

Wavelet smoothing of functional magnetic resonance images: A preliminary report

Bradley J. Lucier¹
Department of Mathematics, Purdue University
150 N. University Avenue, W. Lafayette, IN USA 47907-2067

ABSTRACT

Functional (time-dependent) Magnetic Resonance Imaging can be used to determine which parts of the brain are active during various limited activities; these parts of the brain are called activation regions. In this preliminary study we describe some experiments that are suggested from the following questions: Does one get improved results by analyzing the complex image data rather than just the real magnitude image data? Does wavelet shrinkage smoothing improve images? Should one smooth in time as well as within and between slices? If so, how should one model the relationship between time smoothness (or correlations) and spatial smoothness (or correlations). The measured data is really the Fourier coefficients of the complex image—should we remove noise in the Fourier domain before computing the complex images? In this preliminary study we describe some experiments related to these questions.

Keywords: Functional magnetic resonance imaging; complex wavelets; Daubechies wavelets; wavelet shrinkage; wavelet smoothing; activation regions

1. INTRODUCTION

Functional Magnetic Resonance Imaging (fMRI) is a technique used to determine which parts of the brain are active when doing certain tasks. Images of slices of the brain are taken at regular intervals, some images when the subject is at rest and some when the subject is performing a certain task (tapping a finger, listening to music, etc.). The two types of images are then compared; the regions where there are significant differences between the images (called *activation regions*) are presumed to be (some of) those parts of the brain involved in processing that task. Various statistical tests can be used to determine these activation regions.

While regular MRI images that reveal the anatomical structure of the body are fusions of many MRI measurements and hence have high resolution and low noise, fMRI images, which are not fused, are typically of low resolution (64×64 pixels) and have a fair amount of noise. Thus there is scope for algorithmic processing to improve the determination of activation regions.

The data collected for fMRI images are the complex Fourier coefficients of what we shall call the *complex images*, which are determined by applying the inverse Discrete Fourier Transform (iDFT) to the measured data. (The measured electromagnetic fMRI data have both magnitude and phase, similar to radar data.) Usually, the magnitude of the complex image data is taken to form what we shall call *magnitude images*. One can apply the statistical analyses or noise reduction techniques at various stages in the production of the magnitude images and these choices, among other things, concern us in this paper.

At this point we must apologize for not giving any references to the literature in this area. This is a preliminary report; we may well find that answers to our questions (either positive or negative) have already been given by previous authors. We give enough references to let the reader know the background of the techniques we try and to guide the reproduction of our results.

¹E-mail: lucier@math.purdue.edu. This work was supported in part by the Office of Naval Research, Contract N00014-91-J-1152.

2. FUNCTIONAL MAGNETIC RESONANCE IMAGING DATA

In this section we describe the form of the data with which we work.

Fundamentally, each measurement in an fMRI image consists of two-dimensional complex Fourier coefficients of the complex image data taken at a given horizontal *slice* of the brain. Each slice has 64×64 measurements; the 64×64 complex image is formed by computing the iDFT of the slice data. The complex magnitude is then applied to each complex pixel in the complex image to form a 64×64 magnitude image; this magnitude image is what is usually displayed on the screen.

In our data, 15 slices are taken to image the entire brain. Each slice has about twice the thickness of the width of each pixel within a slice. These slices are collected into what we call a *volume*, which images the entire brain.

A series of volume images are taken over time, some when the subject is at rest, and some when the patient is performing an assigned task. There is a relaxation time involved in MRI imaging and one cannot reduce the time between volume images below a certain amount. For our data, the time increment between volume images was two seconds. In our tests we had a series of 160 volume images; we ignored the first 4 images and otherwise had 64 active volume images (when the subject was performing the test) and 92 inactive volume images (when the subject was resting). The subject was asked to alternate resting or performing the specified task for runs of 16 volume images at a time, and there were extra inactive volumes at the beginning and the end of the testing period.

A typical magnitude image (of volume 50) is shown in Figure 1. The image is normalized so that the maximum magnitude is assigned 255 grey levels. The horizontal slices, ordered from the bottom of the brain to the top, are arranged starting at the upper left corner and proceeding across, then down. The front of the brain is pointing to the bottom in Figure 1.




FIGURE 1. A typical magnitude fMRI image of the brain.

3. SHOULD ONE APPLY STATISTICS TO THE COMPLEX OR THE MAGNITUDE IMAGE?

One has two choices when determining the active regions: applying statistical tests of differences to pixels over time either in the complex images or the magnitude images. The complex images contain phase information as well as magnitude information. The absolute phase of the complex images is relatively arbitrary, and is determined when the experiment is started. Once the experiment is started, however, procedures are put into place so that the phase does not drift over time (the polarity of the measurements is reversed when moving from one row in a slice to the next, for example, and is also reversed at the beginning of half the slices), so *differences* of measurement data should have a fixed, not an arbitrary, phase. The question is whether this phase information is useful.

When one applies statistical tests to the magnitude images, one treats all complex image pixels that are the same distance from the origin in the complex plane as the same. Because taking the complex magnitude is a *nonlinear* operation, statistical estimates of the average magnitude at a given pixel and other statistical parameters are *biased* when they are produced from the magnitude image data rather than the complex image data. Thus, keeping the phase information when determining active regions may give better results.

We examined this hypothesis by using Hotelling's T^2 test [5, page 54] at the 99.9% significance level to compute the active region pixels for both the complex images and the magnitude images. (When applied to real data, Hotelling's T^2 test is equivalent to a two-sided t -test.) The results for the magnitude images are shown in Figure 2, and the results for the complex images are shown in Figure 3. The top half of each image shows the active pixels in white (with a pixel value of 1) and the inactive pixels in black (with a value of 0). The active region information is often used to color fMRI volume images, so on the bottom half of each image we show the active regions image *multiplied by* the volume image in Figure 1. (One could say that we choose which pixels from Figure 1 to display depending on the activation regions.) This has the pleasing visual effect of de-emphasizing the active pixels in regions outside the corpus of the brain. (These appear to arise especially around the optic nerves.)

One notes that more active pixels are specified by the tests on complex data (4.04%) than by tests on the real magnitude data (2.51%). This is to be expected—tests on the real and imaginary components of the complex image pixels have more power than tests on the magnitude alone. We were surprised by the high number of active pixels in each image even at the .999 statistical significance level. While there are differences in the active regions as determined by the two methods, it is not clear that these differences are significant.

4. COMPLEX WAVELET DECOMPOSITIONS AND SHRINKAGE

Our main interest in wavelet smoothing of fMRI images is in studying inter-slice and inter-volume smoothing effects on the complex images. Here we describe the complex wavelet decompositions that we use.

Because the measurements and the complex image pixels are complex, it is natural to use complex wavelet transforms to decompose the data. We use the 4th order accurate complex Daubechies-type wavelets developed by Jean-Marc Lina [4]. These wavelets have several useful properties—they are orthogonal, symmetric, and have compact support. Because we implicitly denoise in L_2 , orthogonality helps. Symmetry reduces the visual effects of shrinkage artifacts. And complex wavelets are natural for complex data like MRI or radar data.

There is the practical question of how to deal with the boundaries of the data. Our images consist of four-dimensional arrays: a time series of volumes, consisting of slices, made up of rows of columns of pixels. So we have four types of boundaries, two along the edges of each slice; at the top and bottom of each volume; and at the beginning and end of the time series.

The MRI data within a slice, being Fourier coefficients of the complex image, are naturally periodic, so we extended the images periodically across the boundaries of each slice. When we applied the wavelet transform to slices within a volume, we first reproduced the last (top) slice to give us 16 slices per volume, and then extended the data above and below the volume by reflection. We dealt with the boundaries in the final (time) dimension by reflection.

We use tensor products of one-dimensional wavelets for our multi-dimensional wavelet decompositions. Because the slice pixels are about twice as deep as they are wide, the question of scale arises when decomposing




FIGURE 2. Active regions determined from the real magnitude images. Active pixels are shown in white at the top; at the bottom the active pixels are multiplied by the grey scales of the image in Figure 1.




FIGURE 3. Active regions determined from the complex images. Active pixels are shown in white at the top; at the bottom the active pixels are multiplied by the grey scales of the image in Figure 1.

three-dimensional volume images. We applied the analysis wavelet filters in the following groups when decomposing an entire volume: first one group of joint row and column filters, then three groups each of joint row, column, and slice filters. The question of differences between time and spatial scales is considered in a later section.

We used shrinkage parameters determined by the SUREShrink procedure of David Donoho and Iain Johnstone [2]. When shrinking complex wavelet coefficients we shrink the magnitude (possibly to zero) and keep the phase of the coefficients.

5. THE EFFECTS OF SLICE AND VOLUME SMOOTHING

In Figure 4 we show the activation regions of the complex images after applying wavelet decomposition and shrinkage independently within each slice. The SUREShrink shrinkage parameter was 8.661238508211351; active pixels make up 4.27% of the image. In Figure 5 we show the activation regions of the complex images after applying wavelet decomposition and shrinkage independently within each volume. The SUREShrink shrinkage parameter was 8.645521068584078; active pixels make up 4.34% of the image. Again, using the active pixels to choose pixels from the magnitude images makes a more pleasing display.

6. INCORPORATING SMOOTHING IN TIME

While the natural relationship between the width of pixels within each slice and the presumed thickness of each slice allows us to choose fairly confidently a series of directions in which to apply the wavelet analysis filter, we know of no such obvious relationship connecting the temporal and spatial scales. In this section we discuss some of the issues.

At the smallest (macro) spatial scales, there are natural discontinuities in brain (and other) images caused by interfaces, membranes, occlusions, etc. Despite these discontinuities, the spatial images can effectively be modeled as having smoothness in various Besov spaces that determine rates of compression or noise removal by wavelet methods [1].

On the other hand, at the smallest (macro) temporal scales, physiological processes are smooth, or at least continuous. Fundamentally, the time progression of physiological processes should be *smoother* than the spatial variation one finds in an image, and it may be possible to improve our data analysis by using techniques and theories that take advantage of this difference in the spatial and temporal smoothness.

A theory of wavelet decompositions for data that have different smoothness properties in different coordinate directions was developed by Christopher Leisner [3]. If one knows, or can predict, the relationship between the smoothness in each coordinate direction (expressed in terms of anisotropic Besov spaces), then one can adapt the multi-dimensional wavelet decompositions to take advantage of this difference in smoothness. Even if one mispredicts the relationships between the various smoothness parameters, one can still obtain some benefit in compression, noise removal, etc.

As a practical matter, one must consider the actual spatial and temporal scales of the data one measures. At a pixel spatial scale of a few millimeters, the data are again (obviously) discontinuous. One sees that at the temporal scale of two seconds, the time separating each pair of volume images in our time series, physiological processes are *also* basically discontinuous; it is not clear that one can, in fact, exploit this smoothness difference in time and space for fMRI images.

We give here the results of two experiments incorporating smoothing in time. In each experiment, we assume that the data have fundamentally twice the smoothness in time as in space (ignoring questions of absolute scales, or, equivalently, ignoring the absolute size of the norms in the relevant Besov spaces). Following Leisner's theory, this means that we should apply the analysis wavelet filters twice as often in space as in time.

To take into account absolute scales, we must decide when specifically to apply the five temporal analysis filters we use to decompose our 160 volumes in the times series. We first choose to apply the time analysis filters early in the analysis process. Specifically we apply the analysis filters in the following order: three times




FIGURE 4. Active regions determined from the complex images after independent smoothing of slices.




FIGURE 5. Active regions determined from the complex images after independent smoothing of volumes.

volume filters alone (in time); row and column filters together; row, column, slice, and volume filters together; row, column, and slice filters together; and finally row, column, slice, and volume filters together. (Remember that we apply row and column filters once without applying the slice filters since the slices are twice as thick as the pixels are wide.) This order assumes that in absolute scales, the variation in time is much smoother than the spatial variation. The SUREShrink parameter for this data set was 7.556510478822085; fully 23% of the pixels were flagged as active.

Secondly, we applied the time analysis filters late in the analysis process. Specifically we apply the analysis filters in the following order: row and column filters together; row, column, slice, and volume filters together; row, column, and slice filters together; row, column, slice, and volume filters together; and finally three applications of the volume filters alone. This order assumes that in absolute scales, the variation in time is similar to the spatial variation, while the fundamental smoothness in time is still twice the fundamental smoothness in space. The SUREShrink parameter for this data set was 7.641056914986314; 5.47% of the pixels were flagged as active.

The results of the first experiment are shown in Figure 6 and the results of the second experiment are seen in Figure 7. Our opinion is that early time-smoothing leads to worse results, while late time-smoothing does not improve results. Thus, with two seconds between each volume image, it appears that one cannot exploit possible difference in spatio-temporal smoothness with our algorithms.

7. SMOOTHING IN THE FOURIER DOMAIN

The data measured by an MRI machine are the Fourier coefficients of the complex image data. If noise enters the data during the measurement process, it manifests itself in the original data in the Fourier domain. One can attempt to remove this noise *before* converting the measured data to the complex image data with the iDFT.

Most methods of noise removal cause as a side effect the data to be smoothed. So to remove noise from the measured data would generally result in smoothing the data. So one must ask what effect smoothing the measured data will have on the complex image data after taking an iDFT.

While we know that smoothing the measured data (in the Fourier domain) will tend to make the complex images decay more quickly at the boundary of each slice, we do not know in detail how smoothing the measured data in the Fourier domain will affect the smoothness in Besov spaces of the complex images. Because the efficacy of further noise removal methods applied to the complex image data is determined by the Besov smoothness of that data, we cannot predict in detail whether to expect an improvement in further processing of the complex image data. On the other hand, noise removal is “good”, at least philosophically, so we conducted an experiment to test its effects.

Figure 8 shows the results of that experiment. Here we applied shrinkage to the wavelet coefficients *of the measured data*, the Fourier coefficients of the complex images, before applying the iDFT. Slices were smoothed independently of each other; no spatial or temporal smoothing was applied after converting the measured data to complex image data using the iDFT.

As you can see, the fraction of pixels in active regions increased to 7.17%, larger than the number of active pixels using only spatio-temporal smoothing. Furthermore, the SUREShrink parameter for this smoothing was 33.466054953629445, which is much larger than the SUREShrink parameters for smoothing in the complex image domain. This difference in size suggests that when modeled in terms of Besov spaces embedded in L_2 , the measured Fourier coefficients of the data are in only a low-smoothness Besov space. This also suggests that one should search for a better method to model the measured data, which would imply a need to choose a different denoising method (since the effectiveness of the wavelet shrinkage method depends solely on the smoothness of the data in the relevant Besov spaces).

8. ACKNOWLEDGEMENTS

The author has benefited from many discussions with Professor Mark J. Lowe, the Director of the Research MRI Facility, the Department of Radiology/Imaging Sciences, Indiana University, who also provided the data for this study.




FIGURE 6. Active regions determined from the complex images after anisotropic spatio-temporal smoothing, with early smoothing in time.




FIGURE 7. Active regions determined from the complex images after anisotropic spatio-temporal smoothing, with late smoothing in time.




FIGURE 8. Active regions determined from the complex images after smoothing of slices in the Fourier domain.

REFERENCES

1. A. Chambolle, R. A. DeVore, N.-Y. Lee, and B. J. Lucier, *Nonlinear wavelet image processing: Variational problems, compression, and noise removal through wavelet shrinkage*, IEEE Trans. Image Processing **7** (1998), 319–335.
2. D. Donoho and I. Johnstone, *Adapting to unknown smoothness via wavelet shrinkage*, J. Amer. Statist. Assoc. **90** (1995), 1200–1224.
3. C. Leisner, *Nonlinear wavelet approximation in anisotropic Besov spaces*, Indiana Univ. Math. J. **52** (2003), 437–455.
4. J.-M. Lina, *Image processing with complex Daubechies wavelets*, J. Mathematical Imaging and Vision **7** (1997), 211–233.
5. B. J. Winer, *Statistical Principles in Experimental Design*, McGraw-Hill, New York, 1971.