ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 23 (2004) 752–763

Measuring functional connectivity during distinct stages of a cognitive task

Jesse Rissman,[a,b,*,1] Adam Gazzaley,[a,c,1] and Mark D'Esposito[a,b,c]

[a]*Henry H. Wheeler Jr. Brain Imaging Center, University of California, Berkeley, Berkeley, CA 94720, USA*
[b]*Department of Psychology, University of California, Berkeley, Berkeley, CA 94720, USA*
[c]*Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, CA 94720, USA*

Received 12 April 2004; revised 8 June 2004; accepted 24 June 2004

The inherently multivariate nature of functional brain imaging data affords the unique opportunity to explore how anatomically disparate brain areas interact during cognitive tasks. We introduce a new method for characterizing inter-regional interactions using event-related functional magnetic resonance imaging (fMRI) data. This method's principle advantage over existing analytical techniques is its ability to model the functional connectivity between brain regions during distinct stages of a cognitive task. The method is implemented by using separate covariates to model the activity evoked during each stage of each individual trial in the context of the general linear model (GLM). The resulting parameter estimates (*beta values*) are sorted according to the stage from which they were derived to form a set of stage-specific *beta series*. Regions whose beta series are correlated during a given stage are inferred to be functionally interacting during that stage. To validate the assumption that correlated fluctuations in trial-to-trial beta values imply functional connectivity, we applied the method to an event-related fMRI data set in which subjects performed two sequence-tapping tasks. In concordance with previous electrophysiological and fMRI coherence studies, we found that the task requiring greater bimanual coordination induced stronger correlations between motor regions of the two hemispheres. The method was then applied to an event-related fMRI data set in which subjects performed a delayed recognition task. Distinct functional connectivity maps were generated during the component stages of this task, illustrating how important and novel observations of neural networks within the isolated stages of a cognitive task can be obtained.

Keywords: Functional connectivity; Network; Correlation; Working memory; Delay period

Introduction

The vast majority of functional magnetic resonance imaging (fMRI) studies have utilized univariate statistical analyses in an attempt to localize activity to brain regions involved in specific cognitive operations. However, brain regions do not act in isolation (Mesulam, 1990) and there is a rapidly growing interest in the field to use fMRI to explore how regions of the brain communicate with one another during cognitive tasks. The ability of fMRI to rapidly sample blood oxygenation level dependent (BOLD) activity throughout the entire functioning brain makes it an ideal tool for studying inter-regional interactions. Such interactions have typically been characterized by identifying regions that show correlated fluctuations in their fMRI time series data, with the belief that temporal correlations in BOLD signal might reflect synchronous neural firing in the communicating regions. The term "functional connectivity" was introduced by Friston et al. (1993) to describe the "temporal correlation between spatially remote neurophysiological events", as assessed by functional neuroimaging data.

Several early explorations of functional connectivity in fMRI data focused on inter-regional interactions during the resting state, when the subject was not explicitly engaged in a cognitive task. These studies demonstrated that sensorimotor regions exhibit correlated low frequency fluctuations in their fMRI time series during the resting state and that these correlations could not be attributed to higher frequency physiological noise (Biswal et al., 1995, 1997; Lowe et al., 1998; Stein et al., 2000; Xiong et al., 1999). The use of resting state data was motivated by the notion that spontaneous firing of functionally connected neurons was all that was needed to induce correlations in the BOLD signal. Later studies began to search for regions whose time series were correlated during the continuous performance of a cognitive task as assessed in a blocked design experiment. For example, Lowe et al. (2000) identified a network of regions that showed stronger correlations with the dorsolateral prefrontal cortex (DLPFC) during a two-back spatial working memory task than during a zero-back motor control task. Hampson et al. (2002) showed that the low frequency correlation that exists between Broca's and Wernicke's

* Corresponding author. Department of Psychology, University of California, Berkeley, 132 Barker Hall, MC# 3190, Berkeley, CA 94720-3190.
E-mail address: jrissman@berkeley.edu (J. Rissman).
[1] These authors contributed equally to this work.
Available online on ScienceDirect (www.sciencedirect.com.)

1053-8119/$ - see front matter © 2004 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2004.06.035

area at rest increases when the language system is actively engaged by a listening task. By having subjects engage in a cognitive task for an extended block of time (~45-s epochs in these two studies), the inter-regional correlations observed in the fMRI time series data collected during these epochs can be attributed to the ongoing cognitive processes taking place during these task periods.

The use of a blocked design in fMRI experiments has the advantage of allowing the functional data collected during the performance of one task to be clearly separated from that collected during another task or rest period. It also serves to provide a large number of temporally contiguous data points, which makes correlation computations robust and sensitive to lower frequency fluctuations. However, a major shortcoming of using blocked design experiments to model brain activity or functional connectivity is that there is no way to estimate what proportion of the evoked BOLD response is attributable to each of the cognitive subcomponents of the task (D'Esposito et al., 1999b). For example, in a two-back working memory task the subject must continuously encode new stimuli, update and maintain mnemonic representations of past stimuli, and decide and respond whether each new stimulus matches the one presented two stimuli previously. Because all of these cognitive operations take place virtually simultaneously, it is impossible to tease them apart temporally. Thus, any statements made about the neural basis of working memory would apply to all processes associated with this task.

One manner of overcoming this drawback of the blocked design is by adopting an event-related design (Postle et al., 2000; Zarahn et al., 1997). With an event-related design it is possible to study the subcomponents of a cognitive process, such as working memory, by employing a task with temporally distinguishable stages. For example, in a typical delayed recognition task, the subject must first encode a visually presented cue stimulus, then hold this percept in mind across an extended delay interval, and finally make a decision as to whether or not the maintained percept matches a probe stimulus. Event-related designs have been extensively applied to investigate the stages of working memory using the delayed recognition task (Courtney et al., 1997; D'Esposito et al., 1999a; Rama et al., 2001; Ranganath and D'Esposito, 2001; Rowe et al., 2000; Sakai et al., 2002; Sala et al., 2003; Zarahn et al., 1999). However, while univariate statistical methods exist for obtaining separate estimates of brain activation during each of these distinct processing stages, currently available multivariate methods for modeling functional connectivity, such as time series correlations, are not capable of evaluating inter-regional interactions within closely spaced stages of a task. In this paper, we present a new method for characterizing functional connectivity in an event-related fMRI experiment that is capable of measuring inter-regional correlations during distinct stages of cognitive task. The method employs a standard general linear model (GLM) approach for estimating stage-specific activity, but adapts the model such that separate parameter estimates are computed for each individual trial and then used as the dependent data in a correlation analysis. We will first validate this method by applying it to a recently published event-related fMRI data set utilizing a simple motor task where there are expected results based on the findings of a bivariate analysis employing coherence (Sun et al., 2004), as well as electrophysiological evidence. Next, we will apply this method to an fMRI data set employing a more complex multistage cognitive task, a delayed recognition working memory task.

Methods

Beta series correlations

To understand how this method is capable of characterizing stage-specific functional interactions, it is important to review how individual stages of a cognitive task can be modeled to obtain estimates of the underlying brain activity. If the fMRI BOLD signal were a direct measure of neural activity, then determining which portions of the measured signal were evoked by different stages, such as the individual stages of a trial of a delayed recognition task—cue, delay, and probe—would be relatively straightforward. Much as a neurophysiologist would analyze single unit data, one could simply subdivide the signal into three discrete trial periods. However, our ability to correctly attribute the measured BOLD signal to the each of the subcomponents of a trial is complicated by the fact that the hemodynamic response acts as a low-pass filter on the neural activity, peaking 4–6 s after the onset of neural activity and then decaying slowly back to baseline. Any attempt to study regional activity or functional connectivity during the individual stages of a cognitive task using fMRI must take this basic principle into account.

One effective and widely used method of determining how much of the measured BOLD signal is attributable to each stage of a multistage trial involves modeling the data with a set of covariates in the context of the general linear model (GLM) (Postle et al., 2000; Zarahn et al., 1997). For example, when applied to a delayed recognition task, separate covariates are constructed to represent neural activity associated with the cue, delay, and probe components of the task. To convert these covariates from models of predicted neural activity to models of predicted BOLD activity, they are convolved with an estimate of the hemodynamic response function (HRF). Since the transition from cue to delay period in a delayed recognition task is instantaneous in both cognitive and neural terms, it might seem that the delay period should be modeled as a block of activity beginning immediately after the cue and continuing until the presentation of the probe. However, such a model would result in a high level of collinearity among the adjacent cue, delay, and probe covariates, and the delay period covariate would invariably capture some of the residual encoding activity, contaminating its estimate of maintenance-related activity. Simulations by Zarahn et al. (1997) suggest that the onsets of temporally adjacent covariates should be spaced at least 4 s apart to minimize this collinearity. By placing the onset of the delay covariate in the middle of the delay interval, each covariate can explain a largely unique source of variance in the signal. This approach has been used to successfully model delay period activity in numerous published studies (Barde and Thompson-Schill, 2002; Druzgal and D'Esposito, 2003; Pessoa et al., 2002; Postle and D'Esposito, 1999; Rypma and D'Esposito, 1999). In estimating the parameters of the GLM, the cue, delay, and probe covariates are scaled to best fit the observed fMRI time series of each voxel, such that their linear combination minimizes the sum of the squared differences between the observed data and the model predictions. This yields a set of parameter estimates (i.e., beta values) for each voxel that reflect how much of that voxel's activity can be attributed to the individual stages of a task (i.e., the cue, delay, and probe phases). In most fMRI analyses, these parameter estimates would be scaled by their error term (unexplained variance), and turned into statistical parametric maps of brain activity. Conclusions could then be made about the brain

regions involved in the various subcomponents of a working memory task.

The classic univariate method for estimating activity during the individual stages of a delayed recognition task uses a single cue, delay, and probe covariate to model the data obtained from many experimental trials, essentially treating any trial-to-trial variability as noise. In contrast, the multivariate method we propose for modeling functional connectivity during the temporally adjacent stages of a multistage cognitive task, although relying on the same principles used to model stage-specific univariate activity, capitalizes on this trial-to-trial variability and uses it to characterize dynamic inter-regional interactions. The premise of this method is that if two areas of the brain are functionally interacting with each other during a particular stage of a cognitive task, then the amount of activity that the two areas exhibit during that stage should be correlated across trials. The goal is to obtain a reasonable measurement of the magnitude of stage-specific activity that each voxel exhibits on each of many task trials, and then search for other voxels in the brain that show correlated fluctuations across trials. This is accomplished by constructing a GLM in which every stage of every trial is modeled with a separate covariate, so that trial-to-trial parameter estimates of stage-specific activity can be obtained. These parameter estimates (*beta values*) can then be sorted according to the stage from which they were derived (what we will refer to as a *beta series*) and correlated across regions to obtain a measure of functional connectivity (*beta series correlation*) during each of the individual task components.

Validation of the beta series correlation method

Before illustrating how this method of obtaining stage-specific estimates of inter-regional correlations is implemented in the analysis of multistage cognitive tasks, such as the delayed recognition task, for which it was specifically designed, it is necessary to demonstrate that it can provide a neurophysiologically plausible measure of functional connectivity between brain regions. To accomplish this, we first demonstrate that the method can be validly applied to single-stage experimental tasks; that is, tasks in which the processing that occurs during each trial can be conceptualized and modeled as a single event. An ideal task is one in which the degree of functional connectivity between two regions is expected to change as a function of task condition in an intuitive way. To this end, we applied the beta series correlation method to a simple motor task to test a straightforward hypothesis about functional connectivity between motor areas of the two cerebral hemispheres.

The goal of this validation study was to replicate the results of Sun et al. (2004). In their event-related fMRI study, they used coherence analysis[2] to assess the strength of interhemispheric interactions in two sequence tapping tasks requiring varying degrees of bimanual coordination. Based on evidence from EEG studies and experiments on acollosal patients (Andres et al., 1999; Gerloff and Andres, 2002; Jeeves et al., 1988; Ohara et al., 2001; Serrien and Brown, 2002), they predicted that a task demanding greater bimanual coordination would induce greater functional connectivity between the hemispheres. Indeed, their results revealed greater coherence between right and left motor cortices when subjects performed a task requiring more bimanual coordination, despite the two tasks showing no significant differences in their univariate activity.

Here, using the fMRI data set collected by Sun et al. (2004), we sought to determine if the same effect is obtained using the beta series correlation method. The specific details of the data collection methods and behavioral tasks were described in Sun et al. (2004), and are summarized below.

Subjects and experimental task

Twelve right-handed subjects participated in the study. Before scanning, the subjects learned two motor sequences. In the *Right-then-Left* sequence, subjects played a sequence of four keystrokes with the fingers of their right hand and then a different sequence of four keystrokes with the fingers of their left hand. In the *Interleaved* sequence, subjects also played a sequence of eight keystrokes, but alternated back-and-forth between hands with each keystroke. Subjects were trained until they could accurately play each sequence in less than 2500 ms. In the scanner, subjects performed a total of 36 trials of each type. On each trial, subjects received a visual cue instructing them which sequence to play. Trials were presented in a randomized fashion and spaced 16–20 s apart. In addition to this task, subjects performed a visuomotor response task to empirically derive a hemodynamic response function (HRF) (Aguirre et al., 1998).

fMRI data acquisition and processing

Functional images were acquired on a 4.0 T Varian scanner using a two-shot gradient-echo EPI sequence (TR = 543 ms per half *k*-space, TE = 28 ms, matrix size = 64 × 64, FOV = 22.4 cm^2). Each volume consisted of ten 5-mm-thick axial slices with a 1-mm gap between slices. fMRI data processing included a linear time-interpolation algorithm to double the effective sampling rate, temporal sync interpolation to correct for between-slice timing differences in image acquisition, motion correction using a six-parameter rigid-body transformation algorithm (Friston et al., 1995), and spatial smoothing with a 7-mm FWHM Gaussian kernel. For the univariate and beta series correlation analyses, the time series of each voxel was normalized by its mean signal value.

Univariate analysis

The reference functions used to model task related activity in each of the two conditions were constructed by convolving the subject's empirically derived HRF (Aguirre et al., 1998) with a function specifying the onset of each trial of that condition, which in turn was convolved with a 2-s block (the average time it took subjects to play a sequence) to account for the extended duration of neural processing. These covariates were entered into the general linear model for analysis with VoxBo (http://www.voxbo.org).

Region-of-interest ("seed") selection

A commonly used approach for identifying networks of functionally connected brain regions involves defining a cluster of activated voxels in a region of interest (ROI) and using this "seed" to determine which voxels throughout the entire brain are functionally interacting with the seed (Cordes et al., 2000; Della-Maggiore et al., 2000; Greicius et al., 2003; Lowe et al., 2000; Quigley et al., 2001; Stein et al., 2000). The seed region is generally chosen based on its known involvement in the behavioral task and the researcher's interest in characterizing its interactions with other regions of brain. In this case, the right and left primary

[2] Coherence measures the linear time-invariant relationship between two time series and is essentially the spectral analog of correlation.

motor cortices (right M1 and left M1) were used as seeds because the experiment was designed to test a specific hypothesis regarding interhemispheric interactions between cortical motor regions. These seeds were identified separately for each subject in their native space (i.e., not spatially normalized) by selecting the 10 voxels with the most task-related activity as determined by an F test, collapsed across task conditions, in the right and left hemisphere primary motor cortices.

Coherence analysis

Since the trials of the two task conditions occurred in a randomly intermixed fashion, the time series data were first reorganized into condition-specific time series by segmenting the data set into 16-s blocks (32 TRs) beginning with the instruction cue of each trial. Each segment was mean-centered, windowed using a 4-point split-cosine bell, and then concatenated with segments of the same condition. Coherence maps were constructed for each subject in their native space by calculating the band-averaged low frequency coherence (0–0.15 Hz) of the condition-specific time series of the seed (averaged across the 10 seed voxels) with that of all other voxels in the brain. Since coherence maps were generated separately for each of the two conditions and with each of the two seeds, four maps were produced for each subject.

Beta series correlation analysis

To implement the correlation analysis, the univariate analysis described above was adapted such that the magnitude of the task-related BOLD response was estimated separately for each of the 72 experimental trials. This was implemented in VoxBo by constructing a design matrix with 72 covariates of interest, each of which modeled the activity evoked during a single trial as a 2-s block (time-locked to the onset of the trial) convolved with the subject's HRF. The model also included covariates of no interest to model the effects of shifting signal levels across runs. A band pass filter was used to attenuate frequencies above 1 Hz and below 0.01 Hz. The least-squares solution of the GLM yielded a unique parameter estimate (beta value) for each trial. These beta values were then sorted according to whether they were derived from a *Right-then-Left* trial or an *Interleaved* trial, yielding a set of 36 beta values for each condition for every voxel in the brain. We refer to such a set of condition-specific beta values as a "beta series". Under the assumptions of this method, the extent to which two brain voxels interact during a given task condition is quantified by the extent to which their respective beta series from that condition are correlated. To attain a parallel analysis with the coherence analysis described above, we used the same right and left M1 seeds to produce maps of functional connectivity. By computing the correlation of the seed's beta series (averaged across the 10 seed voxels) with the beta series of all other voxels in the brain, condition-specific seed correlation maps were generated. All correlation analyses were conducted using Matlab 6.5 (http://www.mathworks.com).

To allow statistical conclusions to be made based on the correlation magnitudes, we applied an arc-hyperbolic tangent transform (Fisher, 1921) to the correlation coefficients of all brain voxels. Since the correlation coefficient is inherently restricted to range from -1 to $+1$, this transformation serves to make its null hypothesis sampling distribution approach that of the normal distribution. The transformed correlation coefficients were then divided by their known standard deviation ($1/\sqrt{N-3}$, where N is the number of data points used to compute the correlation coefficient) to yield z scores.

To assess the map-wise significance of the correlation findings at the group level, the z-transformed correlation maps of the individual subjects were spatially normalized into standard MNI atlas space using routines from SPM99 (http://www.fil.ion.ucl.ac.uk/spm). Group level random-effects t tests were then conducted using the fidl analysis package (http://www.nil.wustl.edu/~fidl) to identify voxels for which the mean of the individual subjects' transformed correlation coefficients was reliably greater than zero. All statistical maps are displayed using MRIcro (http://www.mricro.com).

Delayed recognition task

The bimanual motor task described above provides a strong test of the validity of the beta series correlation method's primary assumption; namely, that correlated fluctuations in trial-to-trial beta values between voxels are indicative of functional interactions between brain regions. However, other methods of assessing functional connectivity, such as coherence analysis or even standard time series correlation can be validly applied to such a data set because the fMRI signal evoked by the two different conditions of interest can be separated temporally. In fact, for the coherence analysis, the data set was essentially treated as a blocked design. Despite the fact that the subject only engaged in the motor task for the first 2 s of the trial, 16 s (or 32 TRs) of fMRI data were taken from each trial, and these miniblocks of single trial activity were concatenated into a larger block from which inter-regional coherence in the low frequency bands was computed.

Below, we illustrate how the beta series correlation method can be applied to obtain separate maps of functional connectivity during the individual subcomponent stages of a multistage cognitive task. We applied this method to an event-related fMRI data set employing a fairly standard visual delayed recognition task. This data set has been previously analyzed with univariate methods, and the results have been described elsewhere (Ranganath and D'Esposito, 2001; Ranganath et al., 2003).

Subjects and experimental task

Eight right-handed subjects participated in this study. Subjects performed three runs of 18 delayed recognition trials for a total of 54 trials. On each trial, subjects were presented with a single grayscale face stimulus for a duration of 1 s, followed by a fixation cross for 7 s, followed by a probe face for 1 s. Subjects were instructed to pay careful attention to the first face in each trial and maintain a mental image of that face throughout the delay period. When the probe face appeared, subjects made a key press with the left index finger if it matched the first face (50%) and the right index finger if it did not (50%). Each trial was followed by a 13-s intertrial interval (ITI).

In addition to this task, subjects performed a visuomotor response task to empirically derive a hemodynamic response function (HRF) (Aguirre et al., 1998) and a "localizer" task in which they passively viewed blocks of face and object stimuli to identify face-sensitive regions of the fusiform cortex (Kanwisher et al., 1997).

fMRI data acquisition and processing

Functional images were acquired on a 1.5 T General Electric scanner with a gradient-echo EPI sequence (TR = 2000 ms, TE = 50 ms, matrix size = 64 × 64, FOV = 24 cm^2). Each functional volume consisted of 21 contiguous 5-mm-thick axial slices. FMRI data processing included sync interpolation in time to correct for

between-slice timing differences in image acquisition, motion correction using a six-parameter, rigid-body transformation algorithm (Friston et al., 1995), normalization of the time series of each voxel by its mean signal value, and spatial smoothing with an 8-mm FWHM Gaussian kernel.

Univariate analysis

BOLD responses during the cue, delay, and probe stages of the task were modeled as impulses of activity convolved with the individually derived HRF. To minimize collinearity between temporally adjacent covariates, onsets were spaced 4 s apart, such that the cue covariate was placed at the start of the trial, the delay covariate was positioned 4 s into the trial, and the probe covariate was positioned 8 s into the trial. These three covariates of interest were entered into the general linear model for analysis with VoxBo. The model included a set of separate covariates to model the cue, delay, and probe stages of those trials for which the subject responded incorrectly. The model also included covariates of no interest to model the effects of shifting signal levels across runs. A band pass filter was used to attenuate frequencies above 0.25 Hz and below 0.01 Hz.

Individual subject activation maps were spatially normalized into standard MNI atlas space using routines from SPM2. Group level random-effects analyses were performed separately for the cue, delay, and probe stages of the task to test whether the mean of the individual subjects' parameter estimates at each voxel was reliably greater than 0.

Seed selection

The seven contiguous voxels in each subject's right fusiform gyrus that exhibited the strongest response preference to faces versus objects in the localizer task, as assessed by a *t* test, were defined as that subject's fusiform face area (FFA) (Kanwisher et al., 1997) and used as a seed in the subsequent correlation analyses. The FFA has been identified as a region of the visual association cortex that is selective for viewing faces (Kanwisher et al., 1997; Puce et al., 1995) and has been widely applied to fMRI analyses of stimulus-specific visual association cortex (Druzgal and D'Esposito, 2003; Lehmann et al., 2004; O'Craven and Kanwisher, 2000; Rossion et al., 2003b; Wojciulik et al., 1998). While face-selective voxels can be found in the fusiform gyri bilaterally in most individuals, we chose to use a right-lateralized FFA seed since lesion, electrophysiological, neuroimaging, and behavioral studies have shown the right hemisphere to play a dominant role in the perceptual analysis and recognition of faces (Bentin et al., 1996; Hillger and Koenig, 1991; Kanwisher et al., 1997; Landis et al., 1988; Rossion et al., 2003a,b).

Beta series correlation analysis

The univariate analysis described above was adapted such that every stage of every trial was modeled with a separate covariate (Fig. 1). Since there were 54 delayed recognition trials, each with three separate task stages, there were a total of 162 covariates of interest entered into the GLM. The model also included covariates of no interest to model the effects of shifting signal levels across runs. A band pass filter was used to attenuate frequencies above 0.25 Hz and below 0.01 Hz. The resulting parameter estimates (from correct trials only) were sorted according to the stage from which they were derived to form a beta series for each stage, reflecting the estimated activity of each voxel in each of the experimental trials that the subject performed correctly. While separate beta series could theoretically be constructed for the cue, delay, and probe stages of the incorrect trials, this was not done in this case due to the insufficient number of error trials (mean accuracy on this task was 97.5%). Stage-specific whole brain correlation maps were obtained by calculating the correlation of the FFA seed's beta series with that of all brain voxels. This was done separately for each of the task stages. The arc hyperbolic tangent transform was then implemented as described above (see Validation study methods), and the transformed correlation coefficients were divided by their standard deviation to produce a map of *z* scores.

To construct group correlation maps, the *z*-transformed correlation maps of the individual subjects were spatially normalized into standard MNI atlas space using routines from SPM2. Group level random-effects *t* tests were then conducted separately for the cue, delay, and probe stages of the task to identify voxels for which the mean of the individual subjects' transformed correlation coefficients was reliably greater than 0.

Results

Validation study

The Right-then-Left and Interleaved task conditions showed highly similar profiles of BOLD activity when analyzed with standard univariate analysis procedures (Sun et al., 2004). Both tasks produced activations bilaterally in primary motor cortex (M1), premotor cortex, supplementary motor area (SMA), and posterior parietal cortex (PPC), and a group level contrast revealed

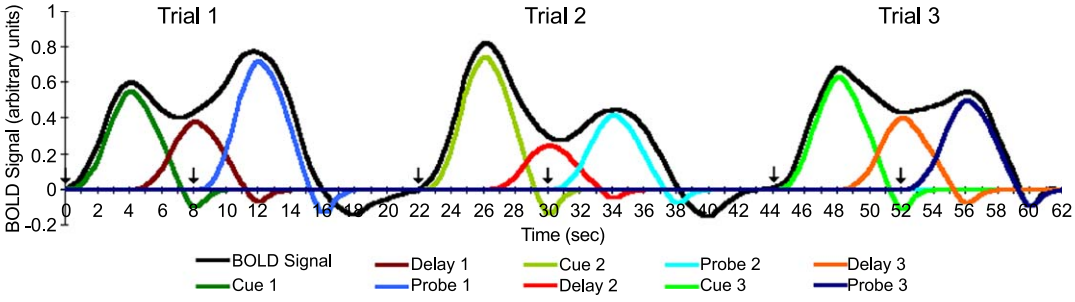

Fig. 1. Schematic of temporally shifted covariates modeling stage-specific activity for each trial of a delayed recognition task. The black line represents the hypothetical BOLD response in a task-related voxel across three trials. Arrows indicate the onsets of the cue and probe stimuli. While this voxel consistently responds more strongly to the cue and probe stages of the task than the delay period, its activity profile varies from trial to trial. This variability is captured by fitting a set of hemodynamic response functions to the data from each trial in the context of the general linear model. Each covariate (colored lines) yields a parameter estimate, or *beta value*, reflecting the amount that the covariate scales to best fit the BOLD data from a single stage of a single trial.

no significant activation differences between the two conditions. Group analysis of the coherence data, also presented in Sun et al. (2004), revealed that the left hemisphere M1 seed showed greater coherence with the contralateral motor/premotor cortex in the Interleaved condition than in the Right-then-Left condition. Similarly, the right hemisphere M1 seed showed greater coherence with the left motor/premotor cortex in the Interleaved condition. Thus, coherence analysis revealed a difference between the two tasks conditions that was not apparent from examining the univariate analysis.

Here, we first present the results of univariate, coherence, and beta series correlation analysis from a single subject to demonstrate the strength of individual subject analysis in native space (Fig. 2a–c). Next, to show that the results of the correlation analysis are robust at the group level, we present group level random-effects *t*-maps of the correlation data for each of the conditions (Fig. 2d).

Closely mirroring the results of the group level univariate analysis described by Sun et al. (2004), the single subject univariate analysis presented here (Fig. 2a) also yielded highly similar activation profiles for the Right-then-Left and Interleaved task conditions. This subject's coherence data were also representative of the group level coherence analysis reported by Sun et al. (2004). Coherence analysis using a left M1 seed revealed greater coherence in motor/premotor areas of the contralateral hemisphere in the Interleaved condition than in the Right-then-Left condition (Fig. 2b). When the same left M1 seed was used in the correlation analysis, the difference in interhemispheric functional connectivity between the two conditions was even more dramatic (Fig. 2c). As with the coherence analysis, right hemisphere motor/premotor areas revealed greater correlation with the left M1 seed in the Interleaved condition. A similar profile of effects emerged when coherence and correlation analyses were conducted using a right M1 seed (data not shown).

As an additional demonstration that the beta series correlation method is sensitive to the increase in functional connectivity between the motor cortices of the two hemispheres in the Interleaved condition, we computed the correlation between the beta series of this subject's right and left M1 seeds in each of the two conditions (Fig. 3). As these scatter plots show, the correlation between the right and left M1 seeds in the Interleaved condition ($r = 0.710$) is nearly double that seen in the Right-then-Left condition ($r = 0.378$). Thus, in the Interleaved condition, the magnitude of activation evoked in right M1 on any given trial (as indexed by its beta value) is tightly coupled to the magnitude of activation evoked in left M1 on that same trial. The parameter estimates of single trial activity obtained for the Right-then-Left trials are much more weakly correlated between these primary motor areas.

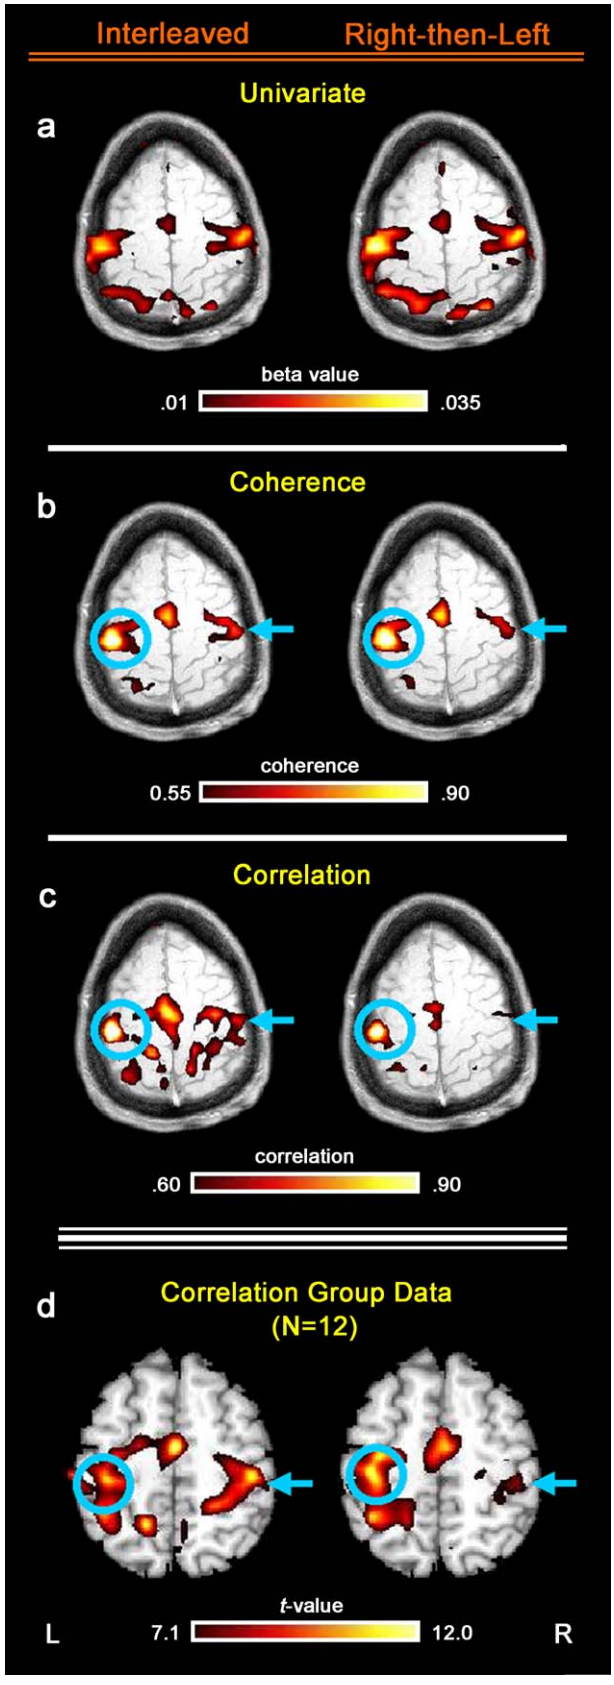

Fig. 2. Comparison of univariate, coherence, and beta series correlation data in the bimanual motor task. (a) A representative univariate activation map from a single subject reveals a very similar profile of brain activity in the Interleaved and Right-then-Left conditions. (b and c) Coherence and correlation analyses using a left M1 seed (located within the blue circle) both reveal stronger functional connectivity with the primary motor cortex of the contralateral hemisphere (blue arrow) in the Interleaved condition than the Right-then-Left condition. (d) Results of the beta series correlation analysis at the group level, using a left M1 seed. The left M1 seed correlates more strongly with contralateral motor regions in the Interleaved condition. These group *t*-maps are shown thresholded at $P < 0.01$ (Bonferroni corrected) and overlaid on an MNI-normalized template brain. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

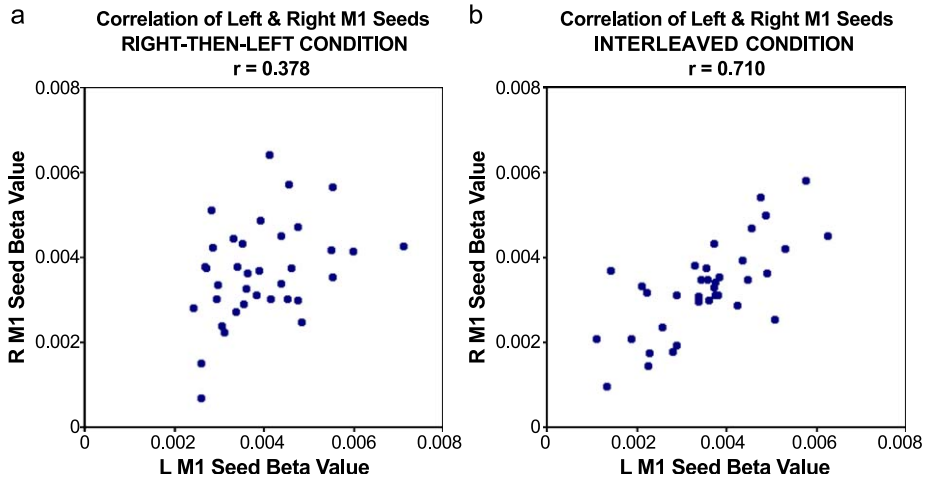

Fig. 3. The correlation between the beta series of the right and left hemisphere M1 seeds for a single subject in each task condition. This figure illustrates the stronger coupling between single trial beta values in the right and left M1 seeds during the performance of Interleaved trials (a) than during the performance of Right-then-Left trials (b). Each dot represents the data from a single trial.

When the correlation maps of each of the 12 subjects were spatially normalized and submitted to a group level random-effects analysis, a pattern of results similar to that seen in the single subject analysis emerged. Fig. 2d shows the group correlation *t*-maps for the Interleaved and Right-then-Left conditions when left M1 was used as the seed. Motor, premotor, and parietal regions contralateral to the left M1 seed were strongly correlated with the seed in the Interleaved condition. In the Right-then-Left condition, both the magnitude and extent of these interhemispheric correlations were substantially diminished. Motor and premotor regions ipsilateral to the left M1 seed also appear to be more strongly correlated with the seed in the Right-then-Left condition than the Interleaved condition. A similar effect was effect was noted in the group coherence data reported by Sun et al. (2004). When the right M1 was used as a seed, the group level correlation analysis similarly revealed increased correlations with motor, premotor, and parietal regions contralateral to the seed in the Interleaved condition relative to the Right-then-Left condition (data not shown).

Delayed recognition task

Beta series correlation analysis was performed on a data set of eight subjects performing a delayed recognition task as described in the methods. Here, we highlight several important observations regarding differences and similarities between univariate and beta series correlation data. The findings reported here are not comprehensive surveys of the correlation data obtained from this group analysis. They are intended only to illustrate the power of this

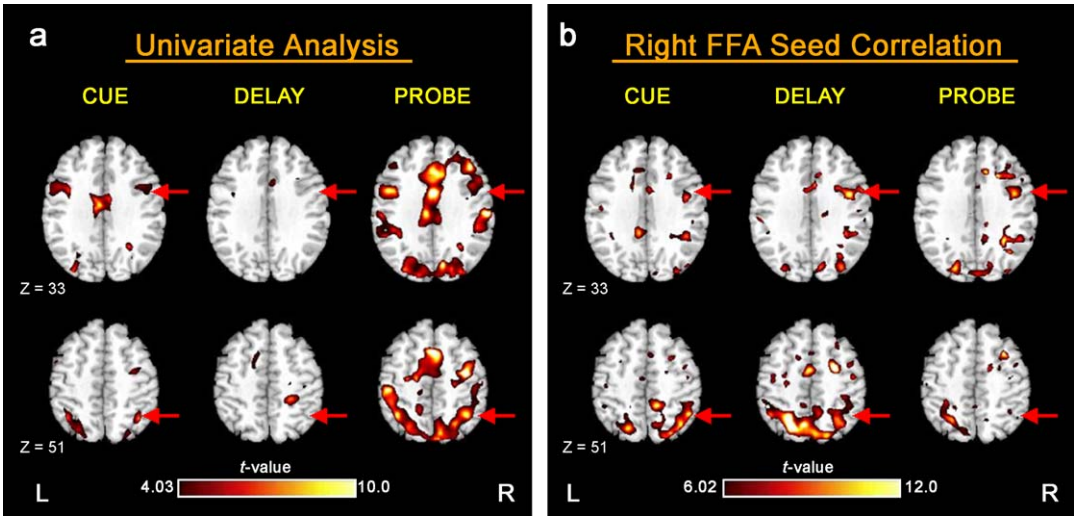

Fig. 4. Group-level *t*-maps of the univariate and beta series correlation data in the delayed recognition task. (a) The univariate data reveal activations in the right prefrontal and parietal cortices during the cue and probe stages of the task, but not during the delay period (red arrows). (b) In contrast, the correlation analysis using the right FFA as a seed reveals that these prefrontal and parietal regions are significantly correlated with the seed during *all three stages* of the task, although during the probe stage the correlation with the right parietal cortex is diminished. For visualization purposes, the univariate activation *t*-map is thresholded at $P < 0.005$ (two-tailed, uncorrected). The beta series correlation *t*-map is thresholded more stringently at $P < 0.0005$ (two-tailed, uncorrected), which meets a corrected threshold of $P < 0.05$ using Gaussian random fields (Friston et al., 1994) with a cluster extent requirement of 34 voxels. Data are shown overlaid on an MNI-normalized template brain. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

method in offering a unique set of data capable of generating novel interpretations distinct from the classic univariate analysis. The results of the group level analysis of the univariate data are depicted in Fig. 4a. During the cue stage of the delayed recognition task, we highlight significant activations (red arrows) that can be observed in the right posterior inferior frontal gyrus (IFG) bordering the premotor area, as well as the right intraparietal sulcus (IPS). During the delay period, no significant activation is seen in these same regions of the right IFG and right IPS. During the probe stage, widespread activations can again be seen in the right IFG and IPS in the same regions activated during the cue. Fig. 4b shows the results of the group level analysis of the correlation data using a right fusiform face area (FFA) seed. Since the FFA is considered to be an important brain region for face processing, it is of interest to determine which cortical regions interact with the FFA to facilitate the active maintenance of face stimuli in this working memory task. Note, that unlike the univariate analysis, regions within the right IFG bordering the premotor area, and the right IPS have significant correlations with the right FFA during the cue, delay, and probe stages of the task, although the FFA's correlation with the right IPS during the probe period has decreased substantially (red arrows). It is worth noting that the region of peak univariate activation in the right IFG is located slightly anterior to the region of peak correlation, suggesting that the IFG voxels that are most activated during the task are not necessarily the ones that interact strongest with the FFA. These maps serve to illustrate both similarities and differences between data obtained from these different methods on the same data set.

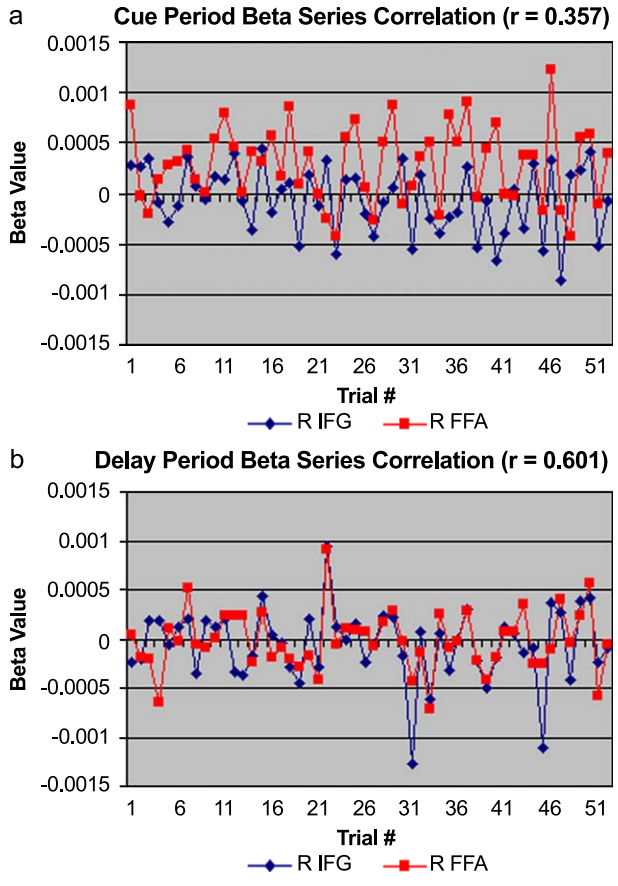

Fig. 5. Example of beta series correlations across the stages of the delayed recognition task in a single subject. The beta series from the right FFA seed (red) is compared with that of a single voxel in the right IFG (blue) during the cue period (a) and the delay period (b). This particular voxel correlates more strongly with the FFA seed during the delay period. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Observations of the group correlation data in Fig. 4b reveal an increase in the correlation between the right FFA and the right IFG from cue to delay. For illustrative purposes, we extracted the beta series from a single voxel in the right IFG from one of the subjects and demonstrate an increase in its correlation with the subject's FFA seed across these stages (Fig. 5).

It is important to evaluate if the beta values obtained from sequential stages of a task are independent from one another to determine if the correlations at each stage reflect interactions specific to that stage. In other words, if the beta series of sequential stages were highly correlated, it might suggest that the beta values were "contaminated" by residual activity of the preceding stage, perhaps secondary to the sluggish hemodynamic response. Under these circumstances, it would be difficult to make the case that the correlation maps generated for these separate task stages provide unique information. To address this issue, we extracted the cue, delay, and probe beta series from each subject's FFA seed, and then calculated the correlation of the cue and delay beta series, as well as the delay and probe beta series for each subject. The average correlation coefficient was 0.055 for the cue/delay correlation and 0.040 for the delay/probe correlation, indicating that the magnitude of the beta values within each stage of any given trial did not have any consistent relationship to the magnitude of the beta values in the subsequent stage derived from that same trial. Based on this near independence, we feel confident that the beta series correlation method is capable of detecting regions correlated with a seed during distinct stages, and if a region does remain correlated with a seed across stages, then the correlation is sustained due neurophysiological factors rather than contamination from the preceding stage.

Discussion

In this paper, we have introduced a new method for modeling functional connectivity in an event-related fMRI design, capable of characterizing network interactions in the individual subcomponents of a multistage cognitive task. To demonstrate the validity of the beta series correlation method, we applied it to an event-related fMRI data set (Sun et al., 2004) in which the subjects performed a simple motor task and showed that a greater demand for bimanual coordination resulted in increased correlation between motor areas of the two hemispheres. We then illustrated how this method could be used to analyze inter-regional interactions in a more complex multistage cognitive task, the delayed recognition task.

The validation study successfully demonstrated that the beta series correlation method is sensitive to changes in functional connectivity. Our results yielded the pattern expected based on extant EEG data and replicated those found with coherence analysis conducted on the same data set (Sun et al., 2004). Specifically, the Interleaved condition, which required increased bimanual coordination, induced greater correlation between right and left hemisphere motor regions.

Despite the fact that coherence analysis and beta series correlation analysis produced similar results when applied to this data set, there are a number of important differences between the methods. Most notably, coherence analysis operates on Fourier

transformed estimates of time series data whereas beta series correlations are based on parameter estimates derived from fitting a statistical model to the time series data. Because coherence analysis is a spectral measure most sensitive to low frequency BOLD fluctuations (<0.15 Hz), relatively long segments of continuous time series data are required to robustly measure inter-regional coherence. Thus, as noted in Sun et al. (2004), a block design consisting of entire runs of one condition or another would be the most optimal for coherence analysis, since it best preserves the spectral information contained in the time series data. The application of coherence analysis to event-related designs is limited to situations where only a single event occurs on each trial and a sufficiently long intertrial interval (ITI) is allowed to elapse between trials to allow the BOLD response to return to baseline. This is necessary because the time series data from all the trials of given condition need to be extracted and concatenated into a condition-specific time series. If segments much shorter than 16 s were concatenated, artificial frequencies could be introduced into the time series that could not be adequately removed by the windowing procedure used by Sun et al. (2004). Thus, coherence analysis could not be adapted for use on an experimental task in which several sequential stages of cognitive processing occur, such as the delayed recognition task, because it would not be possible to extract sufficient time series data from each of the task stages.

The capability of the beta series correlation method to model functional connectivity in such a multistage task represents its primary advantage over existing multivariate methods. However, we would not necessarily advocate using the beta series correlation method over coherence analysis for analyzing the data from a single-stage task with widely spaced trials, such as that of Sun et al. (2004). When applied to such a data set, the beta series correlation method computes the best fit of an HRF-derived reference waveform to the time series data from each trial, distilling 16 s (32 TRs) of data to a single beta value. Thus, it may unnecessarily filter out some of the spectral information contained in the time series data that a method such as coherence analysis could exploit to derive a more robust estimate of functional connectivity. However, the fact that the beta series correlation method was able to closely replicate the results of the coherence analysis despite its inherent reduction of 2304 time points worth of data into 72 beta values indicates that the most critical information contained in the time series data was not discarded.

To demonstrate how the beta series correlation method can be applied to the type of multistage cognitive task for which it was intended, we applied the method to analyze functional connectivity during the individual subcomponents of a fairly standard event-related visual working memory task. We chose the delayed recognition task because it provides an excellent example of a cognitive task that is frequently studied to evaluate component processes in cognitive psychology, single cell neurophysiology, and functional neuroimaging. Many theories of working memory postulate that prefrontal and parietal cortices interact with posterior visual association areas during the delay period to keep behaviorally relevant perceptual representations active when they are no longer present in the environment (Curtis and D'Esposito, 2003a; Fuster, 1997; Miller and Cohen, 2001). Evidence suggestive of such interactions comes from single unit recordings in monkeys performing working memory tasks (Fuster et al., 1985; Tomita et al., 1999) as well as human studies demonstrating the effect that lesions of the prefrontal cortex have on evoked potentials in posterior brain regions (Chao and Knight, 1998). fMRI technology could provide a potentially useful tool for further characterizing these interactions, since it can simultaneously measure activity levels throughout the entire functioning brain with relatively high spatial resolution.

While existing multivariate fMRI methods have succeeded in modeling functional connectivity during single-stage tasks (Buchel et al., 1999; Cabeza et al., 1997; Maguire et al., 2000; Rowe et al., 2002; Sun et al., 2004) these methods cannot be validly applied to multistage tasks. The principal reason for this is that unlike electrophysiological recordings, fMRI can only measure neural activity indirectly through the sluggish BOLD response, making it difficult to determine which aspect of the measured signal is attributable to each of the sequential stages of the task. For example, in the delayed recognition task presented in this paper, the subject encodes a face stimulus for 1 s, maintains the face across a 7-s delay period and then makes a same/different response to a second face stimulus. One cannot simply divide the time series data from each trial into discrete segments, sort by stage, and submit these "stage-specific" time series to correlation or coherence analysis. Such an approach would disregard the fact that the delay period time points are contaminated by residual cue period activity, and the probe period time points are contaminated by both residual cue and delay activity. Even more sophisticated approaches of modeling functional connectivity during the delay period fail to take residual encoding activity into account (e.g., Pessoa et al., 2002).

The beta series correlation method offers a novel approach to modeling functional connectivity during the individual stages of such a task. The method utilizes a general linear model approach similar to that employed by many univariate studies of the delayed recognition task (see D'Esposito et al., 2000 for review). By parameterizing the activity level during each of the individual task stages with a set of time-shifted hemodyamically derived covariates that compete for variance, estimates can be made of the activity specific to each stage that take into account the influence of adjacent stages. The beta series correlation method departs from traditional univariate methods by modeling stage-specific activity separately for every trial in the experiment. The variability inherent in the resulting stage-specific beta values is used to our advantage, with the premise that regions whose beta values fluctuate with the same variability for a given stage are functionally interacting during that stage.

Functional connectivity in the delayed recognition data set was explored by calculating the correlation between the fusiform face area (FFA) seed's beta series and that of every other voxel in the brain. Given the importance of the FFA in face processing, we focused our analysis on characterizing the network of cortical regions that interact with the FFA during the cue, delay, and probe stages of the task. While the network of regions exhibiting significant correlations with the FFA seed had considerable overlap with those regions exhibiting significant univariate activations, interesting differences between the correlation data and the univariate data emerged (Fig. 4). Whereas the univariate data revealed bilateral activations in the posterior IFG only during the cue and probe stages of the task and not during delay, the correlation analysis revealed a right-sided region in the posterior IFG that remained correlated with the right FFA seed throughout all three stages of the task. Univariate activity in the right IPS was also only seen during the cue and probe stages of the task. Like the right IFG, this region was correlated with the right FFA seed during both cue and delay, but unlike the univariate data, its correlation with the seed diminished during the probe stage. The presence of

significant delay period correlations between the FFA and regions of the prefrontal and parietal cortex supports models of working memory that suggest that higher order association cortices interact with posterior sensory regions to facilitate the active maintenance of a sensory percept across the delay period. The correlation between the FFA and these high-order regions appears be initially established during the encoding of the cue stimulus, and these correlations are largely sustained during the delay period, despite a dramatic decrease in the level of univariate activity.

While a detailed description of all the similarities and differences between the univariate and correlation results is beyond the scope of this report, some comments on the relationship between activations and interactions are warranted. Our results showed that the IFG was activated bilaterally during the cue and probe stages of the task, but only the right IFG was correlated with the right FFA seed. One might speculate that the right IFG exhibits high BOLD activity during the cue and probe stages because neurons in this region are actively firing to facilitate communication with the right FFA, as is suggested by the correlation data. The strongly ipsilateral nature of the correlation between this ventral temporal visual association area and the prefrontal cortex fits in with our knowledge of the anatomical connections between these regions (Petrides and Pandya, 2002; Webster et al., 1994). However, the left IFG also showed high BOLD activity during these stages, which raises the question as to what function these left IFG neurons subserve. A reasonable speculation is that this region is interacting with brain regions other than the right FFA, perhaps the left fusiform gyrus, or alternatively its activation reflects the implementation of a verbal strategy involved in the encoding and recognition of the face stimuli. One way to determine what brain regions are interacting with the left IFG is to define it as a seed and run a new beta series correlation analysis to identify the brain regions that are highly correlated with it. By performing correlation analyses with several different seeds, a more detailed profile of the multiple functional networks subserving the individual processing stages of a cognitive task can be revealed.

Our results also revealed regions that did not show high BOLD activity in the univariate analysis, and yet appeared to be strongly correlated with the seed. For example, the right IFG did not show any suprathreshold activation during the delay period when assessed with univariate analysis, but its correlation with the right FFA seed remained robust during the delay period. We have ruled out the possibility that this sustained correlation across stages is due to an artifact of the method; that is, the delay period beta values were simply mirroring the cue period beta values. We demonstrated that the FFA seed's cue period beta series is uncorrelated with its delay period beta series, and the same is true between delay and probe. Thus, correlation maps produced for individual task stages provide unique information about stage-specific functional interactions. What then might it mean for a region to be correlated with the seed while failing to show significant univariate activity? One possibility is that a region that does not exhibit BOLD signal above a specified significance threshold may still contain a population of neurons within those voxels that are communicating with the seed region. While the firing of these neurons might not be enough to drive the BOLD signal above threshold, it may be enough to synchronize the trial-to-trial fluctuations in the BOLD signal between two regions that are communicating. The ability of the beta series correlation method to detect high correlations in the presence of low univariate activity suggests that there is much more information contained in fMRI data than is typically revealed by univariate methods that simply record isolated activity levels exceeding a set threshold.

Although we have focused our attention on the application of the beta series correlation method to the delayed recognition task, the method is not restricted to this particular task design. This method could theoretically be applied to the analysis of functional connectivity during any task in which a series of processing stages occur in succession. It is important that the tasks are designed such that the stages can be adequately modeled by covariates with onsets spaced approximately 4 s apart to minimize collinearity (Zarahn et al., 1997). The design should also include a sufficiently long ITI at the end of every trial to ensure that the BOLD response is allowed to return to baseline levels. Thus, the beta series correlation method cannot be applied to rapid event-related designs. In such designs, the hemodynamic responses of adjacent trials are allowed to overlap in time, and by jittering the degree of this overlap across many trials, the activity evoked by each task condition can be uniquely estimated. However, since the beta series correlation method involves modeling each trial separately, there is no way to systematically account for this hemodynamic overlap when deriving condition-specific single-trial activity estimates. Commonly used experimental paradigms that would be good candidates for beta series correlation analysis include the oculomotor delayed response task (e.g., Curtis and D'Esposito, 2003b), cued response preparation (e.g., Pochon et al., 2001) or anticipation tasks (e.g., Sakai and Passingham, 2003), and tasks involving shifting or sustaining attention.

The example applications of the beta series correlation method provided in this paper have taken a seed-based approach to characterizing functional connectivity. Such an approach is most useful when a particular brain region is known to be involved in a cognitive process and the researcher wishes to determine what other brain areas are interacting with this "seed" region. However, the use of the beta series correlation method is not limited to the generation of exploratory whole-brain seed correlation maps. Once the researcher has characterized the key set of brain regions, or "nodes", involved in a cognitive task, a more explicit model of the inter-regional connections can be constructed. This model can be used to assess how the "effective connectivity" between nodes in the network changes across the stages of the cognitive task. Effective connectivity is defined as the influence one neural element has on another (Friston et al., 1993). The strength of the connections between regions is determined by analyzing the observed covariance structure in the context of an anatomically constrained model. Existing approaches to modeling effective connectivity have computed inter-regional covariances across subjects, as is necessary with positron emission tomography (PET) data (Cabeza et al., 1997; McIntosh et al., 1996, 1994), or from within-subject correlations in the fMRI time series data, which can only be done for blocked design experiments (Buchel and Friston, 1997; Buchel et al., 1999; Friston and Buchel, 2000; Goncalves et al., 2001; Maguire et al., 2000; Rowe et al., 2002; Toni et al., 2002). However, the beta series correlation method can easily be used to calculate the correlations between nodes in event-related fMRI experiments (an example of the correlation between two nodes can be seen in Fig. 3), and these correlations can then be entered into an explicit structural model to characterize the patterns of effective connectivity across the individual stages of a multi-stage cognitive task.

The beta series correlation method presented in this paper should provide a useful tool for researchers interested in exploring

network interactions in multistage cognitive tasks using event-related fMRI designs. Such designs cannot effectively be analyzed with currently available methods for modeling functional connectivity in a way that would allow inter-regional interactions to be assessed independently for the individual stages of the task. We expect that this method will aid in the generation of novel perspectives on brain function that are not offered with standard univariate analyses.

Acknowledgments

The authors thank Felice Sun and Charan Ranganath for the use of their fMRI data sets, Kevin McEvoy and Jon Kelley for assistance with data processing, and Felice Sun for helpful comments on an earlier version of this manuscript. This work was supported by a National Science Foundation Graduate Research Fellowship (J.R.) and grants from the National Institutes of Health and American Federation for Aging Research (A.G. and M.D.).

References

Aguirre, G.K., Zarahn, E., D'Esposito, M., 1998. The variability of human. BOLD hemodynamic responses. NeuroImage 8 (4), 360–369.

Andres, F.G., Mima, T., Schulman, A.E., Dichgans, J., et al., 1999. Functional coupling of human cortical sensorimotor areas during bimanual skill acquisition. Brain 122 (Pt. 5), 855–870.

Barde, L.H., Thompson-Schill, S.L., 2002. Models of functional organization of the lateral prefrontal cortex in verbal working memory: evidence in favor of the process model. J. Cogn. Neurosci. 14 (7), 1054–1063.

Bentin, S., Allison, T., Puce, A., Perez, E., et al., 1996. Electrophysiological studies of face perception in humans. J. Cogn. Neurosci. 8 (6), 551–565.

Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. 34 (4), 537–541.

Biswal, B.B., Van Kylen, J., Hyde, J.S., 1997. Simultaneous assessment of flow and BOLD signals in resting-state functional connectivity maps. NMR Biomed. 10 (4–5), 165–170.

Buchel, C., Friston, K.J., 1997. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modelling and fMRI. Cereb. Cortex 7 (8), 768–778.

Buchel, C., Coull, J.T., Friston, K.J., 1999. The predictive value of changes in effective connectivity for human learning. Science 283 (5407), 1538–1541.

Cabeza, R., McIntosh, A.R., Tulving, E., Nyberg, L., et al., 1997. Age-related differences in effective neural connectivity during encoding and recall. NeuroReport 8 (16), 3479–3483.

Chao, L.L., Knight, R.T., 1998. Contribution of human prefrontal cortex to delay performance. J. Cogn. Neurosci. 10 (2), 167–177.

Cordes, D., Haughton, V.M., Arfanakis, K., Wendt, G.J., et al., 2000. Mapping functionally related regions of brain with functional connectivity MR imaging. Am. J. Neuroradiol. 21 (9), 1636–1644.

Courtney, S.M., Ungerleider, L.G., Keil, K., Haxby, J.V., 1997. Transient and sustained activity in a distributed neural system for human working memory. Nature 386 (6625), 608–611.

Curtis, C.E., D'Esposito, M., 2003. Persistent activity in the prefrontal cortex during working memory. Trends Cogn. Sci. 7 (9), 415–423.

Curtis, C.E., D'Esposito, M., 2003. Success and failure suppressing reflexive behavior. J. Cogn. Neurosci. 15 (3), 409–418.

Della-Maggiore, V., Sekuler, A.B., Grady, C.L., Bennett, P.J., et al., 2000. Corticolimbic interactions associated with performance on a short-term memory task are modified by age. J. Neurosci. 20 (22), 8410–8416.

D'Esposito, M., Postle, B.R., Ballard, D., Lease, J., 1999. Maintenance versus manipulation of information held in working memory: an event-related fMRI study. Brain Cogn. 41 (1), 66–86.

D'Esposito, M., Zarahn, E., Aguirre, G.K., 1999. Event-related functional MRI: implications for cognitive psychology. Psychol. Bull. 125 (1), 155–164.

D'Esposito, M., Postle, B.R., Rypma, B., 2000. Prefrontal cortical contributions to working memory: evidence from event-related fMRI studies. Exp. Brain Res. 133 (1), 3–11.

Druzgal, T.J., D'Esposito, M., 2003. Dissecting contributions of prefrontal cortex and fusiform face area to face working memory. J. Cogn. Neurosci. 15 (6), 771–784.

Fisher, R.A., 1921. On the "probable error" of a coefficient of correlation deduced from a small sample. Metron 1, 3–32.

Friston, K.J., Buchel, C., 2000. Attentional modulation of effective connectivity from V2 to V5/MT in humans. Proc. Natl. Acad. Sci. U. S. A. 97 (13), 7591–7596.

Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S., 1993. Functional connectivity: the principal-component analysis of large (PET) data sets. J. Cereb. Blood Flow Metab. 13 (1), 5–14.

Friston, K.J., Worsley, K.J., Frackowiak, R.S.J., Mazziotta, J.C., et al., 1994. Assessing the significance of focal activations using their spatial extent. Hum. Brain Mapp. 1, 214–220.

Friston, K.J., Ashburner, J., Frith, C.D., Poline, J.-B., et al., 1995. Spatial registration and normalization of images. Hum. Brain Mapp. 2, 165–189.

Fuster, J.M., 1997. Network memory. Trends Neurosci. 20 (10), 451–459.

Fuster, J.M., Bauer, R.H., Jervey, J.P., 1985. Functional interactions between inferotemporal and prefrontal cortex in a cognitive task. Brain Res. 330 (2), 299–307.

Gerloff, C., Andres, F.G., 2002. Bimanual coordination and interhemispheric interaction. Acta Psychol. (Amst.) 110 (2–3), 161–186.

Goncalves, M.S., Hall, D.A., Johnsrude, I.S., Haggard, M.P., 2001. Can meaningful effective connectivities be obtained between auditory cortical regions? NeuroImage 14 (6), 1353–1360.

Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proc. Natl. Acad. Sci. U. S. A. 100 (1), 253–258.

Hampson, M., Peterson, B.S., Skudlarski, P., Gatenby, J.C., et al., 2002. Detection of functional connectivity using temporal correlations in MR images. Hum. Brain Mapp. 15 (4), 247–262.

Hillger, L.A., Koenig, O., 1991. Separable mechanisms in face processing: evidence from hemispheric specialization. J. Cogn. Neurosci. 3, 42–58.

Jeeves, M.A., Silver, P.H., Jacobson, I., 1988. Bimanual co-ordination in callosal agenesis and partial commissurotomy. Neuropsychologia 26 (6), 833–850.

Kanwisher, N., McDermott, J., Chun, M.M., 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. J. Neurosci. 17 (11), 4302–4311.

Landis, T., Regard, M., Bliestle, A., Kleihues, P., 1988. Prosopagnosia and agnosia for noncanonical views. An autopsied case. Brain 111 (Pt. 6), 1287–1297.

Lehmann, C., Mueller, T., Federspiel, A., Hubl, D., et al., 2004. Dissociation between overt and unconscious face processing in fusiform face area. NeuroImage 21 (1), 75–83.

Lowe, M.J., Mock, B.J., Sorenson, J.A., 1998. Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. NeuroImage 7 (2), 119–132.

Lowe, M.J., Dzemidzic, M., Lurito, J.T., Mathews, V.P., et al., 2000. Correlations in low-frequency BOLD fluctuations reflect cortico-cortical connections. NeuroImage 12 (5), 582–587.

Maguire, E.A., Mummery, C.J., Buchel, C., 2000. Patterns of hippocampal–cortical interaction dissociate temporal lobe memory subsystems. Hippocampus 10 (4), 475–482.

McIntosh, A.R., Grady, C.L., Ungerleider, L.G., Haxby, J.V., et al., 1994. Network analysis of cortical visual pathways mapped with PET. J. Neurosci. 14 (2), 655–666.

McIntosh, A.R., Grady, C.L., Haxby, J.V., Ungerleider, L.G., et al., 1996. Changes in limbic and prefrontal functional interactions in a working memory task for faces. Cereb. Cortex 6 (4), 571–584.

Mesulam, M.M., 1990. Large-scale neurocognitive networks and distributed processing for attention, language, and memory. Ann. Neurol. 28 (5), 597–613.

Miller, E.K., Cohen, J.D., 2001. An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. 24, 167–202.

O'Craven, K.M., Kanwisher, N., 2000. Mental imagery of faces and places activates corresponding stiimulus-specific brain regions. J. Cogn. Neurosci. 12 (6).

Ohara, S., Mima, T., Baba, K., Ikeda, A., 2001. Increased synchronization of cortical oscillatory activities between human supplementary motor and primary sensorimotor areas during voluntary movements. J. Neurosci. 21 (23), 9377–9386.

Pessoa, L., Gutierrez, E., Bandettini, P., Ungerleider, L., 2002. Neural correlates of visual working memory: fMRI amplitude predicts task performance. Neuron 35 (5), 975–987.

Petrides, M., Pandya, D.N., 2002. Comparative cytoarchitectonic analysis of the human and the macaque ventrolateral prefrontal cortex and corticocortical connection patterns in the monkey. Eur. J. Neurosci. 16 (2), 291–310.

Pochon, J.B., Levy, R., Poline, J.B., Crozier, S., et al., 2001. The role of dorsolateral prefrontal cortex in the preparation of forthcoming actions: an fMRI study. Cereb. Cortex 11 (3), 260–266.

Postle, B.R., D'Esposito, M., 1999. "What-Then-Where" in visual working memory: an event-related fMRI study. J. Cogn. Neurosci. 11 (6), 585–597.

Postle, B.R., Zarahn, E., D'Esposito, M., 2000. Using event-related fMRI to assess delay-period activity during performance of spatial and nonspatial working memory tasks. Brain Res. Brain Res. Protoc. 5 (1), 57–66.

Puce, A., Allison, T., Gore, J.C., McCarthy, G., 1995. Face-sensitive regions in human extrastriate cortex studied by functional MRI. J. Neurophysiol. 74 (3), 1192–1199.

Quigley, M., Cordes, D., Wendt, G., Turski, P., et al., 2001. Effect of focal and nonfocal cerebral lesions on functional connectivity studied with MR imaging. Am. J. Neuroradiol. 22 (2), 294–300.

Rama, P., Sala, J.B., Gillen, J.S., Pekar, J.J., 2001. Dissociation of the neural systems for working memory maintenance of verbal and nonspatial visual information. Cognit. Affect. Behav. Neurosci. 1 (2), 161–171.

Ranganath, C., D'Esposito, M., 2001. Medial temporal lobe activity associated with active maintenance of novel information. Neuron 31 (5), 865–873.

Ranganath, C., Johnson, M.K., D'Esposito, M., 2003. Prefrontal activity associated with working memory and episodic long-term memory. Neuropsychologia 41 (3), 378–389.

Rossion, B., Joyce, C.A., Cottrell, G.W., Tarr, M.J., 2003. Early lateralization and orientation tuning for face, word, and object processing in the visual cortex. NeuroImage 20 (3), 1609–1624.

Rossion, B., Schiltz, C., Crommelinck, M., 2003. "The functionally defined right occipital and fusiform "face areas" discriminate novel from visually familiar faces. NeuroImage 19 (3), 877–883.

Rowe, J.B., Toni, I., Josephs, O., Frackowiak, R.S., et al., 2000. The prefrontal cortex: response selection or maintenance within working memory? Science 288 (5471), 1656–1660.

Rowe, J., Friston, K., Frackowiak, R., Passingham, R., 2002. Attention to action: specific modulation of corticocortical interactions in humans. NeuroImage 17 (2), 988.

Rypma, B., D'Esposito, M., 1999. The roles of prefrontal brain regions in components of working memory: effects of memory load and individual differences. Proc. Natl. Acad. Sci. U. S. A. 96 (11), 6558–6563.

Sakai, K., Passingham, R.E., 2003. Prefrontal interactions reflect future task operations. Nat. Neurosci. 6 (1), 75–81.

Sakai, K., Rowe, J.B., Passingham, R.E., 2002. Parahippocampal reactivation signal at retrieval after interruption of rehearsal. J. Neurosci. 22 (15), 6315–6320.

Sala, J.B., Rama, P., Courtney, S.M., 2003. Functional topography of a distributed neural system for spatial and nonspatial information maintenance in working memory. Neuropsychologia 41 (3), 341–356.

Serrien, D.J., Brown, P., 2002. The functional role of interhemispheric synchronization in the control of bimanual timing tasks. Exp. Brain Res. 147 (2), 268–272.

Stein, T., Moritz, C., Quigley, M., Cordes, D., et al., 2000. Functional connectivity in the thalamus and hippocampus studied with functional MR imaging. Am. J. Neuroradiol. 21 (8), 1397–1401.

Sun, F.T., Miller, L.M., D'Esposito, M., 2004. Measuring inter-regional functional connectivity using coherence and partial coherence analyses of fMRI data. NeuroImage 21 (2), 647–658.

Tomita, H., Ohbayashi, M., Nakahara, K., Hasegawa, I., et al., 1999. Top-down signal from prefrontal cortex in executive control of memory retrieval. Nature 401 (6754), 699–703.

Toni, I., Rowe, J., Stephan, K.E., Passingham, R.E., 2002. Changes of cortico-striatal effective connectivity during visuomotor learning. Cereb. Cortex 12 (10), 1040–1047.

Webster, M.J., Bachevalier, J., Ungerleider, L.G., 1994. Connections of inferior temporal areas TEO and TE with parietal and frontal cortex in macaque monkeys. Cereb. Cortex 4 (5), 470–483.

Wojciulik, E., Kanwisher, N., Driver, J., 1998. Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study. J. Neurophysiol. 79 (3), 1574–1578.

Xiong, J., Parsons, L.M., Gao, J.H., Fox, P.T., 1999. Interregional connectivity to primary motor cortex revealed using MRI resting state images. Hum. Brain Mapp. 8 (2–3), 151–156.

Zarahn, E., Aguirre, G., D'Esposito, M., 1997. A trial-based experimental design for fMRI. NeuroImage 6 (2), 122–138.

Zarahn, E., Aguirre, G.K., D'Esposito, M., 1999. Temporal isolation of the neural correlates of spatial mnemonic processing with fMRI. Brain Res Cogn Brain Res 7 (3), 55–68.