DUAL-TREE WAVELETS FOR ESTIMATION OF LOCALLY VARYING AND ANISOTROPIC FRACTAL DIMENSION

J. D. B. Nelson and N. G. Kingsbury

Signal Processing and Communications Laboratory
Dept. of Engineering, University of Cambridge, UK
{jdbn2, ngk}@eng.cam.ac.uk

ABSTRACT

The dual-tree wavelet transform is here applied to the problem of fractal dimension estimation. The Hurst parameter of fractional Brownian surfaces is estimated using various wavelet bases. Results are given for global, local, anisotropic, and both local and anisotropic Hurst parameters. It is shown that the directional selectivity of the dual-tree wavelets can be exploited effectively to compute and distinguish Hurst parameters that vary non-trivially with direction and space.

1. INTRODUCTION

Fractal dimension is a generalisation of the familiar Euclidean dimension. It can be regarded as a measure of how much area a curve fills or how much volume a surface fills. It succinctly describes the irregular, fragmented, and often self-similar shapes that occur in natural structures and is known to be well correlated with human perception of texture smoothness. In 2-D data analysis, estimation of fractal dimension has wide ranging applications that include fluid dynamics, surface inspection, adaptive Bayesian denoising, and texture classification of synthetic aperture radar, sidescan sonar, and natural imagery.

Wavelets, which perform a multiresolution decomposition of data, are intimately related to the idea of fractals. By measuring energy with respect to scale, fractal dimension can be extracted very efficiently using wavelets. Wavelet methods are faster and more tolerant to noise and affine transformations than box counting methods [1]. Since they are local, wavelets can also be used to estimate local fractal dimension. This is a significant advantage over methods based on the Fourier transform. However, most fast wavelet transforms have poor directional selectivity. As a result they are not well suited to anisotropic dimension estimation. Vidakovic et al [2] estimated anisotropic dimension in the horizontal, vertical and diagonal directions using discrete wavelets. For finer directionality, one alternative is to implement approximate versions of continuous wavelets such as the Gabor wavelet [3]. However, this comes at a significant computational cost.

This work was supported by the EPSRC (grant number EP/H012834/1) and the UK MoD University Defence Research Centre on Signal Processing.

We propose the use of the dual-tree complex wavelet transform (DTCWT) [4, 5] to estimate local and directional fractal dimension. The DTCWT is fast and, with 6 strongly directional subbands, it has good directional selectivity. Unlike other fast wavelet transforms, the DTCWT also has good shift invariance. This ensures that the magnitude of its complex coefficients remain stable and large near any singularities. It is therefore better at extracting local energy from signals that oscillate rapidly. Since dual-tree wavelets are relatively new they have yet to be applied to the problem of fractal dimension estimation. The main focus of this paper is to establish results on the accuracy and robustness of dual-tree wavelets for fractal dimension estimation. Since ground truth is required, fractional Brownian surfaces (fBs) are synthesised. Of particular interest is how the DTCWT compares to other wavelet methods when the fractal dimension is allowed to vary with respect to space and/or direction.

A surface with pointwise varying fractal dimension is sometimes referred to as multifractal [6]. The distribution of fractal dimension values can be described, globally, by a multifractal (or singularity) spectrum which in turn is related to the Rényi (or generalised fractal) dimension via the Legendre transform [7]. Although dual-tree wavelets could also be used for such analysis, we focus here on estimating the fractal dimension, or related Hurst exponent, in a specific location and direction.

The next section introduces fractional Brownian surfaces. The wavelet analysis method for global, local, and directional fractal dimension estimation is described in Section 3. Experiments are presented in Section 4. The first experiment compares wavelet methods for both global and local fractal dimension estimation. The second experiment tests the methods on locally varying and anisotropic fractal dimension. Finally, a classification problem is used to test how well each method can segment data based purely on the directionality of fractal dimension.

2. FRACTIONAL BROWNIAN SURFACES

For $H \in (0,1)$, an isotropic fractional Brownian surface $B_H \colon \mathbb{R}^2 \mapsto \mathbb{R}$ is a Gaussian zero-mean field with $B_H(\mathbf{0}) = 0$

such that the increments $(\Delta B_H)(\mathbf{x}) \triangleq B_H(\mathbf{x}+\Delta\mathbf{x}) - B_H(\mathbf{x})$ are stationary Gaussian random fields with variance

$$\mathbb{E}\left[\left|(\Delta B_H)(\mathbf{x})\right|^2\right] = \sigma_H^2 \left\|\Delta\mathbf{x}\right\|^{2H} . \quad (1)$$

Putting $\Delta\mathbf{x} = \boldsymbol{\xi} - \mathbf{x}$ into (1) gives the autocorrelation function $(RB_H)(\mathbf{x}, \boldsymbol{\xi}) = \mathbb{E}\left[B_H(\mathbf{x}) B_H(\boldsymbol{\xi})\right]$

$$= \frac{\sigma_H}{2}\left(\left\|\mathbf{x}\right\|^{2H} + \left\|\boldsymbol{\xi}\right\|^{2H} - \left\|\mathbf{x}-\boldsymbol{\xi}\right\|^{2H}\right) . \quad (2)$$

The generalised, or average, power spectrum of B_H is

$$(\mathcal{P}B_H)(\boldsymbol{\omega}) \propto \left\|\boldsymbol{\omega}\right\|^{-2H-2} . \quad (3)$$

The parameter H is referred to as either the Hölder parameter/regularity, Lipschitz regularity, or Hurst parameter. We shall favour the latter term. Values of H close to zero give very rough looking surfaces. Values close to unity will result in relatively smooth surfaces. For 2-D spaces, the parameter is directly related to fractal dimension D via $D = 3 - H$. Given B_H, the problem of fractal dimension estimation is therefore equivalent to the estimation of the Hurst parameter.

2.1. Synthesis of fBs

Fractional Brownian surfaces can be simulated in a variety of ways. Perhaps the most simple is to construct a function in the Fourier domain such that the absolute value is the product of $\left\|\boldsymbol{\omega}\right\|^{-H-1}$ and a normally distributed random variable with zero mean and unit variance, and where the phase is drawn from a uniform distribution. However, this Fourier synthesis method fails to reproduce some key properties of the fBs. In particular, it produces a stationary process.

Kaplan and Kuo [8] made a significant improvement to this technique. Their incremental Fourier synthesis method exploits the fact the increments of the fBs are stationary. The autocorrelations of the first and second order increments are transformed into the Fourier domain, square-rooted, and scaled by white complex Gaussian noise before an inverse Fourier transform is applied. The fBs is then computed by recursively adding the increments together. The result is non-stationary and gives closer approximations to the theoretical values of fBs correlation and variance than the original Fourier synthesis method.

2.2. Synthesis of local and anisotropic fBs

Surfaces with locally varying Hurst can be constructed as follows. First we split up the domain X of the surface into a disjoint covering. That is $\bigcup_n X_n = X$, and $X_n \cap X_m = \emptyset$, for $n \neq m$. Construct the masks $A_n(\mathbf{x}) \triangleq 1$, for $\mathbf{x} \in X_n$ and 0 otherwise. Let $\{\eta(n)\}$ determine the set of different Hurst parameters required. Next generate a set of surfaces $B_{\eta(n)}$ with the incremental Fourier synthesis method, using the same white noise for each surface. Then, an fBs with locally varying H can be constructed with $B_H = \sum_n A_n B_{\eta(n)}$. This surface will have Hurst parameter $\eta(n)$ at $\mathbf{x} \in X_n$.

An anisotropic fBs is synthesised by applying a set of directional binary masks to the set of fBs in Fourier space, summing, and then taking the inverse Fourier transform. First, define a disjoint set of intervals Θ_m that cover $[0, \pi)$. That is $\bigcup_m \Theta_m = [0, \pi)$, and $\Theta_m \cap \Theta_n = \emptyset$, for $m \neq n$. In polar Fourier space, construct the masks $\chi_m(\rho, \theta) \triangleq 1$, for $\theta, \theta - \pi \in \Theta_m$ and 0 otherwise. Again, let $\{\eta(m)\}$ determine the set of Hurst parameters required. Then the anistropic surface can be constructed via $B_H = \left(\sum_m \chi_m B^{\wedge}_{\eta(m)}\right)^{\vee}$, where $^{\wedge}$ and $^{\vee}$ denote forward and inverse Fourier transformation. This surface will have a Hurst parameter $\eta(m)$ for directions $\theta \in \Theta_m$. Both methods above can be combined to construct an fBs with a locally varying and anisotropic Hurst parameter. Although inefficient, this suffices to provide precise ground-truthed test data.

3. WAVELET ANALYSIS OF FBS

The wavelet transform of a fractional Brownian surface is:

$$\left(\mathcal{W}_{\psi;m} B_H\right)(k, \mathbf{x}) = 2^k \int_{\mathbb{R}^2} B_H(\boldsymbol{\xi}) \overline{\psi_m(2^k(\mathbf{x} - \boldsymbol{\xi}))}\, \mathrm{d}\boldsymbol{\xi} ,$$

where ψ_m is a mother wavelet with orientation indexed by m. As described in [2], the variance can be simplified to

$$\mathbb{E}\left[\left|\left(\mathcal{W}_{\psi;m} B_H\right)(k, \mathbf{x})\right|^2\right] = \frac{\sigma_H^2}{2} 2^{-2k(H+1)} c_{\psi_m, H} . \quad (4)$$

Taking the log of both sides of (4) gives

$$\log_2 \mathbb{E}\left[\left|\left(\mathcal{W}_{\psi;m} B_H\right)(k, \mathbf{x})\right|^2\right] = -2k(H+1) + C_{\psi_m, H} \quad (5)$$

The right-hand-side is now a linear function of scale k. Hence, the exponent H can be computed by measuring the average slope, over each of the directional subbands m, of (5) via least squares. In practice the variance is approximated by the sample variance $\frac{1}{N}\sum_n \left|\left(\mathcal{W}_{\psi;m} B_H\right)(k, \mathbf{x}_n)\right|^2$. When the Hurst parameter is allowed to vary as a function of space, $H = H(\mathbf{x})$ and the power spectrum from (3) becomes $(\mathcal{P}B_H)(\boldsymbol{\omega}, \mathbf{x}) \propto \left\|\boldsymbol{\omega}\right\|^{-2H(\mathbf{x})-2}$. In this case, we compute an estimate of local energy around the point $\mathbf{x}$, namely

$$E_{k,m}(\mathbf{x}) \triangleq \left|\left(\mathcal{W}_{\psi;m} B_H\right)(k, \mathbf{x})\right|^2 . \quad (6)$$

For each $\mathbf{x}$, the slope of (6) against k is now computed and averaged over the subbands denoted by m. Typically, better results can be obtained by ignoring the coarsest and finest scale levels. The coarsest levels tend not to be local enough and the finest tend to be more sensitive to noise. For anisotropic fractal dimension, we have $(\mathcal{P}B_H)(\rho, \theta) \propto \rho^{-2H(\theta)-2}$, and for both local and anisotropic, we have $H = H(\mathbf{x}, \theta)$. In this less studied case, the log energies $\log E_{k,m}(\mathbf{x})$ are once again computed for each locality and subband. However, the slopes from all the subbands are not averaged. Instead, each slope describes an estimate of H at $\mathbf{x}$ in the mth subband direction.

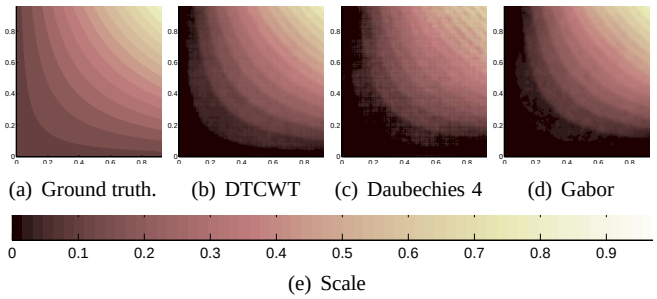

(a) Ground truth. (b) DTCWT (c) Daubechies 4 (d) Gabor

(e) Scale

Fig. 1. Mean Hurst estimates over 100 instances of fractional Brownian surface with locally varying fractal dimension.

4. EXPERIMENTS

4.1. Local Hurst estimation

In our first experiment, the Hurst parameter was allowed to vary with respect to space $H(x, y) = \text{round}(2 + 16xy)/20$, for $\mathbf{x} = (x, y) \in [0, 1]^2$. The surface was realised by the method outlined in Section 2.2. Fig. 1, which shows the mean Hurst estimates, gives a qualitative view using different wavelets. The mean error is plotted in Fig. 2 (a) for various wavelets. Fig. 2 (b) shows results for the case where the Hurst parameter is kept constant throughout the image. The best performing order was chosen for the Coiflet, Daubechies, and symlet wavelets. For a fair comparison, the Gabor wavelet was chosen to have the same orientations $\{(15 + 30m)^\circ\}_{m=0}^{5}$ as the DTCWT. If required, at extra computational cost, it is possible to perform the DTCWT with extra orientations and scales. Double dyadic scale levels were used for all Gabor results in this paper. The single dyadic version resulted in marginally worse results. The benefits of shift invariance and directionality can be seen by the low error in the DTCWT and Gabor results. The DTCWT results are at least comparable to Gabor. However, the DTCWT is faster[1]; it only requires two order-n decimated fast wavelet transforms applied to the rows of the data followed by another two for the columns.

4.2. Local and anisotropic Hurst estimation

The case where the Hurst parameter varies with respect to both space and orientation was also considered. Fig. 3 (b) shows the spatial sub-regions referred to in the following description. Let $H_0 \colon [0, 2\pi) \mapsto (0, 1)$, with $H_0^0(\cdot - \pi) \equiv H_0^0(\cdot)$ describe how the Hurst parameter in region X_0^0 varies with respect to orientation. Using the construction described in Section 2.2, the Hurst parameters in region X_0^0 are given the values $0.2, 0.4, 0.6, 0.8, 0.6, 0.4$ over angular intervals of width 30° centred at $15^\circ, 45^\circ, 75^\circ, 105^\circ, 135^\circ, 165^\circ$. For region X_0^1, the exponents are defined as $H_0^0 - 0.1$. The regions X_n^0 are $30n^\circ$ rotated versions of X_0^0. Likewise, X_n^1 are $30n^\circ$ rotated versions of X_0^1. Since the orientation centres coincide with the subbands of the Gabor and DTCWT, this is a somewhat idealised case. Fig. 4 reports the mean error with respect to the Hurst parameter. The results show that the DTCWT is at least comparable to the slower Gabor method. Similar results were found when H was held constant over space and allowed to vary over direction.

[1] In Matlab experiments on 1024×1024 images, the DTCWT computed local fractal dimension in 4.5 sec.; a spatial domain implementation of Gabor took 42.3 sec.

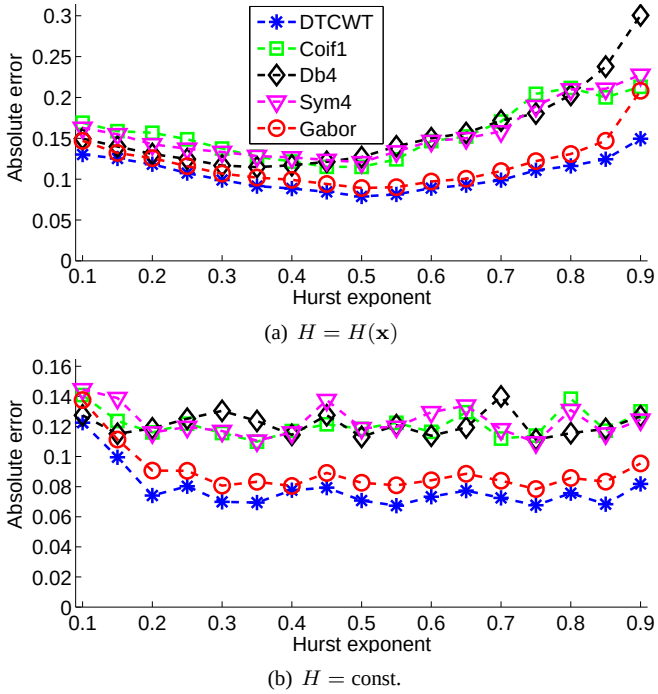

(a) $H = H(\mathbf{x})$

(b) $H = \text{const.}$

Fig. 2. Mean Hurst estimation error over 100 instances of Brownian surfaces: (a) locally varying Hurst and (b) constant Hurst.

4.3. Directionality classification

A final experiment was designed to test how the methods performed at distinguishing between a background and target object, with the same fractal dimension, but different directionality. First we generated a 180° periodic interval Θ with a width of 30° and a centre randomly chosen from the set $\{5m^\circ\}_0^{35}$. The background 256×256 was generated by synthesising an fBs with Hurst H restricted to Θ. For the target, a rectangle with size 10×14 was randomly rotated and randomly placed on the background image. An fBs, restricted to the rectangle, was generated using the same H but different randomly selected centre angle from $\{5m^\circ\}_0^{35}$. The Hurst value was estimated at each orientation using the DTCWT, Gabor, and Daubechies wavelets. The resulting 256×256, 6-element (3-element for Daubechies) vectors were then clustered with the k-means method. This was carried out 100 times for each value of H. Fig. 5 reports the percentage of pixels misclassified. Fig. 6 highlights a significant weakness of other discrete wavelet methods, namely that they cannot distinguish between angles of $\pm 45^\circ$. The results show that the DTCWT is once again comparable to the slower Gabor. Further experiments showed similar relative results for other target sizes.

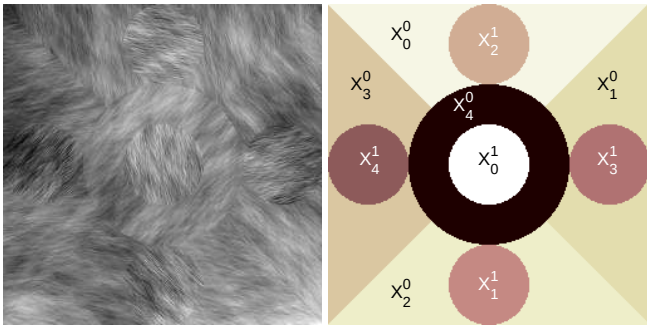

(a) Anisotropic fractional Brownian surface with $H = H(\mathbf{x}, \theta)$.

(b) Spatial sub-regions used.

Fig. 3. Fractional Brownian surface with anisotropic and locally varying fractal dimension.

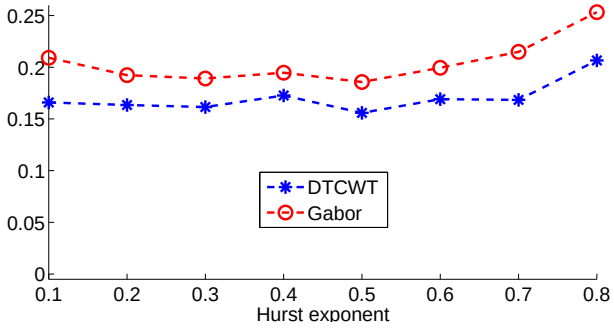

Fig. 4. Hurst error for anisotropic and locally varying H.

5. CONCLUSION

It has been shown that the DTCWT outperforms other fast wavelet transforms, and is at least as good as the slower Gabor transform, for Hurst estimation of fractional Brownian surfaces. The directional selectivity of the DTCWT can be exploited to compute Hurst parameters that vary non-trivially with direction and space. Thus, the DTCWT has been identified as a fast and effective means to compute local and anisotropic fractal dimension.

6. REFERENCES

[1] P. Soilee, "On the validty of fractal dimension measurements in image analysis," *Journal of visual communication and image representation*, vol. 7, no. 3, pp. 217–229, 1996.

[2] B. Vidakovic, O. Nicolis, and C. Garutti, "2-d wavelet-based spectra with applications in analysis of geophysical images," *ASA Joint Statistical Meetings*, 2007.

[3] B. J. Super and A. C. Bovik, "Localized measurement of image fractal dimension using gabor filters," *Journal of visual communication and image representation*, vol. 2, no. 2, pp. 114–128, 1991.

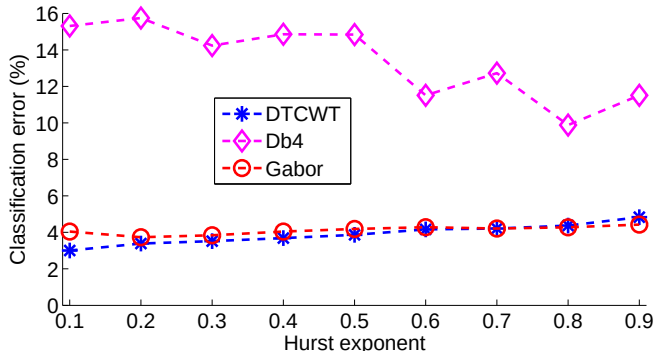

Fig. 5. Classification error (%). Target and background have same fractal dimension but different directionality.

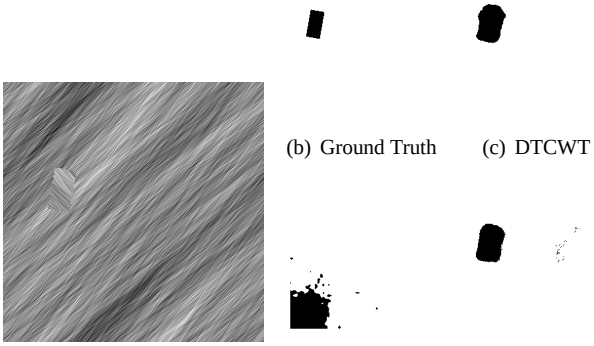

(a) Input image

(b) Ground Truth

(c) DTCWT

(d) Daubechies 4

(e) Gabor

Fig. 6. Example of classification results for anisotropic fractional Brownian surface. The target and background have a Hurst parameter of $H = 0.1$ in directions $+45°$ and $-45°$ respectively.

[4] N. G. Kingsbury, "Complex Wavelets for Shift Invariant Analysis and Filtering of Signals," *Journal of Applied and Computational Harmonic Analysis*, vol. 10, pp. 234–253, May 2001.

[5] I. W. Selesnick, R. G. Baraniuk, and N. G. Kingsbury, "The dual-tree complex wavelet transform," *IEEE Signal Processing Magazine*, vol. 22, no. 6, pp. 123–151, 2005.

[6] B. Pesquet-Popescu and J. L. Vehel, "Stochastic fractal models for image processing," *IEEE Signal Processing Magazine*, vol. 19, no. 5, pp. 48–62, 2002.

[7] B. Lashermes, S. Jaffard, and P. Abry, "Wavelet leader based multifractal analysis," *ICASSP '05*, vol. 4, pp. 161–164, 2005.

[8] L. M. Kaplan and C. C. J. Kuo, "An improved method for 2-d self-similar image synthesis," *IEEE Transactions on Image Processing*, vol. 5, no. 5, pp. 754–761, 1996.