Optimizing Wavelets for the Analysis of fMRI data

Manuela Feilner, Thierry Blu and Michael Unser

Biomedical Imaging Group, DMT/IOA
Swiss Federal Institute of Technology Lausanne
CH-1015 Lausanne EPFL
Switzerland
e-mail: manuela.feilner@epfl.ch

ABSTRACT

Ruttimann *et al.* have proposed to use the wavelet transform for the detection and localization of activation patterns in functional magnetic resonance imaging (fMRI). Their main idea was to apply a statistical test in the wavelet domain to detect the coefficients that are significantly different from zero. Here, we improve the original method in the case of non-stationary Gaussian noise by replacing the original z-test by a t-test that takes into account the variability of each wavelet coefficient separately. The application of a threshold that is proportional to the residual noise level, after the reconstruction by an inverse wavelet transform, further improves the localization of the activation pattern in the spatial domain.

A key issue is to find out which wavelet and which type of decomposition is best suited for the detection of a given activation pattern. In particular, we want to investigate the applicability of alternative wavelet bases that are not necessarily orthogonal. For this purpose, we consider the various brands of fractional spline wavelets (orthonormal, B-spline, and dual) which are indexed by a continuously-varying order parameter α. We perform an extensive series of tests using simulated data and compare the various transforms based on their false detection rate (type I + type II errors). In each case, we observe that there is a strongly optimal value of α and a best number of scales that minimizes the error. We also find that splines generally outperform Daubechies wavelets and that they are quite competitive with SPM (the standard analysis method used in the field), although it uses much simpler statistics. An interesting practical finding is that performance is strongly correlated with the number of coefficients detected in the wavelet domain, at least in the orthonormal and B-spline cases. This suggest that it is possible to optimize the structural wavelet parameters simply by maximizing the number of wavelet counts, without any prior knowledge of the activation pattern. Some examples of analysis of real data are also presented.

Keywords: functional imaging, fMRI, statistical analysis, medical imaging, wavelets, fractional splines

1. INTRODUCTION

fMRI (functional Magnetic Resonance Imaging) is a new and powerful technique for observing the activity of the brain while a subject is performing some cognitive or motor task.[1] The neuronal activity is detected indirectly by a blood oxygen-level dependent (BOLD) contrast signal.[2,3] The basic mechanism responsible for the signal is the local reduction in deoxyhemoglobin which is caused by an increase in cerebral blood flow (CBF). The main advantages of fMRI over other functional imaging techniques such as PET (positron emission tomography) are threefold: 1) fMRI is non-invasive and does not require injection of radioactive isotopes; 2) its time resolution is much better (a few seconds) as compared to PET (several minutes); and 3) the in-plane resolution of the functional images is excellent (better than 1 mm).

Presently, fMRI is primarily used as a research tool in the neurosciences for localizing brain function and for identifying specific neural pathways. It is also gaining importance in medical applications, particularly in the context of neurosurgical planning. For example, it can be used for improving the localization of epileptic seizures, for understanding their progression and for investigating their interaction with various functional areas of the brain.

The analysis of fMRI data is difficult because the signal changes due to activation are small and hidden in large amounts of noise. The noise in fMRI-data has different origins, some of them being

- Subject motion during image acquisition is a major source of artifacts,[4] even considering the fact that head restraints and/or image registration techniques can mitigate this problem somehow[5];

- Fluctuations and vibrations due to the scanner introduce additional motion artifacts[6];
- Physiological effects due to respiration, to cardio-vascular functions, and to the overall background activity of the brain interact with the signal of interest.[7,8]

The use of the wavelet transform for the analysis of functional images of the brain was pioneered by Ruttimann *et al.*, first for PET,[9] and then for fMRI.[10] The advantages of the method over the more conventional spatial approaches are threefold. First, the wavelet transform has a compression effect because it tends to pack the signal of interest into a relatively few number of large coefficients. Since the noise gets diluted evenly in the wavelet domain (provided the wavelet is close to orthogonal), the signal-to-noise ratio is improved for those coefficients where the signal is concentrated. Second, this increase in the signal-to-noise ratio improves the detection rate. Third, the wavelet transform is a one-to-one representation without any redundancies. Thus, because of the decoupling effect in the wavelet domain, it is possible to apply a conservative decision strategy to keep the false-detection rate well under check (Bonferroni correction for multiple testing). This also means that the procedure remains valid—albeit sub-optimal—even when the basic hypothesis of noise independence breaks down.

We apply the wavelet transform to the difference images $I_{\mathrm{D}_\tau}(x,y) = I_{A_\tau}(x,y) - I_{B_\tau}(x,y)$, where the symbols A and B correspond to the two distinct experimental conditions (block paradigm), and where τ is the time index (sequence of realizations). Because of the noise present in the data, it is almost impossible to detect any activation without resorting to statistical analysis. This analysis requires a model of noise[11]; here, we assume that $I_{\mathrm{D}_\tau}(x,y)$ follows a spatio-temporal Gaussian distribution. Thanks to the wavelet transform, we can apply simple statistics, which makes the analysis fast and straightforward.

Ruttimann *et al.*[10] explored the orthonormal Battle-Lemarié wavelets of integer degree; they reported best results using cubic splines. Here, we consider a much larger class of wavelets by exploring the various brands—causal (+) vs. symmetric (∗); orthonormal (`ortho`) vs. B-spline (`bspline`) vs. dual (`dual`)—of the newly-designed fractional spline wavelet transforms (FSWT).[12] Those are indexed by a *continuously-varying* degree parameter α. In principle, this offers the potential for a much finer tuning of the transform for a given detection task, but also raises the delicate issue of how to select the best wavelet parameters in an objective way.

The method of Ruttimann *et al.* was restricted to stationary noise. We propose here to relax this constraint by replacing their z-score calculation in the wavelet domain by a series of t-tests.

The main goal of this paper is to address the important issue of the selection of the "best" wavelet transform (within an infinite—continuous—family of candidates) for the analysis of fMRI data. Our strategy is experimental: to find out the influence of the type of wavelet transform, we perform numerous test experiments using known activation patterns. This requires the definition of an objective performance measure (cf. Section 3.3) that can be used to rank the various transforms or to optimize the structure parameters (i.e., the type of transform, the fractional order α, and the depth of the transform: J=number of scales).

We will also compare our wavelet techniques with SPM* (Statistical Parameter Mapping[13]) which is the reference method used in the field. The SPM package encompasses a broad range of cases and offers many options to tune; nevertheless, its basic working principle is to apply a spatial Gaussian filter to the data before performing statistical testing. This entails a loss of image resolution and also complicates the statistical analysis because the noise can no longer be assumed to be independent.[14,15]

The paper is organized as follows: In Section 2, we present our general procedure. We also describe our new class of fractional spline wavelet transforms and discuss their main properties. The general detection and localization procedures are considered in Sections 2.2 and 2.3, respectively. Our experimental setting is described in Section 3. This includes the description of the fMRI and test data as well as the definition of quality measures for the comparison of wavelet transforms. Experimental data and interpretation of the results are presented in Section 4.

2. DETECTION OF ACTIVATION IN THE WAVELET DOMAIN

The overview of the whole system is shown in Figure 1. The input is a series of time measurements (slices or volumes) falling within the A (activation) or B (rest) states (standard block paradigm); these image sequences are denoted by $I_{A_\tau}(x,y)$ and $I_{B_\tau}(x,y)))$, respectively, where τ is a time index. During the preprocessing step, the data are aligned to

*the latest version, SPM'99 was released in January 2000

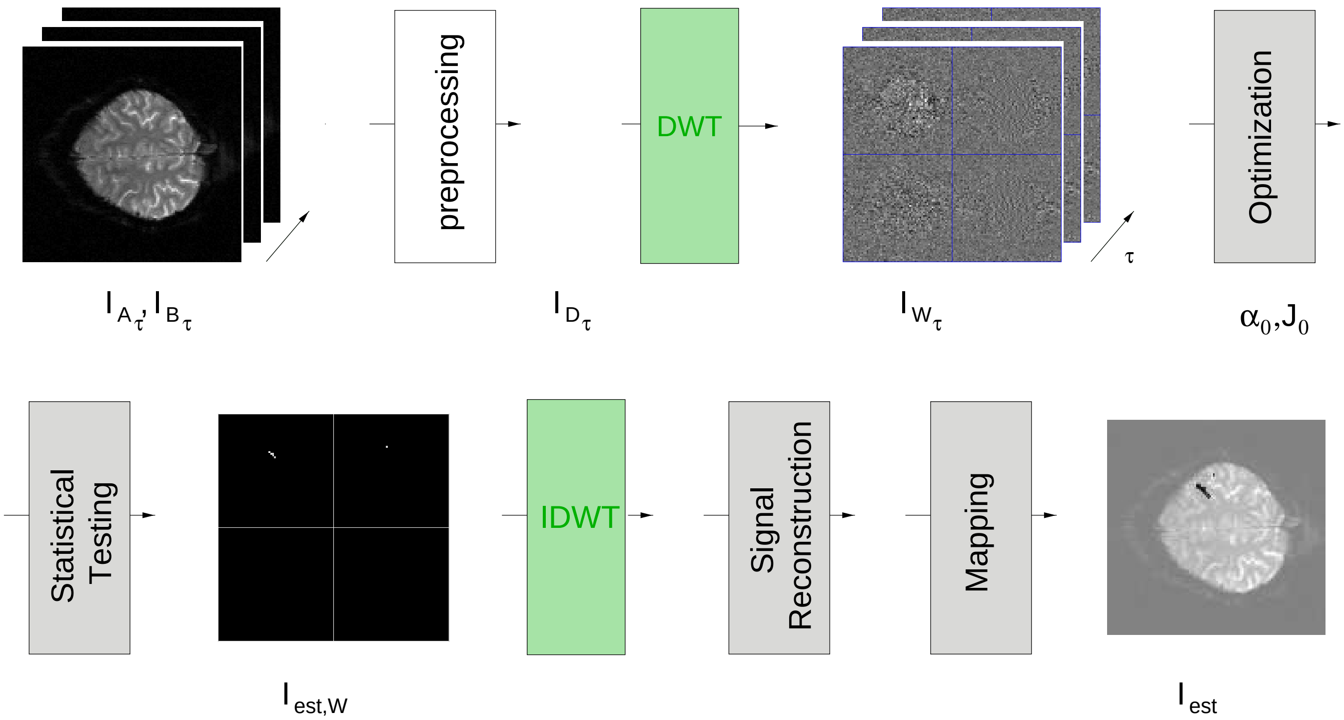

Figure 1. Block diagram of the algorithm

a reference volume using a 3D rigid-body transformation with 6 parameters.[16,17] Note that the motion correction is usually small with respect to the spatial resolution. Drifts in time are also removed using the standard procedure in SPM99. Then, we compute pairwise difference between activation and rest states to produce the sequence difference images $I_{\mathrm{D}_\tau}(x,y) = I_{A_\tau}(x,y) - I_{B_\tau}(x,y)$.

The next step is to apply the Discrete Wavelet Transform (DWT) to $I_{\mathrm{D}_\tau}(x,y)$ to get the decomposed fMRI-data $I_{\mathrm{W}_\tau}(x,y)$, where "W" stands for "wavelet domain", as described below. Note that, due to the linearity of the transform, the Gaussian hypothesis for the noise is preserved in the wavelet domain.

Then, we apply a classical statistical testing on each coefficient of $I_{\mathrm{W}_\tau}(x,y)$ to determine whether it is activated or not. We use a Student t-test as described in Subsection 2.2. At this stage, we zero out all non-significant wavelet coefficients, keeping only those that have been labeled as "activated". This yields the W-estimated activation $I_{\mathrm{est,W}}(x,y)$. Next, we apply the inverse DWT to $I_{\mathrm{est,W}}(x,y)$. Finally, to improve the localization of the activation pattern in the spatial domain, we apply a threshold proportional to the residual noise, which reduces ringing artifacts due to the oscillating nature of the wavelet basis functions. The result is the estimated activation pattern $I_{\mathrm{est}}(x,y)$.

The advantage of a wavelet transform for analyzing fMRI data is that, not only the low-pass contents is processed for the detection of the activation, but also the band-pass components which contribute to increasing the resolution of the detected regions.

In the present study, we are interested in comparing various wavelet transforms. We perform the wavelet analysis repeatedly using different transforms parameters. Then, we evaluate the performance according to our quality criterion.

2.1. Discrete Wavelet Transform

We propose to transform the data using a discrete wavelet transform implemented through an iterated filterbank,[18] whose analysis and synthesis parts are depicted in Figure 2. Even though, in full generality, the synthesis and analysis filters are biorthonormal, they are mostly chosen orthonormal in the literature, i.e., $\tilde{H}(z) = H(z^{-1})$ and $\tilde{G}(z) = G(z^{-1})$; this is in particular the case of Daubechies filters.[19] For such filters that are automatically half-band, there remains only a tradeoff between the sharpness of the frequency cutoff and spatial localization.

Because we want to optimize our detection scheme with respect to orthonormal and biorthonormal transforms, we use a new class of scaling functions and wavelets that provide this flexibility: the fractional splines.[12]

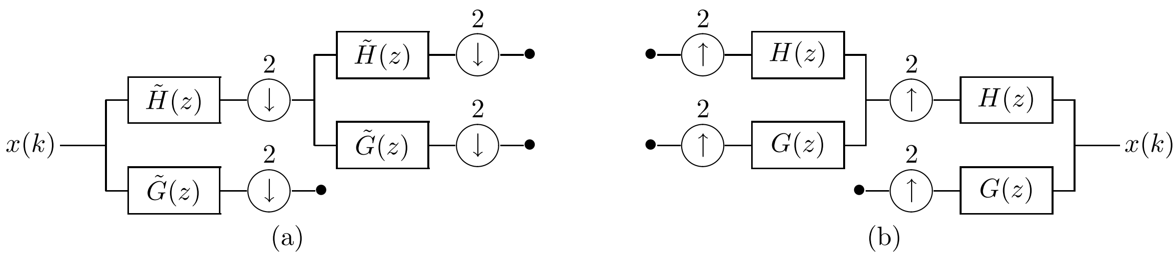

Figure 2. a) Analysis part of the 1D wavelet transform (DWT), where the input $x(k)$ is one line or one row of I_{D_τ}. b) Synthesis part of the 1D wavelet transform (IDWT), applied to $I_{\mathrm{est,W}}$.

2.1.1. The FSWT

We describe here the discrete wavelet transform that arises from a multiresolution analysis generated by fractional splines. There are two flavors of fractional splines: the causal ones, denoted by β_+^α, which are a linear combination of the one-sided power functions $(x-n)_+^\alpha$; and the symmetric ones, denoted by β_*^α, which are a linear combination of the radial power functions $|x-n|^\alpha$ if α is any non-even real number, and of $|x-n|^\alpha \log|x-n|$ if α is even. For each version, the degree α of the splines is assumed to be strictly greater than $-\frac{1}{2}$, in order to ensure a valid multiresolution analysis.[20]

These functions have a simple expression in the frequency domain: $\hat{\beta}_+^\alpha(\omega) = \left(\frac{1-e^{-j\omega}}{j\omega}\right)^{\alpha+1}$ and $\hat{\beta}_*^\alpha(\omega) = |\hat{\beta}_+^\alpha(\omega)|$. They satisfy two-scale difference equations whose scaling filters are given, respectively, by $B_+^\alpha(z) = \sqrt{2}\left(\frac{1+z^{-1}}{2}\right)^{\alpha+1}$ for the causal spline, and by $B_*^\alpha(z) = \sqrt{2}\left|\frac{1+z^{-1}}{2}\right|^{\alpha+1}$ for the symmetric one.

The autocorrelation filter of a B-spline of degree α is given by $A^\alpha(e^{j\omega}) = \sum_n \left|\hat{\beta}_*^\alpha(\omega+2n\pi)\right|^2$. Letting $B^\alpha = B_+^\alpha$ or B_*^α, the discrete FSWT is implemented using the analysis and synthesis filters given in Table 1.

filter \ type	`bspline`	`ortho`	`dual`
$H(z)$	$B^\alpha(z)$	$B^\alpha(z)\sqrt{\dfrac{A^\alpha(z)}{A^\alpha(z^2)}}$	$B^\alpha(z^{-1})\dfrac{A^\alpha(z)}{A^\alpha(z^2)}$
$G(z)$	$-z^{-1}B^\alpha(z^{-1})A^\alpha(-z)$	$-z^{-1}B^\alpha(-z^{-1})\sqrt{\dfrac{A^\alpha(-z)}{A^\alpha(z^2)}}$	$\dfrac{-zB^\alpha(-z)}{A^\alpha(z^2)}$
$\widetilde{H}(z)$	$B^\alpha(z^{-1})\dfrac{A^\alpha(z)}{A^\alpha(z^2)}$	$B^\alpha(z^{-1})\sqrt{\dfrac{A^\alpha(z)}{A^\alpha(z^2)}}$	$B^\alpha(z)$
$\widetilde{G}(z)$	$\dfrac{-zB^\alpha(-z)}{A^\alpha(z^2)}$	$-zB^\alpha(-z)\sqrt{\dfrac{A^\alpha(-z)}{A^\alpha(z^2)}}$	$-z^{-1}B^\alpha(-z^{-1})A^\alpha(-z)$

Table 1. Analysis and synthesis filters for the wavelet transform.

Note that the `dual` type is obtained by exchanging the analysis- and synthesis-part of the `bspline` case. We have worked out a fast and exact implementation of the FSWT[21]: the MATLAB files are freely available on our web at http://bigwww.epfl.ch/blu/fractsplinewavelets/.

2.1.2. Directions for use of the FSWT

We have the choice between three types of analysis-synthesis systems, whose filters are exemplified in Figure 3:

- the `dual` type, characterized on the analysis side by a strongly selective scaling filter, which ensures a higher pixelwise signal-to-noise ratio in the wavelet domain, and thus fewer missed detections: the filter is all the more selective as α is larger. However, on the reconstruction side, the badly localized scaling filter (which even worsens when α increases) is responsible for a higher number of false detections. Because it is more robust to noise, we expect this type to be the best performer when the noise is high.

- the `bspline` type, characterized by a strongly localized synthesis scaling filter, which ensures fewer false detections: this behavior is enhanced as α increases. However, on the analysis side, the badly selective scaling filter (the overshooting indicated by an "$*$" in Figure 3 gets even more prominent as α increases) is responsible for a bad signal-to-noise ratio, and thus misses more detections. Because it has a better localization property and because it is more sensitive to noise, we expect this type to be the best performer when the noise is low.

- the `ortho` type, characterized by a similar space-frequency behavior on the analysis and synthesis sides. By adjusting α, we can control the frequency selectivity and the spatial localization of the filter as summarized below:

$$\alpha \left\langle \begin{array}{l} \text{large } \alpha \left\langle \begin{array}{l} \text{good frequency selectivity} \\ \text{bad spatial localization} \end{array} \right. \\ \text{small } \alpha \left\langle \begin{array}{l} \text{bad frequency selectivity} \\ \text{good spatial localization.} \end{array} \right. \end{array} \right.$$

  We expect the performance of the `ortho` type of transform to be optimal for intermediary noise levels.

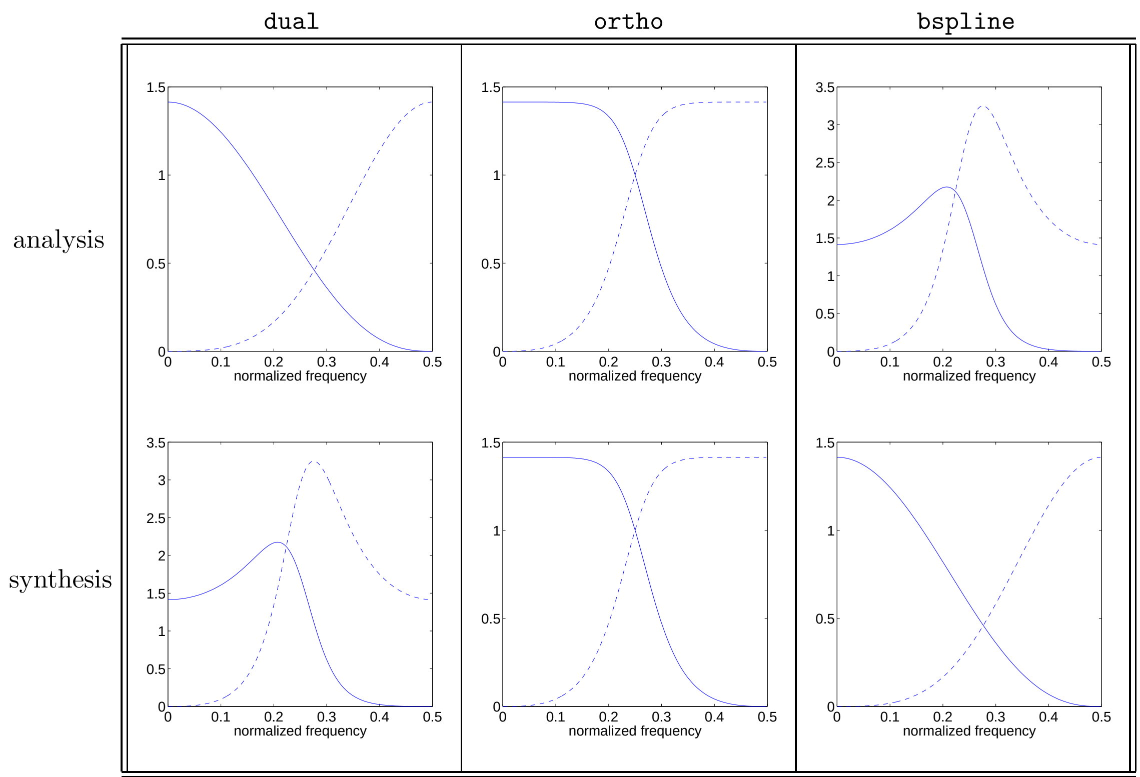

Figure 3. Example of frequency response of the low-pass (plain line) and high-pass (dashed line) filters for the three types of FSWT.

In addition to the type and degree parameter, the FSWT can also be optimized for the iteration depth J: when the noise gets higher, a larger value of J will yield a better pixelwise signal-to-noise ratio and thus more detections of wavelet coefficients. However, the coarser space reconstruction will also yield more falsely detected pixels.

Far less clear is the benefit that the detection process could derive from the symmetric/causal flexibility. However, the fact that the causal splines seem to have a better spatial localization (the integer degree causal splines are even compactly supported) than their symmetric counterparts (which coincid with the regular integer splines only for odd degrees) might explain why, in most of our experiments, the causal ones performed better.

2.2. Activation Detection by Statistical Testing on the Wavelet Coefficients

The statistical analysis is performed at the wavelet coefficient level. We denote by I_{W_τ} the DWT of I_{D_τ} at time (or realization) τ. To decide whether a coefficient is activated or not, we test the activation hypothesis with a predefined confidence level. Under the null hypothesis—the pixel under consideration is not activated—we have that $\mathrm{E}\{I_W(x,y)\} = 0$, where $I_W(x,y) = \frac{1}{n}\sum_\tau I_{W_\tau}(x,y)$ (n is the total number of realizations). If $I_W(x,y)$ exceeds some expected bound which is computed for each (x,y), the null hypothesis is rejected and the alternative (i.e., $\mathrm{E}\{I_W(x,y)\} \neq 0$) is accepted: we conclude that the wavelet pixel is activated.

The decision is made by setting a threshold T chosen as follows: assuming that the pixel is not activated, the probability that some normalized version of $I_W(x,y)$—of known probability density—exceeds T is equal to $1-p$, where the confidence level, p, is usually chosen to be 95%.

It had initially be assumed in the literature[10] that the variance of the noise that corrupt the data I_τ is white and Gaussian, which allowed to apply a z-test to the orthonormal wavelet coefficients. However, we showed that the noise is significantly non-stationary.[11] Moreover, because we will be using a non-orthonormal wavelet transform, we have to rule out definitely the simpler z-test. Instead, we consider a t-test and, accordingly, we normalize I_W by

$$t(x,y) = \sqrt{n}\,\frac{I_W(x,y)}{s_W(x,y)}, \tag{1}$$

where $s_W(x,y)^2 = \frac{1}{n-1}\sum_\tau \left|I_{W_\tau(\mathrm{x,y})} - I_W(x,y)\right|^2$. These t-values are known to follow Student's t-distribution.

We also have to correct for the fact that we will have to make as many independent tests as there are pixels in I_W: indeed, the confidence level decreases after each additional test. This is why if we want to have a confidence level of p for an N-pixel test, we will have to set the pixelwise confidence level substantially higher at $1-\frac{1-p}{N}$. This is Bonferroni-correction which is known to be all the more conservative as the pixels are more correlated.

2.3. Signal Reconstruction

Once the wavelet coefficients of I_W have been flagged either as significantly activated or not, the estimation of the activation pattern is achieved by zeroing out all non significant samples and performing an inverse wavelet transform on I_W.

Since the fractional spline synthesis filters are in general of infinite length, we have to make a decision as to what constitutes a valid activation pattern. Specifically, we apply a threshold to the reconstructed activation pattern, that takes into account the amount of noise: when this noise is very low, we simply choose the quantization level of the initial data; if the noise in every realization of I_D has a non-negligible standard deviation σ, we set the threshold level at $\sigma/\sqrt{n}$, i.e., the noise level of the average image $I_D = \frac{1}{n}\sum_\tau I_{D_\tau}$.

3. OPTIMIZATION

In order to find the structure parameters of the FSWT—its degree α, iteration depth J, and type (causal or symmetric, `dual`, `ortho` or `bspline`)—that are are optimal for a given set of data, we must first define an error measure (see Section 3.3) between the true activation and that which is detected with our method. We thus need synthetic data (see Section 3.2) for which we know the activation a priori. We then have to work out another criterion based on the measured data only (see Section 3.4), that does not take this knowledge into account, and that still can predict the a priori error measure.

3.1. fMRI data

In this study, we investigate simple fMRI block paradigms. Our set of fMRI data consists of 8 repetitions of alternating blocks, according to distinct experimental conditions: brain volume with activation (here, a specific motor task), and brain volume in rest state. Each block contains 4 acquisitions. A full volume of 30 slices of 128×128 pixels each is typically acquired within 6 seconds. If we consider only one slice of this volume, we thus have $n = 32$ realizations consisting of $N = 128 \times 128$ images I_{D_τ} each. These data are corrupted with high noise in particular due to the random activity of the brain, and to the alignment mismatch between different realizations, as explained in Introduction. It is usual to model this noise by a joint Gaussian statistics. Since any volontary movement of the body is mapped contralaterally into the motor cortex, we expect the activation zone to appear in the upper part of the right hemisphere of the brain. We can control the size of the estimated activation pattern, knowing that it shouldn't overlap the visible sulci of each side of the motor cortex.

3.2. Test data

To optimize for the structural wavelet parameters, we used various test data $I(x, y)$, with known activated area A, defined as $\{(x, y)|I(x, y) \neq 0\}$. The complementary region A^c is defined by $\{(x, y)|I(x, y) = 0\}$. We used thresholded ellipse-shaped Gaussian activation patterns. The test-data were generated by adding $n = 20$ realizations of white Gaussian noise: even though this is a further simplification of the fMRI noise model, we believe that this restriction is not likely to alter our final conclusions. In particular, we changed the scale and ellipticity of the pattern, its orientation and position, and the variance of the noise. But we did not restrict ourselves to elliptic shapes, we also tested activations with more high frequency contents.

3.3. Error measures with prior knowledge of the activation pattern

We considered two kinds of errors:

- the number of false detections, i.e., errors of type I, relative to the number of activated pixels

$$E1 = \frac{\mathrm{Card}\big(\{(x, y) : I(x, y) = 0 \text{ and } I_{\mathrm{est}}(x, y) \neq 0\}\big)}{\mathrm{Card}(A)}$$

- the number of missed detections, i.e., errors of type II, relative to the number of activated pixels

$$E2 = \frac{\mathrm{Card}\big(\{(x, y) : I(x, y) \neq 0 \text{ and } I_{\mathrm{est}}(x, y) = 0\}\big)}{\mathrm{Card}(A)}$$

Our global error criterion to optimize is, by definition, the addition of the two types of errors: $E = E1 + E2$. Note that even though this quantity provides a fair control on the localization of the activation zone, it cannot give clues on the amplitude mismatch. Obviously, we need to know the activation pattern to compute E.

3.4. Other measures with unknown activation pattern

We tested three different quantities which, unlike the error measure E defined above, are applied directly on the noisy data: for them, the "true" activation pattern has not to be known a priori.

- Wavelet count: the number of wavelet coefficients that are detected as activated in our statistical procedure. It is thus defined by $\mathrm{C} = \mathrm{Card}\big(\mathrm{I}_{\mathrm{est,W}} \neq 0\big)$.
- Information measure: the average information content of the detected wavelet coefficients, defined by

$$I = -\frac{1}{\mathrm{Card}\big(I_{\mathrm{est,W}}\big)} \sum_{(x,y)\in I_{\mathrm{est,W}}} \log_2 p\big(t(x, y)\big),$$

  where $t(x, y)$ is defined by (1) and $p(t)$ is Student's distribution density.

- Kullback-Leibler information: the map $t(x, y)$ follows the experimental density (histogram) $p_{\mathrm{data}}(t)$, while the statistical model—in the absence of activation—is described by Student's t-distribution, $p(t)$. We measure the similarity between these two densities with Kullback-Leibler information

$$I(p_{\mathrm{data}}, p) = \int_{-\infty}^{\infty} \log_2 \left\{\frac{p_{\mathrm{data}}(t)}{p(t)}\right\} p_{\mathrm{data}}(t)\, dt \qquad (2)$$

  Even though it is not a true distance between probability densities, the Kullback-Leibler information has the important property that $I(p_1, p_2)$ is always strictly positive, except when $p_1 = p_2$, in which case it vanishes.

4. RESULTS

We have thus tested our algorithm on synthetized data for various activation shapes, orientations and noise levels. On each set of realizations, we have optimized for the structure parameters of the FSWT. We also used Daubechies scaling filters (of length between 2 and 16) and finally compared the performance of our algorithm with that of SPM.

We have then verified the good behavior of our algorithm on real fMRI data, the result of which we once again compared with that provided by the SPM package.

4.1. Results: synthetic data

As far as test data are considered, we found that the error criterion E always had a sharp minimum for a particular configuration of the structure parameters: the gain can exceed 40% between the best and the worse configuration, which shows that this optimization is worth the pain. Moreover, in our experiments, the best FSWT always outperformed the wavelet transform based on Daubechies scaling functions: this might be caused by their mediocre space/frequency characteristics, as compared to the orthonormal fractional splines. More specifically, we have found the following trends:

→ **Symmetry.** In most cases, the causal FSWT has fewer errors than its symmetric counterpart. This may be due to a better space localization for the causal splines.

→ **Type.** The `ortho` and `dual` splines in general outperform the `bspline` type FSWT, due to more missed detections (see Figure 4). As a matter of fact, we always tested our algorithm with large noise levels, which are typically observed in practical data. But we surmise that with lower levels, the `bspline` FSWT might be more competitive.

→ **Degree.** The optimal spline degree varies widely in the `ortho` type. However, for the `dual` FSWT, the optimal α is always close to 1.

→ **Iteration depth.** Less surprisingly, when the noise is higher, we need more iterations to detect activated pixels in the wavelet domain. Indeed, we observe that, for lower noise levels the optimal iteration depth is invariably 1, whereas when the variance of the noise gets larger, this optimum is 2. This effect is counterbalanced by the size of the activation: a smaller size can hardly be detected with more than one iteration. For this reason, we never encountered values of J_{opt} exceeding 2.

→ **Correlation *E*–*C*.** For the `bspline` and `ortho` FSWT, the error criterion E is always tightly (inverse) correlated with the number of detected wavelet coefficients, i.e., the wavelet count C. This is extremely useful because, unlike the false+missed detection error count E, the wavelet count does not require any prior knowledge of the activation pattern. This suggest a direct application of our FSWT optimization algorithm to real data: instead of minimizing E, we propose to maximize the wavelet count for the structure parameters.

In more details, we found that for all types of FSWT (i.e., including `dual`) the error type II is strongly correlated with C, i.e., the product $C \times E2$ remains approximately constant. It is thus because, as far as the `bspline` and `ortho` types are concerned, the missed detections count is always so much larger than the false detections $E1$ that we finally have $E \propto 1/C$.

→ **High-pass detections.** Most of the band-pass coefficients of the wavelet transform are usually left undetected, which question the utility of testing for them. However, for activations that have a high frequency contents, some of these coefficients can be detected even with a low SNR. The gain over low-pass detection alone is then clear as exemplified in Figure 6.

The relation between E and C is evidenced in Figure 5. For this specific case, the optimal value of α is 4.2 in the causal `ortho` type and 0.2 in the causal `bspline`; both are reached at a depth of two iterations. The correlation coefficient between C and $1/E$ is close to 0.94, out of a maximum of 1 for perfectly proportional functions. The bad SNR does not favor `bspline`, indeed: we count 80% of mainly missed detections, while `ortho` performs best among all the other wavelets with a little more than 60% detection errors. Here, for $\alpha = 1.6$ and 2 iteration steps, the causal `dual` FSWT is only slightly worse (by 1%) than `ortho`. However, the correlation between $1/E$ and C is so loose (below 0.3) that it makes the prediction of E out of C impossible.

We also applied the SPM software to the synthetic data. In order to be fair, we slightly optimized the width of the Gaussian window that is used for SNR enhancement: the standard value, i.e., twice the space resolution, did not give good results at all. Even with a more suitable window, the SPM method still had over 70% detection errors and was thus outperformed by our FSWT algorithm.

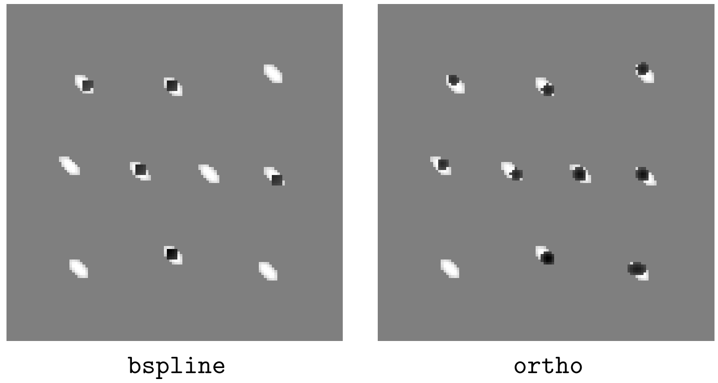

Figure 4. Test data corrupted with a high amount of noise SNR $= -9.15$: the detected pixels are shown in black. Left: detection with the best causal `bspline` FSWT algorithm; right: detection with the best causal `ortho` FSWT algorithm.

4.2. Example of detection on true fMRI data

We applied our algorithm on one slice of the fMRI data described in Section 3.1. It is known that, for the kind of motor task that generates the activation in these data, the signal-to-noise ratio is relatively high compared to that generated by other cognitive tasks.

We first applied SPM (2D version, corrected threshold) which provided the top right image of Figure 7. Then we applied our optimization algorithm to the `ortho` and `bspline` FSWT, replacing the apriori error criterion E, by the wavelet count C: the best result was obtained for the causal `ortho` flavor $\alpha = 5.6$ and for one iteration step. The result is shown in the bottom right image of Figure 7. Finally, we applied our algorithm using a causal `dual` FSWT with $\alpha = 1$ and a depth of one iteration, which provides the bottom left image of Figure 7.

While it is not easy to tell which result is the best, they are all plausible based on the physiology of the brain. We observe that the `ortho` method finds an additional two isolated pixels along another sulcus than that detected by the three methods: it turns out that such a detection is in good agreement with our knowledge of the activation loci of tasks in the motor cortex. We would also like to point out that the zone detected by our algorithm spreads longer along the main sulcus than that of SPM: this is plausible once again, and should it be confirmed by other independent experiments, it would mean that, in this special case, our method outperforms SPM. However, we would rather say that it behaves comparably to SPM, while requiring much simpler statistics (Bonferroni correction).

5. CONCLUSION

We have described a wavelet-based algorithm for detecting the activation in fMRI data. For this purpose, we used a new wavelet transform: the fractional spline wavelet transform (FSWT), for which we have a fast and exact implementation. The characteristic of this transform is that it has a free continuous parameter (the spline degree) and several flavors that make it either orthonormal or biorthonormal, either symmetric or not. We found a way to optimize for these degrees of freedom when the activation pattern is not a priori known.

Our optimized tests on synthetized data have indicated that our method is likely to improve over SPM, the standard technique in the field. This potential gain is obtained using very simple statistics, and indeed, a more sophisticated detection method could easily be worked out: this leaves still a substantial margin of improvement.

We can detect different shapes of different sizes and this detection is robust with respect to noise. Because of its multiresolution nature, our method is especially suited for detecting shapes with various frequency contents, such as mixtures of large and tiny activation zones. This is to be compared with SPM where the choice of the width of the Gaussian filter implicitly sets the size of the detecting shape.

We restricted the presentation to 2-D data but the 3-D extension is straightforward: it is needed to take the volumetric nature of the fMRI data into account.

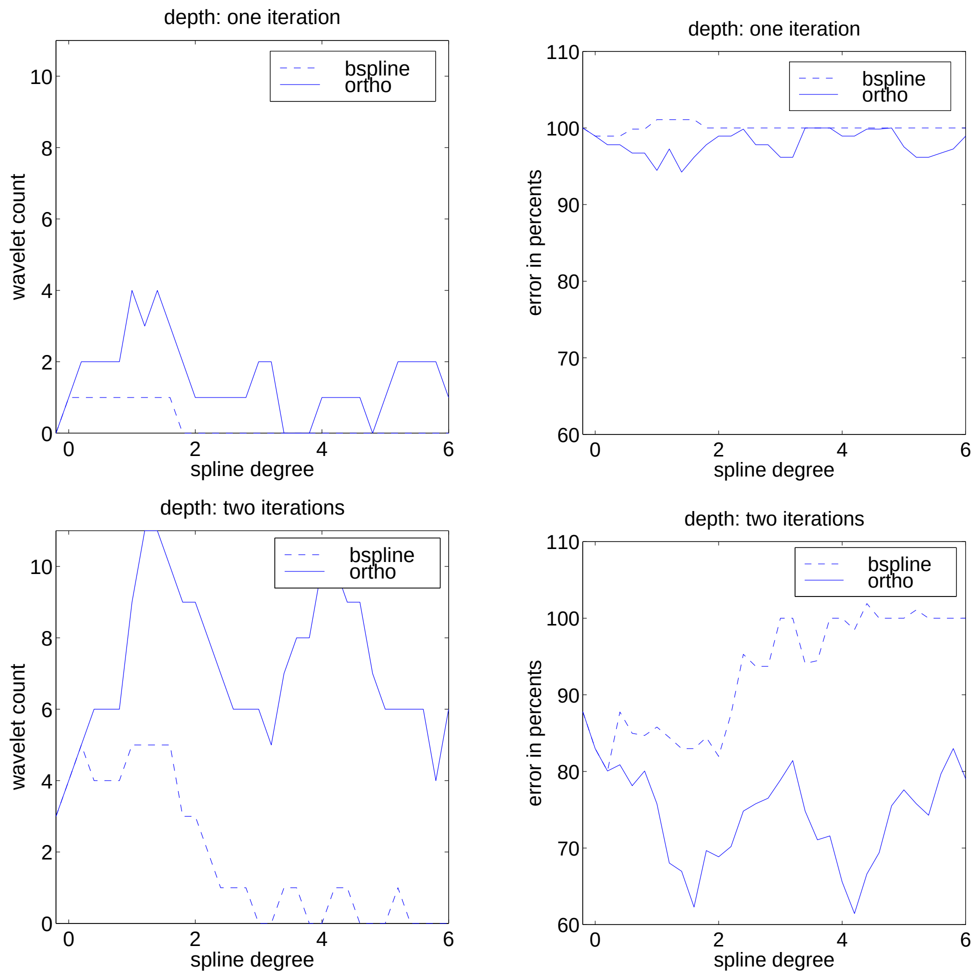

Figure 5. Test data of Figure 4: variations of the detection error count E and the wavelet count C in function of α (causal `ortho` and `bspline` types). Note the sharp minima of E, especially for `ortho`, and the inverse correlation between E and C.

ACKNOWLEDGMENTS

We would like to thank Arto Nirkko of Inselspital Bern, for providing the fMRI data and for fruitful discussions.

REFERENCES

1. Z. Liang and P. Lauterbur, *Principles of Magnetic Resonance Imaging*, IEEE Press, 2000.
2. S. Ogawa, R. Menon, D. Tank, S. Kim, H. Merkle, J. Ellerman, and K. Ugurbil, "Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging," *Biophysical J.* **64**, pp. 803–812, 1993.
3. S. Ogawa, R. Menon, S. Kim, and K. Ugurbil, "On the characteristic of functional magnetic resonance imaging of the brain," *Annu. Rev. Biophys.* **27**, pp. 447–474, 1998.
4. R. Cox, "A two day workshop on functional MRI," in *Proc. Colloque de Physique C2*, June 1996.
5. P. Thévenaz, U. Ruttimann, and M.Unser, "A pyramid approach to subpixel registration based on intensity," *IEEE Trans. Image Processing* **7**(1), pp. 27–41, 1998.
6. A. Macovski, "Noise in MRI," *Magnetic Resonance in Medicine* **36**, pp. 494–497, 1996.
7. M. Dagli, J. Ingeholm, and J. Haxby, "Localization of cardiac-induced signal change in fMRI," *NeuroImage* **9**, pp. 407–415, 1999.

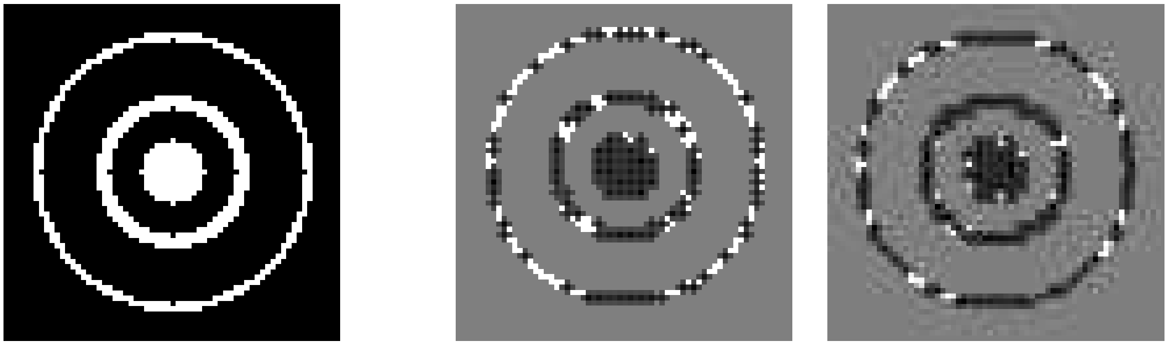

a) b)

Figure 6. Detection of the high frequency activation on the left using: a) only the low-pass coefficients of the FSWT; b) both the low- and high-pass coefficients.

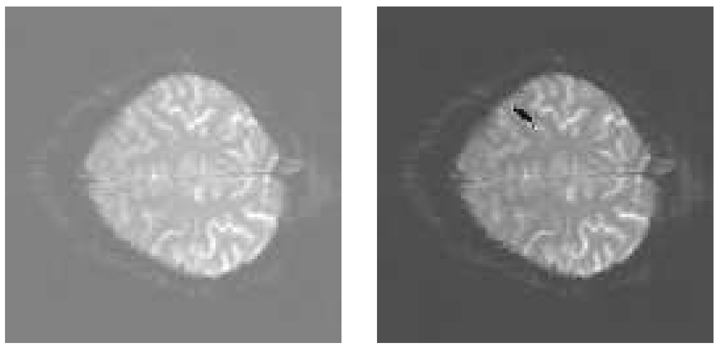

average fMRI image SPM

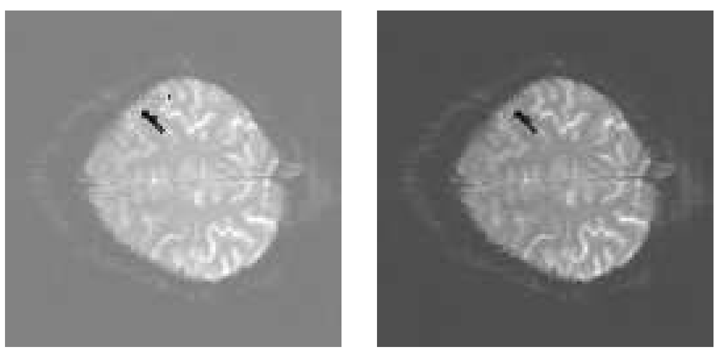

`ortho` `dual`

Figure 7. Real fMRI data: detection of the activation (black pixels) using SPM and using our algorithm based on the causal FSWT of types `ortho` and `dual`.

8. E. Bullmore, M. Brammer, S. Rabe-Hesketh, V. Curtis, R. Morris, S. Williams, T. Sharma, and P. McGuire, "Methods for diagnosis and treatment of stimulus-correlated motion in generic brain activation studies using fMRI," *Human Brain Mapping* **7**, pp. 38–48, 1999.
9. M. Unser, P. Thévenaz, C. Lee, and U. Ruttimann, "Registration and statistical analysis of PET images using the wavelet transform," *IEEE Engineering in Medecine and Biology* **145**, pp. 603–611, September/October 1995.
10. U. Ruttimann, M. Unser, R. Rawlings, D. Rio, N. Ramsey, V. Mattay, D. Hommer, J. Frank, and D. Weinberger, "Statistical analysis of functional MRI data in the wavelet domain," *IEEE Trans. Medical Imaging* **17**(2), pp. 142–154, 1998.

11. M. Feilner, T. Blu, and M. Unser, "Statistical analysis of fMRI data using orthogonal filterbanks," in *SPIE, Wavelet Applications in Signal and Image Processing VII*, vol. 3813, (Denver, CO), July 19–23 1999.
12. M. Unser and T. Blu, "Fractional splines and wavelets," *SIAM Review* **42**, pp. 43–67, 2000.
13. R. Frackowiak, K. Friston, C. Frith, R. Dolan, and J. Mazziotta, *Human Brain Function*, Academic Press, 1997.
14. K. Worsley, "Local maxima and the expected Euler characteristic of excursion sets of ξ^2, t and F fields," *Journal of Applied Probability* **26**, pp. 13–42, 1994.
15. K. Worsley, A. Evens, S. Marrett, and P. Neelin, "A three dimensional statistical analysis for CBF activation studies in human brain," *J. Cerebral Blood Flow and Metabolism* **12**, pp. 900–918, 1992.
16. K. Friston, J. Ashburner, J. Poline, C. Frith, J. Heather, and R. Frackowiak, "Spatial registration and normalization of images," *Human Brain Mapping* **2**, pp. 165–189, 1995.
17. K. Friston, S. Williams, R. Howard, R. Frackowiak, and R. Turner, "Movement-related effects in fMRI time series," *Magnetic Resonance in Medicine* **35**, pp. 346–355, 1996.
18. S. Mallat, "A theory for multiresolution signal decomposition: The wavelet decomposition," *IEEE Trans. Pattern Anal. Mach. Intell.* **11**, pp. 674–693, July 1989.
19. I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Comm. Pure Appl. Math.* **XLI**, pp. 909–996, November 1988.
20. G. Strang and T. Nguyen, *Wavelets and Filter Banks*, Wellesley-Cambridge Press, Cambridge MA, 1996.
21. T. Blu and M. Unser, "The fractional spline wavelet transform: Definition and implementation," in *Proc. ICASSP*, vol. 1, pp. 512–515, (Istanbul, Turkey), June 2000.