



Determining significant connectivity by 4D spatiotemporal wavelet packet resampling of functional neuroimaging data

Rajan S. Patel,^{a,*} Dimitri Van De Ville,^b and F. DuBois Bowman^a

^aDepartment of Biostatistics, The Rollins School of Public Health, Emory University, 1518 Clifton Road, Suite 323, Atlanta, GA 30322, USA

^bBiomedical Imaging Group, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Received 15 August 2005; revised 5 January 2006; accepted 12 January 2006

An active area of neuroimaging research involves examining functional relationships between spatially remote brain regions. When determining whether two brain regions exhibit significant correlation due to true functional connectivity, one must account for the background spatial correlation inherent in neuroimaging data. We define background correlation as spatiotemporal correlation in the data caused by factors other than neurophysiologically based functional associations such as scanner induced correlations and image preprocessing. We develop a 4D spatiotemporal wavelet packet resampling method which generates surrogate data that preserves only the average background spatial correlation within an axial slice, across axial slices, and through each voxel time series, while excluding the specific correlations due to true functional relationships. We also extend an amplitude adjustment algorithm which adjusts our surrogate data to closely match the amplitude distribution of the original data. Our method improves upon existing wavelet-based methods and extends them to 4D. We apply our resampling technique to determine significant functional connectivity from resting state and motor task fMRI datasets.

© 2006 Elsevier Inc. All rights reserved.

Keywords: fMRI; Functional connectivity; Wavelet transform; Permutation test; Resampling

Introduction

The central nervous system consists of billions of interconnected neurons and neuronal ensembles. These intra- and interregional neuronal connections form the basis of neural processing in the human brain. An active area of neuroimaging research involves examining neurophysiologically based functional associations between spatially remote brain regions. We use the term “functional connectivity”, defined by Friston et al. (1993) as the “temporal correlations between spatially remote neurophysiological events”, when describing these functional associations. Traditional activation studies focus on determining distributed

patterns of brain activity associated with specific tasks. However, we can more thoroughly understand brain function by additionally studying the interaction of distinct brain regions, as a great deal of neural processing is performed by an integrated network of several regions of the brain.

Functional neuroimaging methods such as functional magnetic resonance imaging (fMRI) allow us to examine relationships between spatially distinct regions of the human brain. A common measure used to examine functional connectivity is the temporal correlation between brain voxels (Salvador et al., 2005; Hampson et al., 2002; Lowe et al., 1998; Grecius et al., 2003; Xiong et al., 1999). However, correlation in functional neuroimaging data cannot only be attributed to a potential neurophysiologically based association, but also to a host of other factors such as head movement, spatial realignment and normalization, scanner induced correlations, and partial voluming. These non-neurophysiological induced correlations are defined as “background correlations” by Breakspear et al. (2004). We develop a 4D spatiotemporal wavelet packet resampling method which allows us to test the null hypothesis that the correlations between two spatially remote brain regions are due to only the background correlation present in neuroimaging data. Our method generates surrogate data that preserves only the average background spatial correlation within an axial slice, across axial slices, and through each voxel time series, while excluding the specific correlations due to true functional relationships.

These surrogate datasets provide a null distribution of activation and correlation of voxels and between voxel pairs, respectively, against which we can test for significant correlations between voxel pairs and activations due to a stimulus or stimuli of interest for individual voxels. However, constructing surrogate data by simply permuting the original data will destroy the background correlations in the data which we wish to retain within our null distribution. Bullmore et al. (2001) show that resampling the detail coefficients of discrete wavelet transformed data and reconstructing by the inverse discrete wavelet transform preserves much of the correlations of the original data while spatially rearranging the exact correlations. Permuting a subset of the wavelet coefficients corresponding to the intracranial region of the brain allows us to

* Corresponding author. Fax: +1 404 727 1370.

E-mail address: rspate2@emory.edu (R.S. Patel).

Available online on ScienceDirect (www.sciencedirect.com).

construct surrogate data for which a large portion of the energy of the original data is retained in the intracranial space of the surrogate data (Breakspear et al., 2004).

Breakspear et al. (2004) developed a 2D spatiotemporal wavelet resampling method for neuroimaging data which addresses the problem of differentiating correlations that represent neurophysiologically based functional associations from background correlations. They perform a 2D spatiotemporal wavelet decomposition of preselected fMRI slices and randomly permute the detail coefficients of several levels of the wavelet decomposition and subsequently perform the inverse decomposition with the goal of matching the within slice horizontal, vertical, and temporal power spectral densities of the reconstructed surrogate data to that of the temporally demeaned original data. From hereon, we refer to the resampling method for fMRI data developed by Breakspear et al. (2004) as the “wavelet permutation method”. Power spectral densities are a function of correlograms and thus are used to represent the horizontal and vertical axial background correlations and temporal background correlations. However, these densities may not well represent the average background correlation structure of an axial fMRI slice as interhemispheric symmetry due to neurophysiological co-activation (Salvador et al., 2005) selectively influences these spectra. First, we show that the wavelet permutation method generates surrogate datasets whose spectral do not adequately match that of simulated fMRI data and thus construct a poor estimate of null non-neurophysiological correlation inherent in fMRI data. Secondly, the approach is limited in its practical use as it relies on permutation of 2D slices and thus is only able to determine significant connectivity within each slice. Recent connectivity studies (Salvador et al., 2005; Hampson et al., 2002; Lowe et al., 1998; Grecius et al., 2003; Xiong et al., 1999) focus on correlations throughout the entire intra-cranial space in 3D.

Novel extensions in this paper include the first 4D spatiotemporal resampling technique which generates surrogate data that preserves the average background spatial correlation within an axial slice, across axial slices, and through each voxel time series, while excluding the specific correlations due to true functional relationships; the first use of 4D resampling with wavelet packets, which offer a more complex and flexible analysis than traditional wavelet decompositions; the resampling of the exchangeable signs of wavelet coefficients, which allows for surrogate datasets which both adequately match the power spectra of the original while sufficiently randomizing specific activations and correlations without any tradeoff. Additionally, we are the first to employ 2D non-separable polyharmonic B-spline wavelets for resampling, which appropriately do not introduce preferential (horizontal and vertical) directions in the wavelet decomposition within axial slices like their separable counterparts.

Although the spatiotemporally resampled surrogate data has the same overall variability as the original, the entire energy of the original data is constrained within the intra-cranial space, whereas the energy of the surrogate data disseminates into extra-cranial voxels as well as into the temporally padded volumes. Thus, the variability of each intra-cranial voxel in the surrogate data is on average smaller than that of the original data as the number of voxels for which the energy is spread is much larger. We extend a 1D amplitude adjustment algorithm proposed by Theiler et al. (1992) to a 4D lattice and address when it is appropriate to use such an algorithm. Finally, we apply our resampling technique to determine significant functional connectivity from resting state and motor task fMRI datasets.

Materials and methods

Resting state fMRI data

A single healthy right-handed 35-year-old male subject was scanned at rest. The subject gave informed consent in accordance with a protocol reviewed and approved by the Emory University Institutional Review Board. The subject laid in a 3 T scanner with eyes open staring at a crosshair which was projected onto a mirror approximately two feet away. The subject was instructed to let his mind wander and to try not to think about anything specific. Following the acquisition of anatomical reference images, a functional imaging run was acquired. A T2-sensitive, single-shot echo-planar pulse sequence was used for acquisition of 210 functional images of 27 slices each (TR = 2000 ms, TE = 34 ms, flip angle = 90, FOV = 220 mm, and a 64×64 acquisition matrix that resulted in a voxel resolution of $3.44 \text{ mm} \times 3.44 \text{ mm} \times 4 \text{ mm}$). The 210 functional images were accrued only after the signal had reached a steady state. We utilized only the first 180 functional images as the subject experienced a sudden and permanent shift larger than 1 mm after the 180th scan.

All data were motion corrected to the first functional scan and subsequently spatially normalized to the MNI152 template using SPM2 (<http://www.fil.ion.ucl.ac.uk/spm/>). Spatial normalization resulted in a resampled image resolution of $3 \text{ mm} \times 3 \text{ mm} \times 4 \text{ mm}$. No spatial smoothing was done to minimize the amount of preprocessing that induces spatial correlation. We removed the effect of the underlying signal strength by temporally demeaning each voxel time series. For the 180 functional images we analyzed, the subject exhibited translation in any direction of less than 0.5 mm and rotation in any direction of less than 1° .

For our analysis, we zero padded the data so that each intracranial volume exists within a $80 \times 80 \times 48$ lattice. We also zero pad the time series to 192 volumes so that our spatiotemporally resampled data has dimension $80 \times 80 \times 48 \times 192$. These zero padding operations are required to guarantee a number of iterations for the wavelet packet decompositions.

Motor task fMRI data

A single healthy right-handed 22-year-old male subject was scanned during an experimental motor task. The subject laid in a Philips 3 T scanner viewing either ‘L’ or ‘R’ or a fixation cross, ‘+’, in successive blocks of 18 seconds that run in the pattern ‘L + R + L + R + L + R + L + R’. The first block coincides with the first volume. When viewing either ‘L’ or ‘R’, the subject is instructed to repeatedly tap the corresponding left or right index finger, respectively. A T2-sensitive, single-shot echo-planar pulse sequence was used for acquisition of 90 functional images of 30 slices each (TR = 3000 ms, TE = 30 ms, flip angle = 90, FOV = 240 mm, and a 128×128 acquisition matrix. The 90 functional images were accrued only after the signal had reached a steady state.

All data were motion corrected to the first functional scan and subsequently spatially normalized to the MNI152 template using SPM2 (<http://www.fil.ion.ucl.ac.uk/spm/>). Spatial normalization resulted in a resampled image resolution of $3 \text{ mm} \times 3 \text{ mm} \times 4 \text{ mm}$. No spatial smoothing was done to minimize the amount of preprocessing induced spatial correlation. We removed the effect of the underlying signal strength by temporally demeaning each voxel time series. For the 90 functional images we analyzed, the subject

exhibited translation in any direction of less than 0.5 mm and rotation in any direction of less than one degree.

For our analysis we zero padded the data so that each intracranial volume exists within a $80 \times 80 \times 48$ lattice. We also zero pad the time series to 96 volumes so that our spatiotemporally resampled data has dimension $80 \times 80 \times 48 \times 96$.

Wavelets

The discrete wavelet transform provides a multi-resolution, multi-frequency decomposition of a signal across a hierarchy of spatial (temporal) scales. The signal is iteratively decomposed at each scale into a set of approximation coefficients $a_{j,k}$ which provide low frequency signal information at the spatial resolution corresponding to that scale, and a set of detail coefficients $d_{j,k}$, which provide high-frequency information of the signal at the same spatial resolution (Mallat, 1989; Daubechies, 1990). As the scale increases, the spatial resolution and frequency range of information within the corresponding coefficients decrease. For an overview in neuroimaging applications, we refer the reader to Bullmore et al. (2004) and Van De Ville et al. (in press).

The discrete wavelet transform involves a family of wavelet functions $\psi_{j,k}$ and scaling functions $\phi_{j,k}$ generated through dilation (at scale j) and translation k of basis functions, the “mother” wavelet ψ and “father” wavelet ϕ , respectively. The convolution of $\psi_{j,k}$ and $\phi_{j,k}$ with the original signal produces $d_{j,k}$ and $a_{j,k}$, respectively. A signal S can be represented as a wavelet decomposition at scale J by

$$S = \sum_{k \in \mathbb{Z}} a_{J,k} \psi_{J,k} + \sum_{j \leq J} \sum_{k \in \mathbb{Z}} d_{j,k} \phi_{j,k}. \quad (1)$$

For a signal of length M , the number of detail coefficients at scale j is $N_j = 2^{-j} M$ when we use periodic interpolation at the boundaries of the finite signal and M is divisible by 2^j (Breakspear et al., 2004). Fig. 1a illustrates a one level wavelet decomposition and subsequent recomposition.

One of the most fundamental properties of wavelets is their approximate behavior as a multiscale differentiator. This property

is directly coupled to the number of vanishing moments of the wavelet; i.e., the wavelet (locally) removes polynomial trends up to a polynomial degree of $|\gamma| - 1$. In the context of stochastic processes, this feature translates into the highly praised decorrelating or whitening property, which has made wavelets a popular tool for statistical estimation of long-memory processes such as fractional Brownian motion (fBM) (Flandrin, 1992; Dijkerman and Mazumdar, 1994; Percival and Walden, 2000; Fan, 2003). For this kind of process, the correlation between wavelet coefficients rapidly decays as an inverse power of their distance within and between scales.

In the case of orthogonal wavelet transforms, the basis functions are chosen to be orthogonal. The major advantage of orthogonal transforms is their energy conservation property, which is critical for permutation or resampling methods. Here, we choose the popular orthogonal symmetric B-spline wavelets, introduced by Battle (1987). Recently, it has been shown that the essential properties of any wavelet transform, such as multi-scale differentiation, are only due to the convolutional B-spline factor within the scaling function (Unser and Blu, 2003). In the case of B-spline wavelets, the degree α of the B-spline scaling function automatically translates into its order of approximation $\gamma = \alpha + 1$, which imposes, at its turn, the differentiation order γ of the wavelet.

Resampling techniques which utilize the wavelet transform rely on the “exchangeability” property of detail coefficients by decomposing the signal, manipulating exchangeable detail coefficients, and subsequently reconstructing the signal. The resulting signal has similar correlation to the original, while the specific activations and correlations are randomized across time (space). Although the wavelet decomposition decorrelates detail coefficients as the decomposition scale increases yielding them more and more exchangeable, the coefficients at lower scales may not be sufficiently exchangeable in strongly correlated signals. The wavelet packet decomposition addresses the issue of insufficiently decorrelated detail coefficients by the decomposing detail coefficients at each scale as well as approximation coefficients.

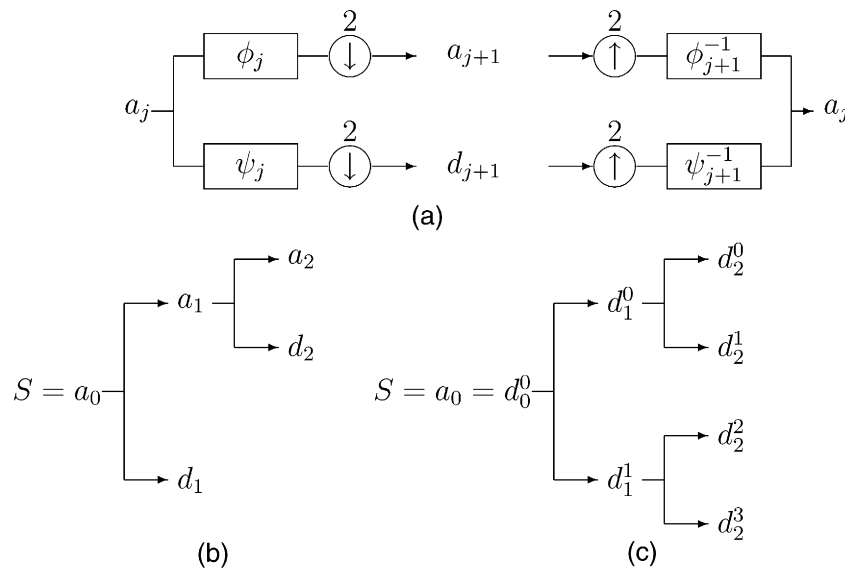


Fig. 1. (a) Single iteration of the analysis-synthesis filterbank for the wavelet transform. (b) Structure of the representation for the wavelet transform, 2 iterations. (c) Structure of the representation for the wavelet packet transform, 2 iterations.




Fig. 2. Frequency tiling for the 2D polyharmonic wavelet packet decomposition (3 iterations, $\gamma = 3$).

Wavelet packets

Wavelet packet decompositions, introduced by Coifman et al. (1992), recursively decompose both approximation and detail coefficients, resulting in a binary tree representation of the original signal (Fig. 1c). Since each set of detail coefficients corresponds to a specific frequency range, their additional recursive decomposition allows for a greater frequency resolution representation of the underlying signal. Furthermore, the additional decomposition of detail coefficients serves to further decorrelate coefficients at smaller scales permitting the manipulation of wavelet packet coefficients for resampling techniques.

Wavelet packets provide us with a very flexible way to represent the signal. From one side, wavelet packets still inherit the “differentiation property” of the normal wavelets; i.e., they still behave as a γ^{th} order differentiator. From the other side, they represent information on a unique and narrow frequency interval. Whitcher (in press) is the first to utilize wavelet packets to decompose 2D fMRI slices before permuting detail coefficients similarly to the wavelet permutation method.

The 2D polyharmonic wavelet transform

The wavelet transform can be extended to multiple dimensions in a straightforward way using the tensor product. Unfortunately, this approach introduces preferential (horizontal and vertical) directions by its sequential row and column-wise series of 1D decompositions while creating a “diagonal” cross-term that does not have a straightforward interpretation.

Here, we opt for the polyharmonic wavelet transform with quincunx subsampling (Van De Ville et al., 2005a), which is a non-separable 2D wavelet transform that is particularly suited to perform an isotropic treatment of the data. Quincunx subsampling provides a more isotropic treatment of the data and a slower progression through scale than dyadic subsampling. We refer to Appendix A for some fundamental characteristics of the polyharmonic wavelets.

In Fig. 2, we show the frequency tiling of the polyharmonic wavelet packet decomposition for three iterations. The contours of the essential support in the frequency domain for each subband are

indicated. Frequency zones that belong to the same subband are indicated in the same grayscale color. Each label refers to the sequence of scaling (“L”) and wavelet (“H”) filters that was used. Notice that the quincunx subsampling scheme creates a ring-like structure for each subband. In Fig. 3A, we show an example of a polyharmonic wavelet packet decomposition for a typical slice of our dataset. In Fig. 3B, the radial power spectrum densities of the 15 wavelet packet subbands (after four iterations) are plotted. Clearly, the spectra are well flattened, which indicates good decorrelation of the coefficients.

Wavelets for 4D spatiotemporal data

In the next section, we apply wavelet packet decompositions for the spatiotemporal resampling of fMRI data (3D volumes over time). The resulting surrogate datasets serve to estimate the null distribution of the spatiotemporal correlation inherent in fMRI. We utilize the 2D polyharmonic wavelet packets for each axial slice to obtain an isotropic treatment, and in the transverse and temporal directions, we apply a 1D B-spline wavelet packet decomposition.

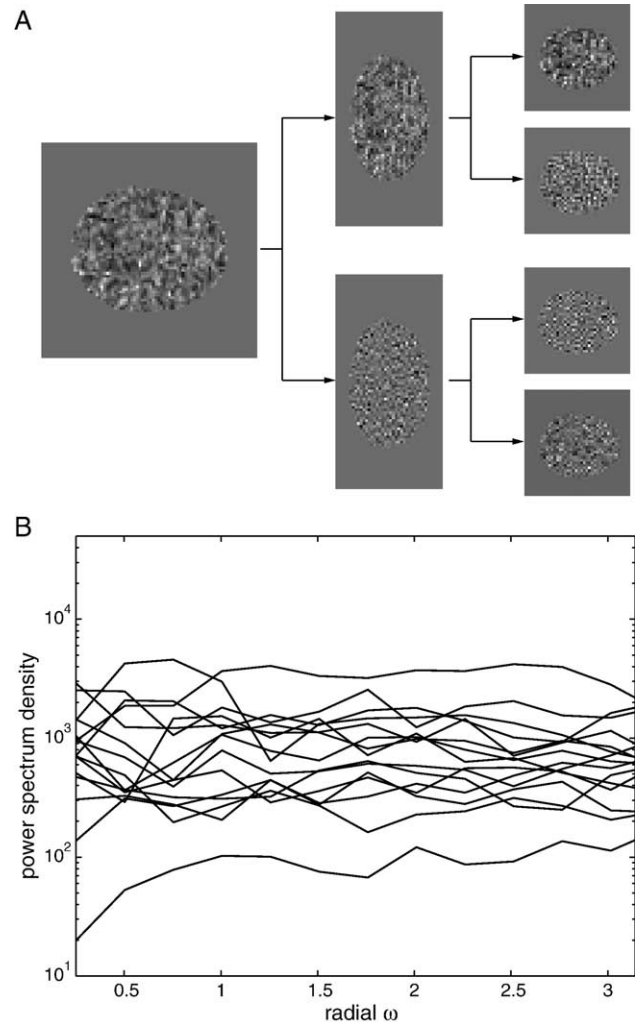


Fig. 3. (A) Polyharmonic wavelet packet decomposition (two iterations, $\gamma = 3$) of a typical slice from our resting state dataset. To better show the coefficients within each subband, the dynamic range has been adapted. The coefficients on the quincunx lattice after the first iteration are represented in a “squeezed” way. (B) Radial power spectrum densities of the 15 wavelet packet subbands after four iterations for the same test image.

Wavelet packet resampling of fMRI data

Axial, inter-axial, and temporal power spectra

For any signal, $S(x)$, the power spectrum of that signal gives the portion of the signal's power, or energy per unit of time, that occurs at a particular frequency. In contrast to a frequency spectrum, the power spectrum contains no spatial or phase angle information. A common method of calculating, the power spectrum for a given signal utilizes a Fourier transform:

$$\text{PSD}_{S(x)}(f) = F(S(x)) \times \overline{F(S(x))}, \quad (2)$$

where $F(S(x))$ is the Fourier transform of $S(x)$ and $\overline{F(S(x))}$ is its conjugate. Breakspear et al. (2004) give a detailed mathematical description of power spectra.

We can determine the average horizontal, vertical, inter-axial, and temporal power spectra for a fMRI volume, $Y(r, c, a, t)$, where r and c are the row and column in axial slice, a , at time

point, t , by taking a 4D Fourier transform, $F_{f_x}(Y)$, where $f_x = \{f_r, f_c, f_a, f_t\}$ is a 4D spatiotemporal frequency vector, multiplying by its conjugate and integrating over the remaining directions:

$$\text{PSD}_H(f_h) = \iiint F(Y) \times \overline{F(Y)} df_c df_a df_t, \quad (3)$$

$$\text{PSD}_V(f_v) = \iiint F(Y) \times \overline{F(Y)} df_r df_a df_t, \quad (4)$$

$$\text{PSD}_A(f_a) = \iiint F(Y) \times \overline{F(Y)} df_r df_c df_t, \quad (5)$$

$$\text{PSD}_T(f_t) = \iiint F(Y) \times \overline{F(Y)} df_r df_c df_a. \quad (6)$$

The average inter-axial and temporal spectra, PSD_A and PSD_T , summarize the background spatial correlation across axial slices

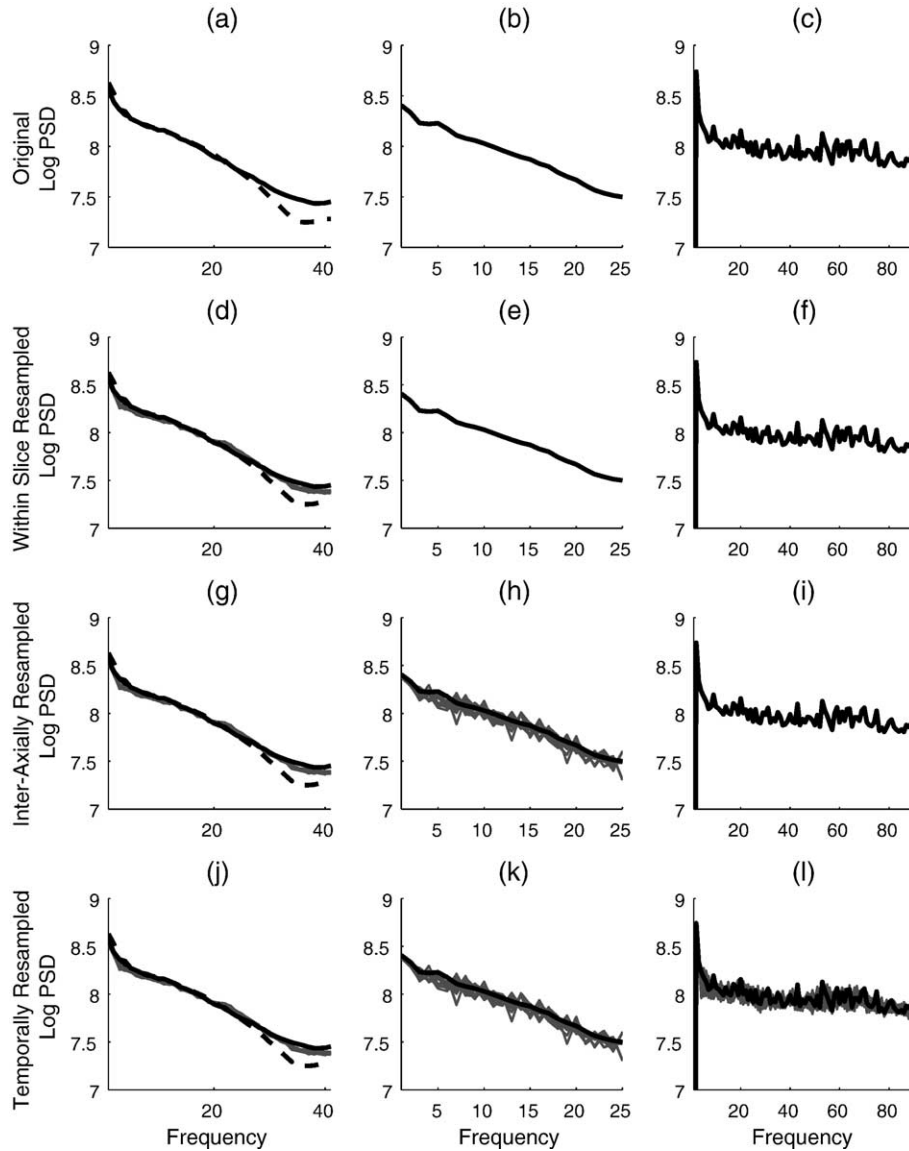


Fig. 4. (a–c) The original (a) horizontal (solid)/vertical (dashed), (b) inter-axial, and (c) temporal power spectral densities for the temporally demeaned resting state data. (d–f) The corresponding power spectra for 19 surrogate datasets (gray) after axial resampling. (g–i) Power spectra for 19 surrogate datasets after inter-axial resampling. (j–l) Power spectra for 19 surrogate datasets after temporal resampling.

and background temporal correlation over the entire voxel time series, respectively. The average horizontal and vertical spectra, PSD_H and PSD_V , give estimates of the inter-axial spatial correlation. We are interested in preserving the average correlation within an axial slice and not necessarily the horizontal and vertical correlations as non-neurophysiological spatial correlation should have no directional bias within an axial slice. It is evident from Fig. 4a that there is a distinct difference in the high frequency range between the horizontal (perpendicular to the coronal plane) and vertical (perpendicular to the sagittal plane) power spectra within an axial slice. A larger percentage of the power of the vertical spectra lies in the low frequency range, suggesting a greater long range correlation along the sagittal axis than the coronal axis. This may suggest significant inter-hemispheric correlation due to neurophysiological co-activation as suggested by Salvador et al. (2005).

Construction and interpretation of axial, inter-axial, and temporal spectra require the assumption of stationarity within an axial slice, across axial slices, and over time, respectively. Stationarity exists in a stochastic process when the probability density function of the elements of the signal does not change over time or position.

4D spatiotemporal wavelet resampling

The square of the detail coefficients of the wavelet packet decomposition represent the energy of the corresponding frequency subband at their respective location, either temporally or spatially. The sign of the detail coefficient, however, gives phase or activation information at its corresponding location. We aim to generate surrogate data which has similar variance and background correlation as the original data, while randomizing specific activations and correlations spatiotemporally. Thus, we develop a novel wavelet packet resampling method which involves resampling, with replacement, the signs of the detail coefficients. However, before resampling, we must ensure that the signs of the detail coefficients are “exchangeable”.

The detail coefficients decorrelate as the scale of the decomposition increases. We utilize this property of the 2D and 1D wavelet decompositions to iteratively build wavelet packet trees, further decomposing each scale until each sub-band is sufficiently decorrelated. We determine whether each subband is sufficiently decorrelated by conducting a spatial or temporal join count test (Cliff and Ord, 1973). Through a join count test (see Appendix B for details), we are able to test the null hypothesis that positive coefficients occur randomly throughout 1D or 2D time or space. Although the join count test yields evidence of spatial or temporal dependence, it cannot be treated as a test for proof of independence. Critical assumptions required of the resampling method are the spatial/temporal independence of the signs of the detail coefficients within each sub-band and the distribution of the detail coefficients is symmetric and has zero median.

Resampling axial slices. The specific spatial correlations within axial slices can be randomized throughout the data while retaining the average background spatial correlation. This is done in the following manner:

- (1) Determine the level of the wavelet packet decomposition, L_s , by iteratively decomposing each axial slice of each volume of each session and conducting a join test for spatial independence at each level. L_s is the first scale at which the average P value of the join test for each slice is greater than $\alpha = 0.15$.

- (2) Perform the L_s -level 2 dimensional wavelet packet decomposition of each slice of each volume of the entire dataset such that the detail coefficients are sufficiently spatially decorrelated at the L_s^{th} level as indicated by the spatial join test.
- (3) For each slice, resample the signs of the detail coefficients by an element-wise multiplication of the i detail coefficient subbands, where $i = 1, \dots, 4^{L_s-1}$, by a matrix $A^{(i)}$, where $A_{r,c}^{(i)} = 2 \times B_{r,c}^{(i)} - 1$ where $B_{r,c}^{(i)}$ is drawn from a Bernoulli(0.5) distribution. We multiply the wavelet packet decomposition of each slice of each volume by the same set of resampling matrices, A .
- (4) Perform the L_s -level inverse wavelet packet recomposition to reconstruct our axially resampled surrogate data.

Since the signs of the detail coefficients of each slice of a volume are resampled by the same resampling matrices, A , the correlations across slices remain unchanged (Breakspear et al., 2004; Prichard and Theiler, 1994). Similarly, since each slice of each volume of the entire time series is resampled by A , the temporal correlation of each voxel time series remains unchanged. Fig. 4d gives the horizontal/vertical power spectra for the surrogate data. For our data, we find that the detail coefficients at $L_s = 4$ are sufficiently decorrelated via the join test described above.

Unlike the original data, the surrogate data has similar horizontal and vertical spectra thus removing directional correlation bias which may be induced by symmetric neurophysiological activation (Salvador et al., 2005). The utilization of 2D non-separable wavelets provide this horizontal and vertical correlation averaging that we desire, as under the null hypothesis, there should be no difference in correlation between the horizontal and vertical directions within an axial slice. A comparison of 2D non-separable and separable wavelets using our resampling technique is shown in Fig. 5. The separable wavelets retain the horizontal/vertical difference in the surrogate data, and thus does not provide an accurate depiction of our null hypothesis.




Fig. 5. The original (black), surrogate using non-separable wavelets (blue), and surrogate using separable wavelets (red), horizontal (solid)/vertical (dashed) power spectral densities for the temporally demeaned resting state data. The surrogate spectra represent one typical surrogate realization.

Resampling across axial slices. The specific spatial correlations across axial slices can be randomized throughout the data while retaining the average background correlation across slices via a 1D wavelet packet decomposition and resampling (Fig. 6). We perform the L_a -level 1D wavelet packet decomposition for each transverse vector through the axial slices of each volume such that the detail coefficients are sufficiently decorrelated at the L_a^{th} level as indicated by the join test. We resample the detail coefficients 1 dimensionally in a similar manner to that described above and reconstruct our inter-axially resampled surrogate data by subsequently performing an inverse wavelet packet recombination. We find that the detail coefficients at $L_a = 4$ are sufficiently decorrelated at each subband. Fig. 4h gives the inter-axial power spectra for the 19 surrogate datasets. The horizontal/vertical and temporal spectra do not change during this step.

Resampling time series. We randomize the specific temporal correlations by performing a 1D wavelet packet resampling for each voxel's time series. Since the polyharmonic wavelet decompositions inherently utilize a periodic extension at the edges, we zero pad the time series to alleviate wrap around effects of the wavelet coefficients and subsequently perform a $L_t = 5$ level decomposition of the zero padded time series. The signs of the detail coefficients of each of the $2^5 - 1$ subbands are sufficiently decorrelated at $L_t = 5$ according to the join test. Using the same set of resampling multiplication vectors for each time series, we are able to keep the spatial correlations the same while randomizing the specific temporal correlations alone. Fig. 4l gives the temporal power spectra for the 19 surrogate datasets. Neither of the spatial spectral densities are affected by this step.

After performing the inverse wavelet packet recombination, we truncate our time series to the original, non-padded length (Fig. 6). Since the surrogate data has non-zero values in the truncated scans, we diminish the overall energy of our voxel time series when compared to the original. We propose an amplitude adjustment step which addresses this issue of energy loss for the surrogate data.

Choice of order γ

Both the 1D B-spline wavelets and 2D polyharmonic wavelets have a tuning parameter, which is their order γ . To select this parameter, we follow the common approach that considers fMRI data as a realization of a fBM, for which the power spectrum can be considered as proportional to $1/\|\omega\|^{2H + N_d}$, where H is the Hurst exponent and N_d is the dimensionality. It has been shown that the order should be chosen such that $\gamma > 2H + N_d$ (Dijkerman and Mazumdar, 1994). We choose $\gamma = 4$ for each decomposition (axial, inter-axial, and temporal).

Amplitude-variance adjustment step

Although the surrogate data has the same overall variability, or energy, as the original, the entire energy of the original data is constrained within the intra-cranial space, whereas the energy of the surrogate data disseminates into extra-cranial voxels as well as into the temporal padding volumes. Ideally, the number of extra-cranial voxels and number of temporal padding volumes should be minimized, thus reducing the likelihood of energy loss into extra-cranial space or into the padded temporal volumes. We pad the data so that the dimensionality of the data allows a wavelet packet decomposition for which the detail coefficients are sufficiently spatially (or temporally) decorrelated. Thus, the variability of each intra-cranial voxel in the surrogate data is on average smaller than

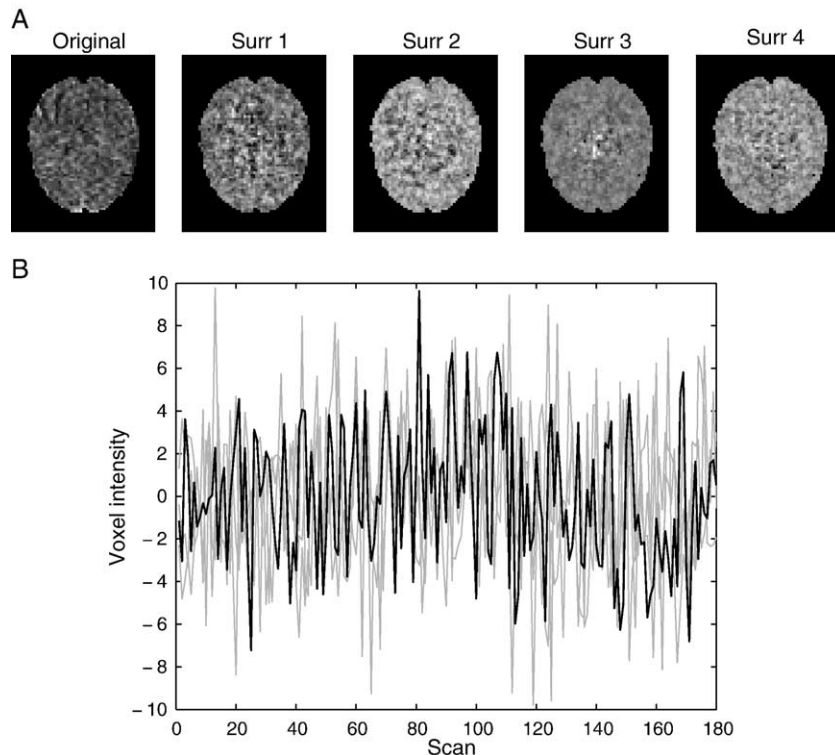


Fig. 6. (A) Example axial slices of the original and 4 surrogate datasets after all 3 resampling steps. (B) Example voxel time series for randomly selected intracranial voxel for original (black and bold) for surrogate data (gray).

that of the original data as the number of voxels for which the energy is spread is larger. The distribution of the variances of each intra-cranial voxel time series is given in Fig. 7.

If the measure of connectivity is not influenced by the large difference in the variance distribution between the original and surrogate data, then an amplitude-variance adjustment step may not be necessary. Otherwise, we propose an extension to an amplitude adjustment step developed by Theiler et al. (1992).

This proceeds by generating V time series, \mathbf{G}_i , from a Gaussian distribution with mean 0 and variance $\sigma_{(i)}^2$ for $i = 1, \dots, V$, where V is the number of intra-cranial voxels and $\sigma_{(i)}^2$ is the i^{th} smallest sample variance of the time series of the V intra-cranial voxels. Let $G_{i(j)}$ be the j^{th} smallest element in \mathbf{G}_i . Then, assign $G_{i(j)}$ to the j^{th} smallest element of the surrogate data voxel time series with the i^{th} smallest variance of the intra-cranial voxels of the surrogate data. The resulting adjusted surrogate data will have approximately the same variance distribution as the original data, while approximately retaining the spatial and temporal power spectral densities of the unadjusted surrogate data. A critical assumption of this adjustment step is that each voxel time series takes a Gaussian distribution, although each distribution can have its own distinct variance.

We choose the smallest possible level of decomposition at each step of the wavelet packet resampling algorithm at which the evidence of dependence among the wavelet coefficients is below some threshold in order to minimize the energy of the surrogate data disseminates into extra-cranial space. If amount of energy disseminates into extra-cranial space is small, and amplitude adjustment step may not be necessary, however if it is large, the variance of each surrogate voxel time series will be significantly smaller than the original data. Due to the additional assumptions required by the amplitude adjustment step, we recommend its use only when the measure of connectivity is influenced by the variance of each time series.

Results

Simulation study

Simulated dataset

The simulated fMRI dataset contains simulated neurophysiological correlation between three sets of brain regions as well as




Fig. 8. Axial slice at $z = +18$ with labeled regions of simulated neurophysiological correlation. Voxels within similarly colored regions are neurophysiologically correlated.

simulated non-neurophysiological background spatial and temporal correlation inherent in fMRI. Our wavelet packet sign resampling method is used to estimate the background correlations and determine significant neurophysiological correlations. We compare our method to the wavelet permutation method which has only been developed for 2D fMRI axial slices.

We simulate 192 functional images of a single fMRI slice (Talairach $z = +18$ mm) in the following manner. Let voxels belonging to the left or right superior temporal gyrus (*LSTG*, *RSTG*) comprise the set of voxels Ω_a . Let voxels belonging to the left and right caudate nucleus (*LCN*, *RCN*) comprise the set of voxels Ω_b . Let voxels belonging to the left frontal middle gyrus or left middle occipital gyrus (*LFMG*, *LMOG*) comprise the set of voxels Ω_c . Fig. 8 illustrates the location of the voxels within Ω_a , Ω_b , and Ω_c .

(1) Simulate neurophysiological correlations within 3 sets of voxels, Ω_a , Ω_b , and Ω_c . For each functional image v , simulate $N(0, 1)$ data with correlation = 0.5 for all voxels in Ω_a , simulate $N(0, 1)$ data with correlation = 0.3 for all voxels in Ω_b , and simulate $N(0, 1)$ data with correlation = 0.1 for all voxels in Ω_c . For all remaining intra-cranial voxels, simulate independent $N(0, 1)$ data.




Fig. 7. Time-series variance distribution of intra-cranial voxels of the original data (black), unadjusted surrogate data (white), and adjusted surrogate data (gray).

(2) Induce spatial correlation inherent in fMRI by smoothing each image with a Gaussian kernel with full width half max (FWHM) of 1.5 voxels. We choose a FWHM of 1.5 to closely match the background spatial correlation in our resting state data.

(3) Induce $1/f$ -like temporal correlation inherent in fMRI (due to *BOLD* response and other factors) by taking Fourier transform of each intra-cranial voxel time series and multiplying the frequency component of each by c/f , where f is the frequency associated with the frequency component and c is a constant. We choose $c = 10$ to closely match the background temporal correlation in our resting state data. The real component of the subsequent inverse Fourier transform is the temporally correlated fMRI signal.

The simulated dataset consists of only one slice so that the method introduced in this paper can be compared to the wavelet permutation method, which only considers 2D functional slices.

Simulation methods

The wavelet permutation method involves random permutation of intra-cranial detail coefficients at each scale of a wavelet decomposition. For our simulated data, each slice is decomposed spatially with a wavelet transform through 4 scales and each time series is decomposed temporally through 6 scales. Breakspear et al. (2004) suggest the Daubechies (Daubechies, 1990) family of wavelets to decompose the data. A Daubechies wavelet of order 6 is chosen to decompose the data based on the most adequate matching of spatial and temporal spectra.

We examine significant correlations from a voxel within *LSTG*, *LCN*, and *LFG* to each other intra-cranial voxel. For the simulation study, 19 surrogate datasets are generated using each method. For each voxel pair, a one-tailed test of significance is performed by comparing the correlation of the voxel pair in the

original simulated data r_{orig} against the rank ordered measures among the 19 surrogate datasets r_{sur} . The null hypothesis is rejected is $r_{\text{orig}} > \max(r_{\text{sur}})$.

Simulation results

The spectra of the surrogate datasets of both algorithms is given in Fig. 9. Our method generates data with surrogate spectra similar to those of the original data both spatially and temporally, however the horizontal and vertical spectra of the surrogate data generated by the wavelet permutation method fail to adequately match those of the original data. The reason for this is likely that the wavelet detail coefficients are not sufficiently decorrelated at lower scales, and thus permutation of these coefficients destroys higher frequency correlations. Breakspear et al. (2004) begin to address this problem by introducing block permutation of wavelet coefficients, however, they do not suggest a way to implement such a method for images where there is an irregular sub-domain (intra-cranial voxels within an axial slice). Block permutation of the wavelet coefficients with an adequate block size would address the issue of matching power spectra, however as the block size increases (allowing for a more adequate match of the power spectra), the number of possible surrogate sets decrease. Consequently, block permutation induces a trade-off between adequate matching of the power spectra and sufficient randomization of the specific neurophysiological activation and correlations in the data Breakspear et al. (2004).

Results in the form of thresholded connectivity maps are given for both our wavelet packet sign resampling algorithm (Figs. 10A–C) and the wavelet permutation method (Figs. 10D–F). Results between the two methods are similar for connectivity within strongly connected regions (Figs. 10A, B, D, E), where




Fig. 9. Power spectral densities of original (black) and surrogate (gray) data. (a) gives the horizontal and vertical spectra of the surrogate data generated by the method developed in this paper. (b) gives the corresponding temporal spectra. (c) gives the horizontal and vertical spectra of the surrogate data generated by the method developed in Breakspear et al. (2004). (d) gives the corresponding temporal spectra.




Fig. 10. Thresholded connectivity results for our method (A–C) and wavelet permutation method (D–F) where $r_{orig} > \max(r_{surv})$ for the corresponding voxel seed ((A, D)-LSTG, (B, E)-LCN, (C, F)-LFMG) and intra-cranial voxel.

the simulated Pearson correlation before inducing background spatial correlation was 0.5 for Ω_a and 0.3 for Ω_b . However, for weakly correlated regions (Figs. 10C, F), where the simulated Pearson correlation before inducing background spatial correlation was 0.1 for Ω_c , the wavelet packet sign resampling method is able to still correctly extract areas of statistically significant correlation, whereas the wavelet permutation method fails to do so.

Resting state study

Functional imaging studies have shown that certain brain regions, such as the posterior cingulate cortex (PCC) and the ventral anterior cingulate cortex (vACC) show a greater activity during the resting state than during cognitive tasks (Shulman et al., 1997; Mazoyer et al., 2001). The finding that certain brain regions exhibit a decreased level of activity during




Fig. 11. Map of significant neurophysiological motor task correlation to the PCC. The numbers above each image correspond the z plane coordinates of Talairach and Tournoux. The left hemisphere of the brain corresponds to the left side of the image. Four significant clusters are labeled A–D: A—Superior frontal gyrus (BA 6), B—Precuneus (BA 7), C—Superior occipital gyrus (BA 19), D—Superior frontal gyrus (BA 10).

cognitive tasks leads to the hypothesis that these regions are components of a default mode of cognitive processing (Raichle et al., 2001).

To illustrate our method, we examine the functional connectivity of the PCC by determining significant functional correlations between the average time series of a contiguous four voxel cluster in the PCC (Talairach coordinates: $-12, -47, 28$) and the set of all other intra-cranial voxels given the single subject resting state fMRI dataset described in Section 2.1. We construct 19 spatiotemporally resampled surrogate datasets and determine significance of a pairwise correlation if the correlation of the observed data exceeds the maximum correlation of the corresponding pair among the 19 surrogate datasets. This results in a pairwise Type I error rate of $\alpha = 0.05$ not corrected for multiple comparisons. We use a low pass filter on each voxel time series to consider only low frequency correlations (<0.1 Hz) as Cordes et al. (2002) suggest that only the “synchronicity of low frequency fluctuations in functionally related regions suggests the existence of neuronal connections characterizing a widespread cortical network”. Greicius et al. (2003) address a similar problem without however addressing and adjusting for non-neurophysiologically based background correlations. We do not perform the amplitude adjustment step for these datasets as correlation is not influenced by variance bias.

Fig. 11 shows 6 axial slices of the brain for which there were clusters of significant correlations to a seed of the average time series of four contiguous voxels in the PCC. Four significant clusters include the superior frontal gyrus (Brodmann’s area (BA) 6 and 10), precuneus (BA 7), and the superior occipital gyrus (BA 19). The connectivity of the PCC with the precuneus corresponds well with the results of Greicius et al. (2003). Further supporting our results, Jiang et al. (2004) conclude in their resting state functional connectivity study that BA 7 and BA 6 are important nodes in the resting state functional network.

Motor task study

To illustrate our method with an activation study, we examine the connectivity from the average time series of a contiguous four voxel seed in the primary motor cortex

(Talairach coordinates: $35, -10, 36$). We again construct 19 spatiotemporally resampled surrogate datasets and determine significance of a pairwise correlation if the correlation (<0.1 Hz) of the observed data exceeds the maximum correlation of the corresponding pair among the 19 surrogate datasets.

Fig. 12 shows 4 axial slices of the brain for which there were clusters of significant correlations to a seed in the primary motor cortex (BA 4) in the right hemisphere. Marker A in Fig. 12 labels a significant positive connectivity in the right cingulate gyrus (BA 24) and a corresponding negative connectivity with the left cingulate gyrus. Marker B labels a significant positive connectivity with the supramarginal gyrus (BA 40) in the parietal cortex. Also evident in the $+32$ mm axial slice is an interesting pattern of connectivity where the right anterior and left posterior (BA 7 and BA 9) intra-cranial regions exhibit a strong negative connectivity to the primary motor cortex while the left anterior and right posterior intra-cranial regions exhibit a strong positive connectivity with the primary motor cortex. Finally, marker C labels a strong positive connectivity with the left middle frontal gyrus (BA 10).

Discussion

We develop a 4D spatiotemporal wavelet packet resampling technique for testing the null hypothesis of no neurophysiological functional relationship between a given voxel pair. The method generates surrogate data that preserves the average background spatial correlation within an axial slice, across axial slices, and through each voxel time series, while excluding the specific correlations due to true functional relationships. Wavelet packets offer a more complex and flexible analysis than traditional wavelet decompositions in that they decompose the frequency spectra into smaller subbands which can be resampled independently. Multi-scale wavelet packet decompositions allow us to resample the exchangeable signs of wavelet coefficients without the need for resampling within a subdomain or block resampling as developed by Breakspear et al. (2003). Our method utilizes two-dimensional isotropic polyharmonic b-spline wavelets to preserve the average background correlation within each slice and subsequent one-dimensional




Fig. 12. Map of significant neurophysiological motor task correlation to the primary motor cortex (Talairach coordinates: $35, -10, 36$) in the right hemisphere. The numbers above each image correspond the z plane coordinates of Talairach and Tournoux. The left hemisphere of the brain corresponds to the left side of the image. Three significant clusters are labeled A–C: A—Cingulate gyrus (BA 24), B—Supramarginal gyrus (BA 40), C—middle frontal gyrus (BA 10).

decompositions to preserve the average transverse and temporal correlations.

An important property of surrogate datasets for any resampling or permutation technique is that they are sufficiently independent. We check the independence of our surrogate datasets by determining the average temporal correlation between each corresponding intra-cranial voxel pair ($n = 44906$) of two randomly generated surrogate sets (95% confidence interval: 0.0079–0.0095). Although the mean correlation is significantly greater than 0 at $\alpha = 0.05$, the mean correlation is so close to 0 that we can consider each dataset as sufficiently independent for the purpose of determining significance. However, one should proceed with caution as any positive correlation among surrogate datasets results in an overall underestimation of Type I Error.

A second caveat which remains from existing wavelet resampling and permutation techniques (Breakspear et al., 2004) is the impact of the spatiotemporal resampling on the distribution of the variance of each voxel time series. Since the variance distribution is not necessarily the same as that of the original data, we propose an extension to the Theiler et al. (1992) amplitude adjustment step which adjusts the data to approximate the variance distribution of the original. However, use of this extension requires the additional assumption of normality of each voxel time series. The variance distribution of the surrogate data is only a concern if the connectivity measure is influenced by changes in the variance distribution without distortion of the power spectra, as is the case with the simple correlation test statistic.

Thirdly, we do not adjust for multiple comparisons in this paper. The use of 19 surrogate datasets allows us to determine voxel pair specific P values. The number of correlations assessed from an individual seed totals the number of intra-cranial voxels. In this paper, no P value adjustment is made to adjust for these comparisons, as this is a topic of future research. Conducting a Bonferroni adjustment would require the generation of a much larger number of surrogate datasets, however there are significant limitations in doing so. The computational requirements of generating and storing surrogate datasets is significant. For a single subject, it takes approximately five hours and two gigabytes of hard drive space to generate 19 surrogate datasets using MATLAB 7.0 on a 12 processor (750 MHz) Sun compute server running the Solaris operating system. Among other variables, processing time and disk space depends on the size of each dataset, for which ours (after padding) is $80 \times 80 \times 48 \times 192$ (for resting state data), the number of levels for which to construct each wavelet packet, and the speed and multitude of the processors. Ideally, we would like to construct hundreds if not thousands of surrogate datasets from which we can more accurately and precisely determine P values of the correlations of voxel pairs, however processing speed and disk space currently limit us to many fewer realizations.

Acknowledgments

The authors would like to acknowledge Dr. Michael Breakspear for his helpful thoughts and comments. We also thank Dr. Erica Duncan of the Emory University School of Medicine and Kathy Hughes of the Mayne Clinical Research Imaging Centre for providing the resting state and motor task fMRI data, respectively. This research was partially supported by the National Institutes of

Health, grant K25-MH65473, and the Centre d'Imagerie Biomédicale (CIBM).

Appendix A. Essential characteristics of the 2D polyharmonic wavelet decomposition

Two important characteristics deserve our attention. First, as a scaling function, the polyharmonic B-spline (Rabut, 1992) is deployed. These basis functions can be regarded as the true multi-dimensional extension of 1D B-splines, spanning the same space as radial basis functions; i.e., functions of the form $\rho(x_1, x_2) = (x_1^2 + x_2^2)^{(\gamma-2)/2}$, where γ is the order of the polyharmonic B-spline. Second, the quincunx subsampling scheme is used, which is characterized by the dilation matrix

$$D = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \quad (7)$$

This scheme provides a more isotropic treatment of the data and a slower progression through scale than dyadic subsampling. The decomposition formula for a 2D signal S at scale J equals

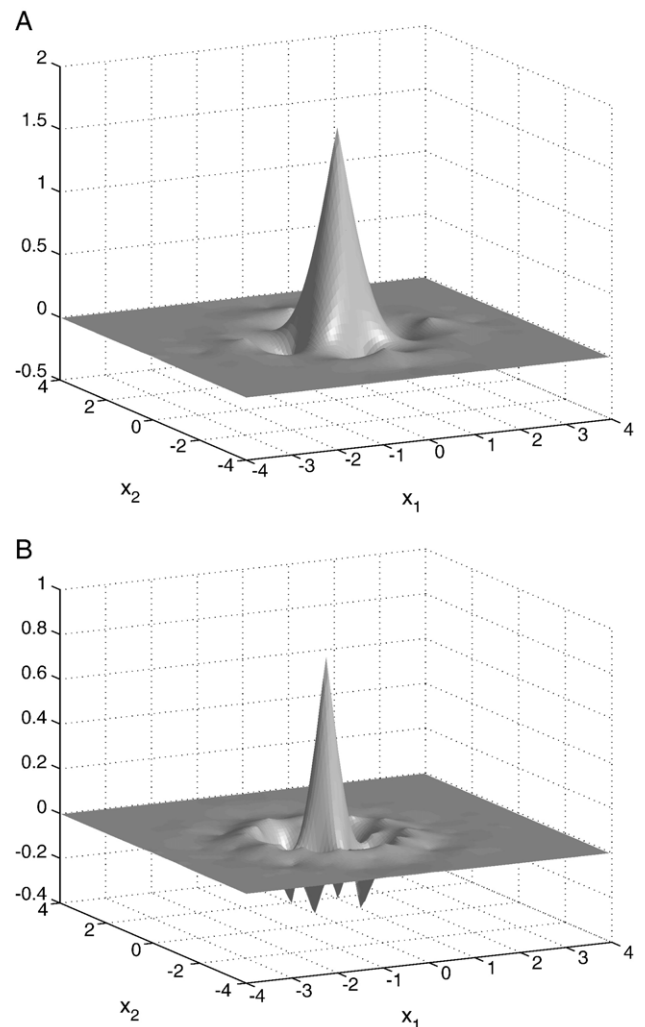


Fig. 13. Orthogonal polyharmonic scaling function (A) and wavelet (B) for order $\gamma = 3$.

$$S(\mathbf{x}) = 2^{J/2} \sum_{k \in \mathbb{Z}^{2 \times 1}} a_{J,k} \phi(D^{-J} \mathbf{x} - k) + \sum_{j=1}^J 2^{j/2} \sum_{k \in \mathbb{Z}^{2 \times 1}} d_{j,k} \psi(D^{-j} \mathbf{x} - k), \quad (8)$$

where $\mathbf{x} = [x_1 \ x_2]^T$.

The polyharmonic wavelet behaves (approximately) as a multiscale version of the $\gamma/2$ -th iterated Laplacian operator; in the Fourier domain we have

$$\psi(\omega) \propto \|\omega\|^\gamma = (\omega_1^2 + \omega_2^2)^{\gamma/2}, \text{ when } \omega \rightarrow 0. \quad (9)$$

Therefore, the polyharmonic wavelet transform has a decorrelating property for isotropic processes with a $1/\|\omega\|$ -like power spectrum density, such as fractional Brownian fields. In Fig. 13, we show both the scaling function and wavelet for the orthogonal version of the transform.

In the 1D case, the number of iterations J_w of the wavelet packet decomposition corresponds to its level L . However, in 2D, we define the level L as the half of the number of iterations. This way, the number of subbands for a one-level 2D packet decomposition coincides with the separable case. At level L of the 2D wavelet packet decomposition of a $R \times C$ matrix, we obtain $4^L - 1$ sets of detail coefficients of size $R2^{-L} \times C2^{-L}$. Similarly, for a 1D wavelet packet decomposition of a vector of length T , we obtain $2^L - 1$ sets of detail coefficients of size $T2^{-L}$.

Appendix B. Join test to determine spatial/temporal independence

Before resampling the signs of the detail coefficients, we must ensure some degree of independence among the signs. We can treat each subband as a binary lattice (or vector) indicating whether the detail coefficient at that location was positive. We can conduct a test for independence of the detail coefficients at level L of the wavelet packet decomposition in the following manner.

If two detail coefficients are adjacent to each other (including diagonally adjacent), they are said to be linked by a “join” (Cliff and Ord, 1973). To test for independence among the coefficients, we see whether the number of positive–positive (PP) joins, or joins between two positive coefficients, negative–negative (NN) joins, and positive–negative (PN) joins are close to what is expected under independence. Let δ be a connection matrix in which $\delta_{ij} = 1$ if coefficients i and j in the L^{th} level are adjacent to each other and $\delta_{ij} = 0$ otherwise. Let $d_i = 1$ if the i^{th} detail coefficient is positive, and $d_i = 0$ if the i^{th} detail coefficient is negative. Detail coefficients equal to 0 are not included in the analysis. The observed number of PP joins is given by $PP = \sum_{i=1}^n \sum_{j=i+1}^n \delta_{ij} d_i d_j$, the observed number of PN joins is given by $PN = \sum_{i=1}^n \sum_{j=i+1}^n \delta_{ij} (d_i - d_j)^2$, and the observed number of NN joins is given by $NN = K - PP - PN$, where n is the number of non-zero detail coefficients at level L and K is the total number of joins between non-zero detail coefficients at level L . K is given by $K = \sum_{i=1}^n \sum_{j=i+1}^n \delta_{ij}$.

Under the assumption that the sign of each detail coefficient is an independent drawing from a Bernoulli(p) distribution, then $E[PP] = Kp^2$ and $VAR[PP] = Kp^2 + 2Dp^3 - (K + 2D)p^4$, where $D = \frac{1}{2} \sum_{i=1}^n H_i(H_i - 1)$ and H_i is the number of positive coefficients adjacent to i (Cliff and Ord, 1973).

Since PP is asymptotically normally distributed (Cliff and Ord, 1973), we can conduct a hypothesis test for the independence of the signs of the detail coefficients by checking $\frac{|PP - E[PP]|}{\sqrt{VAR[PP]}} > Z_{\alpha/2}$, in which case we would conclude that there is sufficient evidence to suggest dependence. For our study, we use $\alpha = 0.05$.

References

- Battle, G., 1987. A block spin construction of ondelettes. Part I: Lemarié functions. Commun. Math. Phys. 110, 601–615.
- Breakspear, M., Brammer, M., Robinson, P., 2003. Construction of multivariate surrogate sets from nonlinear data using the wavelet transform. Physica D 182, 1–22.
- Breakspear, M., Brammer, M., Bullmore, E., Das, P., Williams, L., 2004. Spatiotemporal wavelet resampling for functional neuroimaging data. Hum. Brain Mapp. 23, 1–25.
- Bullmore, E., Long, C., Suckling, J., Fadili, J., Calvert, G., Zelaya, F., 2001. Colored noise and computational inference in neurophysiological time series analysis: resampling methods in time and wavelet domains. Hum. Brain Mapp. 12, 61–78.
- Bullmore, E., Fadili, J., Maxim, V., Sendur, L., Whitcher, B., Suckling, J., Brammer, M., Breakspear, M., 2004. Wavelets and functional magnetic resonance imaging of the human brain. NeuroImage 23, S234–S249.
- Cliff, A., Ord, J., 1973. Spatial Autocorrelation. Pion Limited.
- Coifman, R.R., Meyer, Y., Wickerhauser, M.V., 1992. Wavelets and their Applications. Jones and Barlett, Boston, Ch. Wavelet Anal. Signal Process., 153–178.
- Cordes, D., Haughton, V., Carew, J., Arfanakis, K., Maravilla, K., 2002. Hierarchical clustering to measure connectivity in fMRI resting-state data. Magn. Reson. Imaging 20, 305–317.
- Daubechies, I., 1990. The wavelet transform, time-frequency localization and signal analysis. IEEE Trans. Inf. Theory 36 (5), 961–1005.
- Dijkerman, R.W., Mazumdar, R.R., 1994. On the correlation structure of the wavelet coefficients of fractional Brownian motion. IEEE Trans. Inf. Theory 40, 1609–1612.
- Fan, Y., 2003. On the approximate decorrelation property of the discrete wavelet transform for fractionally differences processes. IEEE Trans. Inf. Theory 49, 516–521.
- Flandrin, P., 1992. Wavelet analysis and synthesis of fractional Brownian motion. IEEE Trans. Inf. Theory 38, 910–917.
- Friston, K., Frith, C., Liddle, P., Frackowiak, R., 1993. Functional connectivity: the principal component analysis of large (pet) data sets. J. Cereb. Blood Flow Metab. 13, 5–14.
- Grecius, M., Krasnow, B., Reiss, A., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proc. Natl. Acad. Sci. 100, 253–258.
- Hampson, M., Peterson, B., Skudlarski, P., Gatenby, J., Gore, J., 2002. Detection of functional connectivity using temporal correlations in mr images. Hum. Brain Mapp. 15, 247–262.
- Jiang, T., He, Y., Zang, Y., Weng, Z., 2004. Modulation of functional connectivity during the resting state and the motor task. Hum. Brain Mapp. 22, 63–71.
- Lowe, M., Mock, B., Sorenson, J., 1998. Functional connectivity in single and multislice echoplanar imaging using resting state fluctuations. NeuroImage 7, 119–132.
- Mallat, S., 1989. A theory for multiresolution signal decomposition: the wavelet decomposition. IEEE Trans. Pattern Anal. Mach. Intell. 11, 674–693.
- Mazoyer, B., Zago, L., Mellet, E., Bricogne, S., Etard, O., Houde, O., Crivello, F., Joliot, M., Petit, L., Tzourio-Mazoyer, N., 2001. Cortical networks for working memory and executive functions sustain conscious resting state in man. Brain Res. Bull. 54, 287–298.
- Percival, D.B., Walden, A.T., 2000. Wavelet Methods for Time Series Analysis. Cambridge Univ. Press, Cambridge.

- Prichard, D., Theiler, J., 1994. Generating surrogate data for time series with several simultaneously measured variables. *Phys. Rev. Lett.* 73, 951–954.
- Rabut, C., 1992. Elementary m -harmonic cardinal B-splines. *Numer. Algorithms* 2, 39–62.
- Raichle, M., MacLeod, A., Snyder, A., Powers, W., Gusnard, D., Shulman, G., 2001. A default mode of brain function. *Proc. Natl. Acad. Sci.* 98, 676–682.
- Salvador, R., Suckling, J., Coleman, M., Pickard, J., Menon, D., Bullmore, E., 2005. Neurophysiological architecture of functional magnetic resonance images of human brain. *Cereb. Cortex* 15, 1332–1342.
- Shulman, G., Fiez, J., Corbetta, M., Buckner, R., Miezin, F., Raichle, M., Petersen, S., 1997. Common blood flow changes across visual tasks: II. decreases in cerebral cortex. *J. Cogn. Neurosci.* 9, 648–663.
- Theiler, J., Eubank, S., Longtin, A., Galdrikian, J., Farmer, J., 1992. Testing for nonlinearity: the method of surrogate data. *Physica D* 58, 77–94.
- Unser, M., Blu, T., 2003. Wavelet theory demystified. *IEEE Trans. Signal Process.* 51 (2), 470–483 (Feb.).
- Van De Ville, D., Blu, T., Unser, M., 2005. Isotropic polyharmonic B-splines: Scaling functions and wavelets. *IEEE Trans. Image Process.* 14, 1798–1813.
- Van De Ville, D., Blu, T., Unser, M., in press. Surfing the brain: An overview of wavelet-based techniques for fMRI data analysis. *IEEE Engineering in Medicine and Biology Magazine*.
- Whitcher, B., in press. Wavelet-based bootstrapping of spatial patterns on a finite lattice. *Computational Statistics and Data Analysis*.
- Xiong, J., Parsons, L., Gao, J., Fox, P., 1999. Interregional connectivity to primary motor cortex revealed using mri resting state images. *Hum. Brain Mapp.* 8, 151–156.