1 ed. MacCormac W JW Kolkman: London pp 218-228
- Golub GH & Van Loan CF (1991) Matrix computations (2nd ed) The Johns Hopkins University Press, Baltimore and London. p241-248
- Grimmett G and Welsh D (1986) Probability and introduction Clarendon Press, Oxford.
- Hasofer AM (1978) Upcrossings of random fields. *Suppl. Adv Appl. Prob.* 10:14-21
- Keys RG (1981) Cubic convolution interpolation for digital image processing *IEEE Trans. Acoustic, Speech and Signal Processing* 29:1153-1160
- Kwong KK, Belliveau JW Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, Cheng HM, Brady TJ & Rosen BR (1992) Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation *Proc. Natl. Acad. Sci.* 89:5675-5679
- More JJ (1977) The Levenberg-Marquardt Algorithm: Implementation and theory, *Numerical Analysis*, ed Watson GA, Lecture notes in mathematics 630, Springer Verlag, p105-116

Nelson JI Salin PA Munk MHJ Arzi M & Bullier J (1992) Spatial and temporal coherence in cortico-cortical connections: A cross-correlation study in areas 17 and 18 in the cat. *Visual Neurosci.* 9:21-37

Ogawa S Tank DW, Menon R, Ellermann JM, Kim SG, Merkle H & Ugurbil K. (1992) Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc. Natl. Acad. Sci.* 89:5951-5955

Phillips CG, Zeki S, Barlow HB (1984) Localization of function in the cerebral cortex. Past present and future. *Brain* 107:327-361

Talairach P & Tournoux J (1988) *A Stereotactic coplanar atlas of the human brain.* Stuttgart:Thieme.

Worsley KJ Evans AC Marrett S & Neelin P (1992) A three-dimensional statistical analysis for rCBF activation studies in human brain. *J. Cereb. Blood. Flow Metab.* 12:900-918

Worsley KJ & Vandal (1994) Tests for changes in random fields with applications to medical images - submitted

Legends for Figures

Figure 1 *The analysis of a simulated data set*

Top left: The assumed hemodynamic response function $h(\tau)$ which has a Poisson distribution with parameter 7.37 seconds. Top right: The original square wave contrast [broken line - $c(t)$], the contrast convolved with assumed response function [solid line - $y(t)$] and the response of the simulated voxel data $x(t)$ that best correlated with the convolved contrast (dotted line). Bottom left: The distribution (over voxels and time) of $\zeta(\tau)$ calculated according to eqn(9) (dotted line). This distribution is indistinguishable from that predicted theoretically (solid line). The broken vertical line is the threshold that any value would have to reach to be considered significant at $p = 0.05$ according to eqn(14). Bottom right: The corresponding SPM(z), a statistical parametric map of $\zeta(0)$. The gray scale is arbitrary and has been scaled to the maximum and minimum of the SPM.

Figure 2 *Actual and assumed intrinsic autocorrelation and spectral densities*

Empirical estimates of the intrinsic autocorrelation function and spectral density compared with those based on temporal smoothness. Left: The autocorrelation function of the difference between the MRI time-series from one stimulus cycle and another (phase-shifted differences), averaged over all voxels and all pairs of cycles (broken line). The autocorrelation function obtained by convolving an uncorrelated innovation with a Poisson function based on the measured temporal smoothness is also shown (solid line). Right: The corresponding empirical (dots) and assumed (solid line) spectral densities of the phase-shifted differences.

Figure 3 *The analysis of a real data set*

As for Figure 1 but using real data. In this case the distribution (lower left) is for $\zeta(0)$ (dotted line) and is dramatically skewed with respect to the distribution predicted under the null hypothesis (solid line).

Figure 4 *Thresholding the SPM{z}*

Left: The SPM{z} from Figure 3 is presented as a surface rendering, with the corrected threshold [corrected for multiple non-independent comparisons according to eqn(14)]. The resulting 'excursion set' is highlighted on the right and includes V1 and several extrastriate regions according to the atlas of Talairach and Tournoux (1988).

Figure 5 *A null analysis of real data*

Exactly as for Figure 3 but with a change to the contrast (the 'on-off' cycle period has been reduced from 20 scans to 16 scans). The impact on the distribution of $\zeta(0)$ (lower left) is clearly evident.

Figure 6 *Estimating the response function directly*

Top left: The first spatial mode or eigenimage of the real data. The gray scale is arbitrary and has been scaled to the image maximum. Darker means a higher loading. Only positive loadings are shown. Note the correspondence between this profile and the significantly 'activated' regions in figure 3. Top right: The assumed input $c(t)$ (solid line) and the response of the first mode or eigenimage (dots). Lower left: The response function based only on the intrinsic temporal smoothness (a function of the variance of the first derivative) of the MRI data (broken line) and the response function obtained by least squares deconvolution (dots), based on an assumed input $c(t)$ and the first mode's response. Lower right: The actual response (dots) and the response predicted

using the Poisson response function (broken line) and the least squares deconvolution estimate (solid line).

Figure 7 *Transfer functions*

Left: The modulation transfer function expressed as the square of the frequency representation of a hemodynamic response function. Right: Relative activation expressed as a function of 'on-off' cycle length. The data are copied from the graphical data in Bandettini (1993b) and the solid line represents that predicted with a hemodynamic response function of Poisson form and parameter 5 seconds.

Figure 8 *Phase differences and delay*

The distribution of phase differences between the contrast and the response at each voxel. The phase difference was computed for the largest Fourier component of the contrast (input). The vertical broken line is the delay estimated on the basis of temporal smoothness.

Figure 9 *The effects of changing the response function*

Upper left: The exceedence proportion (number of voxels above threshold $p = 0.001$) as a function of λ . Lower left: The SPM $\{z\}$ thresholded at a descriptive (non-corrected) level of $p = 0.001$ (one tailed). In this instance the voxel values are the largest $\zeta(0)$ obtained over a range of λ . Lower right: The optimal value of λ for all voxels surviving a $p = 0.001$ threshold (voxels highlighted on the Lower left). The optimal value of λ is defined as that which gives the greatest $\zeta(0)$. The gray scale is arbitrary with small values of λ (quicker and more acute responses) being brighter. Note that V1 (centro-medial dark areas) appear to be subject to the greatest delay and dispersion.

