Multiresolution Analysis in fMRI: Sensitivity and Specificity in the Detection of Brain Activation

M. Desco,[1]* J.A. Hernandez,[1,2] A. Santos,[3] and M. Brammer[4]

[1]*Medicina Experimental, Hospital General Universitario "Gregorio Marañón", Madrid, Spain*
[2]*Grupo de Bioingeniería y Telemedicina, ETSI Telecomunicación, UPM, Madrid, Spain*
[3]*Department of Ingeniería Electrónica, ETSI Telecomunicación, UPM, Madrid, Spain*
[4]*Department of Biostatistics and Computing, Institute of Psychiatry, London, United Kingdom*

Abstract: Multiresolution analysis of fMRI studies using wavelets is a new approach, previously reported to yield higher sensitivity in the detection of activation areas. No data are available, however, in the literature on the analytic approach and wavelet bases that produce optimum results. The present study was undertaken to assess the performance of different wavelet decomposition schemes by making use of a "gold standard," a realistic computer-simulated phantom. As activation areas are then known 'a priori,' accurate assessments of sensitivity, specificity, ROC curve area and spatial resolution can be obtained. This approach has allowed us to study the effect of different factors: the size of the activation area, activity level, signal-to-noise ratio (SNR), use of pre-smoothing, wavelet base function and order and resolution level depth. Activations were detected by performing *t*-tests in the wavelet domain and constructing the final image from those coefficients that passed the significance test at a given *P*-value threshold. In contrast to previously reported data, our simulation study shows that lower wavelet orders and resolution depths should be used to obtain optimum results (in terms of ROC curve area). The Gabor decomposition offers the maximum fidelity in preserving activation area shapes. No major differences were found between other wavelet bases functions. Data pre-smoothing increases ROC area for all but very small activation region sizes. *Hum. Brain Mapping 14:16–27, 2001.*

Key words: SPM; wavelet; Gabor; phantom; ROC

INTRODUCTION

Since the introduction of BOLD (blood oxygen level-dependent) sequences [Ogawa et al., 1992], the number of applications of functional magnetic resonance imaging (fMRI) on neurological research and clinical studies is continuously increasing [Barnes et al., 2000; Bullmore et al., 1999a; Curtis et al., 1998, 1999; Honey et al., 1999; Phillips et al., 1999; Ring et al., 1999].

The fact that fMRI is non-invasive and can be used repeatedly in individuals has led to its widespread application, both in the understanding of normal cerebral functioning and in disease states. The number of areas in which the technique is being applied is clear from examination of the proceedings of the latest (6th) annual meeting of the Organization for Human Brain Mapping in San Antonio (See Neuroimage, vol. 11 no. 5, 2000). The number of centers using fMRI worldwide has grown from a handful in the early to mid-1990s to hundreds today.

Contract grant sponsor: FIS; Contract grant number: 00/36; Contract grant sponsor: Comunidad de Madrid; Contract grant number: 8.1/49/1998, III PRYCIT; Contract grant sponsor: CICYT; Contract grant number: TIC-99-1085-C02.

*Correspondance to: Manuel Desco, PhD, Medicina Experimental. Hospital G.U. "Gregorio Marañón", Dr. Esquerdo 46. E-28007. Madrid. Spain. E-mail: desco@mce.hggm.es

Received 10 January 2001; Accepted 18 April 2001

fMRI studies, however, are difficult to analyze due to the relatively small difference (1–4%) between activation and non-activation areas and to the high noise level intrinsic to the fast sequences involved [Brammer, 1998; Ruttimann et al., 1998; Vlaardingenbroek and den Boer, 1999]. As this noise level depends on voxel size, the signal-to-noise ratio (SNR) can be increased by using larger voxels at the expense of a lower spatial resolution. A typical figure for 1.5 T scanners and 256 × 256 images is a 20% noise level for a voxel size of 1 mm × 1 mm × 4 mm [Vlaardingenbroek and den Boer, 1999]. For these reasons, it is usually necessary to perform a long series of repeated experiments (rest and activation interlaced) and to use statistical tests to detect the existence of significant differences (activation areas).

Results are commonly presented as activation maps depicting those voxels that show significant changes between the different brain states, usually rest and activation. Different statistical approaches have been proposed, including Student's *t*-test, Kolmogorov-Smirnov test, correlation analysis and Markov random fields [Aguirre et al., 1998; Kuppusamy et al., 1997; Worsley et al., 1995, 1996a, 1997, 1998, 1999; Worsley and Friston, 1995] and non-parametric approaches [Bullmore et al., 1996]. These statistical tests are normally applied independently to each and every voxel in the study. Many activation mapping methods do not take advantage of neighborhood information (activated regions usually span over several voxels) to increase their statistical power. When this information is used, it is typically in a post-processing step, to enhance the shape of the activation areas detected or to correct the significance levels according to the size of the area [Bullmore et al., 1999b; Holmes, 1994; McColl et al., 1994; Poline et al., 1997].

Neighborhood information, however, can also be exploited through the use of multiresolution decomposition based on **discrete** wavelet transformation (DWT). This approach has the potential advantage of analyzing the images at different spatial scales, concentrating the information from voxels with a high spatial correlation into a few wavelet coefficients [Antonini et al., 1992; Mallat, 1989]. Statistical tests can then be applied in the wavelet domain, taking benefit from this spatial correlation, in an adaptive way without "a priori" assumptions about the size or shape of the activation.

Other alternatives to the discrete wavelet transform have also been proposed, such as the use of scale-space transforms [Poline and Mazoyer, 1994; Worsley et al., 1996b], but these are computationally more expensive than the DWT.

A key point when applying the wavelet decomposition is the selection of the appropriate bases functions and their order. Different bases functions offer different compromises between scale/frequency and spatial resolution, the maximum joint resolution being limited by Heisenberg's uncertainty principle [Daubechies, 1998]. The order of the wavelet also engenders a compromise between spatial and frequency resolution both resolutions: higher order wavelets offer better frequency resolution (higher sensitivity in our case), but lower spatial resolution (localization accuracy) [Banham et al., 1994; Cohen and Kovacevic, 1996; Daubechies, 1998].

Application of wavelet analysis to functional brain studies has been proposed by several authors, mainly for filtering noise either in fMRI [LaConte et al., 2000; Millet et al., 2000; Ngan et al., 2000; Zaroubi and Goelman, 2000] or in PET/SPECT studies [Turkheimer et al., 1999, 2000].

Two recent studies [Brammer, 1998; Ruttimann et al., 1998] presenting a multiresolution analysis of fMRI using wavelets reported a higher sensitivity in the detection of activated brain regions compared with more established methods of analysis. Brammer used Daubechies wavelets and reported a maximum sensitivity when using 12 coefficients at the fourth resolution level (that contained the experimental condition alternation frequency). Ruttimann made use of Lemarie-Battle wavelets, decomposing up to the fourth resolution level (neither wavelet order nor resolution being explicitly indicated in the paper). Neither of these studies provided comparisons of the results using different wavelet bases families, nor they show clear reasons to choose a particular wavelet order and number of decomposition levels. Ruttimann justifies the selection of the Lemarie wavelet family on the basis of its good theoretical properties (orthogonality, symmetry, reduced spectral overlap) but no additional theoretical justification is provided, either for the particular wavelet order nor for the resolution depth actually chosen.

As indicated, both authors reported good results compared with other established methods of analysis, mainly reflected as increased sensitivity. Nevertheless, it is noteworthy that any improvement in sensitivity of a medical test cannot be reliably assessed without also rating its specificity, using receiver operating characteristic (ROC) curves or similar methodologies [Constable et al., 1995].

A problem that both the above mentioned studies share is the lack of a 'gold standard,' necessary to assess objectively the quality of the results. In patient studies neither the exact size and position of activated

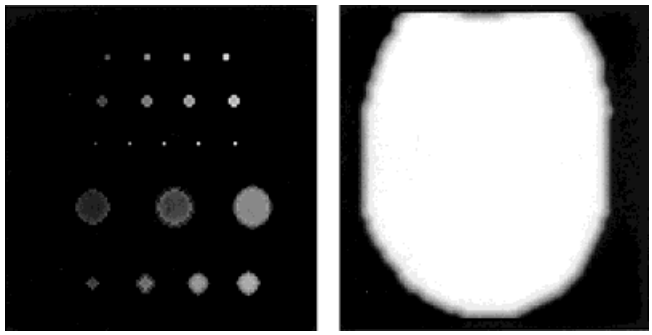

Figure 1.
Images used to build the software phantom: (left) activation regions; (right) intracranial mask.

areas nor the activation levels are known. This implies that sensitivity and specificity of the proposed methods cannot be reliably measured without using additional information.

The present study has assessed the performance of the different wavelet decomposition schemes by making use of a gold standard that is a realistic computer-simulated phantom with activation areas introduced by the operator. In this way, activation areas are known 'a priori', and accurate assessments of sensitivity, specificity and spatial resolution can be easily obtained.

This approach has allowed us to study the effect of different factors on activation mapping, namely, size of the activation area, activity level, signal-to-noise ratio (SNR), use of pre-smoothing, wavelet bases function and order and resolution level depth. Other limitations found when using the multiresolution method are also discussed.

MATERIALS AND METHODS

Algorithm

The implementation of the algorithm used in this work follows in general terms the approach proposed by [Ruttimann et al., 1998].

TABLE I.

	No pre-smoothing				Pre-smoothing (FWHM = 3 pixels)			
	Order	Resolution	*P*-value	% Area	Order	Resolution	*P*-value	% Area
Noise 5%								
t-Test		0	0.20	55%		0	0.15	85%
Gabor		2	0.0005	87%		1	0.25	88%
Daubechies	2	3	0.05	69%	2	1	0.05	86%
Lemarie	4	1	0.005	79%	2	1	0.05	86%
Symlet	2	1	0.005	79%	2	1	0.05	86%
Noise 10%								
t-Test		0	0.20	38%			0.15	72%
Gabor		2	0.05	81%		2	0.005	80%
Daubechies	2	3	0.20	60%	2	2	0.0005	75%
Lemarie	4	2	0.005	62%	2	2	0.005	74%
Symlet	2	3	0.20	60%	2	2	0.0005	75%
Noise 15%								
t-Test		0	0.20	29%		0	0.15	63%
Gabor		2	0.10	76%		2	0.10	76%
Daubechies	4	3	0.05	49%	2	2	0.005	67%
Lemarie	4	2	0.25	51%	2	2	0.005	65%
Symlet	2	3	0.25	47%	2	2	0.005	67%
Noise 20%								
t-Test			0.20	25%		0	0.15	56%
Gabor		2	0.15	72%		2	0.15	72%
Daubechies	2	3	0.25	40%	2	2	0.005	62%
Lemarie	4	3	0.10	45%	4	2	0.000005	61%
Symlet	4	3	0.05	44%	2	2	0.005	62%

Column "% Area" contains the maximum area under ROC curve obtained for the different wavelet families used in the study at noise levels of 5%, 10%, 15% and 20%. Columns order, resolution and *P*-value reflect the settings at which the maximum area was obtained. "*t*-test" rows correspond to the baseline *t*-test. The second half of the table show the same results after applying pre-smoothing (FWHM = 3 pixels). ROC area is calculated up to 95% of specificity.

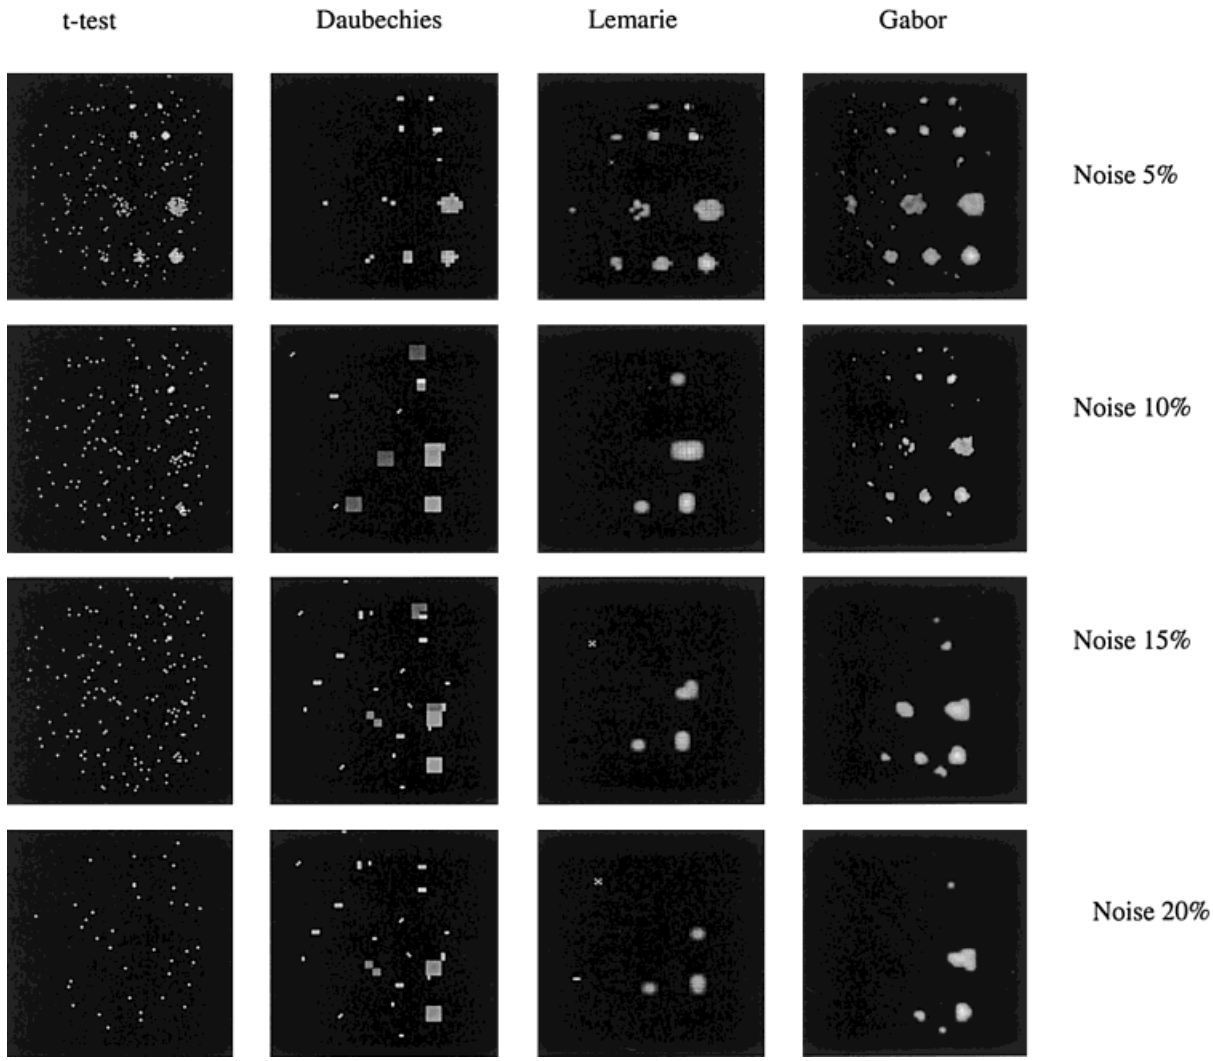

Figure 2.

Result of the baseline *t*-test (first column) and wavelet-based multiresolution analysis ($P < 0.01$) using Daubechies with two coefficients (second column), Lemarie with four coefficients (third column) and Gabor (fourth column), all without pre-smoothing. Each row corresponds to a different noise level: 5% (first row), 10%, 15%, and 20% (last row).

The "level" of multiresolution decomposition corresponds to the number of times the wavelet transform is applied to the original image. The "order" of the wavelet filter is the number of coefficients of the discrete mother wavelet [Mallat, 1989; Ruttimann et al., 1998].

Functional images are analyzed by performing a bidimensional discrete wavelet transform (DWT), decomposing up to the sixth level. At each level, the null-hypothesis is tested on the values of the wavelet coefficients (*t*-test without Bonferroni correction) at different *P*-values (ranging from 5×10^{-9} to 0.99) and the inverse transform is computed using only the significant coefficients that passed the test at the chosen *P*-value threshold. An alternative implementation, using separated *t*-tests for every individual pixel instead of a single pooled variance, was also evaluated and discarded, because in our case (additive Gaussian noise and no artifacts) it produced the same results. For real patient images, however, this latter approach may be preferable. Bonferroni corrections were not applied because in our setting the test is repeated for a whole range of *P*-values and the use of Bonferroni correction would simply shift the results toward higher *P*-values.

Several wavelet bases functions have been included: Daubechies, Lemarie and Symlets (orders 2, 4, 8, 16, 32, 40, 48 and 64), Spline (orders 2,8;4,4) and Maxflat

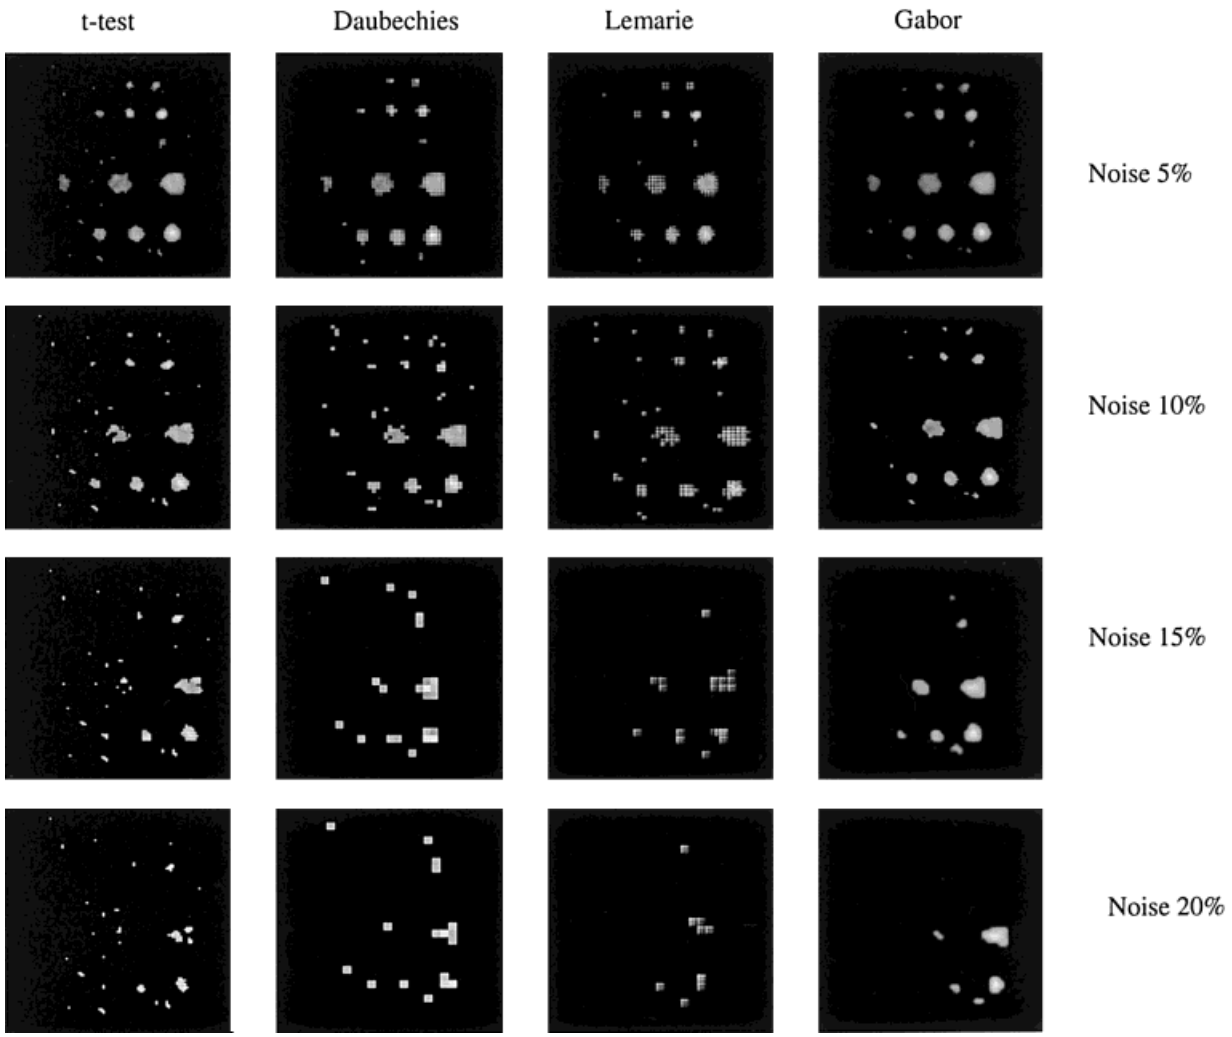

Figure 3.
Same data as in Figure 2, applying pre-smoothing (gaussian filter, FWHM = 3 pixels) before performing the tests.

(orders 3,3;5,5). Implementation of the transforms was performed using the "uvi_wave" toolbox for Matlab (The Mathworks, Inc.) [González and Márquez, 1996; Márquez and González, 1996]. The Gabor transformation [Gabor, 1946; Navarro et al., 1996] was implemented according to the proposed implementation with 11 coefficients described by Nestares et al. [1998]. Presmoothing of the data, when used, was performed using a gaussian filter with an FWHM of 3 pixels.

Software Phantoms

A set of 2D phantom studies was built upon a basal image with a uniform intensity level in a "brain-like" shape. On this image, twenty activation regions with different sizes and activation levels were created (Fig. 1). These regions comprised clusters of 1, 2, 4, 8, and 12 voxels. The intensity levels in the activation regions were 0.5%, 1%, 2% , 3% and 4% higher than the surrounding area. To produce more realistic images, avoiding sharp edges, all the images (activation and rest) were smoothed by applying a gaussian filter (FWHM = 3 pixels). Different levels of gaussian white noise (5%, 10%, 15%, or 20%) were then added.

Each phantom study consisted of four simulated fMRI scans (128 × 128 × 64), each consisting of eight activation-rest cycles of 6 images per epoch.

All the variables used in constructing the phantoms (activation level, noise, number of images) were chosen trying to resemble commonly accepted values for actual

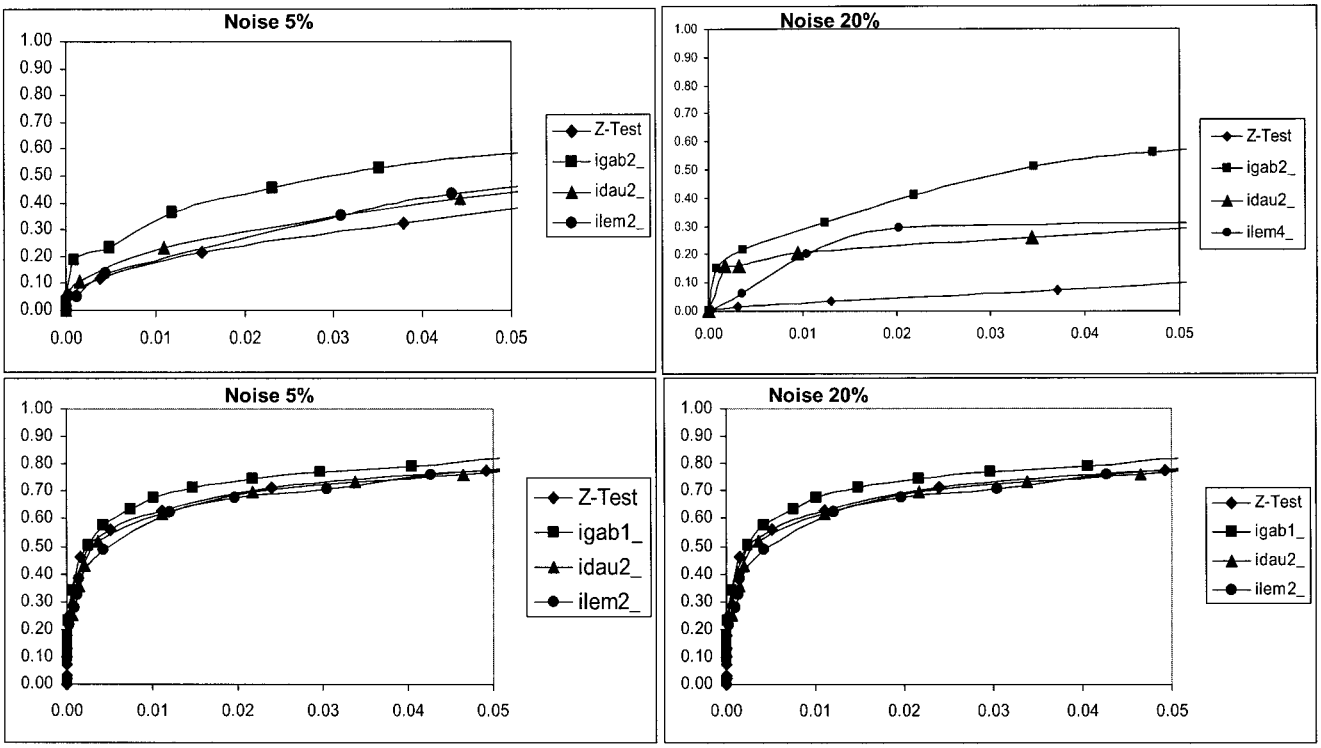

Figure 4.

ROC curves showing the best results at 5% (left) and 20% (right) noise levels. Upper row corresponds to images without pre-smoothing, analyzed up to second resolution for 5% and third resolution for 20% noise levels. Lower row corresponds to images with pre-smoothing, first resolution used at 5% and second resolution at 20% noise levels.

brain activation studies [Brammer, 1998; Ruttimann et al., 1998; Vlaardingenbroek and den Boer, 1999].

Evaluation

According to [Ruttimann et al., 1998], we used as baseline for the comparisons the result of conventional *t*-tests, presented as Z-scores. For each activation image, sensitivity (defined as the percentage of phantom activation voxels that were properly detected as activated) and specificity (defined as the percentage of voxels correctly detected as non activated) were determined. Comparisons between different methods were assessed by using receiver operating characteristic (ROC) curves.

ROC curves were generated by performing the statistical tests at different significance levels, ranging from $P = 5.10^{-9}$ to $P = 0.99$. The relative quality of the results using different methods was assessed by the area of the ROC curve from 0 to 95% specificity.

To test the effect of image pre-smoothing all the experiments were repeated after applying a gaussian filter of FWHM = 3 pixels, chosen according to usual recommendations [Friston et al., 1995; Worsley and Friston, 1995].

RESULTS

Overall performance is rated as the area under the ROC curve up to a 95% of specificity. Table I summarizes the best results obtained for each wavelet family, with or without pre-smoothing.

In Figure 2 the baseline results are compared with three examples of the wavelet-based method using different bases, Daubechies (2 coefficients), Lemarie (4 coefficients) and Gabor. The first two were selected because they apparently produced good results in previous studies [Brammer, 1998; Ruttimann et al., 1998]. The Gabor transform was included because it yielded the best results in our investigation. It is noticeable that all multiresolution results are better than those produced using the baseline *t*-test, especially at higher noise figures. The shape of activated areas detected by Gabor analysis was the closest to their true extent in the phantom.

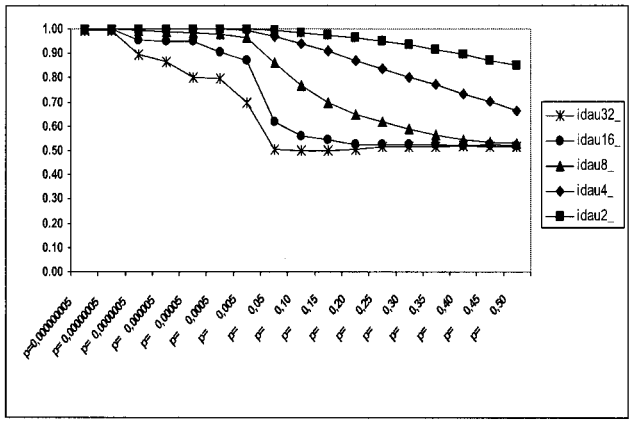

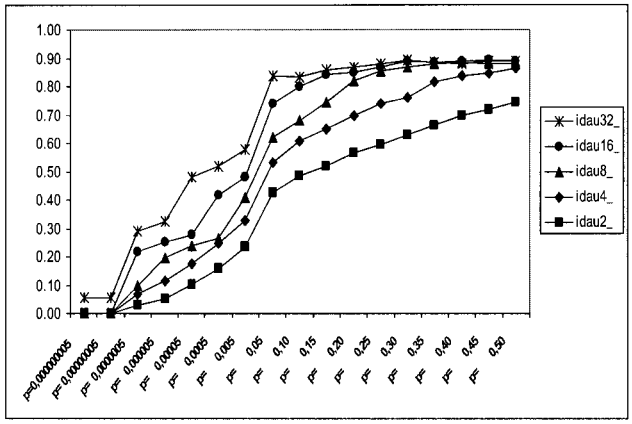

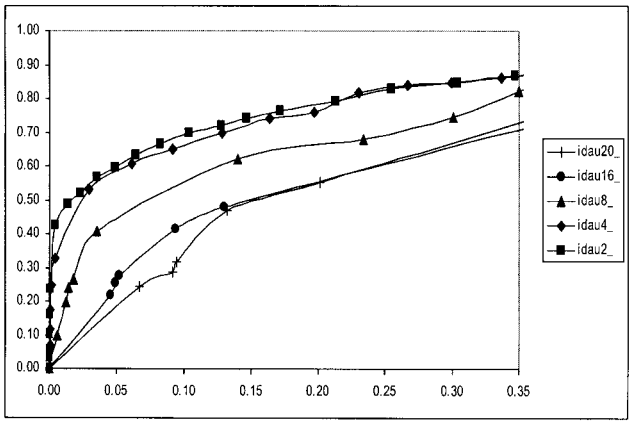

Figure 5.
Sensitivity, specificity, and ROC curves for different orders (32, 16, 8, 4, and 2 coefficients) of the Daubechies wavelet family. Data correspond to resolution 1, 10% noise level, pre-smoothing applied.

Figure 3 shows a similar comparison to that shown in Figure 2, but now obtained after applying pre-smoothing to the original images. In this case little improvement was observed on the un-presmoothed multiresolution results for noise levels of 5% and 10%, though for 15% and 20% better results were obtained. The better performance of the Gabor transform compared with the other approaches in preserving original shapes is evident. Figure 4 shows the ROC curves of the best results at the two extreme noise levels used (5% and 20%), with and without pre-smoothing.

Figure 5 illustrates the effect of the wavelet order on sensitivity and specificity. As an example, results using the Daubechies wavelet family are shown (32, 16, 8, 4 and 2 coefficients), using resolution 1 and 10% noise level with pre-smoothing. Similar results are obtained for other families and combinations of parameters. The ROC curves demonstrate that better results are obtained with lower wavelet orders, the ROC curve for 2 coefficients having the largest area in the 0–95% specificity range.

The resolution level also affects sensitivity and specificity. In Figure 6, ROC area is plotted against resolution level for the two extreme noise figures tested in this study (5% and 20%) in four different cases: baseline *t*-test, Daubechies (2 coefficients), Lemarie (4 coefficients) and Gabor. Figure 7 shows the results obtained after pre-smoothing with a gaussian filter of FWHM = 3 pixels.

The combined effect of resolution level and wavelet order can be seen in Figure 8, which shows the images obtained at resolutions 1 to 6 for high-order (32 coefficients) Daubechies and Lemarie wavelets. It is noticeable how the increase in sensitivity is associated with a substantial loss of specificity. The images corresponding to resolution 4 represent the setting proposed by Brammer [1998] and Ruttimann et al. [1998].

To assess the performance of the method on activation areas of different sizes, the analysis was repeated separately for the these activation regions. Table II summarizes the results obtained with multiresolution

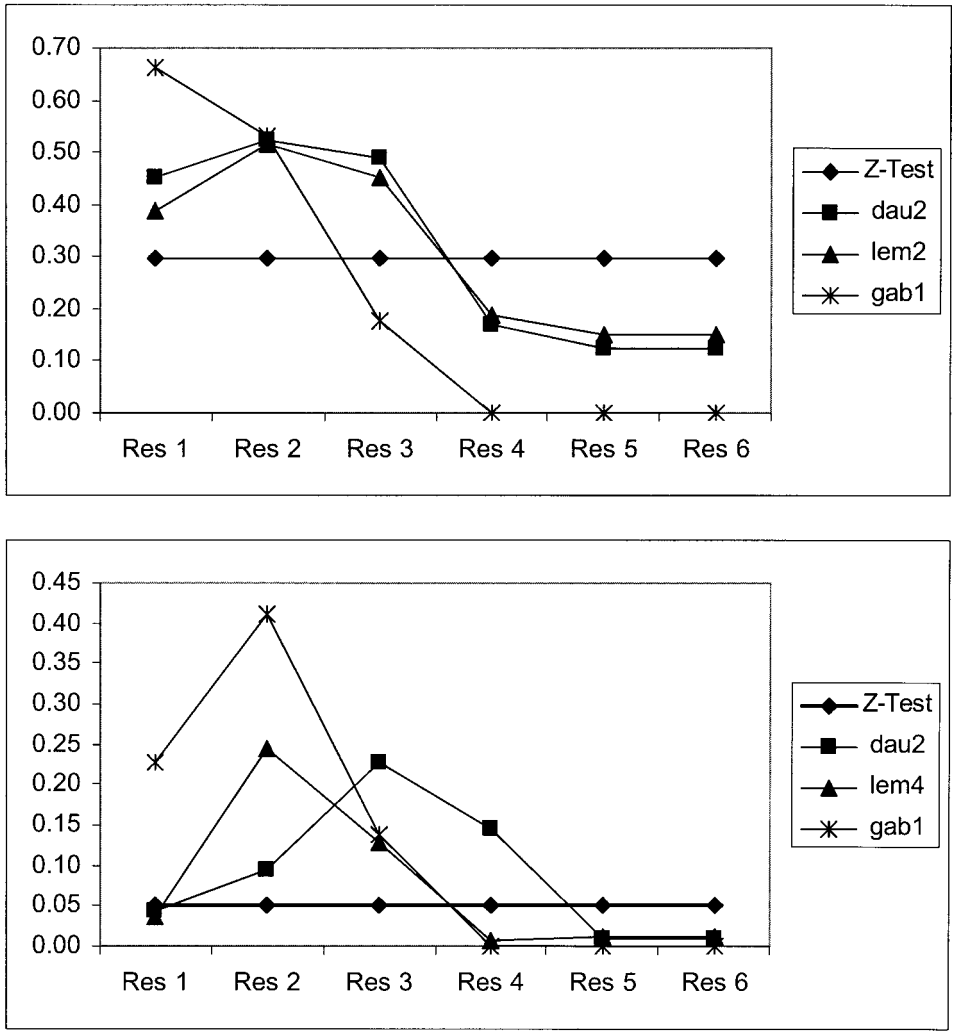

Figure 6.
ROC area vs. resolution level for two different noise figures (upper: 5%, lower: 20%). Curves correspond to baseline z-score, Daubechies, Lemaire, and Gabor. No pre-smoothing applied.

methods for area sizes 1 to 12 voxels and noise figures 5% and 20%.

DISCUSSION

On our phantom data, the baseline approach using conventional *t*-test was able to detect the 4% activation area essentially completely, but only part of the 3% activation area at a 5% noise level. Multiresolution methods clearly outperformed the standard approach, improving sensitivity and being able to locate even the 1% activation area, with different degrees of spatial resolution. The best results were obtained with the multiresolution Gabor decomposition, which provided the best sensitivity/specificity combination and a good spatial localization and region shape preservation.

Sensitivity decreases at higher noise levels both with conventional techniques and multiresolution methods. The advantage of the multiresolution approach is more noticeable, however, with low activation levels on noisy images, a result of some relevance for real activation mapping experiments. This is probably due to the fact that information from areas of noisy voxels with a high spatial correlation is concentrated into fewer wavelet coefficients that pass the statistical tests more easily.

The increase in sensitivity that we have obtained is in accordance with previous findings [Brammer, 1998; Ruttimann et al., 1998]. These previous studies, however, did not provide a quantitative assessment of their achievements, because they did not employ any "gold standard" for comparison. Furthermore, they selected 'a priori' a single wavelet family and order, thus not offering comparative data. Our results show that size and shape of the activation areas are closer to the expected ones when using Gabor transformation than when using the previously proposed Daubechies

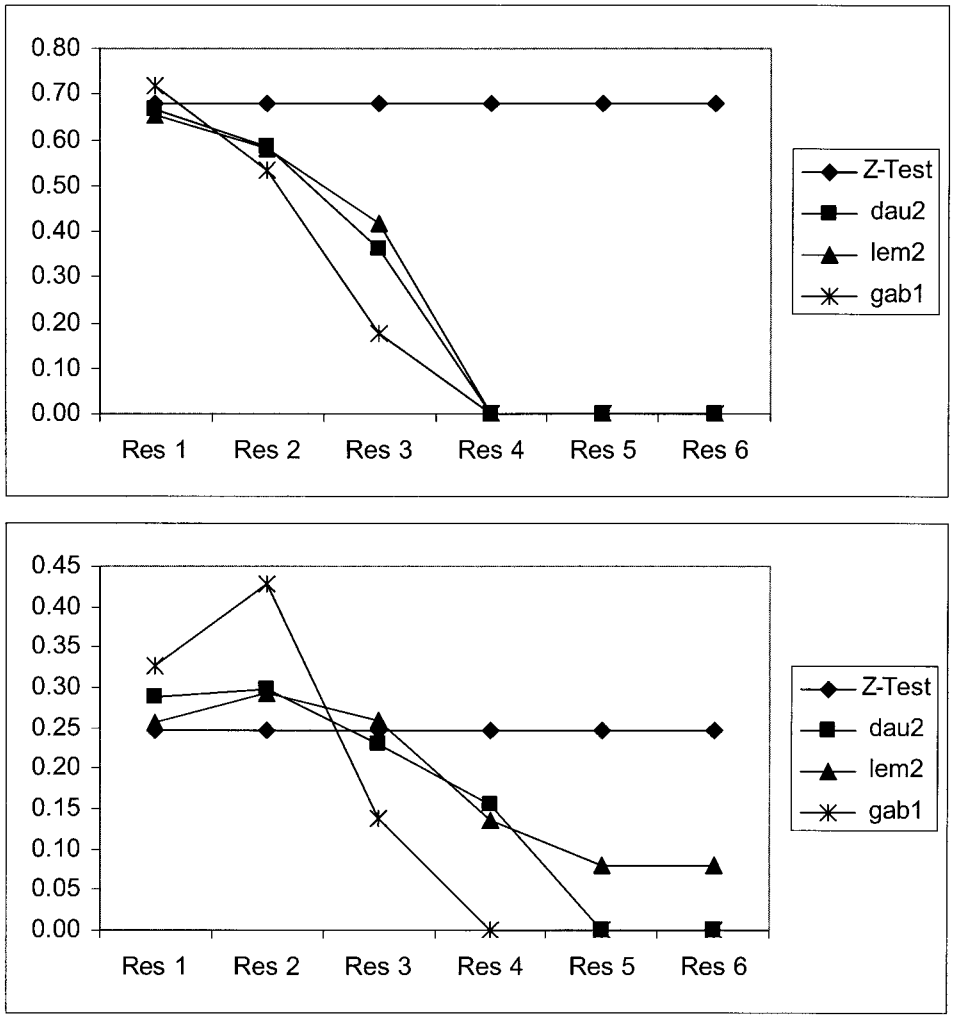

Figure 7.
ROC area vs. resolution level for two different noise figures (upper: 5%, lower: 20%), after applying pre-smoothing (gaussian filter FWHM = 3 pixels). Curves correspond to baseline z-score, Daubechies, Lemarie, and Gabor.

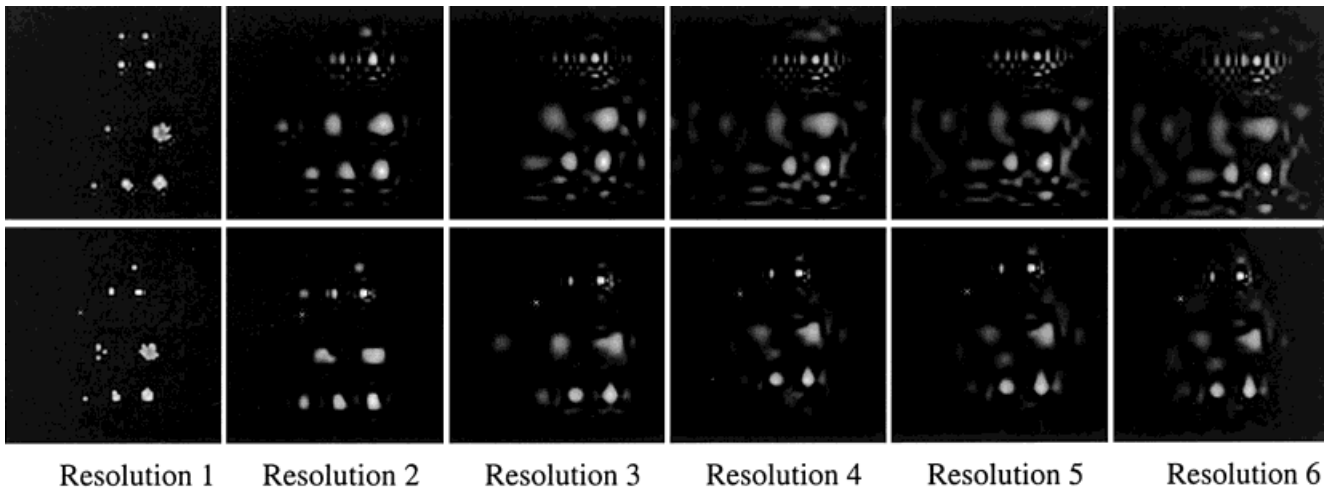

Figure 8.
Images obtained using Daubechies (upper row) and Lemarie (lower row) wavelet decomposition with 32 coefficients, from resolutions 1 to 6 (5% noise level, $P = 0.05$).

TABLE II.

	5% of Noise				20% of Noise			
	Method	Order	Resolution	Pre-smooth	Method	Order	Resolution	Pre-smooth
12 pixels								
2%	Gabor		1	Yes	Gabor		2	Yes
1%	Gabor		1	Yes	Gabor		2	Yes
0.5%	Gabor		2	No	Gabor		2	No
8 pixels								
3%	Gabor		1	Yes	Gabor		1	Yes
2%	Gabor		1	Yes	Gabor		2	Yes
1%	Gabor		1	Yes	Gabor		2	Yes
0.5%	Gabor		2	Yes	Lemarie	2	4	Yes
4 pixels								
3%	Gabor		1	Yes	Symlets	4	1	Yes
2%	Daubechies	2	1	No	Daubechies	4	1	Yes
1%	Gabor		1	No	Lemarie	2	2	Yes
0.5%	Spline	2,8	1	No	Spline	2,8	1	No
2 pixels								
3%	Gabor		1	No	Lemarie	2	1	No
2%	Lemarie	2	1	No	Daubechies	2	1	Yes
1%	Daybechies	2	1	No	Daubechies	2	1	No
0.5%	Symlet	4	1	No	Symlets	4	1	No
1 pixel								
4%	Gabor		1	No	Spline	2,8	1	No
3%	Gabor		1	No	Gabor		1,2	No
2%	No detection				No detection			
1%	No detection				No detection			
0.5%	No detection				No detection			

Best method for detecting activity depending on the size of the activation region and activity level. Data for the two extreme noise figures tested (5% and 20%) are provided. Order, resolution and whether pre-smoothing was applied are also indicated.

or Lemarie wavelets. Regarding the preservation of shape, it must be noted that the Gabor basis is the one that better fits the circular activation areas of our phantom; thus, these results should be cautiously extrapolated to non-circular areas.

When evaluating the performance of an analysis method, specificity is a parameter as important as sensitivity because sensitivity can always be increased at the expense of lowering specificity. The correct shape of the detected activation areas is also clearly dependent on obtaining adequate specificity. Unless a gold standard is available, specificity cannot be determined. The spatial distribution of false positive voxels is different using conventional or multiresolution methods. In the first case, false positive pixels are randomly distributed all over the image. In contrast, wavelet-based analysis tends to concentrate false pixels in the borders of real activation areas, increasing their apparent size ('border effect'). In this way, similar values of specificity may have a different interpretation: in the case of traditional analysis the lack of specificity appears randomly, giving a 'noisy' aspect to the results, whereas in the wavelet-based approach most of the error contributes to increase the apparent size of activation areas. In all the cases, increasing in the resolution level or the wavelet order led to worse spatial resolution and more 'border effect.'

Another interesting issue relates to the use of pre-smoothing. It is usually recommended to pre-smooth images with a gaussian filter of FWHM ranging between 8 mm and 12 mm [Friston et al., 1995; Worsley and Friston, 1995]. The need for this pre-filtering step is debatable when applying multiscale approaches: Ruttimann states that the wavelet decomposition itself should provide an adequate degree of low-pass filtering; Brammer, on the other hand, applied smoothing before proceeding with the wavelet analysis.

From our data on the whole image (including different sizes of activation regions) it seems that pre-smoothing always improves the overall performance,

either with conventional or multiresolution methods. When the analysis is performed separately for each activation area, however, better performance was obtained for the 1 and 2 voxel "activations" when no pre-smoothing was applied. This seems reasonable, because smoothing flattens sharp activity peaks, making their detection more difficult. As can be seen in Table II, Gabor decomposition outperforms any other wavelet decomposition scheme in most cases, and lower order and resolution levels are always preferable.

The results of this study cannot be directly compared with those obtained with traditional analysis when applying 'a posteriori' corrections according to the area size [Friston et al., 1995; Worsley and Friston, 1995]. This comparison was not the aim of this work, and we have not studied to what extent postprocessing techniques applied after the multiresolution analysis might improve the results, in terms of ROC curve area.

A limitation of the sensitivity/specificity and ROC data obtained from the whole image is that performance is measured mixing data that come from regions with different area sizes. For this reason, additional analysis for each region size were also performed.

CONCLUSION

Multiresolution analysis of fMRI studies outperforms more standard techniques, increasing the probability of detecting activation areas (sensitivity) while keeping specificity within reasonable limits. In contrast to previously reported data, our simulation study shows that lower wavelet orders and resolution depths should be used to obtain optimum results (in terms of ROC curve area). The Gabor transformation offers the maximum fidelity in preserving activation area shapes and no major differences were found between other wavelet bases functions. Data presmoothing increases ROC area except for very small activation region sizes.

ACKNOWLEDGMENTS

The authors thank the collaboration of the Instituto de Optica 'Daza Valdés' (CSIC) and the Signal Theory Group of the University of Vigo, whose implementations of the different wavelet transforms were used in this work. This work has been partially supported by grants from FIS (00/36), Comunidad de Madrid (8.1/49/1998 and III PRYCIT) and CICYT TIC-99-1085-C02.

REFERENCES

Aguirre GK, Zarahn E, D'Esposito M (1998): A critique of the use of the Kolmogorov-Smirnov (KS) statistic for the analysis of BOLD fMRI data. Magn Reson Med 39:500–505.

Antonini M, Barlaud M, Mathieu P, Daubechies I (1992): Image coding using wavelet transform. IEEE Trans Imag Proc 1:205–220.

Banham RM, Mark R, Galatsanos NP, González HL, Katsaggelos AK (1994): Multichannel restoration of single channel images using a wavelet-based subband decomposition. IEEE Trans Imag Proc 6:821–833.

Barnes J, Howard RJ, Senior C, Brammer M, Bullmore ET, Simmons A, Woodruff P, David AS (2000): Cortical activity during rotational and linear transformations. Neuropsychologia 38:1148–1156.

Brammer MJ (1998): Multidimensional wavelet analysis of functional magnetic resonance images. Hum Brain Mapp 6:378–382.

Bullmore E, Brammer M, Williams SC, Curtis V, McGuire P, Morris R, Murray R, Sharma T (1999a): Functional MR imaging of confounded hypofrontality. Hum Brain Mapp 8:86–91.

Bullmore E, Brammer M, Williams SC, Rabe Hesketh S, Janot N, David A, Mellers J, Howard R, Sham P (1996): Statistical methods of estimation and inference for functional MR image analysis. Magn Reson Med 35:261–277.

Bullmore ET, Suckling J, Overmeyer S, Rabe Hesketh S, Taylor E, Brammer MJ (1999b): Global, voxel, and cluster tests, by theory and permutation, for a difference between two groups of structural MR images of the brain. IEEE Trans Med Imaging 18:32–42.

Cohen A, Kovacevic J (1996): Wavelets: the mathematical background. Proc IEEE 84:514–522.

Constable RT, Skudlarski P, Gore JC (1995): An ROC approach for evaluating functional brain MR imaging and postprocessing protocols. Magn Reson Med 34:57–64.

Curtis VA, Bullmore ET, Brammer MJ, Wright IC, Williams SC, Morris RG, Sharma TS, Murray RM, McGuire PK (1998): Attenuated frontal activation during a verbal fluency task in patients with schizophrenia. Am J Psychiatry 155:1056–1063.

Curtis VA, Bullmore ET, Morris RG, Brammer MJ, Williams SC, Simmons A, Sharma T, Murray RM, McGuire PK (1999): Attenuated frontal activation in schizophrenia may be task dependent. Schizophr Res 37:35–44.

Daubechies I (1998): Orthonormal bases of compactly supported wavelets. Comm Pure App Math 41:909–996.

Friston KJ, Holmes AP, Worsley KJ, Poline JP, Frith CD, Frackowiak RSJ (1995): Statistical parametric maps in functional imaging: a general linear approach. Hum Brain Mapp 2:189–210.

Gabor D (1946): Theory of communication. Instr Electr Engin 93: 429–457.

González N, Márquez OW (1996): Uvi_wave, the ultimate toolbox for wavelet transforms and filter banks. In: Workshop on intelligent methods in signal processing and communication. Docampo D, Figueiras A, Peréz F, Editors. Bayona, Spain. p 224–227.

Holmes AP (1994): Statistical issues in functional brain mapping. PhD Thesis. Department of Statistics. University of Glasgow, Glasgow.

Honey GD, Bullmore ET, Soni W, Varatheesan M, Williams SC, Sharma T (1999): Differences in frontal cortical activation by a working memory task after substitution of risperidone for typical antipsychotic drugs in patients with schizophrenia. Proc Natl Acad Sci USA 96:13432–13437.

Kuppusamy K, Lin W, Haacke EM (1997): Statistical assessment of cross-correlation and variance methods and the importance of electrocardiogram gating in functional magnetic resonance imaging. Magn Reson Imaging 15:169–181.

LaConte SM, Ngan SC, Hu X (2000): Wavelet transform-based wiener filtering of event-related fMRI data. Magn Reson Med 44: 746–757.

Mallat SG (1989): A theory for multiresolution signal decomposition: the wavelet representation. IEEE Trans Patt Anal Mach Intel 11:399–404.

Márquez OW, González N (1996): An interactive software for a hypertext course in wavelets. In: The Seventh IEEE Digital Signal Processing Workshop. IEEE Press. Loen, Norway. p 105–108.

McColl JH, Holmes AP, Ford I (1994): Statistical methods in neuroimaging with particular application to emission tomography. Stat Methods Med Res 3:63–86.

Millet P, Ibanez V, Delforge J, Pappata S, Guimon J (2000): Wavelet analysis of dynamic PET data: application to the parametric imaging of benzodiazepine receptor concentration. Neuroimage 11:458–472.

Navarro R, Tabernero A, Cristóbal G (1996): Image representation with Gabor wavelets and its applications. In: Advances in imaging and electronic physics. Hawkes P, Ed. Boston, MA. Academic Press. p 1–84.

Nestares O, Navarro R, Portilla J, Tabernero A (1998): Efficient spatial-domain implementation of a multiscale image representation based on Gabor functions. Electronic Imaging 7:166–173.

Ngan SC, LaConte SM, Hu X (2000): Temporal filtering of event-related fMRI data using cross-validation. Neuroimage 11:797–804.

Ogawa S, Tank GW, Menon RS, Ellerman J, Kim S, Merkle H, Ugurbil K (1992): Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. Proc Nat Acad Sci USA 89:5951–5955.

Phillips ML, Williams L, Senior C, Bullmore ET, Brammer MJ, Andrew C, Williams SC, David AS (1999): A differential neural response to threatening and non-threatening negative facial expressions in paranoid and non-paranoid schizophrenics. Psychiatry Res 92:11–31.

Poline JB, Mazoyer BM (1994): Enhanced detection in brain activation maps using a multifiltering approach. J Cereb Blood Flow Metab 14:639–642.

Poline JB, Worsley KJ, Evans AC, Friston KJ (1997): Combining spatial extent and peak intensity to test for activations in functional imaging. Neuroimage 5:83–96.

Ring HA, Baron Cohen S, Wheelwright S, Williams SC, Brammer M, Andrew C, Bullmore ET (1999): Cerebral correlates of preserved cognitive skills in autism: a functional MRI study of embedded figures task performance. Brain 122:1305–1315.

Ruttimann UE, Unser M, Rawlings RR, Rio D, Ramsey NF, Mattay VS, Hommer DW, Frank JA, Weinberger DA (1998): Statistical analysis of functional MRI data in the wavelet domain. IEEE Trans Med Imaging 17:142–154.

Turkheimer FE, Banati RB, Visvikis D, Aston JA, Gunn RN, Cunningham VJ (2000): Modeling dynamic PET-SPECT studies in the wavelet domain. J Cereb Blood Flow Metab 20:879–893.

Turkheimer FE, Brett M, Visvikis D, Cunningham VJ (1999): Multiresolution analysis of emission tomography images in the wavelet domain. J Cereb Blood Flow Metab 19:1189–1208.

Vlaardingenbroek MT, den Boer JA (1999): Magnetic resonance imaging, vol. 1. Heidelberg: Springer-Verlag.

Worsley DF, Johnson RD, Kwong JS (1996a): Bronchogenic cyst causing unilateral ventilation-perfusion mismatch. Clin Nucl Med 21:249–250.

Worsley K, Marret S, Neelin P, Evans A (1996b): Searching scale space for activation in PET images. Human Brain Mapp 4:74–90.

Worsley KJ, Andermann M, Koulis T, MacDonald D, Evans AC (1999): Detecting changes in nonisotropic images. Hum Brain Mapp 8:98–101.

Worsley KJ, Cao J, Paus T, Petrides M, Evans AC (1998): Applications of random field theory to functional connectivity. Hum Brain Mapp 6:364–367.

Worsley KJ, Friston KJ (1995): Analysis of fMRI time-series revisited—again [comment]. Neuroimage 2:173–181.

Worsley KJ, Poline JB, Friston KJ, Evans AC (1997): Characterizing the response of PET and fMRI data using multivariate linear models. Neuroimage 6:305–319.

Worsley KJ, Poline JB, Vandal AC, Friston KJ (1995): Tests for distributed, nonfocal brain activations. Neuroimage 2:183–194.

Zaroubi S, Goelman G (2000): Complex denoising of MR data via wavelet analysis: application for functional MRI. Magn Reson Imaging 18:59–68.