

WSPM: Wavelet-based statistical parametric mapping

Dimitri Van De Ville,^{a,*} Mohamed L. Seghier,^b François Lazeyras,^c
Thierry Blu,^a and Michael Unser^a

^aBiomedical Imaging Group, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

^bWellcome Trust Centre for Neuroimaging, UCL, London, UK

^cDepartment of Radiology and Medical Informatics, University Hospital Geneva, Switzerland

Received 16 February 2007; revised 23 May 2007; accepted 3 June 2007
Available online 19 June 2007

Recently, we have introduced an integrated framework that combines wavelet-based processing with statistical testing in the spatial domain. In this paper, we propose two important enhancements of the framework. First, we revisit the underlying paradigm; i.e., that the effect of the wavelet processing can be considered as an adaptive denoising step to “improve” the parameter map, followed by a statistical detection procedure that takes into account the non-linear processing of the data. With an appropriate modification of the framework, we show that it is possible to reduce the spatial bias of the method with respect to the best linear estimate, providing conservative results that are closer to the original data. Second, we propose an extension of our earlier technique that compensates for the lack of shift-invariance of the wavelet transform. We demonstrate experimentally that both enhancements have a positive effect on performance. In particular, we present a reproducibility study for multi-session data that compares WSPM against SPM with different amounts of smoothing. The full approach is available as a toolbox, named WSPM, for the SPM2 software; it takes advantage of multiple options and features of SPM such as the general linear model.

© 2007 Elsevier Inc. All rights reserved.

Keywords: Wavelets; Wavelet thresholding; Statistical testing; Bias reduction; Shift-invariant transform; Reproducibility study

Introduction

Statistical parametric mapping (SPM) (Friston et al., 1995; Frackowiak et al., 1997) is probably the most popular parametric hypothesis-driven method for the analysis of fMRI data. To control the multiple testing problem, SPM considers the data as a lattice

representation of a continuous Gaussian Random Field (GRF). To conform with this hypothesis, the data are pre-smoothed with a Gaussian filter with fixed size (Worsley et al., 1996; Poline et al., 1997). The user has the option to adjust the smoothing strength for optimal detection (compromise between SNR enhancement and spatial definition).

The discrete wavelet transform (DWT) has also been applied to the analysis of fMRI data for parametric hypothesis testing. Three important properties justify the use of wavelets in this application. First, they typically encode activation patterns with a small set of wavelet coefficients; this is referred to as the DWT’s sparsifying property. Second, an orthogonal DWT leaves the noise evenly distributed in the wavelet domain; it therefore increases the SNR. Third, the DWT acts (approximately) as a decorrelator. Therefore, a conservative Bonferroni correction for multiple hypothesis testing is closer to optimal in the wavelet domain than in the spatial domain. Basically, the standard wavelet approach to parametric hypothesis testing consists in statistical testing the wavelet domain representation of the parameter map (Ruttimann et al., 1998; Turkheimer et al., 2000). The remaining difficulty is how to fully exploit the reconstruction of the parameter map after thresholding in the wavelet domain. Several approaches have been proposed, such as variance reconstruction or estimation of the residual in the spatial domain (Desco et al., 2005; Aston et al., 2005), Bayesian modeling (Turkheimer et al., 2006; Flandin and Penny, 2007), optimizing statistical power while controlling false discovery rate (Sendur et al., 2005; Srikanth et al., 2006), or considering the wavelet processing as an alternative preprocessing step (Wink and Roerdink, 2004). Recently, we have proposed a variation of the wavelet-based framework (WSPM) that performs the model-fitting and processing in the wavelet domain, but transfers the statistical testing back into the spatial domain (Van De Ville et al., 2004).

Most parametric hypothesis-driven approaches that have been proposed so far fall within the general conceptual framework that is summarized in Fig. 1. In SPM, the initial smoothing with a Gaussian filter can be seen as a denoising procedure to increase the SNR of the data. The subsequent detection procedure is applied on

* Corresponding author. Fax: +41 21 6933701.

E-mail addresses: Dimitri.VanDeVille@epfl.ch (D. Van De Ville), m.seghier@fil.ion.ucl.ac.uk (M.L. Seghier), Francois.Lazeyras@sim.hcuge.ch (F. Lazeyras), Thierry.Blu@epfl.ch (T. Blu), Michael.Unser@epfl.ch (M. Unser).

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




Fig. 1. Conceptual view of a parametric hypothesis-driven approach to fMRI data analysis. First, the original time-series data are fitted according to a general linear model, resulting in a so-called parameter map. This operation is voxel-wise. Next, the parameter map is denoised to increase the SNR. This can be done using assumptions for the activation patterns. Finally, an activation map is obtained by statistical testing of the denoised parameter map.

the parameter map obtained from the smoothed data using the GLM; it is based on GRF theory. The temporal processing by the GLM and the spatial processing by the Gaussian filter are both linear and separable (time \times space) operations and therefore can be interchanged; i.e., SPM's parameter map can also be obtained by filtering the parameter map obtained from unsmoothed data. In our wavelet-based framework, we can also separate the different processing steps according to Fig. 1: temporal modeling by the GLM followed by (non-linear) denoising in the wavelet domain. The subsequent statistical testing takes into account the influence of the wavelet processing but remains in the spatial domain, which has obvious advantages.

In this paper, we further investigate and extend the integrated WSPM framework in two important aspects:

- (1) Reduction of spatial bias: Although the framework takes into account the statistical effect of denoising, it can still manifest spatial bias. For instance, weak activations in the parameter map may resemble the underlying pattern poorly if only a limited subset of wavelet coefficients is retained. Notice that a similar effect can be observed with SPM: the data are denoised by the Gaussian filtering and so is the parameter map as well, which potentially introduces spatial bias. We extend the wavelet framework to reduce this effect. Consequently, the final detected parameter map can be considered as more closely related to the measured data with respect to spatial bias.
- (2) Better shift-invariance: The fact that the DWT is shift-variant is often recognized as a major disadvantage. We show how to incorporate results of multiple non-redundant DWTs, for different shifts of the data, which essentially make the analysis shift-invariant.

The proposed framework has been integrated into SPM2 as a “WSPM toolbox”, allowing the user to have the usual SPM-based analysis and the wavelet-based framework side by side. With respect to temporal domain modeling, the full GLM fitting as provided by SPM (including compensation for serial correlation) is used.

In what follows, we briefly review the standard wavelet-based method, and then introduce the enhanced framework with bias reduction and the shift-invariant extension. We illustrate and discuss the concept with several examples. First, we present a 1D synthetic data set to demonstrate the main effect of the proposed enhancements. Then, we analyze an experimental multi-session data set using a reproducibility study to estimate the sensitivity and specificity of WSPM and SPM at different smoothing settings. We extend the typical receiver-operating-characteristics

(ROC) curves by a third dimension that measures the bias of the methods with respect to the best-linear-unbiased estimated of the non-smoothed data. These tri-variate plots allow us to evaluate the trade-off that is provided by the various methods and settings.

Background

We start this section with a brief review of the DWT. We then present the integrated framework for fMRI analysis.

The discrete wavelet transform

The discrete wavelet transform (DWT) is a powerful tool for multi-resolution signal analysis (Mallat (1989, 1998)). The transform of a multi-dimensional signal $v(\mathbf{x})$ is a decomposition into a sum of shifted and scaled versions of (bandpass) wavelet functions (typically 2^d-1 different wavelets in d dimensions) and shifted versions of a (lowpass) scaling function. We introduce the shorthand notation

$$v(\mathbf{x}) = \sum_{\mathbf{k}} v_w[\mathbf{k}] \psi_{\mathbf{k}}(\mathbf{x}), \quad (1)$$

where $v_w[\mathbf{k}]$ can be both lowpass or wavelet coefficients at all different scales and orientations, for which $\psi_{\mathbf{k}}$ are the corresponding basis functions.

A review of the integrated framework for wavelet-based statistical analysis of fMRI data

We denote an fMRI data set as $v[\mathbf{n}; t]$, $\mathbf{n} \in Z^3, t \in Z$, where \mathbf{n} and $t=1, \dots, N_t$ are the 3D-spatial and temporal indices, respectively. The non-redundant spatial 3D DWT of a volume, $v[\mathbf{n}; t]$, yields the coefficients $v_w[\mathbf{k}; t]$. As in Eq. (1), we compactly denote the wavelet decomposition as

$$v[\mathbf{n}; t] = \sum_{\mathbf{k}} v_w[\mathbf{k}; t] \psi_{\mathbf{k}}(\mathbf{n}). \quad (2)$$

Since each volume $v[\mathbf{n}; t]$ is transformed to the wavelet domain, we can introduce a time-series vector of length N_t for each wavelet coefficient: $v_w[\mathbf{k}] = [v_w[\mathbf{k}; 1] \dots v_w[\mathbf{k}; N_t]]^T$. Next, the temporal behavior of the wavelet coefficient is described by a GLM (Friston et al., 1995; Van De Ville et al., 2003; Mueller et al., 2003), so we put forward

$$v_w[\mathbf{k}] = \mathbf{X}_y v[\mathbf{k}] + \mathbf{e}_w[\mathbf{k}], \quad (3)$$

where \mathbf{X} is the $N_t \times L$ design matrix, and $\mathbf{e}_w[\mathbf{k}]$ a (random) error. The matrix \mathbf{X} contains L signal regressors.

Given the observed data $\mathbf{v}_w[\mathbf{k}]$, the least-squares estimate of the parameters of the GLM is given by $\bar{\mathbf{y}}_w[\mathbf{k}] = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{v}_w[\mathbf{k}]$. This estimate is optimal provided that the error component is independently and identically Gaussian distributed. The corresponding residual is $\mathbf{e}_w[\mathbf{k}] = \mathbf{v}_w[\mathbf{k}] - \mathbf{X} \bar{\mathbf{y}}_w[\mathbf{k}]$. Next, the information of interest is extracted from $\bar{\mathbf{y}}_w[\mathbf{k}]$ by a contrast vector \mathbf{c} . At this stage, we obtain two scalar values for the \mathbf{k} -th wavelet coefficient:

$$u_w[\mathbf{k}] = \mathbf{c}^T \bar{\mathbf{y}}_w[\mathbf{k}], \quad (4)$$

$$s_w^2[\mathbf{k}] = \bar{\mathbf{e}}_w[\mathbf{k}]^T \bar{\mathbf{e}}_w[\mathbf{k}] \mathbf{c}^T (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{c}, \quad (5)$$

where $u_w[\mathbf{k}]$ and $s_w^2[\mathbf{k}]/J$, under the null hypothesis, follow a Gaussian and a χ^2 distribution (with $J = N_t - \text{rank}(\mathbf{X})$ degrees of freedom), respectively (Jezzard et al., 2001; Worsley et al., 2002). At this point, we introduce the unprocessed parameter map as the reconstruction of the raw wavelet coefficients $u_w[\mathbf{k}]$ without any processing:

$$u[\mathbf{n}] = \sum_{\mathbf{k}} u_w[\mathbf{k}] \psi_{\mathbf{k}}(\mathbf{n}). \quad (6)$$

Notice that the same $u[\mathbf{n}]$ would be obtained by fitting the GLM directly in the spatial domain.

The wavelet processing of the integrated framework consists of thresholding the parameter map $u_w[\mathbf{k}]$ based on the corresponding t -values, which are determined as

$$t_w[\mathbf{k}] = \frac{u_w[\mathbf{k}]}{\sqrt{s_w^2[\mathbf{k}]/J}}. \quad (7)$$

This operation requires a threshold τ_w . In WSPM, this threshold is not obtained as the result of a statistical test, but kept as a general parameter of the algorithm. The wavelet coefficients $u_w[\mathbf{k}]$ that survived the thresholding $|t_w[\mathbf{k}]| > \tau_w$ are then reconstructed as

$$\tilde{u}[\mathbf{n}] = \sum_{\mathbf{k}} \frac{H(|t_w[\mathbf{k}]| - \tau_w) u_w[\mathbf{k}] \psi_{\mathbf{k}}(\mathbf{n})}{\tilde{u}_w[\mathbf{k}]}, \quad (8)$$

where $H(t)$ is the Heaviside step function defined as

$$H(t) = \begin{cases} 0, & \text{when } t < 0, \\ 1, & \text{otherwise} \end{cases} \quad (9)$$

In other words, the term $H(|t_w[\mathbf{k}]| - \tau_w)$ in Eq. (8) acts as an indicator function which is equal to 1 for $|t_w[\mathbf{k}]| \geq \tau_w$ and 0 otherwise.

Finally, we perform a hypothesis test in the spatial domain to determine whether the voxel \mathbf{n} is activated and thus has a non-zero mean:

$$\mathcal{H}_0 : E[\tilde{u}[\mathbf{n}]] = 0, \quad (10)$$

$$\mathcal{H}_1 : E[\tilde{u}[\mathbf{n}]] > 0. \quad (11)$$

Notice that we opt for a one-sided test as it is also SPM's convention. The main result of our integrated framework (Van De Ville et al., 2004) states that the statistical test amounts to checking whether or not $\tilde{u}[\mathbf{n}] \geq \tau_s \Lambda[\mathbf{n}]$, where τ_s is the threshold in the

spatial domain and $\Lambda[\mathbf{n}]$ is the reconstruction of the values $s_w[\mathbf{k}]/\sqrt{J}$ by a modified inverse DWT algorithm:

$$\Lambda[\mathbf{n}] = \sum_{\mathbf{k}} \frac{s_w[\mathbf{k}]}{\sqrt{J}} |\psi_{\mathbf{k}}(\mathbf{n})|. \quad (12)$$

The result is the detected parameter map

$$\tilde{u}'[\mathbf{n}] = H(\tilde{u}[\mathbf{n}] - \tau_s \Lambda[\mathbf{n}]) \tilde{u}[\mathbf{n}], \quad (13)$$

for which the probability of a false detection under the null hypothesis can be bound as

$$P[\tilde{u}'[\mathbf{n}] \geq \tau_s \Lambda[\mathbf{n}]] \leq \Upsilon(\tau_w, \tau_s) = \alpha_B, \quad (14)$$

where $\Upsilon(\tau_w, \tau_s) = \min_{a>0} E[(1 + a(\xi_{\tau_w} - \tau_s \varsigma))_+]$ is a data-independent function of the parameters τ_w and τ_s , with ξ_{τ_w} and ς being known reference random variables¹ and $(t)_+ = \max(0, t)$. This bound can be chosen equal to α_B , the desired confidence level after Bonferroni correction; i.e., $\alpha_B = \alpha/N_c$, where N_c is the number of intracranial voxels. However, there is an infinite number of combinations (τ_w, τ_s) that yield the same probability α_B . In Van De Ville et al. (2004), we looked for the solution that minimizes the worst-case error between the unprocessed and detected parameter map. This constraint turned out to be equivalent to minimizing the sum $\tau_w + \tau_s$, subject to $\Upsilon(\tau_w, \tau_s) = \alpha_B$. A closed-form solution can be found for $N_t > 50$, as is typically the case for fMRI:

$$\tau_w = \sqrt{-W_{-1}(-2\pi\alpha_B^2)}, \quad \tau_s = 1/\tau_w, \quad (15)$$

where W_{-1} is the -1 -branch of the Lambert W -function.

In Fig. 2, we show schematically the integrated framework and its different steps.

Methods

Specificity improvement by spatial bias reduction

We revisit the central idea of the integrated framework: the separation of denoising and detection. Let us take a closer look at the schematic overview of Fig. 2. We recognize that $u[\mathbf{n}]$ is the linear estimate of the parameter map, which is the best possible unbiased estimate under the Gaussian assumption. Unfortunately, a direct voxel-wise statistical test on $u[\mathbf{n}]$ with Bonferroni correction would show a disappointingly low sensitivity since spatial correlation is not taken into account. This is the primary reason to exploit the spatial correlation by thresholding in the wavelet domain: noise is suppressed for a good part, while signal-carrying wavelet coefficients are essentially preserved. The framework then relies on a bound for the probability that the denoised parameter map $\tilde{u}[\mathbf{n}]$ gets falsely detected under the null hypothesis:

$$P[\tilde{u}[\mathbf{n}] \geq \tau_s \Lambda[\mathbf{n}]] \leq \Upsilon(\tau_w, \tau_s), \quad (16)$$

where $\tau_s \Lambda[\mathbf{n}]$ is an appropriate non-stationary threshold adapted for testing the denoised map $\tilde{u}[\mathbf{n}]$. By construction, this result is conservative with respect to the denoised parameter map $\tilde{u}[\mathbf{n}]$. Since this parameter map is obtained by an inverse DWT from (typically) a subset of wavelet coefficients, the reconstruction can be spatially biased by the synthesis process.

¹ ξ_{τ_w} follows a truncated normalized Gaussian distribution, and ς and χ -distribution with J degrees of freedom, see Van De Ville et al. (2004).

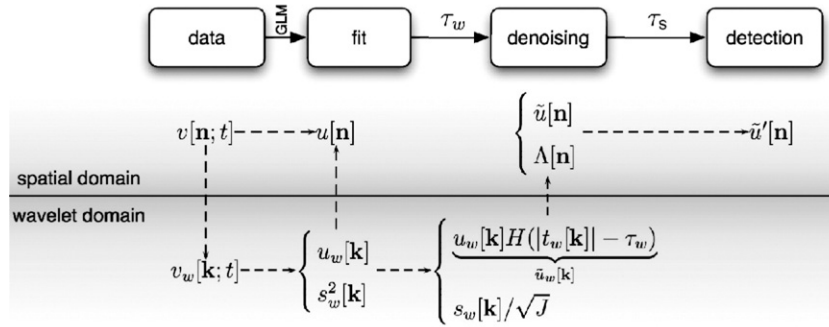


Fig. 2. Schematic overview of the integrated framework.

We first recall that checking Eq. (16) corresponds to a one-sided test; i.e., larger values of $\tilde{u}[\mathbf{n}]$ increase the probability of detection. With this in mind, we propose to construct an improved parameter map $\hat{u}[\mathbf{n}]$ out of $\tilde{u}[\mathbf{n}]$ with reduced spatial bias. We opt for a conservative point-of-view; i.e., we want to correct for the case when the parameter is overestimated by the reconstruction after thresholding in the wavelet domain. This typically arises when the thresholding operation keeps isolated wavelet coefficients in which case the secondary ripples of the wavelet function may survive spatial thresholding (“ringing” effect). Therefore, we compare the denoised map $\tilde{u}[\mathbf{n}]$ against the linear estimate $u[\mathbf{n}]$, and we construct the corrected map as

$$\hat{u}[\mathbf{n}] = \min(u[\mathbf{n}], \tilde{u}[\mathbf{n}]) \leq \tilde{u}[\mathbf{n}]. \quad (17)$$

Basically, this means that we do only follow the denoised map if it is not higher than the linear estimate. Our choice could miss activations that are underestimated by the linear approach (as the denoised map will be corrected accordingly). However, the linear estimate is unbiased and we believe that spatial bias introduced by the basis functions is important to compensate for.

The conservativeness of the spatial bias reduction (17) follows from $\hat{u}[\mathbf{n}]$ being obviously closer to $u[\mathbf{n}]$ than $\tilde{u}[\mathbf{n}]$ is; i.e., $\|\hat{u} - u\| \leq \|\tilde{u} - u\|$. Because $\hat{u}[\mathbf{n}] \leq \tilde{u}[\mathbf{n}]$, we automatically have that $\hat{u}[\mathbf{n}] \geq \tau_s \Lambda[\mathbf{n}]$ also implies $\tilde{u}[\mathbf{n}] \geq \tau_s \Lambda[\mathbf{n}]$. Consequently, we have

$$P[\hat{u}[\mathbf{n}] \geq \tau_s \Lambda[\mathbf{n}]] \leq P[\tilde{u}[\mathbf{n}] \geq \tau_s \Lambda[\mathbf{n}]], \quad (18)$$

which, in turn, is upper bounded by $\Upsilon(\tau_w, \tau_s)$. We can thus perform the same test on $\hat{u}[\mathbf{n}]$ as we did on $\tilde{u}[\mathbf{n}]$, with the same bound. Finally, the detected parameter map becomes

$$\hat{u}'[\mathbf{n}] = H(\hat{u}[\mathbf{n}] - \tau_s \Lambda[\mathbf{n}]) \hat{u}[\mathbf{n}]. \quad (19)$$

Another useful interpretation of the improved parameter map $\hat{u}[\mathbf{n}]$ is illustrated by Fig. 3. Clearly, we would like the denoised version $\tilde{u}[\mathbf{n}]$ to be as close as possible to the underlying true activation $u_0[\mathbf{n}]$. Since $u[\mathbf{n}]$ is the best linear unbiased estimate for $E[u[\mathbf{n}]] = u_0[\mathbf{n}]$, we can interpret $e[\mathbf{n}] = u[\mathbf{n}] - \tilde{u}[\mathbf{n}]$ as an estimate for the difference $u_0[\mathbf{n}] - E[\tilde{u}[\mathbf{n}]]$. When the denoised map $\tilde{u}[\mathbf{n}]$ satisfies $\tilde{u}[\mathbf{n}] > u[\mathbf{n}]$

(where $u[\mathbf{n}]$ is close to $u_0[\mathbf{n}]$, on average), it might survive spatial thresholding while the true map $u_0[\mathbf{n}]$ may not, thus leading to false positives. Therefore, when $e[\mathbf{n}] < 0$, we prefer to stick to the linear estimate by selecting $\hat{u}[\mathbf{n}] = u[\mathbf{n}]$. On the other hand, when $e[\mathbf{n}] \geq 0$, we keep $\hat{u}[\mathbf{n}] = \tilde{u}[\mathbf{n}]$ and take advantage of the powerful denoising properties of the non-linear operation in the wavelet domain.


We recall that the framework is characterized by the two threshold values τ_w and τ_s . For the original framework, we showed that minimizing the worst-case difference between the linear fit $u[\mathbf{n}]$ and the detected parameter map $\hat{u}'[\mathbf{n}]$, corresponds to minimizing the sum $\tau_w + \tau_s$. For the extended framework, we ask a similar question, but this time for the new end-to-end difference $|u[\mathbf{n}] - \hat{u}'[\mathbf{n}]|$. By expressing this difference as

$$\begin{aligned} |u[\mathbf{n}] - \hat{u}'[\mathbf{n}]| &= |u[\mathbf{n}] - \hat{u}[\mathbf{n}] + \hat{u}[\mathbf{n}] - \hat{u}'[\mathbf{n}]| \\ &\leq |u[\mathbf{n}] - \hat{u}[\mathbf{n}]| + |\hat{u}[\mathbf{n}] - \hat{u}'[\mathbf{n}]| \\ &= |u[\mathbf{n}] - \min(u[\mathbf{n}], \tilde{u}[\mathbf{n}])| + |\hat{u}[\mathbf{n}]| (1 - H(\hat{u}[\mathbf{n}] - \tau_s \Lambda[\mathbf{n}])) \\ &\leq |u[\mathbf{n}] - \tilde{u}[\mathbf{n}]| + \tau_s \Lambda[\mathbf{n}] \\ &\leq (\tau_w + \tau_s) \Lambda[\mathbf{n}], \end{aligned}$$

we see that the optimal values of τ_w and τ_s remain unmodified; i.e., they are obtained by minimizing their sum subject to $\Upsilon(\tau_w, \tau_s) = \alpha_B$.

Shift-invariant wavelet processing

One of the major disadvantages of the non-redundant DWT is its shift variance; i.e., a shift of the input signal does not simply translate into a shift of the wavelet coefficients. Consequently, shifting an activation region could turn out to give different detected patterns. The potential of the redundant (translation-invariant) DWT has been suggested before by Turkheimer et al. (2000) in the context of direct statistical testing in the wavelet domain. At first sight, plugging in the redundant DWT into our framework looks very tempting since the threshold values τ_w and τ_s would remain unchanged. Unfortunately, there is a catch. Small activation regions or activations barely detected would be detected by a few coefficients only, whose energy has been decreased by the redundancy factor. Therefore, the denoised parameter map would show lower values than in the non-redundant

Fig. 3. The denoising procedure of the integrated framework has been extended to obtain a parameter map $\hat{u}[\mathbf{n}]$ after denoising and bias reduction.

case (assuming the non-redundant transform is located at the “right” shift), and detection would become less sensitive.

We propose to mitigate this problem by analyzing the data under M different shifts. The data volumes are shifted by a vector \mathbf{d}_m , $m=0, \dots, M-1$ analyzed using the normal DWT, and the results are shifted back by $-\mathbf{d}_m$. That way, we obtain for each shift:

$$P[\hat{u}^{(m)}[\mathbf{n}] \geq \tau_s \Lambda^{(m)}[\mathbf{n}] \leq \Upsilon(\tau_w, \tau_s)]. \quad (20)$$

We combine these results by selecting the one that results into the highest normalized value:

$$P\left[\max_m \left(\frac{\hat{u}^{(m)}[\mathbf{n}]}{\Lambda^{(m)}[\mathbf{n}]}\right) \geq \tau_s\right] = P\left[\bigvee_{m=0}^{M-1} \hat{u}^{(m)}[\mathbf{n}] \geq \tau_s \Lambda^{(m)}[\mathbf{n}]\right] \quad (21)$$

$$\leq \sum_{m=0}^{M-1} P[\hat{u}^{(m)}[\mathbf{n}] \geq \tau_s \Lambda^{(m)}[\mathbf{n}]] \quad (22)$$

$$\leq M \Upsilon(\tau_w, \tau_s). \quad (23)$$

Clearly, the redundancy factor M increases the threshold values with respect to the non-redundant case. From this point of view, the proposed extension is not yet completely optimal, since it does not take into account the correlation between the different shifted versions. Despite the higher threshold values, experimental results will show that the combinations increases sensitivity, although the threshold values are slightly higher.

To be fully shift-invariant with J_w decomposition levels, we need $M=2^{2J_w}$ shifts in 2D, and $M=2^{3J_w}$ in 3D. They can be constructed easily by considering all possible shifts $0, \dots, 2^{J_w}-1$ in every dimension. For example, in 3D we have

$$\mathbf{d}_m = \begin{pmatrix} m \\ \lfloor m/2^{J_w} \rfloor \\ \lfloor m/2^{2J_w} \rfloor \end{pmatrix} \text{ modulo } 2^{J_w}, \quad m = 0, \dots, 2^{3J_w}-1. \quad (24)$$

In practice, it is sufficient to keep M low and to be only shift-invariant at the first decomposition level; i.e., 4 when the wavelet transform is deployed in 2D or 8 for 3D.

Compensation for serial correlation

In practice, fMRI data sets are known to have a non-negligible temporal correlation. Hence, if not properly compensated for, ordinary least-squares estimates are sub-optimal (i.e., the estimates of the standard deviation are biased). For our framework, we choose to incorporate the method deployed by SPM (Worsley and Friston, 1995; Friston et al., 2000, 2002). Basically, the least-squares estimate is properly weighted by the covariance matrix, which is estimated using an ARMA-model and a restricted maximum likelihood (ReML) method, which is then incorporated into the linear model to prewhiten the data. The degrees of freedom are estimated by the Satterthwaite approximation (Worsley and Friston, 1995). Since the temporal model is spatially stationary, it can be transposed without any problem to the time-series of the wavelet coefficients. Whitening and parameter estimation are dealt with by functions available from SPM.

WSPM: a new toolbox for SPM

The extended framework has been implemented as a “WSPM toolbox” (version 1.2) for SPM2. In this way, the user can setup his

experiments as usual using SPM’s extensive features for preprocessing (e.g., registration) and GLM specification, including the HRF modeling. Next to the standard analysis performed by SPM, the toolbox allows to use our framework for spatio-wavelet statistical testing. Its results are added as new “contrasts” to the SPM structure related to the experiment, and they can be explored using SPM’s extensive features for visualization and cluster analysis. This toolbox can be freely downloaded from <http://bigwww.epfl.ch/wspm/>.

Simulation results

For more insight, we illustrate the modification of the denoising step by a 1D example in Fig. 4. We selected the following activation function

$$u_0[n] = \exp\left(-\frac{(n-16)^2}{4}\right) + \frac{1}{3} \exp\left(-\frac{(n-28)^2}{4}\right), \quad n = 1, \dots, 32, \quad (25)$$

and generated $N_t=80$ signal realizations corrupted with additive white Gaussian noise of standard deviation 0.10. In Fig. 4(a), we show the true activation $u_0[n]$ and the linear fit $u[n]$, which is simply obtained as the average over all realizations.² The error bars correspond to the standard deviation. The integrated framework is applied using the orthogonal B-spline DWT (degree 3, $J_w=2$ decomposition levels), which corresponds to the well-known cubic Battle–Lemarié wavelet. The confidence level is $\alpha=1\%$, which yields the thresholds values $\tau_w=4.14$ and $\tau_s=0.24$ after Bonferroni correction. In Fig. 4(b), we observe that the wavelet basis functions contributing to the reconstructed signal of the main activation have a strong weight and induce bias in the denoised parameter map; i.e., the left positive sidelobe is spurious while the right positive sidelobe reinforces the smaller activation. In Fig. 4(c), we see how the systematic overshoot (positive sidelobes) has been suppressed in the bias-reduced parameter map \hat{u} ; the negative lobes are still remaining but they have no incidence on detections, as a one-sided test is used.

In Fig. 5, we take the previous 1D example again. In the non-redundant case, shown in Fig. 5(a) and (b), the influence of a shift becomes apparent. The result in Fig. 5(b) corresponds to shifting the activation by one sample position, which implies that the signal is less efficiently captured in the wavelet domain; so the smaller activation stays undetected. In Fig. 5(c), we see how the redundant DWT is able to combine the result for both shifts, but suffers from a reduced signal intensity. Finally, in Fig. 5(d), we show the proposed combination of Fig. 5(a) and (b). The corresponding threshold values have slightly changed to $\tau_w=4.31$ and $\tau_s=0.23$ due to the redundancy factor of 2.

Experimental results

In this section, we propose an evaluation of WSPM compared to SPM2, based on real multi-session fMRI data. The data set here comes from a carefully conducted experiment with auditory

² This example corresponds to a design matrix with only one column; i.e., the constant value.




Fig. 4. One-dimensional example to illustrate the effect of the spatial bias reduction after the denoising step. (a) Original signal and linear estimate. (b) Non-linear estimate after adaptive denoising in the wavelet domain. (c) Non-linear estimate after bias reduction. Legend: 1: ground truth u_0 ; 2: linear estimate u ; 3: non-linear estimate \hat{u} ; 4: spatial threshold $\tau_{s,A}$; 5: non-linear estimate with bias reduction \hat{u} .

stimulation. All experiments were approved by the Ethical Committee of the Geneva University Hospital.

Paradigm and stimuli

A block paradigm that alternates between “rest” and “stimulation” sequences is applied. One period of the design consists of

24 s of auditory stimulation followed by 12 s of silence. The subject is exposed to single-frequency acoustic stimulation, which are 0.5 s sinusoidal tone bursts at a rate of 1 Hz. These stimuli are delivered binaurally at a comfortable loudness level. Four different frequencies are used: 300 Hz, 1126 Hz, 2729 Hz, and 4690 Hz. A total of 4 sessions is carried out within the same fMRI experiment, and the order of blocks of auditory stimulation at different frequencies is permuted across sessions. In each session, the stimulation/silent block is repeated twice for each of the 4 frequency values for a total acquisition time of about 5 min per session.

MRI acquisition

The MRI data are acquired on a 1.5 T system (Philips Medical Systems, Best, The Netherlands). The multi-slice volume is positioned using sagittal scout images. Before the functional MR scans, an anatomical scan (a GRE T1-weighted sequence, TR/TE/Flip=162 ms/4.47 ms/80°, FOV=230 mm, matrix=256×256, slice-thickness=3 mm) is performed to acquire the same volume as in the functional session. Functional imaging consists of an EPI GRE sequence (TR/TE/Flip=1.2 s/40 ms/80°, FOV=230 mm, matrix=128×128, 14 contiguous 3 mm axial slices, spatial resolution 1.8×1.8×3 mm). The explored volume is measured 20 times during each period of auditory stimulation and 10 times during each silent period. Functional scanning is always preceded by 8 s of dummy scans to ensure tissue steady-state magnetization. The subject’s head is placed within a custom-designed headset with insulation, which filters out most of the MR scanners noise (70 dB attenuation for 250–8000 Hz range).

Data analysis

Data processing is carried out with the Statistical Parametric Mapping SPM2 software package (Wellcome Department of Imaging Neuroscience, London UK, <http://www.fil.ion.ucl.ac.uk/spm/>). All functional volumes of the 4 sessions are spatially realigned to a reference functional volume. Realigned images are then smoothed with an isotropic Gaussian kernel of appropriate size (FWHM). The pre-processed functional volumes of each subject are then submitted to fixed-effects analysis using the general linear model (GLM) applied at each voxel across the whole brain. Each condition of interest (i.e., auditory stimulation and silent blocks) is modeled by a boxcar waveform convolved with a canonical hemodynamic response function (with no dispersion or temporal derivatives) and subjected to a multiple regression analysis with six covariates of no interest representing the head motion parameters) see Friston et al., 1996a; Johnstone et al., 2006). The estimation process also includes high-pass filtering (1/128 Hz cutoff) to remove low-frequency noise and signal drift, as well as AR-modeling to compensate for serial correlation.

The statistical tests with SPM are performed on the smoothed and realigned images. Contrast volumes are computed for the main effect of auditory stimulation (“all frequencies”) versus silent periods over the 4 sessions. Statistical parametric maps of the t statistics (SPMt) are obtained; these are corrected for multiple testing based on SPM’s Gaussian Random Field theory (FWE correction). To evaluate the influence of smoothing, we considered three Gaussian filter settings: 4 mm, 6 mm, and 8 mm FWHM (voxel size=1.8×1.8×3 mm).

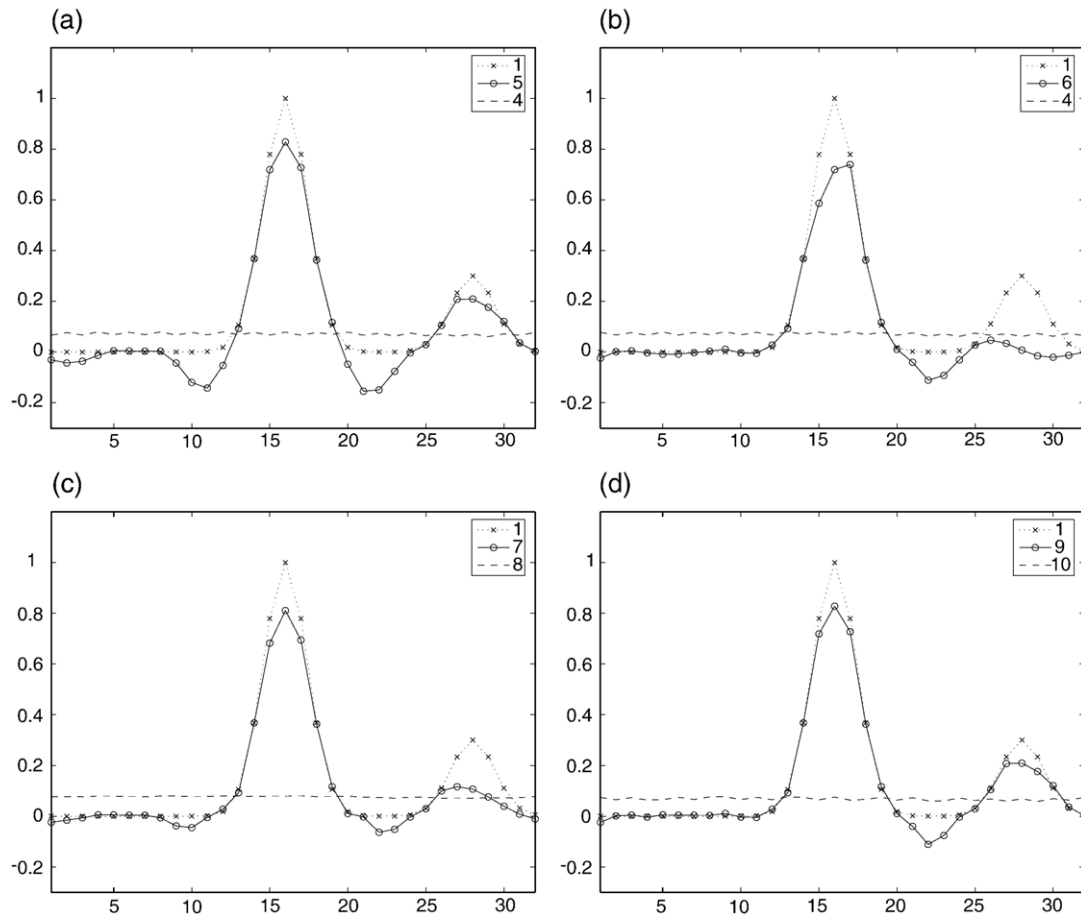


Fig. 5. One-dimensional example to illustrate the effect of translation. (a) Non-linear estimate for non-redundant DWT. (b) Non-linear estimate for non-redundant DWT, but signal shifted by one position. (c) Non-linear estimate for redundant DWT. (d) Combining the estimates using the non-redundant DWT. Legend: 1: ground truth u_0 ; 4: spatial threshold $\tau_s A$; 5: non-linear estimate $\hat{u}^{(1)}$ using non-redundant DWT; 6: non-linear estimate $\hat{u}^{(2)}$ using non-redundant DWT but with shifted signal; 7: non-linear estimate \hat{u}_{red} using the redundant DWT; 8: spatial threshold $\tau_s A_{\text{red}}$ using the redundant DWT; 9: non-linear estimate \hat{u}_{max} combining the non-redundant DWTs; 10: spatial threshold $\tau_s A_{\text{max}}$.

For the WSPM analysis, the same realigned images (but without smoothing) are decomposed using the following wavelet transforms:

- 2D orthogonal B-spline wavelet transform of degree 1.0, number of decomposition levels $J_w=1$ and $J_w=2$.
- 3D orthogonal B-spline wavelet transform of degree 1.0, number of decomposition levels $J_w=1$ and $J_w=2$. WSPM uses simple Bonferroni correction to deal with the multiple testing problem.

We present results by projecting the detected clusters on the T1-weighted SPM2 template brain in the radiological convention. The statistical significance is fixed at 5% corrected for multiple hypothesis testing. The different results are shown in Figs. 6–11 and will be discussed later. As a reference, we also include the result of a voxel-by-voxel t -test on the unsmoothed data.

Reproducibility study

To investigate the performance of SPM and WSPM, we conducted a reproducibility analysis. While the true status (truly active or truly inactive) of each voxel is unknown, we can estimate

the proportion of truly active voxels that are classified correctly based on multi-session data. These estimates can be interpreted as the empirical sensitivity and specificity of the method, which provides us with a way to establish the receiver-operating-characteristics (ROC) curve. Several studies have recommended the ROC approach for validating and comparing statistical methods, e.g., Friston et al. (1996b), Genovese et al. (1997), Skudlarski et al. (1999), Maitra et al. (2002), Liou et al. (2003), Fadili and Bullmore (2004), and Liou et al. (2006).

The main idea is to exploit the consistency in the detections over multiple sessions. To that end, we consider the cumulative sum of the individual binary activation maps, which takes values between 0 and K , where K is the number of sessions. The normalized histogram of the cumulative map is denoted by $q[k]$, $k=0, \dots, K$, and considered to be a realization of a process that follows the binomial mixture law:

$$\binom{K}{k} [\lambda p_A^k (1 - p_A)^{K-k} + (1 - \lambda) p_I^k (1 - p_I)^{K-k}], \quad (26)$$

where p_A and p_I denote the sensitivity and false positive rate, respectively, and λ the mixture parameter that represents the proportion of truly active voxels.

These parameters are estimated from the data using an EM-algorithm (Fadili and Bullmore, 2004). Since the mixture




Fig. 6. Activation maps for the contrast “all frequencies – rest” obtained by SPM at various smoothing, 5% (FWE correction).

parameter is assumed to be constant for the given data set, it is estimated globally. The estimated sensitivity p_A and specificity $1-p_I$ are then used to specify an empirical ROC curve by letting the significance level vary (in our case, from $\alpha=1$ to $\alpha=10^{-11}$ corrected for multiple testing).

Next to this reproducibility analysis, we also measure a bias term that corresponds to the difference between the denoised

contrast image and the best linear unbiased estimate $u[\mathbf{n}]$ (i.e., the voxel-based estimate of the unsmoothed data). For the SPM framework, the parameter map is denoised by the smoothing operation. We define the empirical relative bias as

$$\varepsilon_{\text{SPM}} = \frac{\sum_{\mathbf{n} \in \mathcal{A}} |u_{\text{SPM}}[\mathbf{n}] - u[\mathbf{n}]|}{\sum_{\mathbf{n} \in \mathcal{A}} |u[\mathbf{n}]|}, \quad (27)$$

where $u[\mathbf{n}]$ is the unsmoothed voxel-wise estimate and \mathcal{A} denotes the set of active-declared voxels. For WSPM, the parameter map is denoised by the non-linear operation of thresholding the wavelet coefficients. The corresponding empirical relative bias is

$$\varepsilon_{\text{WSPM}} = \frac{\sum_{\mathbf{n} \in \mathcal{A}} |\hat{u}[\mathbf{n}] - u[\mathbf{n}]|}{\sum_{\mathbf{n} \in \mathcal{A}} |u[\mathbf{n}]|}, \quad (28)$$

where $\hat{u}[\mathbf{n}]$ is given by Eq. (17).

In Fig. 12, we show the tri-variate plot of the estimated specificity, sensitivity, and bias. We also included the voxel-by-voxel spatial t -test that has no bias, but suffers from a very low sensitivity while its specificity is well controlled.

Discussion

Based on the results for the multi-session fMRI data set, we discuss the proposed extensions to WSPM’s framework, and make the comparison with SPM.

We first investigate the effect of the bias reduction and the shift-invariant extension of WSPM for the 2D slice-by-slice wavelet transform (orthogonal B-spline wavelet of degree 1.0, $J_w=2$ decomposition levels). In Fig. 8, we see that the bias reduction effectively makes the detected activation patterns closer to typical activations within the grey matter. As this correction is conservative, the number of detections can only potentially decrease, but there is no apparent sensitivity reduction at the cluster-level. In Fig. 9, we show the activation maps for the data analyzed using various shifts of the original data in the axial plane. The influence of such a shift is clearly not negligible and leads to different detections and shapes. Using the shift-invariant extension of the framework (Fig. 11, top), we see that the results for the various shifts are properly combined. While the redundancy by a factor of 4 slightly increases the threshold values, there is no noticeable loss in sensitivity. From now on, we consider only




Fig. 7. Activation maps for the contrast “all frequencies – rest” obtained by voxel-wise t -test, 5% (Bonferroni corrected).




Fig. 8. Activation maps for the contrast “all frequencies – rest” obtained by WSPM, 5% (Bonferroni corrected), for the 2D wavelet transform without redundancy (2 decomposition levels). Left: with bias reduction; right: without bias reduction.

WSPM results using the bias reduction and the shift-invariant extension (4 shifts for 2D, 8 shifts for 3D).

We now take a closer look at the results for SPM as shown in Fig. 6. For this type of auditory stimulation, we expect that

activated regions are mainly located bilaterally in the auditory cortex with more extent in the left hemisphere (e.g., see [Bilecen et al., 1998](#)). Indeed, we observe different activated foci in the auditory cortex and extended anterior-lateral to posterior-mesial

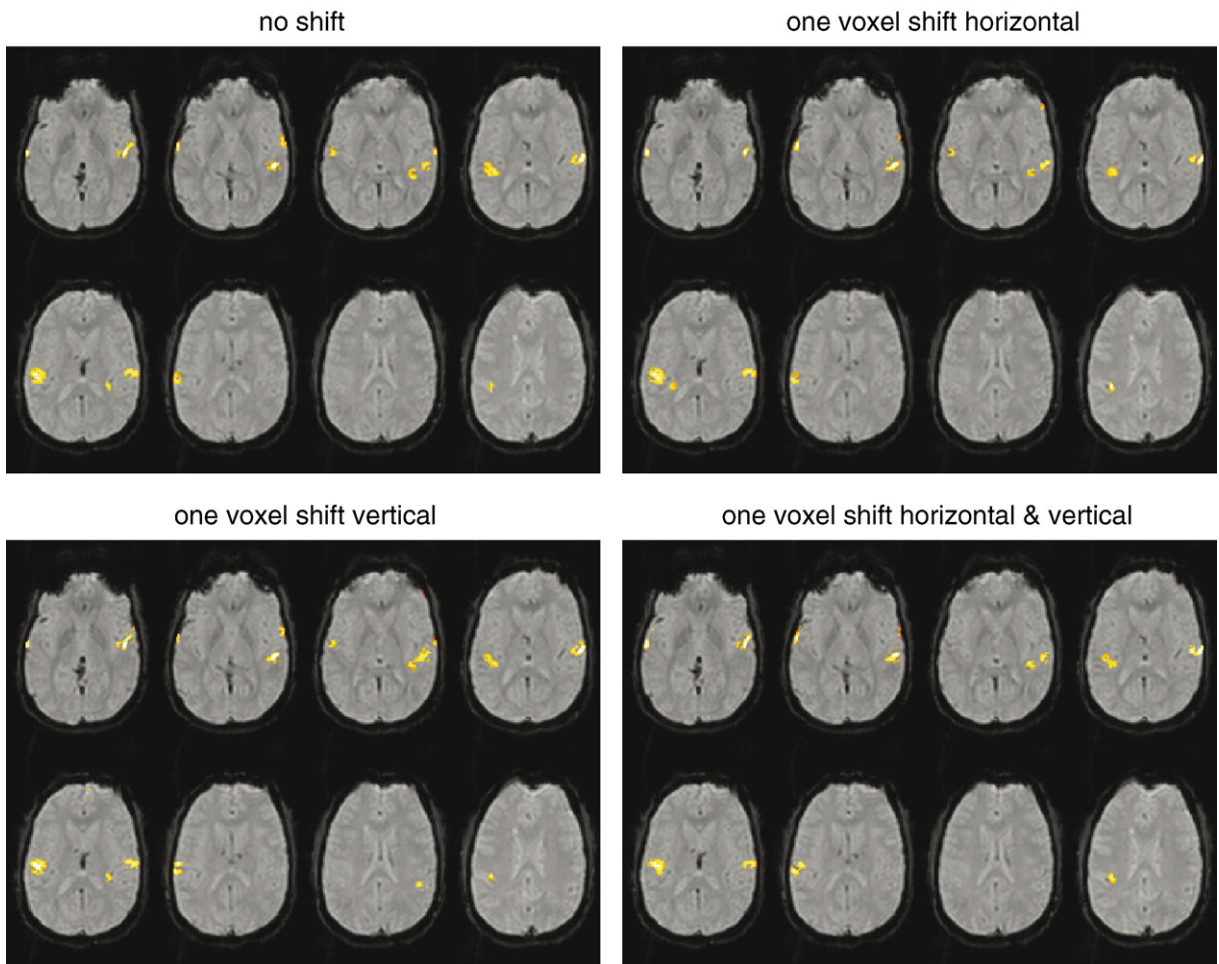


Fig. 9. Activation maps for the contrast “all frequencies – rest” obtained by WSPM, 5% (Bonferroni corrected), for the 2D wavelet transform without redundancy (2 decomposition levels). Shifts are in the axial plane.

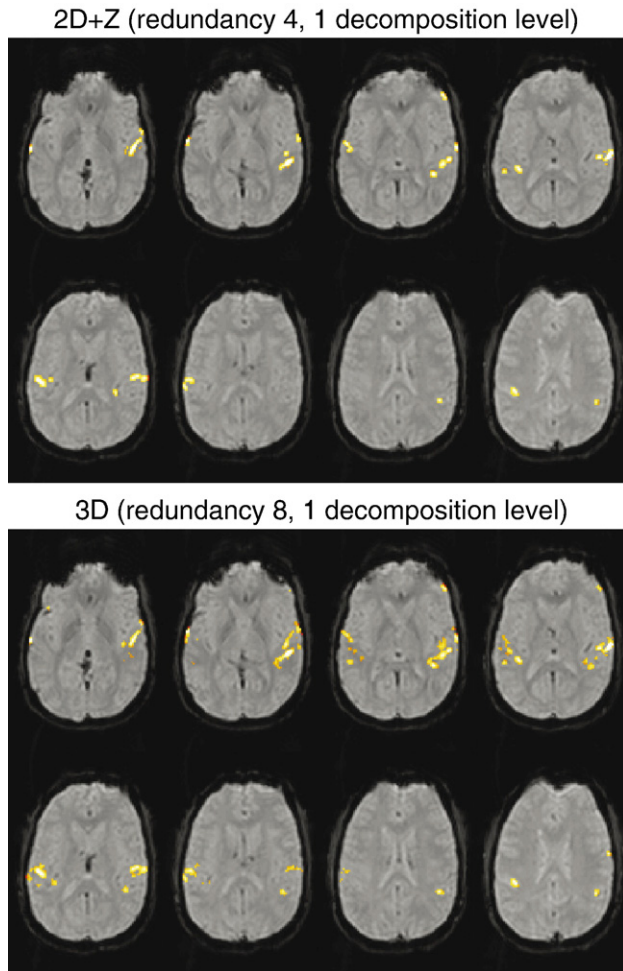


Fig. 10. Activation maps for the contrast “all frequencies – rest” obtained by WSPM, 5% (Bonferroni corrected).

within the Heschl’s gyrus. Increasing the Gaussian smoothing strength results into a grouping of activated clusters and an enlargement of their extent (left auditory cortex in Fig. 6), which is in line with previous assessments of the effects of data smoothing; e.g., Worsley et al. (1996). This increase in the spatial extent of activation is inherently accompanied by a decrease in the spatial definition (i.e., the ability to identify distinct activated foci), which might be relevant when high spatial resolution acquisitions are performed to better delimit subdivisions of the auditory cortex or to visualize tonotopy. Critically, the activated clusters outside the auditory cortex (which are more likely to be false positives) are also increasing with spatial smoothing, as illustrated by frontal or occipital foci in SPM with 8 mm smoothed data (Fig. 6, bottom). Finally, we notice that the sensitivity of SPM has indeed greatly improved with respect to the voxel-wise spatial t -test.

In Figs. 10 and 11, we show the various results for WSPM (2D and 3D wavelet transform, 1 and 2 decomposition levels). In general, functional patterns detected by WSPM are comparable to those by SPM but with some notable differences. Interestingly, significant clusters are localized with high spatial definition, which seems to be fairly independent of the type of wavelet transform. The 3D transform has a higher sensitivity than the 2D one; e.g., a cluster located in the anterior-lateral part of the left auditory cortex is left undetected by the 2D transform. This superior performance

can be explained by inter-slice correlations that are exploited by the additional transform along the Z-direction for the 3D transform. However, the inter-slice correlation is local, and having more than 1 decomposition level starts degrading the result for the 3D transform. Slice timing correction did not influence the results in this case, but it should be considered for higher TR (with non-sequential slice acquisition schemes or event-related paradigms). For data sets with less inter-slice correlation, the use of the 2D transform can be more appropriate since it is computationally faster, and its shift-invariant extension is less redundant.

The ROC curves obtained from the reproducibility study over the multiple sessions are shown in Fig. 12. By definition, the voxel-wise statistical test has no bias, but a very low sensitivity. SPM’s Gaussian smoothing introduces a bias but also significantly increases sensitivity. It is interesting to see that, with WSPM, we obtain a lower bias than SPM 4 mm for all types of wavelet transforms. For very high significance levels, the bias of WSPM approaches the linear unbiased estimate. In addition to a lower bias, we also observe an improved ROC behavior. In particular, the 3D transform with 1 decomposition level yields a higher sensitivity than SPM 4 mm for the same empirical specificity. The only exception is the 3D transform with 2 decomposition levels. Finally, we note that the wiggly curves for WSPM are due to the (non-linear) thresholding operation in the wavelet domain.

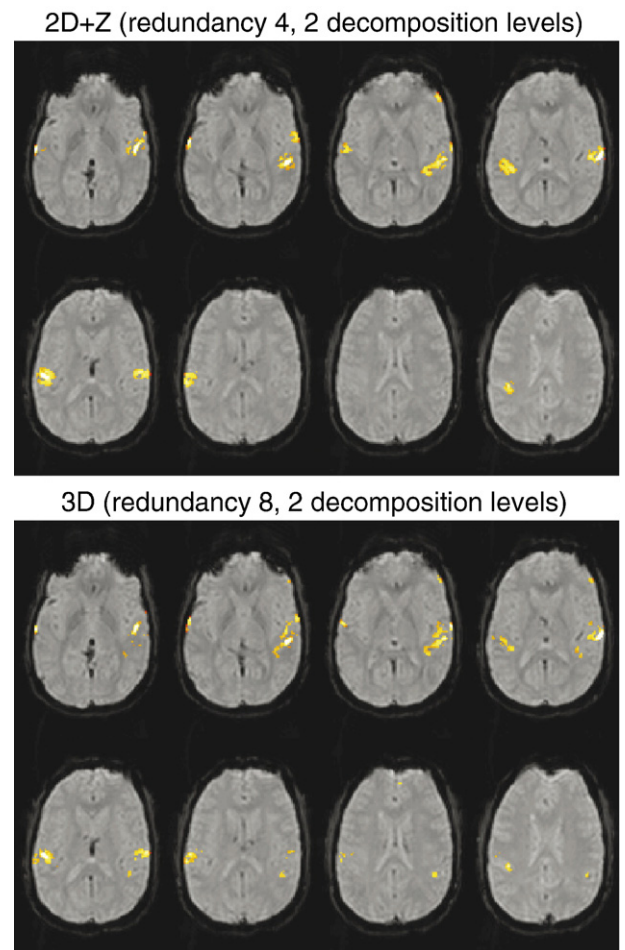


Fig. 11. Activation maps for the contrast “all frequencies – rest” obtained by WSPM, 5% (Bonferroni corrected).




Fig. 12. Empirical ROC curves for the multi-session experiment. (a) Tri-variate plot showing (1-specificity), sensitivity, and fidelity (bias) for the various approaches. (b) View only showing bias and sensitivity. (c) View only showing specificity and sensitivity. (d) Zoom of (c).

These findings suggest that WSPM should be useful for high spatial resolution mapping. As shown above, the false positives are well controlled, while the true positives are precisely localized. The

latter property is a consequence of the spatial adaptivity of the wavelet representation and of the bias reduction procedure that tends to preserve smaller activations. Preserving localization is a

key consideration in many fMRI studies. For example, investigations of the auditory cortex that try to identify cortical subdivisions or visualize tonotopy have analyzed raw fMRI data without spatial smoothing (e.g., Talavage et al., 2000, 2004; Schönwiesner et al., 2002; Formisano et al., 2003; Seghier et al., 2005). Indeed, it is highly desirable to preserve the high spatial resolution provided by the scanner and to avoid integration of signal from different voxels that may have different functional properties. However, better resolution usually comes at the expense of noise and many authors have employed spatial smoothing (e.g., Bilecen et al., 2002; Wessinger et al., 2001) to improve the signal-to-noise ratio and to increase the sensitivity of the statistical analysis. Within this context, WSPM can be seen as an adaptive technique, in which the spatial filtering is performed by wavelets that can preserve edges and can potentially recover high-resolution activation patterns. We have shown that this can be achieved without any compromise in sensitivity. For instance, it is noteworthy that five different foci can be distinguished within the left auditory cortex (Fig. 10, bottom); these may correspond to subdivisions of the primary auditory cortex (Morosan et al., 2001) or to different secondary auditory areas (Rivier and Clarke, 1997; Talavage et al., 2000).

In this study, we opted for the orthogonal B-spline wavelet of degree 1.0 (Battle, 1987; Unser and Blu, 2000). Thanks to orthogonality, the measurement noise remains white in the wavelet domain; therefore, the denoising operation by coefficient-wise thresholding is very efficient. The transform has 2 vanishing moments, which means that constant and linear trends in the data are filtered out, which could compensate for inhomogeneity effects. While the activation maps and ROC curves for the current study show that a single decomposition level yields the better results, we expect a beneficial effect of more decomposition levels for data with higher spatial resolution (leading to relatively larger activation patterns that can be more efficiently represented by larger basis functions). A similar effect has been observed when applying the framework to analyze data from optical imaging, as in Bathellier et al. (2007).

Finally, we recall that WSPM relies on SPM's temporal modeling (GLM and AR-model). One shortcoming of this model is that it may produce overestimates when there is serial correlation in the data that has not been dealt with correctly. One possibility could be to deploy more advanced methods to deal with colored noise; e.g., Bullmore et al. (2001) and Fadili and Bullmore (2002).

Conclusions and outlook

In this paper, we extended WSPM's framework to further improve the results for fixed-effect analysis using the spatial discrete wavelet transform. In particular, we proposed a simple bias reduction method, and an approach to overcome the shift-variance of the wavelet transform.

The beneficial influence of both extensions has been illustrated using synthetic and experimental data, including a comparison against SPM's results for various degrees of smoothing. This evaluation also clarified the trade-off between bias, sensitivity, and specificity that WSPM achieves.

We believe that the use of WSPM is particularly interesting for the analysis of high spatial-resolution fMRI data, such as studies involving sensory (Beauchamp et al., 2004) or visual (Menon et al., 1997; Kim et al., 2000) cortex. Another potential area of application is clinical fMRI. Indeed, the characterization of perilesional activation is highly significant for the assessment of

recovery and plasticity in patients after brain insult (e.g., Breier et al., 2004) and it is of paramount importance to obtain the most accurate spatial mapping.

Acknowledgments

The authors would like to thank Marco Pelizzone from the Centre romand d'implants cochleaires (University Hospital Geneva) for the experimental data. This work was supported in part by the Swiss National Science Foundation under Grant 200020-101821 and Center for Biomedical Imaging (CIBM) of the Geneva – Lausanne Universities and the EPFL, as well as the foundations Leenaards and Louis-Jeantet.

References

- Aston, J., Gunn, R.N., Hinz, R., Turkheimer, F., 2005. Wavelet variance components in image space for spatio-temporal neuroimaging data. *NeuroImage* 25 (1), 159–168.
- Bathellier, B., Van De Ville, D., Blu, T., Unser, M., Carleton, A., 2007. Wavelet-based multi-resolution statistics for optical imaging signals: application to automated detection of odour activated glomeruli in the mouse olfactory bulb. *NeuroImage* 34, 1020–1035.
- Battle, G., 1987. A block spin construction of ondelettes: Part I. Lemarié functions. *Communications in Mathematical Physics* 110, 601–615.
- Beauchamp, M.S., Argall, B., Bodurka, J., Duyn, J., Martin, A., 2004. Unraveling multisensory integration: patchy organization within human STS multisensory cortex. *Nature Neuroscience* 7, 1190–1192.
- Bilecen, D., Scheffler, K., Schmid, N., Tschopp, K., Seelig, J., 1998. Tonotopic organization of the human auditory cortex as detected by BOLD-fMRI. *Hearing Research* 126, 19–27.
- Bilecen, D., S., E., Scheffler, K., Henning, J., Schulte, A.C., 2002. Amplitude of the human auditory cortex: an fMRI study. *NeuroImage* 17, 710–718.
- Breier, J.I., Castillo, E.M., Boake, C., Billingsley, R., Maher, L., Francisco, G., Papanicolaou, A.C., 2004. Spatiotemporal patterns of language-specific brain activity in patients with chronic aphasia after stroke using magnetoencephalography. *NeuroImage* 23, 1308–1316.
- Bullmore, E., Long, C., Suckling, J., Fadili, J., Calvert, G., Zelaya, F., Carpenter, T., Brammer, M., 2001. Colored noise and computational inference in neurophysiological time series analysis: resampling methods in time and wavelet domains. *Human Brain Mapping* 12, 61–78.
- Desco, M., Penedo, M., Gispert, J.D., Vaquero, J.J., Reig, S., Garcia-Barreno, P., 2005. ROC evaluation of statistical wavelet-based analysis of brain activation in [¹⁵O]-H₂O PET scans. *NeuroImage* 24, 763–770.
- Fadili, M.J., Bullmore, E., 2002. Wavelet-generalised least squares: a new BLU estimator of linear regression models with 1/f errors. *NeuroImage* 15, 217–232.
- Fadili, M.J., Bullmore, E.T., 2004. A comparative evaluation of wavelet-based methods for multiple hypothesis testing of brain activation maps. *NeuroImage* 23 (3), 1112–1128.
- Flandin, G., Penny, W.D., 2007. Bayesian fMRI data analysis with sparse spatial basis function priors. *NeuroImage* 34, 1108–1125.
- Formisano, E., Kim, D.S., Di Salle, F., van de Moortele, P.F., Ugurbil, K., Goebel, R., 2003. Mirror-symmetric tonotopic maps in human primary auditory cortex. *Neuron* 40, 859–869.
- Frackowiak, R., Friston, K., Frith, C., Dolan, R., Mazziotta, J., 1997. *Human Brain Function*. Academic Press.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.P., Frith, C.D., Frackowiak, R.S.J., 1995. Statistical parametric maps in functional imaging: a general linear approach. *Human Brain Mapping* 2, 189–210.
- Friston, D.J., Williams, S., Howard, R., Frackowiak, R.S., Turner, R., 1996a. Movement-related effects in fMRI time-series. *Magnetic Resonance in Medicine* 35, 346–355.
- Friston, K.J., Holmes, A., Poline, J.-B., Price, C.J., Frith, C., 1996b.

- Detecting activations in PET and fMRI: levels of inference and power. *NeuroImage* 4, 223–235.
- Friston, K.J., Josephs, O., Zarahn, E., Holmes, A.P., Poline, J.-B., 2000. To smooth or not to smooth? Bias and efficiency in fMRI time series analysis. *NeuroImage* 12, 196–208.
- Friston, K.J., Penny, W., Phillips, C., Kiebel, S., Hinton, G., Ashburner, J., 2002. Classical and Bayesian inference in neuroimaging: theory. *NeuroImage* 16, 465–483.
- Genovese, C.R., Noll, D.C., Eddy, W.F., 1997. Estimating test-retest reliability in functional MR imaging I: statistical methodology. *Magnetic Resonance in Medicine* 38, 497–507.
- Jezzard, P., Matthews, P.M., Smith, S.M., 2001. *Functional MRI an Introduction to Methods*. Oxford University Press.
- Johnstone, T., Ores Walsh, K.S., Greischar, L.L., Alexander, A.L., Fox, A.S., Davidson, R.J., Oakes, T.R., 2006. Motion correction and the use of motion covariates in multiple-subject fMRI analysis. *Human Brain Mapping* 27, 779–788.
- Kim, D.S., Duong, T.Q., Kim, S.G., 2000. High-resolution mapping of isoorientation columns by fMRI. *Nature Neuroscience* 3, 164–169.
- Liou, M., Su, H., Lee, J., Cheng, P., Huang, C., Tsai, C., 2003. Bridging functional MR images and scientific inference: reproducibility maps. *Journal of Cognitive Neuroscience* 15, 935–945.
- Liou, M., Su, H.R., Lee, J., Aston, J., Tsai, A., Cheng, P.E., 2006. A method for generating reproducible evidence in fMRI studies. *NeuroImage* 29, 383–395.
- Maitra, R., Roys, S.R., Gullapalli, R.P., 2002. Test–retest reliability estimation of functional MRI data. *Magnetic Resonance in Medicine* 48, 62–70.
- Mallat, S., 1989. A theory for multiresolution signal decomposition: the wavelet decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 674–693.
- Mallat, S., 1998. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego (CA).
- Menon, R.S., Ogawa, S., Strupp, J.P., Ugurbil, K., 1997. Ocular dominance columns in human V1 demonstrated by functional magnetic imaging. *Journal of Neurophysiology* 77, 2780–2787.
- Morosan, P., Rademacher, J., Schleicher, A., Amunts, K., Schormann, T., Zilles, K., 2001. Human primary auditory cortex: cytoarchitectonic subdivisions and mapping into a spatial reference system. *NeuroImage* 13, 684–701.
- Mueller, K., Lohmann, G., Zysset, S., von Carmon, Y., 2003. Wavelet statistics of functional MRI data and the general linear model. *Journal of Magnetic Resonance Imaging* 17, 20–30.
- Poline, J., Worsley, K., Evans, A., Friston, K., 1997. Combining spatial extent and peak intensity to test for activations in functional imaging. *NeuroImage* 5 (2), 83–96.
- Rivier, F., Clarke, S., 1997. Cytochrome oxidase, acetylcholinesterase, and NADPPH-diaphorase staining in human supratemporal and insular cortex: evidence for multiple auditory areas. *NeuroImage* 6, 288–304.
- Ruttimann, U., Unser, M., Rawlings, R., Rio, D., Ramsey, N., Mattay, V., Hommer, D., Frank, J., Weinberger, D., 1998. Statistical analysis of functional MRI data in the wavelet domain. *IEEE Transactions on Medical Imaging* 17 (2), 142–154.
- Schönwiesner, M., von Cramon, Y.D., RübSamen, R., 2002. Is it tonotopy after all? *NeuroImage* 17, 1144–1161.
- Seghier, M.L., Boëx, C., Lazeyras, F., Sigrist, A., Pelizzone, M., 2005. fMRI evidence for activation of multiple cortical regions in the primary auditory cortex of deaf subjects users of multichannel cochlear implants. *Cerebral Cortex* 15, 40–48.
- Sendur, L., Maxim, V., Whitcher, B., Bullmore, E., Sep. 2005. Multiple hypothesis and mapping of functional MRI data in orthogonal and complex wavelet domains. *IEEE Transactions on Signal Processing* 53 (9), 3413–3426.
- Skudlarski, P., Constable, R.T., Gore, J.C., 1999. ROC analysis of statistical methods used in functional MRI: individual subjects. *NeuroImage* 9, 311–329.
- Srikanth, R., Casanova, R., Laurienti, P.J., Peiffer, A.M., Maldjian, J.A., 2006. Estimation of false discovery rate for wavelet-denoised statistical parametric maps. *NeuroImage* 33, 72–84.
- Talavage, T.M., Ledden, P.J., Benson, R.R., Rosen, B.R., Melcher, J.R., 2000. Frequency-dependent responses exhibited by multiple regions in human auditory cortex. *Hearing Research* 15, 225–244.
- Talavage, T.M., Sereno, M., Melcher, J., Ledden, P.J., Rosen, B.R., Dale, A.M., 2004. Tonotopic organization in human auditory cortex revealed by progressions of frequency sensitivity. *Journal of Neurophysiology* 91, 1282–1296.
- Turkheimer, F.E., Aston, J.A.D., Asselin, M.-C., Hinz, R., 2006. Multi-resolution Bayesian regression in PET dynamic studies using wavelets. *NeuroImage* 32, 111–121.
- Turkheimer, F.E., Brett, M., Aston, J.A.D., Leff, A.P., Sargent, P.A., Wise, R.J., Grasby, P.M., Cunningham, V.J., 2000. Statistical modelling of positron emission tomography images in wavelet space. *Journal of Cerebral Blood Flow and Metabolism* 20, 1610–1618.
- Unser, M., Blu, T., 2000. Fractional splines and wavelets. *SIAM Review* 42, 43–67.
- Van De Ville, D., Blu, T., Unser, M., 2003. Wavelets versus resels in the context of fMRI: establishing the link with SPM. *SPIE's Symposium on Optical Science and Technology: Wavelets X*. Vol. 5207. SPIE, San Diego CA, USA (Aug.).
- Van De Ville, D., Blu, T., Unser, M., 2004. Integrated wavelet processing and spatial statistical testing of fMRI data. *NeuroImage* 23 (4), 1472–1485 (Dec.).
- Wessinger, C.M., VanMeter, J., Tian, B., Van Lare, J., Pekar, J., Rauschecker, J.P., 2001. Hierarchical organization of the human auditory cortex revealed by functional magnetic resonance imaging. *Journal of Cognitive Neuroscience* 13, 1–7.
- Wink, A.M., Roerdink, J.B.T.M., 2004. Denoising functional MR images: a comparison of wavelet denoising and Gaussian smoothing. *IEEE Transactions on Medical Imaging* 23 (3), 374–387 (Jun.).
- Worsley, K.J., Friston, K.J., 1995. Analysis of fMRI time-series revisited—Again. *NeuroImage* 2, 173–181.
- Worsley, K., Marrett, S., Neelin, P., Evans, A., 1996. Searching scale space for activation in PET images. *Human Brain Mapping* 4 (1), 74–90.
- Worsley, K., Liao, C., Aston, J., Petre, V., Duncan, G., Evans, A., 2002. A general statistical analysis for fMRI data. *NeuroImage* 15, 1–15.