

Spatiotemporal wavelet analysis for functional MRI

Chris Long,^{a,*} Emery N. Brown,^{b,c} Dara Manoach,^a and Victor Solo^{a,d}

^aMGH Martinos Center for Biomedical Imaging, Charlestown, MA 02129, USA

^bDepartment of Anesthesia and Critical Care, Massachusetts General Hospital, Boston, MA 02129, USA

^cHarvard-MIT Division of Health Sciences and Technology, Cambridge, MA 02139, USA

^dDepartment of Electrical Engineering, University of Michigan, Ann Arbor, MI 48109, USA

Received 16 October 2003; revised 9 April 2004; accepted 9 April 2004
Available online 7 August 2004

Characterizing the spatiotemporal behavior of the BOLD signal in functional Magnetic Resonance Imaging (fMRI) is a central issue in understanding brain function. While the nature of functional activation clusters is fundamentally heterogeneous, many current analysis approaches use spatially invariant models that can degrade anatomic boundaries and distort the underlying spatiotemporal signal. Furthermore, few analysis approaches use true spatiotemporal continuity in their statistical formulations. To address these issues, we present a novel spatiotemporal wavelet procedure that uses a stimulus-convolved hemodynamic signal plus correlated noise model. The wavelet fits, computed by spatially constrained maximum-likelihood estimation, provide efficient multiscale representations of heterogeneous brain structures and give well-identified, parsimonious spatial activation estimates that are modulated by the temporal fMRI dynamics. In a study of both simulated data and actual fMRI memory task experiments, our new method gave lower mean-squared error and seemed to result in more localized fMRI activation maps compared to models using standard wavelet or smoothing techniques. Our spatiotemporal wavelet framework suggests a useful tool for the analysis of fMRI studies.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Spatiotemporal; Functional MRI; Wavelets

Introduction

Functional magnetic resonance imaging of human subjects aims to associate region-specific brain activity arising from stimulus-dependent neural firing—the BOLD (Blood Oxygenation Level-Dependent) effect (Kwong et al., 1992; Ogawa et al., 1992)—with sensory or cognitive stimuli. These changes in blood flow and volume are caused by the cerebral response induced by

the stimuli and are manifested in the form of spatially localized fluctuations in gray-level contrast across T2*-weighted echo-planar images (EPI). In activated regions of the cortex, the fMRI signal is composed of spatially coherent clustered time series (Forman et al., 1995; Mitra et al., 1997) and is often studied in a framework consisting of spatial smoothing, temporal regression, and hypothesis testing. Functional MRI is capable of acquiring high-resolution images simultaneously in both space and time, maintaining relatively good spatial detail while sampling the experimentally induced hemodynamic and volume effects at an adequate rate. Any spatial presmoothing within the analysis, though potentially improving the signal to noise, will degrade the native resolution of the experiment, increasing the likelihood of losing or at best blurring important detail within the functional map. The spatial activation profiles typically encountered in fMRI range from well-localized high-amplitude features to low-amplitude but relatively large connected structures. The morphology of these functional structures is influenced both by the underlying heterogeneity of the brain anatomy and by the task-related response of different brain locales. To account for this intrinsic inhomogeneity, we have developed a novel spatiotemporal wavelet procedure capable of adapting to the spatially varying nature of the fMRI signal. We draw on the adaptive properties of spatial wavelets to capture this variability in the activation maps and use a non-wavelet convolution model based on our understanding of the BOLD signal to inform the temporal estimation. We use the term *spatiotemporal* in this paper to refer to the sense in which the spatial wavelet estimations are performed; in essence, they are driven by the underlying temporal processes in the fMRI data. Such an interpretation differs to a more conventional definition where spatiotemporal wavelets typically designate families of n -dimensional bases (e.g., tensor products), constructed and applied in both space and time.

Wavelets have attracted recent attention from a diverse range of fields in science and engineering. These have included applications in signal and image processing (Chambolle et al., 1998; Coifman et al., 1992; Devore and Lucier, 1992; Vetterli and Kovačević, 1995), fractals (Abry et al.), vision (Mallat, 1996), meteorology (Fournier, 1996), time series (Nason and von Sachs, 1999; von Sachs, 1998), and statistics (Benjamini and Hochberg, 1995; Donoho and Johnstone, 1995; Johnstone and Silverman, 1997).

* Corresponding author. Department of Radiology, Massachusetts General Hospital-East, 13th Street, Building 149, Charlestown Navy Yard, Charlestown, MA 02129. Fax: +1-617-726-8410.

E-mail address: cjl@nmr.mgh.harvard.edu (C. Long).

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

In the functional MRI literature, Brammer (1998) and Ruttiman et al. (1998) have used the good spatial localization properties of wavelets to better characterize fMRI activation maps. Temporal applications of wavelets in fMRI have included Bullmore et al. (2002) who utilized the decorrelating properties of wavelets to control Type-I error. Also, Fadili and Bullmore (2002) constructed a temporal maximum-likelihood framework for estimating signal in the presence of $1/f$ noise.

The wavelet procedure proposed in the current work efficiently characterizes multiscale variability in the spatial fMRI activations by drawing on the multiresolution properties of wavelets. The main novelty of this technique, however, lies in its spatiotemporal formulation and combined wavelet-regression solution. In effect, the wavelet estimation is informed both by anatomical considerations and by the temporal behavior of the fMRI data. The computations are carried out in four stages. In the first stage, a temporal physiological model is solved using a weighted least-squares approach to partially absorb the influence of correlated noise. In the second step, refined estimates of the noise parameters are calculated using an Expectation–Maximisation (EM) algorithm. In the third, the temporal BOLD signal components are wavelet smoothed using a nonstandard wavelet procedure to generate spatiotemporal parameterization of the activations. These steps are next repeated until the minimum of a spatiotemporal likelihood function is reached. In the sequel, we show that modifications to standard wavelet methods are required to maximize this likelihood leading to a new type of wavelet estimator. We illustrate our method on simulated and real fMRI data, and compare the outcome with standard wavelet estimation and Gaussian smoothing techniques.

Methods

Acronyms and notation

Throughout this paper, we shall assume that all temporal signals x are of dyadic length n . These will be indexed as $(x_j)_{j=0}^{n-1}$, i.e., with n time points, or in the frequency domain as $(\hat{x}_k)_{k=0}^{N-1}$, where N signifies an N -point discrete Fourier transform (DFT) and k is frequency. We designate the DFT of a (zero-padded) signal x_t as $\hat{x}_k = \sum_{t=0}^{N-1} x_t e^{-j\omega_k t}$, where $\omega_k = (2\pi k)/(N)$ represents the normalized Fourier frequencies. Also in future sections, we use the superscript H to denote the conjugate transpose, that is, $\hat{z}_k^H = (\hat{z}_k^*)^T$, where $(\cdot)^T$ signifies the transpose operation. In the spatial domain, we assume that each 2D (two-dimensional) slice is indexed by, $\mathbf{p} = (p_1, p_2)$, ($p_d = 1, \dots, M_d$), ($d = 1, 2$). In the spatiotemporal activation analysis that follows, \mathbf{y}_p is a 2D map used to denote the spatially unregularized voxelwise BOLD signal estimate at voxel \mathbf{p} . $\hat{\mathbf{f}}_p^{(W)}$ is a map used to designate the final regularized spatiotemporal BOLD activations, where W implies that a wavelet processing has taken place. In the wavelet domain, $\hat{f}_{j,\tau}^l$ defines the wavelet coefficient gained from the 2D discrete wavelet transform (DWT) at scale j , location $\tau = (\tau_1, \tau_2)$, and wavelet orientation l , respectively. In this paper, all DWTs are two-dimensional, though the method can easily be extended to higher dimensions, see Appendix C. We use $\hat{\mathbf{f}}_p^{(SM)}$ to indicate that the signal parameters \mathbf{f}_p were recovered using conventionally smoothed or presmoothed data with no wavelet treatment. Finally, the “ ℓ_1 norm” of a vector \mathbf{x} , indicates the sequence norm $\|\mathbf{x}\|_1$, defined as $\sum_i |x_i|$.

Voxelwise modeling

In the following, we assume that the gray-level fluctuations in the acquired T2*-weighted echo-planar images are caused by blood deoxygenation effects governed by the dynamic behavior of the blood flow and volume within activated brain matter. Combining these dynamics with empirical observations about the nature and form of the fMRI noise and including a term to capture low-frequency static magnetic field drift culminates in the following model for the time course at each location in the brain

$$x_{t,\mathbf{p}} = m_{\mathbf{p}} + b_{\mathbf{p}}t + s_{t,\mathbf{p}} + v_{t,\mathbf{p}} \quad (1)$$

Within each voxel at location \mathbf{p} , we assume the time course to consist of (i) a background or dc component $m_{\mathbf{p}}$, (ii) a linear trend component, $b_{\mathbf{p}}t$, (iii) a signal term $s_{t,\mathbf{p}}$, and (iv) a stochastic noise term $v_{t,\mathbf{p}}$. The convolutional signal model $s_{t,\mathbf{p}}$ is constructed to reflect experimentally induced cerebral hemodynamics and is similar to earlier work, i.e., Friston et al. (1994), in which a Poisson-shaped function was used to model the relationship between stimulus and response. We assume that the response of the cerebral cortex is proportional to small contrast changes in the measured T2*-weighted MRI signal and we ascribe these changes to the joint behavior of intracerebral volume- and flow-related blood deoxygenation. Specifically, our BOLD model comprises the product between these components. That is

$$s_{t,\mathbf{p}} = H(t - D_{\mathbf{p}})(\gamma_{0,\mathbf{p}} + \gamma_{1,\mathbf{p}} g_{\mathbf{p},t}^a * c_{t-D_{\mathbf{p}}}) \times (V_{0,\mathbf{p}} + V_{1,\mathbf{p}} g_{\mathbf{p},t}^b * c_{t-D_{\mathbf{p}}}) \quad (2)$$

where $V_{0,\mathbf{p}}$, $\gamma_{0,\mathbf{p}}$ are the baselines of the volume and flow terms, $V_{1,\mathbf{p}}$, $\gamma_{1,\mathbf{p}}$ are the response amplitudes, $D_{\mathbf{p}}$ is the hemodynamic delay, $H(\cdot)$ is the unit-step function present to impose causality upon the system, and $*$ denotes convolution. Also, c_t represents the experimental stimulus, $g_{\mathbf{p},t}^a$, $g_{\mathbf{p},t}^b$ are the blood flow and blood volume impulse response terms whose fundamental shape and rise-decay rates were motivated by empirical work (Mandeville et al., 1996). In response to fixed input stimuli, the hemodynamic behavior is supposed to exhibit classic first-order system features, including a lag time between input and response, rise time, decay to baseline, and response undershoot. Consequently, the hemodynamic response $g_{\mathbf{p},t}^a$ possesses a smooth biphasic form that may, in practice, be closely characterized by a discrete gamma function, $g_{\mathbf{p},t}^a = (1 - e^{-1/\tau_{a,\mathbf{p}}})(t + 1)e^{-t/\tau_{a,\mathbf{p}}}$. This function has its maximum at $t = \tau_{a,\mathbf{p}} + 1$ and has been normalized to ensure unity gain. The blood volume impulse response is also based on empirical studies (Boynton et al., 1996) and has form $g_{\mathbf{p},t}^b = (1 - e^{-1/\tau_{b,\mathbf{p}}})e^{-t/\tau_{b,\mathbf{p}}}$. Ideally, the system time constants would also require estimation since the spatial dispersion characteristics of the respective impulse responses might be clinically or experimentally relevant. However, in the present work, these properties are kept fixed with time constants $\tau_{a,\mathbf{p}} = 1.5$ s for the flow component and $\tau_{b,\mathbf{p}} = 11$ s for the blood volume effect. For generality, the optimization algorithms outlined in the later sections are capable, with appropriate extensions, of performing such time-constant computation under the same maximum likelihood assumptions. But because these parameters are kept fixed in this work, the signal model terms $(g_{\mathbf{p},t}^a * c_t)$, $(g_{\mathbf{p},t}^b * c_t)$, and $(g_{\mathbf{p},t}^a * c_t)(g_{\mathbf{p},t}^b * c_t)$ are not spatially dependent.

Continuing, Eq. (2) can be reorganized to yield

$$s_{t,\mathbf{p}} = V_{0,\mathbf{p}}\gamma_{0,\mathbf{p}}H(t - D_{\mathbf{p}}) + \xi_t \mathbf{f}_{\mathbf{p}} = \text{Flow and Volume Baseline} \\ + \text{Weighted Signal Response}$$

with

$$\xi_t = \begin{bmatrix} (g_{\mathbf{p},t}^a * c_t) & (g_{\mathbf{p},t}^b * c_t) & (g_{\mathbf{p},t}^a * c_t)(g_{\mathbf{p},t}^b * c_t) \end{bmatrix}$$

and

$$\xi_{\mathbf{p}} = \begin{bmatrix} (g_{\mathbf{p},t}^a * c_t)_1 & (g_{\mathbf{p},t}^b * c_t)_1 & (g_{\mathbf{p},t}^a * c_t)_1(g_{\mathbf{p},t}^b * c_t)_1 \\ \vdots & \vdots & \vdots \\ (g_{\mathbf{p},t}^a * c_t)_n & (g_{\mathbf{p},t}^b * c_t)_n & (g_{\mathbf{p},t}^a * c_t)_n(g_{\mathbf{p},t}^b * c_t)_n \end{bmatrix} \quad (3)$$

In this setup, $\mathbf{f}_{\mathbf{p}} = [f_{\mathbf{p}}^{(1)}, f_{\mathbf{p}}^{(2)}, f_{\mathbf{p}}^{(3)}]'$ is the three-vector of weighting or activation coefficients at voxel \mathbf{p} associated with each of the physiological model components, flow, volume, and interaction. In its current form, this model assumes similar response onset times for both the blood flow and volume components, and so contains just a single delay parameter. However, it may be desirable to decouple the response times for each of these components, and in the estimation scheme that follows, this is simple to accomplish if necessary.

The stochastic fMRI noise $v_{t,\mathbf{p}}$ is modeled by additively combining a serially correlated component with an independent white noise (Purdon and Weisskoff, 1998; Solo et al., 2001; Weisskoff et al., 1993). These two noise components primarily arise from respective contributions of physiologic- and machine-induced (Johnson) noise sources. Much of the serial autocorrelation may be explained by the impact of cardiac pulsations or respiratory effects on the brain though physical aspects of the imaging process may also have a bearing upon this correlation, e.g., machine acquisition parameters or repetition time (TR) (Purdon and Weisskoff, 1998; Yoo et al., 2001). Shortfalls in the signal model may also fail to fully capture experimental variance, compounding the correlation problem further. To deal with these issues, our noise model combines these (often low-frequency) physiologic fluctuations within the brain with the scanner noise. We have shown in earlier work (Solo et al., 2001) that such mixtures of independent noises can be compactly characterized in terms of first-order auto-regressive moving-average (ARMA(1,1)) processes $v_{t,\mathbf{p}}$ similar to those first suggested for use in fMRI by Locascio et al. (1997). Often, it is useful to consider such models in the frequency domain, in which case the power spectral density of $v_{t,\mathbf{p}}$ is

$$F_{k,\mathbf{p}} = \frac{\sigma_{\eta,\mathbf{p}}^2}{|1 - a_{\mathbf{p}}e^{-j\omega_k}|^2} + \sigma_{v,\mathbf{p}}^2 \quad (4)$$

Here, $\sigma_{v,\mathbf{p}}^2$ is the machine or background noise, $\sigma_{\eta,\mathbf{p}}^2$ is the correlated noise variance, and $a_{\mathbf{p}}$ is the correlation coefficient. Identifiability of the machine and correlated noise variances will depend on the strength of this correlation coefficient. For example, when $a_{\mathbf{p}} \rightarrow 0$, Eq. (4) just becomes the sum of the two variances. ARMA(1,1) models enable powerful yet parsimonious representations of the noise spectra in fMRI and differ from the

simpler and more widely used AR(1) models first proposed in fMRI analysis by Bullmore et al. (1996). Assuming pure AR processes in fMRI time series amounts to combining the two independent noise sources into a single correlated component and this may lead to problems in the AR estimation. Naturally, an AR(1) model will suffice as $\sigma_{v,\mathbf{p}}^2 \rightarrow 0$, but otherwise, AR(1) will be a biased estimator for an ARMA(1,1) process. For equivalence, one needs an arbitrarily high-order AR(p) model to approximate an ARMA realization. Since the value of p depends upon the behavior of the poles and zeros in the ARMA transfer function which will itself vary across the brain, the AR order will need to be chosen voxelwise—an unwieldy task. Detail on the relationship between these models is given by the Wold Decomposition Theorem (Wold, 1954). Finally, we note that more complex long-term $1/f$ noise processes have been suggested for fMRI noise, for example, Fadili and Bullmore (2002), that may yield equally compact and perhaps more powerful representations of brain noise. Indeed, these methods have been shown in some cases to better constrain Type-I errors in resting data compared to AR(p) models. The spatial wavelet framework outlined in the following sections may be easily modified to incorporate noise spectra derived by these or other means, if desired.

The final piece of the overall model (Eq. (1)) is the drift term $b_{\mathbf{p}}t$ that captures slow changes in image intensity caused, for example, by low-order drift in the static magnetic field during scanning or unresolved head motion. The profile of such drifts is generally monotonic but may be spatially heterogeneous. In most cases, the magnitude of these artifacts is large compared to the surrounding signal and noise. To a reasonable degree, this trend can be modeled using a simple linear drift term though higher order regressors, such as families of sinusoids, may easily be interchanged in our formulation, depending upon the extent of the problem (Friston, 1997). In general, choosing the optimal number of terms in such expansions can become problematic if the stimulus is not a simple blocked design, raising concerns respecting proper determination of drift from task-related activity. Other trend removal strategies such as running-lines smoothers are equally effective, if not preferable (e.g., Marchini and Ripley, 2000) to these regression models but may have similar drawbacks related to the choice of tuning or smoothing parameter. The estimation of $b_{\mathbf{p}}t$ in Eq. (1) is computed such that the presence of both ARMA(1,1) noise and the physiologic transfer function are accounted for.

Voxelwise estimation

In previous work (Purdon et al., 2001; Solo et al., 2001), we estimated the model parameters in Eq. (1) by specifying and solving the following frequency domain negative likelihood function. For completeness, its solution is given in Appendix A.

$$J = \sum_{\mathbf{p}} \sum_k \frac{|\tilde{x}_{k,\mathbf{p}} - m_{\mathbf{p}}\delta_{0,k}N - b_{\mathbf{p}}\tilde{u}_k - e^{-j\omega_k D_{\mathbf{p}}} \tilde{\xi}_k \mathbf{f}_{\mathbf{p}}|}{2F_{k,\mathbf{p}}} \\ + \frac{1}{2} \sum_{\mathbf{p}} \sum_k \log(F_{k,\mathbf{p}}) \quad (5)$$

Once more, x reflects the time series at voxel \mathbf{p} , and $\tilde{x}_{k,\mathbf{p}}$ is its Fourier transform. $\delta_{0,k} (= 0, k \neq 0, = 1 \text{ otherwise})$ is the dirac or impulse function that accompanies the mean $m_{\mathbf{p}}$, singling out, in essence the dc frequency component from the spectrum. \tilde{u}_k is the

DFT of a ramp function with zero mean and b_p is the associated drift coefficient. To estimate the model parameters, we iterate between computing the signal and noise in a cyclic descent procedure. In Purdon et al. (2001), spatial smoothing or regularization of the noise parameters was shown to solve a variant of Eq. (5) while retaining the native resolution of the activation coefficients. In the following sections, we extend the likelihood in Eq. (5) to include a spatial component that imposes nonuniform regularization upon the activation maps.

Spatiotemporal wavelet estimation

We aim to combine previously outlined empirical and physiological observations about BOLD and the fMRI noise with a spatially heterogeneous wavelet procedure. But to correctly incorporate spatial wavelets in this context, we need to revisit some of the basic assumptions associated with wavelet estimation and apply some reformulations. Before investigating this further, we will first review some relevant properties of wavelets.

Wavelet analysis begins by defining a prototype or mother wavelet (such as Daubechies, Haar or Coiflet) whose width and position can vary but whose fundamental shape is kept fixed. Wavelet analysis represents data using a sequence of compact weighted functions (building blocks or basis functions) into which any finite energy signal or image can be decomposed. The respective weightings of these wavelet basis functions are the desired wavelet coefficients that, roughly speaking, are generated from a series of convolutions carried out at different resolutions. If the dilation and translation parameters of these basis functions are sampled appropriately and combined with certain imposed conditions upon the mother wavelet, e.g., regularity, compactness, and vanishing moments, families of orthonormal basis functions can be constructed. Any finite-energy signal or image can be efficiently decomposed and uniquely reconstructed in these orthonormal bases using the Discrete Wavelet Transform (DWT) and its inverse (see Appendix C or e.g., (Daubechies, 1990; Mallat, 1989) for further details). The DWT implements a multiresolutional analysis, compactly describing (or compressing) signal information using a few large amplitude wavelet coefficients. Fig. 1a illustrates these localizing properties in the time–frequency plane. The DWT analyzes high-frequency events with good temporal resolution and low-frequency features with good frequency resolution. In contrast, Fig. 1b illustrates the way in which a classic analysis method, the windowed Fourier transform, behaves in the same context. In this case, all frequencies are analyzed with the same fixed-width window. Fig. 2 shows an example of a synthetic heterogeneous signal consisting of a longer duration, low-frequency event in conjunction with two higher-frequency burst-like features. With the windowed Fourier transform, the window size may be chosen to localize optimally on any one feature, but its invariant nature means it cannot capture all three very efficiently. The wavelet transform, on the other hand, is able to localize comparatively well on all three features.

Wavelet smoothing: Because of the multiresolution property of wavelets, compact representations of signals containing irregular or smooth structures are possible. This fundamental feature gives rise to powerful estimation strategies for discriminating signals of unknown smoothness in the presence of white (Donoho and Johnstone, 1994) or structured noise (Johnstone and Silverman, 1997). Wavelet smoothing is a nonlinear operation that discards (sets to zero or kills) wavelet coefficients representing noise (usually many and small) while compensating for the additive effects of this noise

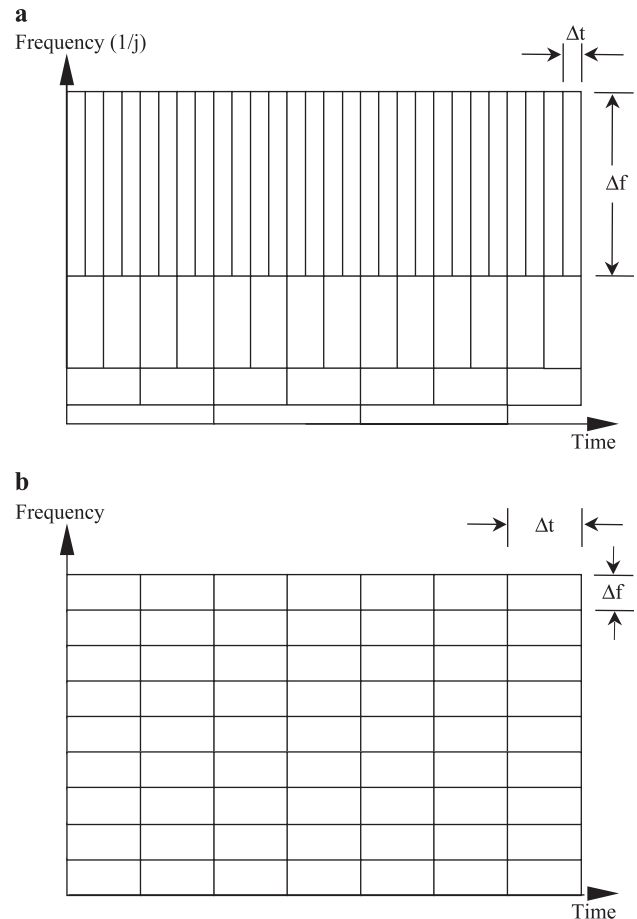


Fig. 1. (a) Tiling of the time–scale plane via the wavelet transform. The window area stays constant: the length halves (doubles) while the width doubles (halves) between scales. In the figure, j is the scale or dilation parameter. (b) Resolution of the time–frequency plane using the short time Fourier transform (STFT). Note that this method can only resolve two impulses in time if at least Δt apart and two spectral peaks if they are more than Δf apart.

on the signal by shrinking those coefficients (generally few and large) toward zero. This mechanism is known as soft thresholding. The exact choice of surviving wavelet coefficients in a given wavelet analysis depends upon the way in which the signal has been coded in that particular basis, and this in turn uniquely reflects the underlying smoothness structure of the signal in question. The choice of threshold depends on the expected long-term behavior of the noise process. If the noise variance is previously known or can be reliably estimated, an appropriate thresholding rule can be constructed. It then remains simply to filter the wavelet coefficients globally (or across a select number of scales) by comparing each against this (normally) fixed threshold. If Gaussian white noise is the major corrupting force, the thresholding prescription reduces to that of Donoho and Johnstone (1994) (D & J).

Wavelet smoothing and ℓ_1 penalization: Thresholding in the wavelet domain is equivalent to solving an ℓ_1 penalized least-squares problem. Such problems specify a data-fidelity term to ensure good correspondence of candidate fits against the data and a penalty term to impose a running cost on the overall minimization, helping preclude unlikely or degenerate solutions while preserving transients and other irregular structures. In contrast,

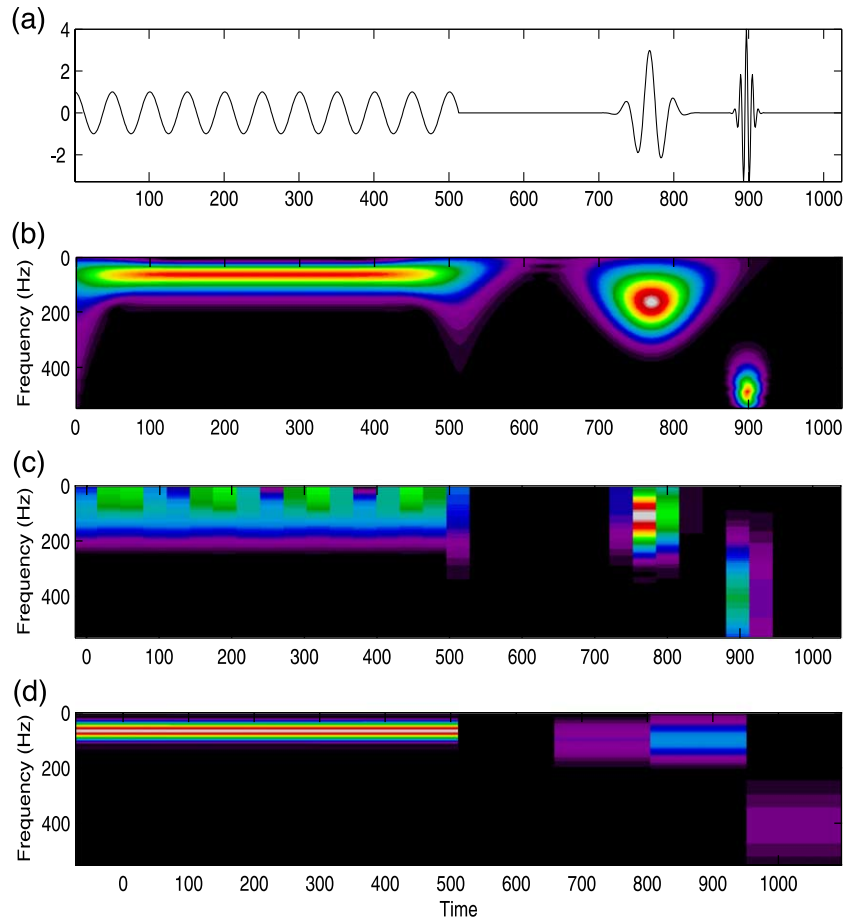


Fig. 2. This figure compares the analysis of an artificial signal (a) using the continuous wavelet transform (CWT) and (b) the STFT with different choices of window sizes (c and d). The signal shown in (a) was constructed by adding a low-frequency (80 Hz) sine wave with two higher-frequency bursts with shorter time duration according to their frequency (120 and 450 Hz, respectively). In (b) with the CWT, one gains good frequency localization at the lower frequencies while improved temporal resolution occurs at higher frequencies. If we try analyzing the same signal with the STFT using a relatively short-duration window (c), discrimination of the shorter-term events is possible, but it is more difficult to localize their frequencies. In particular, it is hard to separate the 80-Hz longer wave from the 100-Hz pulse. Moving on to (d), the frequency localization is improved by increasing the duration of the window function. This choice efficiently localizes the longer term 80 Hz event since it is stationary over time, but the shorter events now become averaged out having sacrificed localization power in frequency.

non- ℓ_1 (e.g., quadratic) penalties attempt to find a trade-off between overfitting the data and oversmoothing potential solutions. While robust, such approaches tend to degrade edges and other potentially informative irregularities in the recovered image. ℓ_1 constraints can facilitate a larger range of candidate fits by including those belonging to a minimally smooth subspace (see Chambolle et al., 1998 of all possible functions). In the current application, use of an ℓ_1 penalty enriches detail in the activation maps while accounting for low-amplitude, smooth connected structures. Significantly, wavelet thresholding techniques afford a practical means of solving such problems. This relationship has been formalized in the work of Devore and Lucier (1992) and also in Donoho and Johnstone (1994).

Spatiotemporal wavelets for fMRI

We now construct a weighted least-squares model for fMRI analysis with this type of spatial penalty. To generate a possible spatiotemporal wavelet solution, we first formulate an appropriate likelihood function with this kind of structure. Consequently, we have incorporated physiological assumptions about the BOLD

contrast mechanisms and noise structure while placing the ℓ_1 constraint on the spatial variability of the BOLD signal. To estimate the model parameters, we maximize the following wavelet-penalized Gaussian-negative log-likelihood function. That is

$$J(\mathbf{f}, \mathbf{D}, \sigma_\eta, \sigma_v, \mathbf{a}) = \sum_{\mathbf{p}} \sum_k \frac{|\langle \tilde{x}_{k,\mathbf{p}} - e^{-j\omega_k D_{\mathbf{p}}} \tilde{\zeta}_k \mathbf{f}_{\mathbf{p}} \rangle|^2}{2F_{k,\mathbf{p}}} + \frac{1}{2} \sum_{\mathbf{p}} \sum_k \log(F_{k,\mathbf{p}}) + \lambda \sum_{j,\tau,l} \|\tilde{\mathbf{f}}_{j,\tau}^l\|_1 \quad (6)$$

Here $\tilde{\mathbf{f}}$ is the spatial wavelet transform of the unregularized activation map. For the first \mathbf{f} component, that is, $\mathbf{f}_{\mathbf{p}}^{(1)}$, $\tilde{\mathbf{f}}$ is just the 2D DWT of the map $\mathbf{f} = \{\mathbf{f}_{\mathbf{p}}^{(1)}\}_{\mathbf{p} \in \mathbb{Z}^2}$. Note that the 2D wavelet coefficients are indexed by scale j , shift τ , and orientation l . For ease of exposé, we have assumed that the baseline and drift terms have been previously computed and subtracted off the original

data. In practice though, all the parameters are estimated within an iterative procedure similar to that described in Appendix A, the difference here just involves a further spatial processing of \mathbf{f} . The first term in Eq. (6) enforces fidelity on the data, comparing candidate parameterizations of the delayed BOLD signal (effected by the term $e^{-j\omega_k D_p}$) with the measured time course. Inclusion of the noise spectrum $F_{k,p}$ recognizes serial correlation in the errors. The final term of this likelihood $\lambda \sum_{j,\tau,l} \|\tilde{f}_{j,\tau}^l\|_1$ is the wavelet domain ℓ_1 spatial constraint on the activation map \mathbf{f} . Due to the flexible nature of this penalty, we can adaptively accommodate heterogeneous spatial activations and can hope to use wavelet thresholding to derive a computationally tractable solution. We begin by multiplying out the least-squares term in Eq. (6) and reorganizing, i.e.,

$$J = \frac{1}{2} \sum_{\mathbf{p}} (\mathbf{y}_{\mathbf{p}} - \mathbf{f}_{\mathbf{p}})' \Omega_{\mathbf{p}} (\mathbf{y}_{\mathbf{p}} - \mathbf{f}_{\mathbf{p}}) + \lambda \sum_{j,\tau,l} \|\tilde{f}_{j,\tau}^l\|_1 \quad (7)$$

which contains the purely spatial components:

$$\mathbf{y}_{\mathbf{p}} = \left(\sum_k \frac{\tilde{\xi}_k^H \tilde{\xi}_k}{F_{k,p}} \right)^{-1} \left(\sum_k \frac{\tilde{\xi}_k^H \tilde{x}_{k,p}}{F_{k,p}} \right) \quad (8)$$

and

$$\Omega_{\mathbf{p}} = \sum_k \frac{\tilde{\xi}_k^H \tilde{\xi}_k}{F_{k,p}} \quad (9)$$

The regression parameters in Eq. (8) constitute a crude version of the desired $\mathbf{f}_{\mathbf{p}}$ parameters. Eq. (9) represents the spatial noise variance and is based upon the normalized power of the hemodynamic model.

The solution of Eq. (7) does not depend solely upon information in the estimated spatial activation map, \mathbf{y} . One possible approach would simply involve calculating the noise standard deviation from these spatial maps—treating them in essence as observed static 2D images—then applying and constructing a conventional thresholder, thus ignoring the temporal basis under which the maps were derived. Instead though, the temporal behavior of the noise processes in the functional data are characterized through Ω and the spatial thresholding informed by this dynamic information. $\Omega_{\mathbf{p}}^{-1}$ may be thought of as a spatial noise covariance matrix that arises as a by-product during the maximization of Eq. (6). Utilizing $\Omega_{\mathbf{p}}^{-1}$ in the spatial thresholding embeds the temporal part of the problem into the overall estimation. The resulting spatiotemporal framework is thereby reduced to that of a more conventional ‘signal plus noise’ recovery problem whose objective is to extract estimates of the true activation $\mathbf{f}_{\mathbf{p}}$ from a spatially white background noise process $\nu_{\mathbf{p}}$ with zero mean, but in this instance, known spatially varying variance $\Omega_{\mathbf{p}}^{-1}$. That is, $\mathbf{y} = \mathbf{f} + \nu$, $\nu = \{\nu_{\mathbf{p}}\}_{\mathbf{p} \in \mathbb{Z}^2}$, $\nu_{\mathbf{p}} \sim \mathcal{N}(0, \Omega_{\mathbf{p}}^{-1})$. A modified spatial wavelet thresholding will now be shown to offer an approximate solution to this problem.

Strictly speaking, the optimization of Eq. (7) does not in fact admit a true thresholding solution and this will be pursued elsewhere. But we can still proceed if several reasonable approximations are made. The most important of these treats each $\{\mathbf{f}_{\mathbf{p}}^{(q)}\}_{q=1}^3$ component separately rather than as a full vector

optimization, i.e., for any choice of q the voxelwise signal plus noise model may just be written

$$\mathbf{y}_{\mathbf{p}} = \mathbf{f}_{\mathbf{p}} + \nu_{\mathbf{p}} \quad (10)$$

Taking 2D wavelet transforms of \mathbf{y} and Ω (denoted by (\cdot)) and ignoring cross-scale correlations gives

$$\tilde{y}_{j,\tau}^l = \tilde{f}_{j,\tau}^l + \tilde{\nu}_{j,\tau}^l \text{ where } \text{var}(\tilde{\nu}_{j,\tau}^l) \approx \frac{1}{\tilde{\Omega}_{j,\tau}^l} = \sum_{j,\tau,l} \psi_{j,\tau}^{(l)2} \frac{1}{\Omega_{\mathbf{p}}}$$

For 2D wavelets, $1 \leq l < 2^2 - 1$ and ψ is the wavelet basis function (see Appendix C for more detail). Eq. (7) can then be rewritten in the wavelet domain as

$$J = \frac{1}{2} \sum_{j,\tau,l} (\tilde{y}_{j,\tau}^l - \tilde{f}_{j,\tau}^l)^2 \tilde{\Omega}_{j,\tau}^l + \lambda \sum_{j,\tau,l} |\tilde{f}_{j,\tau}^l| \quad (11)$$

This expression has the following spatially varying thresholder as its minimizer (see Appendix B for proof) and yields estimates $\hat{\mathbf{f}}$ of the desired wavelet-domain activation coefficients $\hat{\mathbf{f}}$

$$\hat{\mathbf{f}}_{j,\tau}^{(W)l} = \text{sgn}(\tilde{y}_{j,\tau}^l) \left(|\tilde{y}_{j,\tau}^l| - \frac{\lambda}{\sqrt{\tilde{\Omega}_{j,\tau}^l}} \right)_+ \quad (12)$$

In other words, if the value of $|\tilde{y}_{j,\tau}^l|$ is greater than $\frac{\lambda}{\sqrt{\tilde{\Omega}_{j,\tau}^l}}$, then $\tilde{y}_{j,\tau}^l$ shrinks toward zero by this amount. Otherwise, set $|\tilde{y}_{j,\tau}^l| = 0$.

$\tilde{\mathbf{y}}$ is the 2D map of discrete wavelet-transformed regression coefficients derived from Eq. (8) and $\hat{\mathbf{f}}^{(W)}$ are the recovered (thresholded) wavelet coefficients representing the spatiotemporal map of signal parameters. These are inverse wavelet transformed to yield the final signal estimate $\hat{\mathbf{f}}$, which as a consequence of the spatial penalty will have retained inherent spatial detail and variability. λ is a (normally positive) parameter that controls the degree to which the data fidelity and penalty terms interact with one another. Intuitively, we can see that a high value of λ will cause a low energy in the penalty, leading in turn to smoother fits. Conversely, small λ will lead to larger energy in the wavelet coefficients permitting more flexibility in the reconstruction but potentially overfitting the data. In this work, we use the universal thresholder that sets $\lambda = \sqrt{2 \log_e(M_1 M_2)}$, where $M_1 M_2$ is the number of voxels in the brain. This choice quantifies an upper bound on the risk of the wavelet threshold estimator (see e.g., Mallat, 1998 for more detail).

Computational Steps

Ignoring drift and baseline terms for the moment, the parameters remaining to be estimated at each pixel \mathbf{p} are $(\mathbf{f}_{\mathbf{p}}, D_{\mathbf{p}}, \sigma_{\eta,\mathbf{p}}, \sigma_{\nu,\mathbf{p}}, a_{\mathbf{p}})$, i.e., the BOLD model estimates, the hemodynamic delay, and the noise components (white & correlated variances and the correlation coefficient), respectively. The overall minimizer of Eq. (6) consists of the following eight steps (shown graphically in Fig. 3):

- (1) Obtain initial estimates for the noise and construct the power spectrum, $F_{k,p}$.

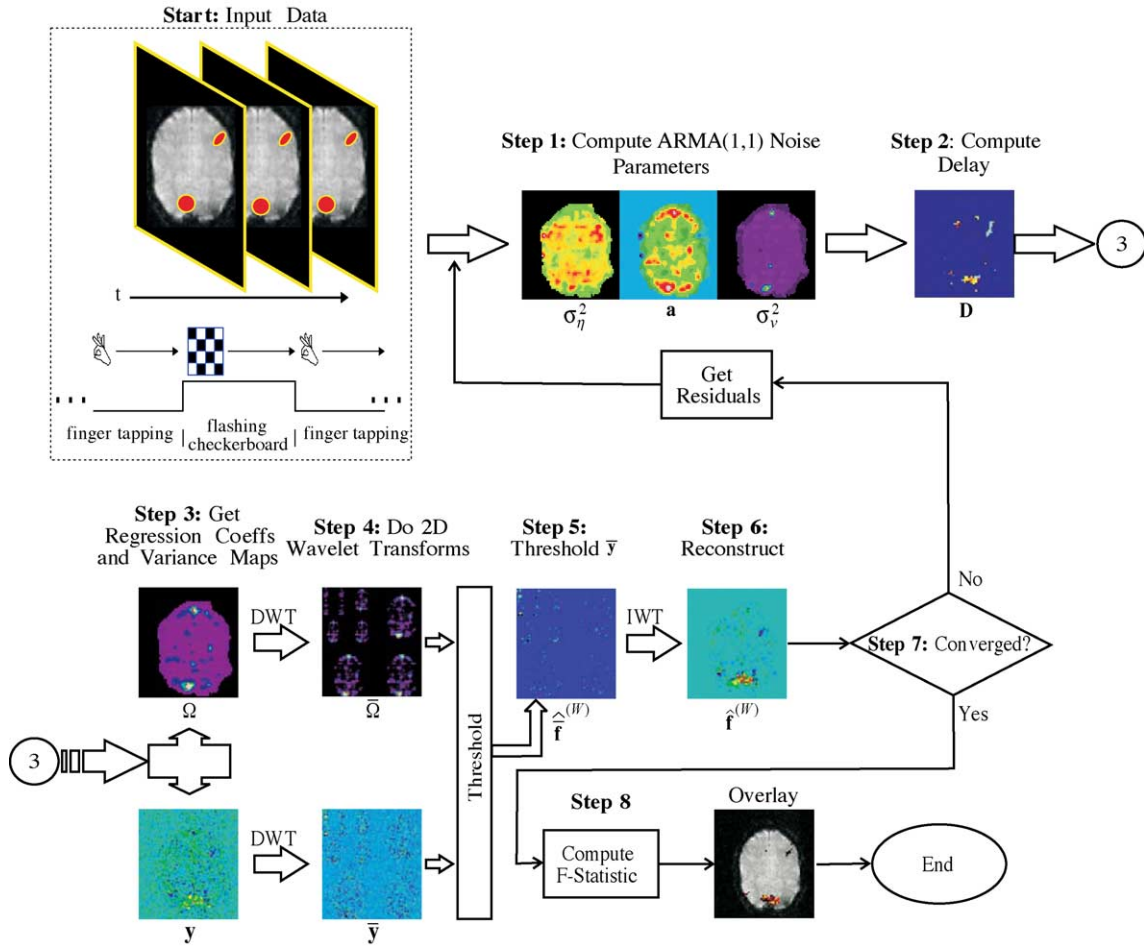


Fig. 3. Graphical schema of the algorithm described below. The example used for illustration was taken from an experiment where the subject was asked to alternately tap their index finger and view a flickering checkerboard. Significant activation is therefore expected in the motor strip and visual cortex. In the figure, DWT and IWT correspond to the forward and inverse wavelet transforms, respectively.

- (2) Compute voxelwise estimates of the hemodynamic delay, that is,

$$\hat{D}_p = \arg.\max \left(\begin{array}{c} \operatorname{Re} \left(\sum_k \frac{e^{-j\omega_k D_p} \tilde{z}_k^H \tilde{x}_{k,p}}{F_{k,p}} \right)' \left(\sum_k \frac{\tilde{z}_k^H \tilde{z}_k}{F_{k,p}} \right)^{-1} \\ \operatorname{Re} \left(\sum_k \frac{e^{-j\omega_k D_p} \tilde{z}_k^H \tilde{x}_{k,p}}{F_{k,p}} \right) \end{array} \right) \quad (13)$$

- (3) Perform weighted regression to obtain raw estimates of the signal parameters y_p , that is,

$$y_p = \left(\sum_k \frac{\tilde{z}_k^H \tilde{z}_k}{F_{k,p}} \right)^{-1} \left(\sum_k \frac{e^{-j\omega_k \hat{D}_p} \tilde{z}_k^H \tilde{x}_{k,p}}{F_{k,p}} \right) \quad (14)$$

- (4) Take 2D DWTs of maps y and Ω to obtain \tilde{y} and $\tilde{\Omega}$.
 (5) Wavelet threshold the activation maps using the modified procedure, Eq. (12) to obtain $\hat{f}_p^{(W)}$.

- (6) Reconstruct $\hat{f}_p^{(W)}$ back into the spatial domain by taking the inverse wavelet transform, yielding $\hat{f}_p^{(W)}$, and compute the new residuals.

- (7) Return to step (1) and repeat until convergence.

- (8) Compute activation maps using the following approximate frequency domain F-statistic, with m and $n - q$ degrees of freedom and overlay.

$$S_p(m, n - q) = \left(\hat{f}_p^{(W)'} \tilde{\xi}' (RSS)^{-1} \operatorname{diag}(\mathbf{F}_p)^{-1} \tilde{\xi} \hat{f}_p^{(W)} \right) / m \quad (15)$$

where RSS is the sum of squares of the whitened residuals.

Under the null hypothesis of no significant activation, $m = q$, where q is the rank of $\tilde{\xi}$ and n is the number of time points in the experiment. We derive the power spectral density of the temporal ARMA(1,1) noise using an EM algorithm (see Solo et al., 2001 for further details), which is incorporated into the complex reweighted least-squares computations that generate D_p and y_p . The predictor $\tilde{\xi}$ is constructed off-line by convolution of the input stimulus $c(t)$ with the impulse response terms representing the blood flow and volume components of the BOLD response. Once D_p is calculated, $\tilde{\xi}_k$ is shifted to $\tilde{\xi}_{k-D_p}$ and the y_p estimates obtained from Eq. (14). The delay computations are based on a simple correlation analysis

between $\tilde{\xi}_{k-D_p}$ and each time course, possessing a structure that reflects the presence of ARMA(1,1) noise and potentially, signal drift (e.g., replace $\tilde{x}_{k,p}$ by $\tilde{x}_{k,p}^0$ in Eq. (13), see Appendix A, step (iv)). This quantity is also determined in the frequency domain utilizing the Fourier Shift Theorem to modulate the phase of each time course. Because of the duality property of the FST, this procedure is equivalent to applying the same latency in continuous time. The D_p calculation can be summarized as a line search, comprising repeated correlation computations applied across a range of candidate delays at arbitrarily fine resolutions.

Finally, the wavelet thresholding rule of Eq. (12) is applied to the temporal maps, an improved estimate of the noise is obtained, and the procedure iterated until convergence.

p value interpretation of wavelet thresholding

To digress slightly, we will describe the relationship between wavelet thresholding and significant testing. Although this particular interpretation of wavelet thresholding is not drawn upon in this paper, we feel it is worth including given its potential for combining estimation and inference in functional MRI.

Not only can wavelets recover functions of unknown smoothness from corrupted sets of observations, but in most cases are near-optimal estimators in situations where the signal is nonuniform or heterogeneous. In terms of minimum mean-square error, wavelets generally outperform or at worst match conventional denoising or estimation techniques depending on the local smoothness properties of the underlying signal. Donoho (1995) noted that with “overwhelming probability”, one can hope to recover a function that is at least as smooth as the true function, and that this observation holds across a wide range of smoothness classes. In the wavelet domain, this is equivalent to stating that noise has a high likelihood of disappearing from the recovered signal, which in turn leads us to have high confidence in the zeroed wavelet coefficients. Abramovich and Benjamini have used this observation in the context of simultaneous inference. They noted that thresholding by the universal thresholder $\sqrt{2\log(n)}$ was close to a Bonferroni correction over a range of probabilities α . As $n \rightarrow \infty$, the tail of the quantile $z(1/n)$ approaches $\sqrt{2\log(n)}$. If d_k is an iid Gaussian process with zero mean and variance σ^2 , then the probability level $\alpha_n = P(\max_{1 \leq k \leq n} |d_k| > \sqrt{2\log(n)} \leq nP(|d_k| > \sqrt{2\log(n)})$. This yields the asymptotic relationship

$$\alpha_n \leq \sqrt{\frac{2}{\pi}} \left(\frac{n}{\sqrt{2\log n}} e^{-\log(n)} \right) = (\pi \log(n))^{-\frac{1}{2}} \quad (16)$$

as n gets large.

Data description and experimental paradigms

Simulated data

1.5 T Siemens Sonata, 3.0 T Siemens Allegra, and 3.0 T Siemens Trio: Data were acquired from five normal volunteers (two at 1.5 T and three at 3.0 T) who were scanned while lying quietly with their eyes closed for 6 min (1.5 T) and 5 min (3.0 T). Images were contiguously acquired at a repetition rate of 2 s (1.5 T) and 5 s (3.0 T) for the duration of each session, yielding a total of 180 (1.5 T), 60 (3.0 T), T2*-weighted 3D images each containing 21 (1.5 T) or 3 (3.0 T) noncontiguous, axially acquired slices. The data were acquired on the systems at the MGH

Martinos Center, Charlestown, MA, with time-to-echo (TE) = 40 ms, TR = 2 s (1.5 T), = 5 s (3.0 T), in-plane resolution = 3.13 mm (1.5 T), = 1.9 mm (3.0 T), and slice thickness = 6 mm (1.5 T), = 4 mm (3.0 T).

Data construction: Clusters of temporal signals were synthesized by constructing multiple realizations of the physiological model $s_{t,p}$ (Eq. (3)), while keeping \mathbf{f}_p spatially constant within each of three different regions. Spatiotemporal activations were created by adding these clusters onto the three areas in respective slices from the T2*-weighted EPIs (see Fig. 4, column (i)). Average SNRs were calculated voxelwise from $\text{SNR} = \frac{1}{n} \sum_t s_{t,p}^2 / \sigma_B^2$ across each region, for each choice of \mathbf{f} . The measure of SNR was used to ascertain detectability of the embedded signal as opposed to, say, using the power of $s_{t,p}$. σ_B^2 is the temporal variance of the background fMRI data. As shown in Fig. 4, the signal regions consisted of two relatively localized heterogeneous features with high SNR and a single larger and smoother region with an average SNR $\sim 60\%$ lower than in regions 1 or 2. The overall goal of this simulation was to recover an estimate $\hat{\mathbf{f}}$ of \mathbf{f} through the wavelet thresholder (Eq. (12)), such that the penalized likelihood (Eq. (11)) was solved. In this way, the relative performance of the proposed variable thresholder could be measured. In all of the simulations, a simple check of algorithm performance was computed from the mapwise mean-squared error, that is,

$$\text{MSE}(\hat{\mathbf{f}}) = \sum_p \sum_k (\xi_{t,p} \mathbf{f}_p - \xi_{t,p} \hat{\mathbf{f}}_p) \quad (17)$$

Sternberg item recognition memory task (3.0 T, event-related)

A single healthy subject performed a numerical Sternberg item recognition paradigm (SIRP) that had been adapted for rapid presentation, event-related fMRI. This version of the SIRP required subjects to encode a set of target digits, maintain them in working memory (WM) over a fixed period, and to respond to a probe digit by indicating whether it was a target. Further details of the data acquisition can be found in Manoach et al. (2003).

Each working memory (WM) trial began with a central fixation cross for 500 ms followed by the presentation of a set of five digits (targets) to be learned (3500 ms). This constituted the encode phase that was directly followed by the Delay (or holding) epoch during which time the screen was blank. During the Probe epoch, subjects were presented with a single digit (probe) for 2000 ms. In half the trials, the probe was a target (a member of the memorized set) and in half the trials the probe was a foil (not a member of the memorized set). Subjects responded by pressing a button box with their right thumb for targets and their left thumb for foils. The three trial types differed only in the length of the delay period that lasted 0, 2, or 4 s. The three trial types randomly alternated with a fixation baseline condition within each run. During the baseline condition, subjects fixated on an asterisk that appeared in the center of the screen. The duration of fixation was randomly varied in increments of 2 s up to a maximum of 12 s. Subjects performed six runs of 4 min 48 s each. Each run contained nine trials of each WM condition and 72 s of fixation.

Functional images were collected on the 3.0 T Allegra system using BOLD contrast and a gradient-echo T2*-weighted sequence (TR/TE/Flip = 2000 ms/30 ms/90°) to measure variations in blood flow and oxygenation. Twenty contiguous horizontal 6-mm slices parallel to the intercommissural plane (voxel size 3.13 × 3.13 × 6 mm) were acquired interleaved. Four images at the beginning of each scan were acquired and discarded to allow longitudinal

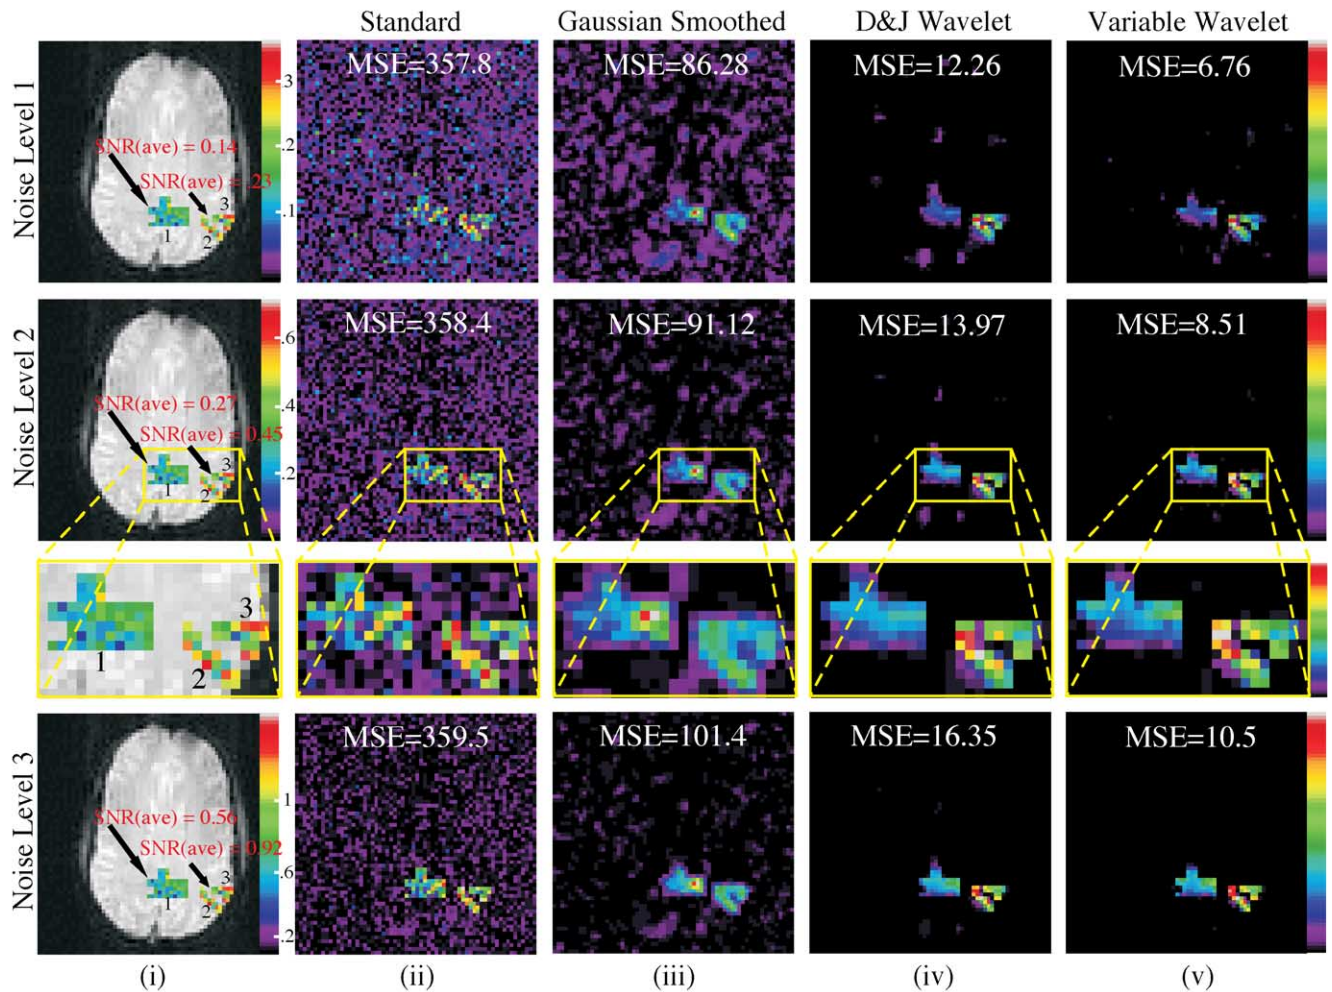


Fig. 4. Simulated data set showing relative MSE performances among various estimation strategies. In column (i), we constructed a synthetic signal by adding spatially constant realizations of the temporal signal described by Eq. (2) to a resting data set. Shown is the relative SNR in the appropriate areas. Shown furthermore in each of columns (ii), (iii), (iv), and (v) are the results of estimating the underlying (known) signal using (ii) no spatial treatment, (iii) Gaussian presmoothing, (iv) Donoho and Johnson wavelet thresholding, and (v) the variable wavelet thresholding procedure proposed in this paper. Also shown in each panel is the MSE measuring the recovery performance of each method (see Eq. (17)). Note that the Gaussian smoothing is applied as a preprocessing step in line with common strategies, whereas both spatial wavelet techniques are invoked within the estimation procedure.

magnetization to reach equilibrium. Regions known to be reliably activated by blocked versions of SIRP include the dorsolateral prefrontal cortex (DLPFC), intraparietal sulcus, lateral premotor cortex, supplementary motor area, and the insula (e.g., Manoach et al., 1997; Rypma and D'Esposito, 1999). This experiment was designed with the intent of investigating the neural circuitry underlying each epoch (encode, delay, probe) of the working memory task (Manoach et al., 2003).

Spatial working memory task (3.0 T, blocked-periodic)

Subjects were presented with a set of target stimuli consisting of either two or three shapes (irregular polygons) that appeared in particular locations on the right and left side of the screen. Subjects were asked to hold in WM either the locations or shapes of the targets. The targets then disappeared and subjects then responded to the appearance of individual probe shapes by indicating whether they appeared in one of the memorized locations regardless of its shape (spatial) or whether the probe was one of the memorized shapes regardless of its location (shape). The control condition used

identical stimuli to both the spatial and the shape tasks, but there was no display of targets and subjects responded to probes by indicating whether they appeared on the right or left side of the screen. This required a visually guided rather than a WM guided response. In the WM conditions, half of the probes were targets (a member of the memorized set) and half were foils (not a member of the memorized set). In all conditions, half of the probes appeared on the right and half appeared on the left. Subjects responded by pressing a button box with their right thumb for targets and their left thumb for foils. Response time (RT) and side (right or left) were recorded. Each run of the task contained two blocks of the low- and high-load WM conditions (consisting of two and three targets), three blocks of the control condition (c), and three periods of fixation (*). The fixation baseline condition consisted of an asterisk that flashed at 2-s intervals to maintain the subjects' visual attention and gaze (see Manoach et al., 2004 for further details).

Images were collected on the 3.0 T Allegra system. Functional images were collected using Blood Oxygen Level Dependent (BOLD) contrast and a gradient-echo T2*-weighted sequence

(TR/TE/Flip = 2000 ms/30 ms/90°) to measure variations in blood flow and oxygenation. Twenty contiguous horizontal 5-mm slices parallel to the intercommissural plane (voxel size: 3.13, 3.13, 5 mm) were acquired interleaved. Four images at the beginning of each scan were acquired and discarded to allow longitudinal magnetization to reach equilibrium. This experiment was designed to investigate whether spatial WM is associated with functional specialization of the right prefrontal cortex (PFC) relative to WM for shapes.

Results

Simulated data

To isolate the effect of the proposed spatially varying wavelet thresholder, four variations of the algorithm were contrasted. The first and simplest aimed to recover the spatiotemporal signals using the purely voxel-based procedure described in Appendix A. After noise estimation, each time course was corrected for the effect of drift at each location in the brain followed by construction and regression of the delayed physiological model Eq. (3) against each time course, thus generating the spatially unregularized maps \mathbf{y} . This process was iterated until the estimate for each model term had settled down to within a predefined tolerance limit. This procedure corresponded to a spatially unregularized version of the technique described in Solo et al. (2001) and Purdon et al. (2001). Fig. 4, column (ii) illustrates the recovered voxelwise signal parameters over a range of SNRs on data from one of the subjects acquired at 1.5 T. The mean-squared error of the recovered signal was calculated from Eq. (17) and is shown on each panel. Fig. 4, column (iii) shows the result of imposing spatial regularization upon the activation parameters by presmoothing the data using a Gaussian kernel with FWHM = 2.6 mm. This choice of kernel size represented a reasonable compromise between the disparate sizes of the synthesized features, maintaining some structure in the detail while capturing to some extent the implicit smoothness of the larger-scale features. As before, a voxelwise cyclic-descent strategy was applied to each time course recovering the corresponding $\hat{\mathbf{f}}^{(SM)}$ maps (shown in the respective rows) at each level of SNR. Fig. 4, column (iv) incorporates the D & J spatial wavelet thresholding method (Donoho and Johnstone, 1994) in place of the variable thresholder to estimate \mathbf{f} . Due to the nature of the D & J thresholder, this amounts to treating the parameter maps \mathbf{y} as purely spatial entities not accounting for the temporal aspects of the data. However, an intuitive invocation of wavelets in this manner could still be considered quite reasonable. The maps in Fig. 4, column (iv) clearly show the “zeroing” property of wavelet estimation. This is a consequence of the nonlinear smoothing characteristics of wavelet threshold estimators that shrink or kill the coefficients. Fig. 4, column (v) illustrates the effect of recovering the synthesized signal by the full-blown spatiotemporal wavelet technique. The wavelet thresholding in this instance guarantees a minimum in the likelihood function Eq. (6). This process was repeated on data from a further four subjects acquired at 1.5- and 3.0-T field strengths across a wider range of SNR. Their average MSEs are summarized in Table 1.

SIRP task

In this example, we analyzed the combined effect of the three trial types. Fig. 5a illustrates the statistical maps and their

Table 1

Average MSEs across five subjects from simulations performed on data acquired at 1.5 and 3.0 T

SNR (Average)	N	P	D & J	VW
0.18	735.08	291.17	74.865	18.637
0.73	735.43	293.97	76.008	20.237
1.67	736.01	298.81	77.207	21.578
3.00	736.68	305.6	78.463	22.989
4.73	737.43	314.48	79.681	24.692
6.84	738.37	325.49	81.061	26.189

N = no spatial treatment, P = presmoothing with a Gaussian kernel, $D \& J$ is the Donoho and Johnson wavelet thresholding, and VW is the variable wavelet thresholding procedure described in this paper.

subsequent anatomical overlay after applying four different analysis procedures with the same temporal modeling structure, but in each case different choices of spatial treatment, namely Gaussian or wavelet smoothing. The slices shown were chosen to transect motor areas and the prefrontal cortex. Before analysis, the data were subjected to rigid-body motion correction to align the scans across time (Cox and Jesmanowicz, 1999). In Fig. 5a, columns (i) and (ii), spatial smoothing was applied using Gaussian kernels of size 4 and 12 mm, respectively, F -statistics and p values computed, and respective Bonferroni corrections applied. To render the standardized maps comparable with the proposed wavelet method, we smoothed the signal maps \mathbf{f} , the noise parameters σ_n^2 , σ_v^2 , and the correlation coefficient a within the estimation procedure rather than as a prior step. Had we been concerned with direct comparisons between the respective \mathbf{f} maps from the different methods (as in the simulations), smoothing in this way or presmoothing would have little affected the final result. However, because of its denoising aspect, presmoothing inflates the test statistic (Eq. (15)) at the expense of spatial resolution, confounding direct comparison with the proposed wavelet thresholder that does not involve preprocessing. To gain insight into the effect of the spatially varying component of our technique on these data, we analyzed the experiment by substituting the spatially varying thresholder with the $D \& J$ wavelet thresholder. Column (iii) illustrates the maps gained from this analysis, and column (iv) shows the result of applying the spatially varying thresholding (Appendix B). For overlay onto the corresponding T1-weighted structural images (top eight panels), the maps were thresholded at the level indicated ($p < 1e - 3$).

In the first row of Fig. 5a, each analysis seems to show activation in several frontal and parietal regions including the supplementary motor area and lateral motor and premotor areas. In the second row, all analyses exhibit activity in the intraparietal sulcus, the cingulate, and lateral prefrontal regions including premotor areas and the dorsolateral prefrontal cortex. The third and fourth rows illustrate the corresponding p value maps before overlay.

SWM task

In this experiment, we examined the response of the subject to the spatial location of each probe versus the control condition. In Fig. 6a, two representative slices are shown after analysis with the same four variations of spatial treatment as those described above. Once again, the data were motion corrected and subjected to identical analyses as outlined in the previous example. p values were once again com-

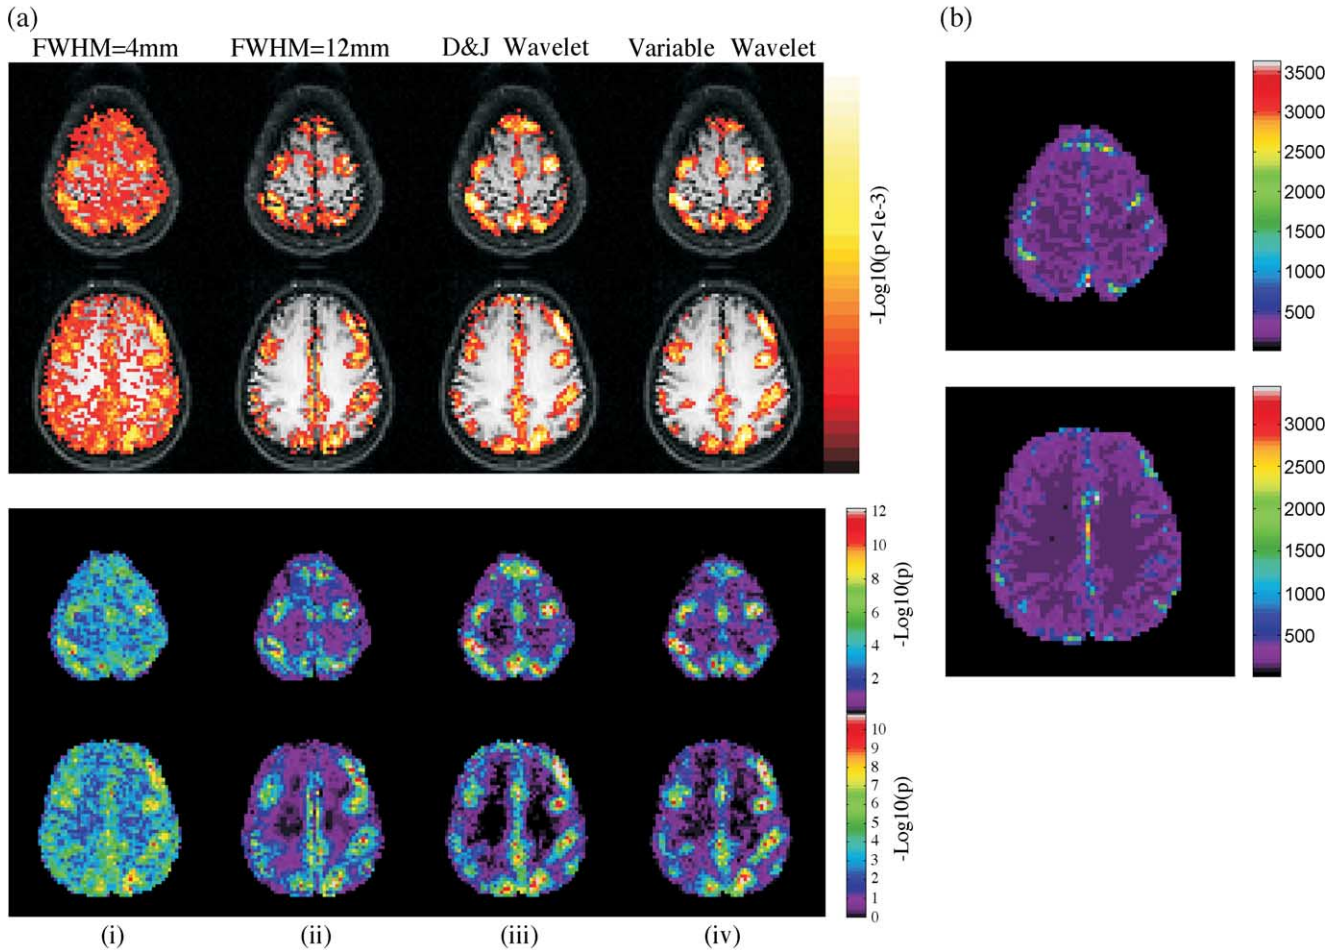


Fig. 5. (a) SIRP task; columns (i) and (ii) show the effect of smoothing the parameters with Gaussian filters, FWHM = 4 and 12 mm, respectively. Illustrated in the bottom two rows are p values derived from the smoothed regression parameters $\hat{\mathbf{f}}_p^{(SM)}$, normalized by the spatially smoothed noise spectra such that $S_p(m, T - k) = (\hat{\mathbf{f}}_p^{(SM)T} \hat{\boldsymbol{\xi}}^T (RSS)^{-1} \text{diag}(\mathbf{F}_p) \hat{\mathbf{f}}_p^{(SM)})/m$. In columns (iii) and (iv), the bottom two rows represent p values obtained in a similar way with the exception that the recovered parameters $\hat{\mathbf{f}}_p^{(W)}$, were incorporated. The top two rows show the result of overlaying and thresholding these maps at $p < 1e - 3$. Note that at FWHM = 4 mm, the activations are well localized but the images are noisy. At FWHM = 12 mm, the images appear cleaner but the activations are less distinct. In both the wavelet approaches, the images are cleaner without sacrificing spatial resolution. The variable thresholder activations appear better localized than the D & J thresholder. (b) These figures show the underlying temporal variance structure Ω employed by the wavelet thresholder. Differences between the two wavelet methods in (a) are largely due to the added uncertainty these maps provide.

puted (bottom eight panels), thresholded, and overlaid onto the structural T1-weighted scan at ($p < 1e - 5$). As expected for this task, lateral, medial frontal, and parietal regions all showed increased activation relative to baseline.

Discussion

In this paper, we have developed a new spatiotemporal model for analyzing functional MRI. Our approach begins by defining a spatially constrained maximum likelihood function that characterizes fMRI activation by combining a model of the BOLD effect into a modified spatial wavelet procedure. The use of spatial wavelets in this context appears to better enable adaptation to the inherent spatial variability of the functional signal. Wavelets are natural candidates for situations involving spatial heterogeneity, but in their standard form may not be appropriate for solving spatiotemporal problems. To overcome this issue, we have shown

the rationale behind the use of a modified wavelet thresholding rule that minimizes a spatiotemporal likelihood function. The results of this approach indicate that smoothing procedures and their associated drawbacks may potentially be avoided in fMRI analyses. The suggested framework should be transparent to the specifics of underlying temporal signal and noise models, enabling experimentation with different and perhaps more complex choices for either component. The resulting spatially varying wavelet thresholding solution is simple and fast to apply, requiring little computational effort beyond that of a standard D & J wavelet estimator.

We comparatively evaluated the performance of the new regime on the simulated data illustrated in Fig. 4 with (1) no spatial treatment, (2) Gaussian presmoothing, (3) standard D & J wavelet thresholding, and (4) the proposed variable wavelet thresholder. In these simulations, we constructed spatiotemporal signals of varying strengths and added them to functional data acquired from five resting subjects. Across a range of SNRs, both wavelet estimators

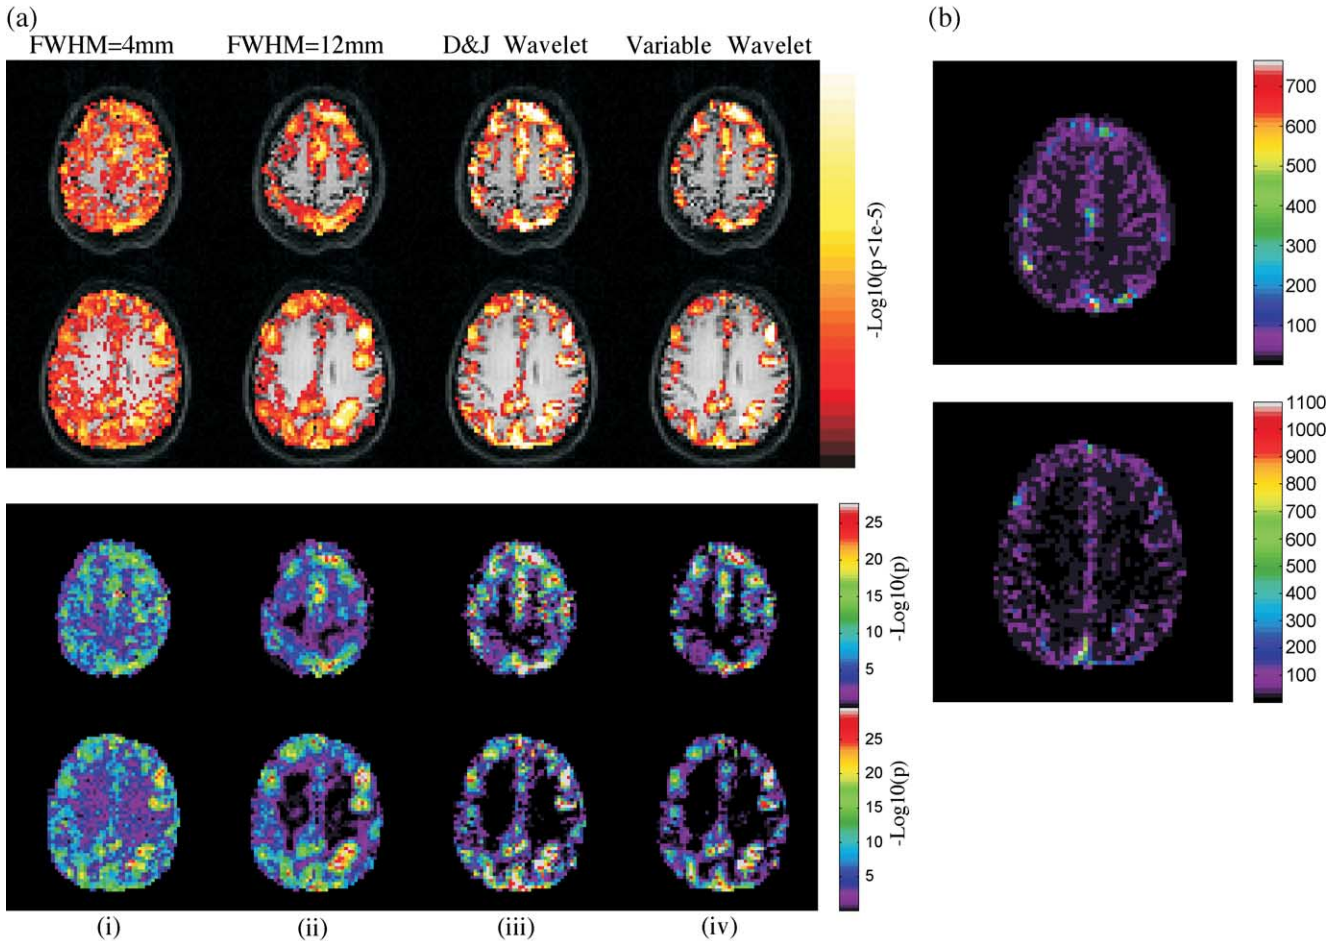


Fig. 6. (a) SWM task; columns (i) and (ii) show the effect of smoothing the parameters with Gaussian filters, FWHM = 4 and 12 mm, respectively. Illustrated in the bottom two rows are p values derived from the smoothed regression parameters $\hat{\mathbf{f}}_p^{(SM)}$ normalized by the spatially smoothed noise spectra such that $S_p(m, T - k) = (\hat{\mathbf{f}}_p^{(SM)'} \xi' (RSS)^{-1} \text{diag}(\mathbf{F}_p)^{-1} \xi \hat{\mathbf{f}}_p^{(SM)})/m$. In columns (iii) and (iv), the bottom two rows represent the p values obtained using the recovered parameters $\hat{\mathbf{f}}_p^{(W)}$ after respective use of the D & J wavelet thresholding and the new variable thresholder. The top two rows show the result of overlaying and thresholding these maps at $p < 1e - 5$. These images show a similar response to the different analysis methods as the previous example. (b) SWM task; these figures show the underlying temporal variance structure Ω employed by the variable wavelet thresholder. Differences between the two wavelet methods in (a) can be mostly ascribed to the added uncertainty offered by these maps.

showed quite distinct demarcation of the activated regions compared to maps derived from the Gaussian-smoothed analyses. In regions containing lower SNR (around .15), some erosion was exhibited after wavelet thresholding, but these areas fared no better under the Gaussian presmoothed or purely voxelwise analyses. In all cases involving presmoothing, discrimination between neighboring features proved challenging with regions (1) & (2) becoming fused together. Also, some loss of detail in region (3) was observed at these SNRs. Table 1 shows the averaged MSEs for each case across the five subjects who had been scanned on either 1.5- or 3.0-T systems. The difference between the competing wavelet approaches became marked across these systems—at higher field strength performance of the D & J wavelet thresholding deteriorated quite significantly, increasing its average MSE. The MSE of the variable thresholder remained, however, relatively stable at these field strengths indicating comparative robustness of the proposed method. Such behavior seems to strengthen the validity of utilizing spatiotemporal assumptions in the variable thresholding model compared to the purely spatial D & J thresholding. While both wavelet approaches are based on universal

thresholding, the D & J estimator employs a single variance gathered from the estimated noise distribution of the derived BOLD parameter map, \mathbf{y} . In contrast, the variable thresholder uses a different noise variance for each voxel to inform its estimation. The D & J approach still, however, appears preferable to smoothing the maps with a fixed-width Gaussian kernel. Like analyses employing Gaussian smoothing, the D & J estimator is also wholly spatial, implicitly ignoring the temporal basis under which the \mathbf{y} coefficients were derived. Recognizing unified spatiotemporal structure in the data allows inclusion of a richer source of variance into the variable wavelet estimation, resulting in this comparative MSE improvement.

In the SIRP and SWM tasks, both wavelet procedures seemed to provide well-localized, physiologically plausible descriptions of the task-related activation (Figs. 5 and 6a, columns (iii) and (iv)) compared to those employing Gaussian smoothing (columns (i) and (ii)). In the two smoothing analyses, the first choice of kernel (FWHM = 4 mm) retained good localization of the activations but yielded noisy images. The second smoothing choice (FWHM = 12 mm) resulted in cleaner maps at the

expense of resolution. In contrast, the wavelet methods in columns (iii) and (iv) were better able to localize on the activations while efficiently preserving larger-connected areas. The variable wavelet thresholding appeared to yield sparser maps, but in the absence of ground truth and from these limited data sets, it is perhaps difficult to definitively choose a preferred technique. Given, however, that the maps obtained through the variable thresholding appeared consistently more conservative than those derived using D & J thresholding (which also showed upward bias in the simulations), it is feasible that D & J may be an under-conservative estimator for fMRI analysis since the temporal covariance structure is ignored. It will be the subject of further work to develop confidence bounds and other diagnostics that accurately quantify the performance of these spatial and spatiotemporal estimators.

Inspection of the variance maps in Figs. 5 and 6b suggests that the temporal process underlying the thresholding is distinctly neurophysiological, with much of the power appearing concentrated in the cortical areas of the brain. This is particularly noticeable in Fig. 5b. The wavelet transform of these variance maps constitutes the added information used in the spatiotemporal wavelet thresholder compared to the D & J thresholder. It is also possible to witness a certain correspondence between areas of high power in the Ω maps and significant regions of activation. Such behavior might either evidence some misspecification in our temporal BOLD model leading to strongly correlated residuals, or simply indicate increased physiological noise activity in these areas.

Choice of wavelet: In the simulations, a Daubechies wavelet with two vanishing moments (the Haar) was chosen to capture the blocky nature of the added signal. In addition to its discontinuous nature, the Haar wavelet seemed appropriate here due to its linear phase. This property improves the spatial localization of the activation clusters compared with asymmetric bases, of which higher-order (>2) Daubechies wavelets are one example. The frequency domain behavior of Haar wavelets is, however, fraught with side lobes that die off gradually. These lead to relatively high inter-scale correlation between the wavelet coefficients, an effect contrary to the assumptions of coordinate-wise thresholding. At the other extreme, use of a sinc-type function in the spatial domain (i.e., a main lobe with oscillations that die out slowly across space or time) would result in almost no interference between frequency channels (being approximately discontinuous over there). But this desirable phenomenon would cost us spatial localization. In the simulations, we tried reducing crosstalk between scales by increasing the wavelet smoothness, but perhaps surprisingly, little difference was observed in the overall MSEs. In this case, the benefit of using discontinuous basis functions for inherently blocky data apparently outweighed any deterioration in performance caused by inter-scale correlations. In the real fMRI experiments, a Symmlet wavelet with four vanishing moments was used. These wavelets are symmetrical and fairly smooth, and seemed to provide reasonable trade-off between spatial localization and small inter-scale correlation. In the activation maps at least, little positive effect was observed on increasing wavelet smoothness beyond this level.

Our approach can be most closely compared to the work of Ruttiman et al. (1998) who developed thresholding methods for fMRI based on hypothesis testing. The wavelet expansion coefficients were classified according to the validity of the null-hypothesis for each, and on this basis either selected for use in the reconstructed activation map or discarded. Also, Brammer

(1998) developed a spatiotemporal wavelet method (using tensor products of wavelets in both space and time) for analyzing fMRI data. This technique performed a linear spatial smoothing on the wavelet domain data by zeroing out prechosen (higher frequency) spatial scales. A temporal wavelet denoising was subsequently performed on the spatial presmoothed data that included prior knowledge about the frequency structure of the experimental response. Temporal applications of wavelets in fMRI have included Bullmore et al. (2002), Fadili and Bullmore (2002), and Long et al. (2001). In Bullmore et al. (2002), temporal wavelets were used to serially decorrelate the data, enabling the use of resampling schemes to accurately ascertain null distributions of activation statistics, and better control Type-I error. In Fadili and Bullmore (2002), a least-squares likelihood function was optimized in the wavelet domain under $1/f$ assumptions about the noise. In earlier work, we (Long et al., 2001) developed wavelet packet time-series analysis in a correlated noise setting that constructed basis expansions using time–frequency properties of the experimental stimulus.

The major difference between these approaches and that of the current work lies in our recognition and explicit modeling of the transfer function component within fMRI data and the way in which this information is combined in the spatial wavelet thresholder. In the present study, we have formulated a hybrid spatiotemporal formulation of wavelets, taking advantage of their adaptive spatial properties and informing the spatial estimation through the temporal covariance structure.

Other non-wavelet techniques have been developed in the past to address the same questions as those considered here. One major class of spatial (or spatiotemporal) procedure that have been applied to fMRI data are Markov Random Fields (or MRFs) that are largely based upon the original work of Besag (1986) from the spatial statistical literature. These approaches also aim to preserve heterogeneous spatial activation patterns, see, for example, Descombes et al. (1998a,b). Originally designed in the light of contemporary image processing problems, e.g., image recovery or denoising (see Geman and Geman, 1984), MRFs have a natural ability to retain anatomical or functional detail in a similar manner to wavelets. An important difference between MRFs and the method proposed here relates to the fact that in Descombes et al. (1998a,b), MRFs are invoked as a preprocessing step and as such can be considered as a spatially adaptive denoising utility. In contrast, our spatial wavelet component is embedded within an estimation procedure that operates directly on signal parameterizations rather than on raw data. A further difficulty relating to the kind of product-form combinations of spatial and temporal covariance functions implied by the use of MRFs in this manner involves their extension from spatial to temporal dimensions. This does not seem to follow naturally for fMRI data or indeed any spatiotemporal data containing transfer function dynamics. In fMRI data, the assumptions required for each dimension suggest to us that a nonproduct form spatiotemporal model such as described in the current work may be fundamentally more appropriate.

In summary, we note that the proposed method possesses a data-based/mechanistic or ‘grey box’ structure, distinguishing it from other fMRI signal processing strategies. These usually fall into either category but rarely combine both, as we have done here. We impart interpretability to the parameters by a physiological transfer function model of the BOLD signal and apply spatial regularization by drawing on underlying characterizations of the

noise component. Because of the data-dependence of the wavelet thresholding, the process may be also considered partly data-driven based to some degree on the statistical properties of a particular experiment.

Conclusions

We conclude that the spatiotemporal model outlined in this paper provides a useful alternative to standard smoothing methods. The structure of our mathematical formulation is intuitively appealing since wavelet smoothing is constructed to guarantee a minimum in a spatiotemporal likelihood function. It is likely that this approach could be further improved by generalizing the standard discrete wavelet transform (DWT) to the spatial wavelet packet transform (WPT) that is yet more adaptive. It is also likely that the complex spatial correlations in fMRI might be better captured using more general thresholding rules than the current universal thresholding approach.

Acknowledgments

We thank Christina Triantafyllou and Doug Greve of the MGH/HMS/MIT Martinos Center for Biomedical Imaging, Charlestown, MA for their help in acquiring the fMRI data sets used in our simulations. The WaveLab (<http://www-stat.stanford.edu/~wavelab/>) and fmristat (<http://www.math.mcgill.ca/~keith/fmristat/>) toolboxes provided some of the utilities used in this work. We also thank two anonymous reviewers for helpful comments on an earlier draft of this paper.

This work was supported by NIH grants NCR P41 RR14705 to the Center for Functional Imaging Technologies at Massachusetts General Hospital, NIMH K02 MH61637, and NIBIB R01 EB0522.

Appendix A. Voxelwise estimation

The fitting of the overall model Eq. (1) is nonlinear in some of the parameters, but a frequency domain optimization can be achieved by minimizing the following negative likelihood function

$$J = \sum_{\mathbf{p}} \sum_k \frac{|\tilde{x}_{k,\mathbf{p}} - m_{\mathbf{p}}\delta_{0,k}N - b_{\mathbf{p}}u_k - e^{-j\omega_k D_{\mathbf{p}}} \tilde{\zeta}_k \mathbf{Y}_{\mathbf{p}}|^2}{2F_{k,\mathbf{p}}} + \frac{1}{2} \sum_{\mathbf{p}} \sum_k \log(F_{k,\mathbf{p}})$$

As before (see Eq. (2)), x reflects the time series at voxel \mathbf{p} and $\tilde{x}_{k,\mathbf{p}} = \sum_0^{N-1} x_{t,\mathbf{p}} e^{-j\omega_k t}$. $u_k = \frac{N}{1-e^{-j\omega_k}}$, $k \neq 0$; $= 0$ otherwise. u_k is the DFT of a ramp function with zero mean. Also, δ is the dirac function, $\delta_{0,k} = 1$, $k = 0$; $= 0$ otherwise. Details of a closely related spatially regularized strategy can be found in Solo et al., 2001, but here we consider just the temporal aspect.

The overall minimization strategy proceeds as follows:

- (i) Initialize $\hat{\mathbf{y}}$ (the spatially unregularized map of the true signal parameters \mathbf{f}). Perform an Expectation–Maximization (EM)

computation (see Solo et al., 2001 for details) to gain estimates of the three noise parameters, $\sigma_{\eta,\mathbf{p}}^2$, $\sigma_{v,\mathbf{p}}^2$, $a_{\mathbf{p}}$.

- (ii) Given these noise estimates, compute the drift $\hat{b}_{\mathbf{p}}$, the signal delay $\hat{D}_{\mathbf{p}}$, the spatially unregularized signal parameters $\hat{y}_{\mathbf{p}}$, and the baseline component $\hat{m}_{\mathbf{p}}$.
- (iii) Given these four parameters, construct updated residuals and re-estimate the noise components (Eq. (4)).
- (iv) Go to (ii) and repeat until convergence.

Step (ii) is divided up into several pieces:

- (a) Initialize the drift $\hat{b}_{\mathbf{p}}$ and subtract this off the raw signal $\tilde{x}_{k,\mathbf{p}}$. Call this new time course $\tilde{x}_{k,\mathbf{p}}^0$.
- (b) Compute the delay using a line search Eq. (13), that is, choose $D_{\mathbf{p}}$ such as to maximize the correlation between the drift-corrected signal and the physiological model over a suitable range of D .
- (c) $\hat{\mathbf{y}}_{\mathbf{p}}$ is computed at the point where Eq. (13) is at its maximum ($=\hat{D}_{\mathbf{p}}$), i.e., from Eq. (14)
- (d) Apply delay to the reconstructed model and subtract this signal off $\tilde{x}_{k,\mathbf{p}}$. Call this new quantity $\tilde{x}_{k,\mathbf{p}}^u$

$$\text{then } \hat{b}_{\mathbf{p}} = \left(\sum_k \frac{\tilde{u}_k^H \tilde{u}_k}{F_{k,\mathbf{p}}} \right)^{-1} \left(\sum_k \frac{\tilde{u}_k^H \tilde{x}_{k,\mathbf{p}}^u}{F_{k,\mathbf{p}}} \right)$$

Finally, Step (iii) is calculated using an EM algorithm. The mean $\hat{m}_{\mathbf{p}}$ is generated from the difference between the frequency domain dc components of the signal and the reconstructed model, that is, $\hat{m}_{\mathbf{p}} = \text{Re}(\tilde{x}_{0,\mathbf{p}} - \zeta_0 \hat{\mathbf{y}}_{\mathbf{p}})/N$.

Computations are done in the frequency domain, mainly for reasons of efficiency. But also, since the colored noise term $v_{k,\mathbf{p}} = \sum_0^{N-1} v_{l,\mathbf{p}} e^{-j\omega_k t}$ is a stationary process, its ordinates $\tilde{v}_{k,\mathbf{p}}$, $\tilde{v}_{l,\mathbf{p}}$ almost decouple for $k \neq l$. This means that the associated covariance matrix is approximately diagonalized in the frequency domain, partitioning the noise from the other model components during convergence.

Appendix B. Spatiotemporal wavelet estimation

The approximate solution to the optimisation of

$$J(f, D, \sigma_{\eta}, \sigma_v, a) = \sum_{\mathbf{p}} \sum_k \frac{|\tilde{x}_{k,\mathbf{p}} - e^{-j\omega_k D_{\mathbf{p}}} \tilde{\zeta}_k f_{\mathbf{p}}|^2}{2F_{k,\mathbf{p}}} + \frac{1}{2} \log(F_{k,\mathbf{p}}) + \lambda \sum_{j,\tau,l} \|\hat{f}_{j,\tau}^l\|_1$$

is given by the following thresholding rule

$$\hat{\mathbf{f}}_{j,\tau}^l = \text{sgn}(\tilde{y}_{j,\tau}^l) \left(|\tilde{y}_{j,\tau}^l| - \frac{\lambda}{\sqrt{\hat{\Omega}_{j,\tau}^l}} \right)_+$$

PROOF: We start by considering two cases for the solution \hat{f} . Ignoring the wavelet indices for the moment, then either

$$(1) \hat{f} = 0 \Rightarrow J(0) = \bar{y}^2/2\bar{\Omega}$$

(2) $\hat{f} \neq 0$. In this case, one can differentiate to find the minimum:

$$\frac{\partial(J(\bar{f}))}{\partial \bar{f}} = -(\bar{y} - \bar{f})\bar{\Omega} + \lambda \text{sgn}(\bar{f}) = 0,$$

$$\Rightarrow \bar{\Omega}\bar{y} = \bar{\Omega}\bar{f} + \lambda \text{sgn}(\bar{f})$$

$$\Rightarrow \bar{\Omega}\text{sgn}(\bar{y})|\bar{y}| = \bar{\Omega}\text{sgn}(\bar{f})\left(|\bar{f}| + \frac{\lambda}{\bar{\Omega}}\right) \Rightarrow \text{sgn}(\bar{y}) = \text{sgn}(\bar{f})$$

And $|\bar{f}| = (|\bar{y}| - \lambda/\bar{\Omega})$ and since $|\bar{f}| = \bar{f}|\text{sgn}(\bar{f})$

$$\bar{f} = \text{sgn}(\bar{y})\left(|\bar{y}| - \frac{\lambda}{\bar{\Omega}}\right) \quad (\text{i})$$

In practice, one would initially compute this soft estimate of \bar{f} . Step 2 then involves ascertaining whether $\hat{f} = 0$, a step that is achieved in practice by comparing the costs for the two cases, that is, $(J(0)) < (J(\text{sgn}(\bar{y})(|\bar{y}| - \lambda/\bar{\Omega})))$ is the condition that would constitute case (1). In this case, a hard thresholding should occur when

$$|\bar{y}| < \frac{\lambda}{\bar{\Omega}} \quad (\text{ii})$$

Combining Eqs. (i) and (ii) yields

$$\bar{f} = \text{sgn}(\bar{y})\left(|\bar{y}| - \frac{\lambda}{\bar{\Omega}}\right)_+ \quad (18)$$

where $t_+ = \max(|t|, 0)$.

This thresholded quantity is applied separately to each of the three components in Eq. (3), before the inverse wavelet transform is applied and the regularized spatiotemporal maps obtained.

Appendix C

C.1. Wavelet computations in higher dimensions

C.1.1. One-dimensional case

We begin with a brief description of 1D wavelets. The basic form of the wavelet functions is $\psi_{j,\tau}(t) = 2^{-j/2}\psi(2^{-j}t - \tau)$, $j, \tau \in \mathbb{Z}$, where τ is the temporal or spatial shift component, j is the scale component governing the dilation or contraction of the wavelet basis function $\psi_{j,\tau}$. These wavelet functions are themselves derived from a low-pass generating function $\phi_{j,\tau}(t) = 2^{-j/2}\phi(2^{-j}t - \tau)$, $j, \tau \in \mathbb{Z}^2$. Dilations of ϕ enable traversals from one scale to the next. This process can alternatively be considered in terms of the sequence of subspaces or signal spaces which they generate (Mallat, 1989).

$$\{0\} \subset \dots \subset V_j \subset V_2 \subset V_1 \subset V_0 \subset V_{-1} \subset V_{-2} \subset \dots \subset L_2\{\mathbb{R}\}$$

Resolution increasing \rightarrow

The basis $\{\phi_{j,\tau}\}_{j = J,\tau}$ and all its translated versions are orthonormal at any scale (or resolution) J and span the subspace V_j . When the resolution becomes arbitrarily fine, $j \rightarrow$

$-\infty$, and $V_{-\infty}$ converges toward the original or native signal space $L_2\{\mathbb{R}\}$.

Conversely, when the resolution becomes very coarse, all detail about the signal is ultimately lost $V_j \cong \{0\}$. The detail absent at a particular level V_j is contained by the wavelet collection $\{\psi_{j,\tau}\}_{j = J,\tau}$, i.e., the orthogonal complement of $\{\phi_{j,\tau}\}_{j = J,\tau}$. The wavelets define a separate, orthogonal subspace W_j . Progression up the resolution ladder can be attained by adding the missing detail at the j th level.

$$V_{j-1} = V_j \oplus W_j \quad (20)$$

A fine-to-coarse progression, in contrast, would just require the iteration $\{\phi_{j,\tau}\}_{j = j+1,\tau}$. The wavelet basis like the scaling function $\{\psi_{j,\tau}\}_{j = J,\tau}$ is also orthonormal at scale J . However, across all scales and shifts, $\{\psi_{j,\tau}\}_{j,\tau \in \mathbb{Z}^2}$ forms an orthonormal basis of $L_2\{\mathbb{R}\}$. Because $V_j, W_j \subset V_{j-1}$, there must be level-wise connections between ϕ and ψ .

$$\phi(x) = \sqrt{2} \sum_{l=0}^{L-1} h_l \phi(2x - l) \quad (21)$$

and for the wavelet

$$\psi(x) = \sqrt{2} \sum_{l=0}^{L-1} g_l \phi(2x - l). \quad (22)$$

These equations provide the link between the scaling function at the next finest scale and its integer translates with the filter h_l . In a similar fashion, the wavelet spanning the detail space at the higher resolution is generated from the same scaling function, except the connection is effected with a different filter sequence g_l . These filter sequences, $H = \{h_l\}_{0 \leq l \leq L-1}$ and $G = \{g_l\}_{0 \leq l \leq L-1}$, also called analysis filters are time-reversed, modulated versions of one another, $g_l = (-1)^l h_{L-l-1}$, $0 \leq l \leq L-1$, and for most wavelet bases also satisfy conditions that render them perfect reconstruction, or quadrature mirror filters.

C.1.2. Higher dimensions

Wavelet filters can also be generated for higher dimensions from the fundamental 1D case. If we define the domain $x = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d$.

Then for 2D, the basis consists of the following three wavelets

$$\begin{aligned} \psi_{j,\tau}^{(1)}(x) &= \psi(x_1)\phi(x_2); \psi_{j,\tau}^{(2)}(x) = \phi(x_1)\psi(x_2); \\ \psi_{j,\tau}^{(3)}(x) &= \psi(x_1)\psi(x_2) \end{aligned} \quad (23)$$

that is, $1 \leq l < 2^d - 2 - 1$ and so on for higher dimensions (see Mallat, 1998).

$$\psi_{j,\tau}^{(l)}(x_1, x_2, x_3) = 2^{-d(j/2)} \psi^l(2^{-j}x_1 - \tau, 2^{-j}x_2 - \tau, 2^{-j}x_3 - \tau) \quad (24)$$

Subsequently, $\Psi = \{\psi^{(1)}, \psi^{(2)}, \psi^{(3)}\}$ is an orthonormal basis of $L^2(\mathbb{R}^2)$, and the wavelet expansion coefficients of the 2D image $U(x)$ are

$$c = \int_{\mathbb{R}^2} U(x)\Psi dx \quad (25)$$

that is,

$$U = \sum_{\substack{\tau \in \mathbb{Z}^2 \\ j, l \in \mathbb{Z}}} c_{j, \tau, l} \psi_{j, \tau}^{(1)}$$

Because we are dealing with finite domains, namely bounded images, τ is constrained to those shifts occurring only in the span of the image. Also, in practical situations, the following restrictions on the scale should be imposed; (1) an upper bound on $j = K$; (2) $K < 0$; (3) $J \ll K$, i.e., $2^J = n$. The wavelet coefficients are normally computed (Eq. (25)) in the context of the efficient dual channel filter banks described by Mallat (1989). The procedure is a realization of Eq. (25) effecting at each scale two filtering steps, one low, associated with H , and one high, associated with G . These filtering steps precede a subsampling stage that captures the resolution-dependent aspect of the transformation. At each stage, the wavelet coefficients are retained and the scaling coefficients are passed on to the next level. This process is repeated until the coarsest scale L is reached wherein both vectors of expansion coefficients are retained. Thus

$$U = \sum_{\substack{j \leq L, \tau \in \mathbb{Z}^2 \\ l \in \mathbb{Z}}} c_{j, \tau}^l \psi_{j, \tau}^l + \sum_{\tau \in \mathbb{Z}^2} d_{L, \tau} \phi_{L, \tau} \quad (26)$$

References

- Abramovich, F., Benjamini, Y., 1995. Thresholding of wavelet coefficients as multiple hypotheses testing procedure. In Antoniadis and Oppenheim (Antoniadis and Oppenheim, 1995), pp. 5–14.
- Abry, P., Gonçalves, P., Flandrin, P., 1995. Wavelets, spectrum analysis and 1/f processes. In Antoniadis and Oppenheim (Antoniadis and Oppenheim, 1995), pp. 15–29.
- Antoniadis, A., Oppenheim, G. (Eds.), 1995. Wavelets and Statistics. Lecture Notes in Statistics, vol. 103. Springer-Verlag, New York.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J.R. Stat. Soc., B* 57 (1), 289–300.
- Besag, J., 1986. On the statistical analysis of dirty pictures. *J. Royal Stat. Soc.* 48 (Ser. B), 259–302.
- Boynton, G.M., Engel, S.A., Glover, G.H., Heeger, D.J., 1996. Linear systems analysis of functional MRI in human v1. *J. Neurosci.* 16, 4207–4221.
- Brammer, M.J., 1998. Multidimensional wavelet analysis of functional magnetic resonance images. *Hum. Brain Mapp.* 6, 378–382.
- Bullmore, E., Brammer, M., Williams, S.C.R., Rabe-Hesketh, S., Janot, N., David, A., Mellors, J., Howard, R., Sham, P., 1996. Statistical methods of estimation and inference for functional MR image analysis. *Magn. Reson. Med.* 35, 261–277.
- Bullmore, E., Long, C.J., Suckling, J., Fadili, J., Calvert, G., Zelaya, F., Carpenter, T.A., Brammer, M., 2002. Coloured noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains. *Hum. Brain Mapp.* 12, 61–78.
- Chambolle, A., Devore, R., Lee, N., Luicier, B., 1998. Nonlinear wavelet image processing: variational problems, compression, and noise removal through wavelet shrinkage. *IEEE Trans. Image Process.* 7 (3), 319–335 (March).
- Coifman, R.R., Meyer, Y., Wickerhauser, M.V., 1992. Wavelet analysis and signal processing. In: Ruskai, M.B., Beylkin, G., Coifman, R., Daubechies, I., Mallat, S., Meyer, Y., Raphael, L. (Eds.), *Wavelets and their Applications*. Jones and Barlett, Boston, pp. 153–178.
- Cox, R.W., Jesmanowicz, A., 1999. Real-time 3d image registration for functional MRI. *Magn. Reson. Med.* 42, 1014–1018.
- Daubechies, I., 1990. The wavelet transform, time–frequency localization and signal analysis. *IEEE Trans. Inf. Theory* 36 (5), 961–1005.
- Descornes, X., Kruggel, F., Von-Cramon, D., 1998a. Spatio-temporal fMRI analysis using Markov random fields. *IEEE Trans. Med. Imag.* 17 (6), 1028–1039.
- Descornes, X., Kruggel, F., Von-Cramon, D.Y., 1998b. fMRI signal restoration using a spatio-temporal Markov random field preserving transitions. *NeuroImage* 8, 340–349.
- Devore, R.A., Lucier, B.J., 1992. Fast wavelet techniques for near-optimal image processing. *IEEE Military Commun. Conf. Rec.*, Piscataway, NJ, pp. 1129–1135.
- Donoho, D.L., 1995. De-noising by soft-thresholding. *IEEE Trans. Inf. Theory* 41 (3), 613–627.
- Donoho, D.L., Johnstone, I.M., 1994. Ideal spatial adaptation by wavelet shrinkage. *Biometrika* 81 (3), 425–455.
- Donoho, D.L., Johnstone, I.M., 1995. Adapting to unknown smoothness by wavelet shrinkage. *J. Am. Stat. Assoc.* 90, 1200–1224.
- Fadili, M.J., Bullmore, E.T., 2002. Wavelet-generalized least squares: a new BLU estimator of linear regression models with 1/f errors. *NeuroImage* 15, 217–232.
- Forman, S., Cohen, J., Fitzgerald, M., Eddy, W., Mintun, M., Noll, D., 1995. Improved assessment of significant activation if fMRI: use of a cluster-size threshold. *Magn. Reson. Med.* 33, 636–647.
- Fournier, A., 1996. Wavelet analysis of observed geopotential and wind: blocking and local energy coupling across scales. In: Unser, M.A., Aldroubi, A., Laine, A.F. (Eds.), *Wavelet Application in Signal and Image Processing IV*. Proceedings of SPIE, vol. 2825, p. 1044. 4–9 August, Denver, CO.
- Friston, K.J., 1997. SPM course notes. <http://www.fil.ion.ucl.ac.uk/spm/course/notes.html>.
- Friston, K.J., Jezzard, P., Turner, R., 1994. Analysis of functional MRI time series. *Hum. Brain Mapp.* 1, 153–171.
- Geman, S., Geman, D., 1984. Stochastic relaxation, gibbs distribution and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* 6 (6), 721–741.
- Johnstone, I.M., Silverman, B.W., 1997. Wavelet threshold estimators for data with correlated noise. *J. R. Stat. Soc.* 59 (2), 319–351.
- Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, P.B., Kennedy, D.M., Hoppel, B.E., Cohen, M.S., Turner, R., Brady, H.M., Rossen, B.R., 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Nat. Acad. Sci. U. S. A.* 89, 5675–5679.
- Locascio, J.J., Jennings, P.J., Moore, C.I., Corkin, S., 1997. Time series analysis in the time domain and resampling methods for studies of functional magnetic resonance brain imaging. *Hum. Brain Mapp.* 5, 168–193.
- Long, C.J., Bullmore, E.T., Brown, E.N., Solo, V., Brammer, M.J., 2001. Regularised fMRI signal response estimation in a library of orthonormal bases. *Proc. Int. Soc. Magn. Reson. Med.*, 1738 (Glasgow).
- Mallat, S., 1989. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (7), 674–693.
- Mallat, S., 1996. Wavelets for a vision. *Proc. IEEE* 84 (4), 604–614.
- Mallat, S., 1998. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego.
- Mandeville, J., Marota, J., Keltner, J., Kosofsky, B., Burke, J., Hyman, L., LaPointe, L., Reese, T., Kwong, K., Rosen, B., Weissleder, R., Weisskoff, R., 1996. CBV functional imaging in rat brain using iron oxide agent at steady-state concentration. *Int. Soc. Magn. Res. Med.*, 292.
- Manoach, D.S., Schlag, G., Siewert, B., et al., 1997. Prefrontal cortex fMRI signal changes are correlated with working memory load. *NeuroReport* 8, 545–549.
- Manoach, D.S., Greve, D.N., Lindgren, K.A., Dale, A.M., 2003. Identifying regional activity associated with temporally separated components

- of working memory using event-related functional MRI. *NeuroImage* 20 (3), 1670–1684.
- Manoach, D.S., White, N.S., Lindgren, K.A., Heckers, S., Coleman, M.J., Dubal, S., Holzman, P.S., 2004. Hemispheric specialization of the lateral prefrontal cortex for strategic processing during spatial and shape working memory. *NeuroImage* 21 (3), 894–903.
- Marchini, J., Ripley, B., 2000. A new statistical approach to detecting-significant activation in functional MRI. *NeuroImage* 12, 366–380.
- Mitra, P.P., Ogawa, S., Hu, X., Ugurbil, K., 1997. The nature of spatio-temporal changes in cerebral haemodynamics as manifested in functional MRI. *Magn. Reson. Med.* 37, 511–518.
- Nason, G.P., von Sachs, R., 1999. Wavelets in time series analysis. *Philos. Trans. R Soc. London, A* 357 (1760), 2511–2526.
- Ogawa, S., Tank, D.W., Menon, R., Ellerman, J.M., Kim, S.G., Merkle, H., Ugurbil, K., 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc. Natl. Acad. Sci. U. S. A.* 89, 5951–5955.
- Purdon, P.L., Weisskoff, R.M., 1998. Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fMRI. *Hum. Brain Mapp.* 6 (4), 239–249.
- Purdon, P.L., Solo, V., Weisskoff, R.M., Brown, E.N., 2001. Locally regularized spatio-temporal modelling and model comparison for functional MRI. *NeuroImage* 14, 912–923.
- Ruttiman, E.E., Unser, M., Dawlings, R.R., Ramsey, N.F., Roi, D., Mattay, V.S., Hommer, D.W., Frank, J.A., Weinberger, D.R., 1998. Statistical analysis of functional MRI data in the wavelet domain. *IEEE Trans. Med. Imag.* 17, 142–152.
- 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, 6558–6563.
- Solo, V., Purdon, P.L., Brown, E.N., Weisskoff, R.M., 2001. A signal estimation approach to functional MRI. *IEEE Trans. Med. Imag.* 20, 26–35.
- Vetterli, M., Kovačević, J., 1995. *Wavelets and Subband Coding*. Prentice Hall, New Jersey.
- von Sachs, R., 1998. *Nonparametric wavelet methods for nonstationary time series*. Tech. Rep. 98/19, Institut de Statistique, Université Catholique de Louvain.
- Weisskoff, R.M., Baker, J., Belliveau, J., Davis, T.L., Kwong, K.K., Cohen, M.S., Rosen, B.R., 1993. Power spectrum analysis of functionally-weighted MR data: what’s in the noise? *Proc. Int. Soc. Magn. Reson. In Med.*, 1–7.
- Wold, H., 1954. *A Study in the Analysis of Stationary Time Series*. Almqvist and Wiksell, Stockholm.
- Yoo, S., Guttman, C.R., Panych, L., 2001. Multiresolution data acquisition and detection in functional MRI. *NeuroImage* 14 (6), 1476–1485.