Wavelet-Based Bootstrapping for MRI Data

Brandon Whitcher

GlaxoSmithKline, Greenford, United Kingdom

1 Introduction

Medical imaging data is acquired in multiple spatial dimensions, usually three, and also over time – thus producing a four-dimensional data set per experiment. Taking into account the dependence structure in more than one dimension is difficult and not found in routine statistical analysis techniques. For example, testing brain activation in functional magnetic resonance imaging (fMRI) typically involves fitting a regression model to each voxel (volume element or three-dimensional pixel) independently, calculating a test statistic and then correcting for multiple comparisons to produce a so-called statistical parametric map (SPM).

The analysis of fMRI data is not limited to detecting brain activation. The concept of functional connectivity, a descriptive measure of spatio-temporal correlations between spatially distinct regions of the cerebral cortex (Friston *et al.*, 1993), is also an active area of research and development. Instead of regression models, cross-correlation sequences are estimated either between pairs of voxels or using a template sequence derived from the experimental design. Voxels exhibit significant correlation when it is determined that they exceed a specified threshold.

Given a two-dimensional finite lattice of observations (i.e., an image), the methodology presented here implements a resampling procedure in order to produce bootstrap realizations of the data and build up a bootstrap distribution for a statistic of interest. This procedure enables the bootstrap to be applied to statistics of interest when the observations show evidence of spatial correlation. The implicit differencing of the DWPT has the potential to handle certain intrinsic random fields, but more work is required to establish this fact. An example is provided here motivated by the notion of functional connectivity in the human brain. Briefly, it has been established that there are associations present between regions of the brain even when a subject is at rest. To establish associations when performing a specific task, it is not enough to test the hypothesis of no spatial association but to test against the presence of a baseline spatial association. We utilize the bootstrap to provide an empirical distribution for the measure of association in a stationary random field.

Brief descriptions of the methods and algorithm are provided in this abstract, the interested reader is encouraged to look at Whitcher (2004) for a more detailed account

2 Methods and Results

2.1 Discrete Wavelet Transforms

We consider an image $\{Z_{x,y} : x = 0, \ldots, M-1; y = 0, \ldots, N-1\}$, where M, N are not necessarily equal but both dimensions must be dyadic in length. The first scale of coefficients from a two-dimensional discrete wavelet transform (2D DWT) of Z are $W^{(h)}_{x,y,1}, W^{(v)}_{x,y,1}, W^{(d)}_{x,y,1}$ – the horizontal, vertical and diagonal wavelet coefficient sub-images, – and $V_{x,y,1}$ is the scaling coefficient sub-image. All four sub-images have dimension $M/2 \times N/2$. Subsequent scales of wavelet sub-images are derived from recursively filtering and downsampling the scaling sub-image $V_{x,y,j-1}$ (Mallat, 1998).

Initial investigations into producing approximately uncorrelated wavelet coefficients in space indicated that the default frequency partition of the 2D DWT was not sufficient. Hence, we also propose to use the two-dimensional discrete wavelet packet transform (2D DWPT) – a generalization of the filtering operations described above, allowing for a more flexible frequency partitioning of the original spatial process (Mallat, 1998). The notation for wavelet coefficient sub-images is more involved and is associated with a quad-tree structure. That is, each iteration of the 2D DWPT decomposes every wavelet sub-image at scale j into four wavelet sub-images at scale $j+1$. If the original image is defined to be $W_0 = Z$, then the first scale of 2D DWPT sub-images the same as the 2D DWT; i.e., $W_1 = \{W_1^{(a,b)} : a, b = 0, 1\}$ where each wavelet sub-image is of dimension $M/2 \times N/2$. The second scale of the 2D DWPT produces four wavelet sub-images for each of the four elements of W_1, such that $W_2 = \{W_2^{(a,b)} : a, b = 0, \dots, 3\}$ with each element having dimension $M/4 \times N/4$, and so on.

The spectral density function (SDF) of a wavelet sub-image from the 2D DWPT may be obtained using a multidimensional downsampling formula. Let $S_j^{(a,b)}(\boldsymbol{f})$ be the SDF associated with the wavelet sub-image at node (j, a, b) in the wavelet packet tree $\mathcal{T}$. Loosely speaking, $S_j^{(a,b)}(\boldsymbol{f})$ is formed by averaging aliased versions of a portion of the original spectrum determined by the squared gain function of the wavelet filter. The consequence of this is for regions of the spectrum of Z that are slowly varying, the wavelet spectrum associated with that region will have a small dynamic range and thus contain approximately uncorrelated wavelet coefficients. A more detailed description of multidimensional filtering and downsampling applied to Matérn covariance functions, and its effect on the spectrum of wavelet sub-images, may be found in Whitcher (2004).

2.2 Wavestrapping Algorithm

Let $Z(x)$ be an stationary random field (SRF) on a finite lattice (e.g., an image) and let $\mathcal{T} = \{(j, a, b) : j = 0, \dots, J; a, b = 0, \dots, 2^j - 1\}$ denote the quad-tree structure of the 2D DWPT. The general algorithm to implement wavelet-based bootstrapping on Z is:

1. Compute the two-dimensional DWPT on the observed image Z.
2. Determine the orthonormal basis $\mathcal{B} \subset \mathcal{T}$ by recursive testing for spatial autocorrelation.
3. Apply the naïve bootstrap within all wavelet sub-images $\mathbf{W}_j^{(a,b)}$, such that $(j, a, b) \in \mathcal{B}$.
4. Apply the inverse two-dimensional DWPT to obtain Z^*.
5. Compute the statistic of interest T^* on Z^*.
6. Repeat steps (3)-(5) a suitable number of times to produce the bootstrap distribution of T^*.

In the example presented here, and simulation studies found in Whitcher (2004), the statistic of interest is the isotropic sample variogram. Thus, we are measuring the ability of an adaptive wavelet-based bootstrapping procedure to produce images that share the same spatial covariance structure with the original image. This has intrinsic value by validating the ability of the 2D DWPT to adapt to arbitrary spatial covariance and produce approximately uncorrelated wavelet sub-images, and may also be used to explore the validity of hypothesis tests from statistics calculated on the image.

Testing for spatial autocorrelation (step 2) is performed via a top-down algorithm where frequency partitions are included in the orthonormal basis $\mathcal{B}$. Starting with the root of the quad-tree at $W_{0,0} = Z$, the first level of four wavelet sub-images $W_1 = \{W_1^{(a,b)} : a, b = 0, 1\}$ is tested for CSR (complete spatial randomness). If CSR cannot be rejected, then the wavelet sub-image is declared uncorrelated and its associated frequency partition is included in the orthonormal basis. If CSR is rejected, the four wavelet sub-images associated with the next level of the 2D DWPT are tested. This recursive testing procedure is run to a pre-specified depth. If any portion of the frequency plane is uncovered, then it is filled in using partitions from the final level of the decomposition.

2.3 Example: Functional MRI Data

The analysis of functional connectivity in magnetic resonance imaging (MRI) data has been around since the early 1990s. Functional connectivity is loosely defined as a measure of spatio-temporal correlations between spatially distinct regions of cerebral cortex (Friston *et al.*, 1993). Originally found in animal models and described as low-frequency oscillations of regional blood flow and oxygenation, recent publications indicate that the underlying biophysics that form the basis for functional connectivity are the same in rats and humans. Functional connectivity is usually characterized in humans as low-frequency physiological fluctuations (Hyde and Biswal, 1999).

Using data obtained from the Brain Mapping Unit, Department of Psychiatry, University of Cambridge, we illustrate how wavelet-based bootstrapping may be applied to fMRI. The data set is identified via 000444-m4_6_EPI and is available for download at

`http://www-bmu.psychiatry.cam.ac.uk/DATA/NULLdata/`.

It consists of a 4D data structure with three spatial dimensions $(X, Y, Z) = (64, 64, 21)$ and one temporal dimension $T = 108$. Each voxel dimension is $3.9 \times 3.9 \times 5.0$mm. For the subsequent analysis here we select the 10th axial slice ($Z = 10$) at time $T = 1$. The image was cropped to 32×32 voxels to minimize the amount of non-brain voxels. To eliminate any anatomical changes in the mean, a 2D wavelet shrinkage procedure was applied to the image. The smoothed version of the image was then subtracted from the original to produce an image representing the background noise structure (Figure 1, upper-left corner).

Figure 1 summarizes the results of two-dimensional wavestrapping using two different wavelet filters, the Haar and LA(8). The first attempt at reproducing the spatial variation used the Haar wavelet filter. Visually, the Haar wavestrap image does not have gross differences from the original version but its variogram is reduced for all lags. This illustrates the need for statistical summaries of the spatial structure since a visual comparison may not be sufficient. The bootstrap image using the LA(8) wavelet filter performs well both visually and as evaluated by the variogram. Although the smoothed curves indicate subtle differences between the simulated and bootstrap images, their correspondence is quite good with exception of the smallest lag where the bootstrap image is elevated.

3 Conclusions

As medical imaging data sets become commonplace, more difficult scientific questions will be asked of them. The methodology presented here is an attempt to acknowledge the fact that spatial autocorrelation exists in medical images. By adapting the recursive filtering operations in the wavelet transform to the estimated spectral density function of the image, the naïve bootstrap may be applied to each collection of coefficients in wavelet space independently – instead

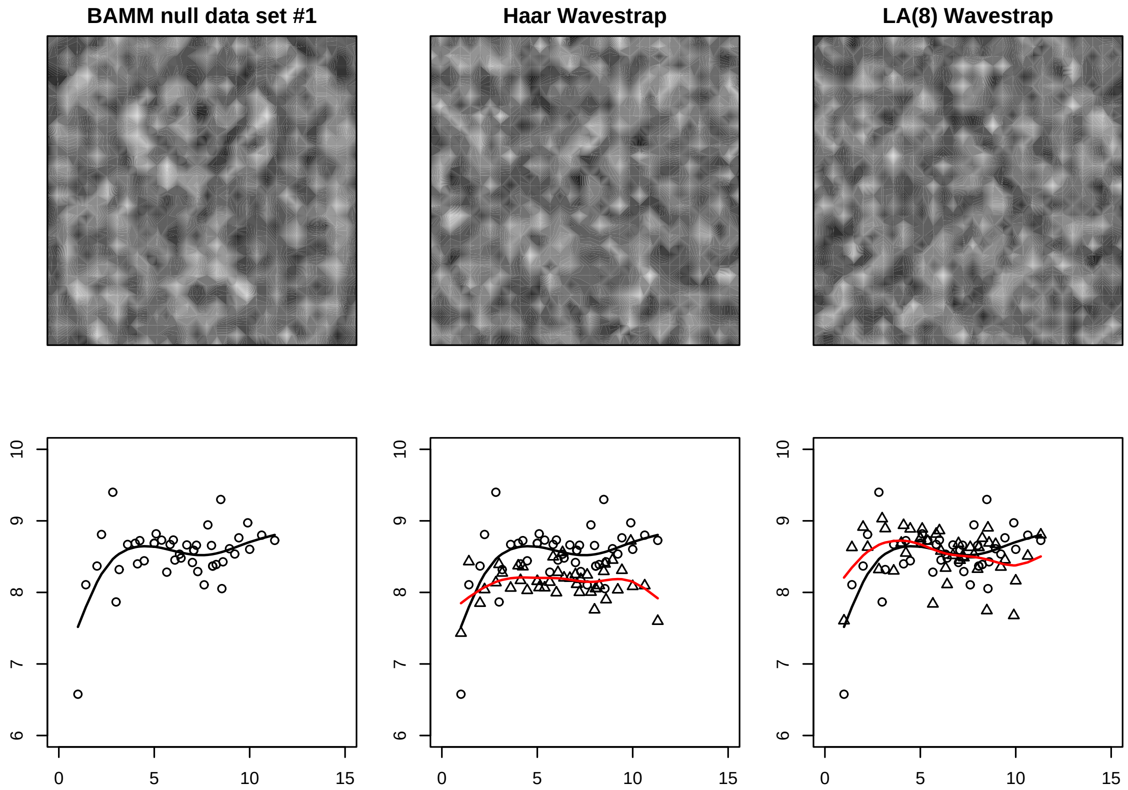

Figure 1: Functional MRI data set and wavestrapped samples using the Haar and LA(8) wavelet filters. The empirical variogram for the data is provided (dots) with a smoothed version (black line), along with the empirical variogram for the bootstrap image (triangles and red line).

of applying a more complicated sampling procedure in image space. Extensions to higher dimensions, along with the analysis of spatio-temporal data, are currently under investigation.

References

Friston, K.J., Frith, C., Liddle, P., Frickowiak, R. (1993). Functional connectivity: the principal components analysis of large (PET) data sets. *Journal of Cerebral Blood Flow and Metabolism* **13**, 5–14.

Hyde, J.S., Biswal, B.B. (1999). Functionally related correlation in the noise. In: Moonen, C.T.W., Bandettini, P.A. (Eds.), *Functional MRI*. Springer-Verlag, Berlin, pp. 263–275.

Mallat, S. (1998). *A Wavelet Tour of Signal Processing* (1st Ed.), Academic Press, San Diego.

Whitcher, B. (2004). Wavelet-based bootstrapping of spatial patterns on a finite lattice, submitted.