

THE DUAL-TREE COMPLEX WAVELET TRANSFORM: A NEW EFFICIENT TOOL FOR IMAGE RESTORATION AND ENHANCEMENT

Nick Kingsbury

Signal Processing Group, Dept. of Engineering
University of Cambridge, Cambridge, CB2 1PZ, UK
e-mail: ngk@eng.cam.ac.uk

ABSTRACT

A new implementation of the Discrete Wavelet Transform is presented for applications such as image restoration and enhancement. It employs a *dual tree* of wavelet filters to obtain the real and imaginary parts of the *complex* wavelet coefficients. This introduces limited redundancy (4 : 1 for 2-dimensional signals) and allows the transform to provide approximate shift invariance and directionally selective filters (properties lacking in the traditional wavelet transform) while preserving the usual properties of perfect reconstruction and computational efficiency. We show how the dual-tree complex wavelet transform can provide a good basis for multi-resolution image denoising and de-blurring.

1 INTRODUCTION

Although the Discrete Wavelet Transform (DWT) in its maximally decimated form (Mallat's dyadic filter tree [1]) has established an impressive reputation as a tool for image compression, its use for other signal analysis and reconstruction tasks, such as image restoration and enhancement, has been hampered by two main disadvantages:

- Lack of *shift invariance*. This means that small shifts in the input signal can cause major variations in the distribution of energy between DWT coefficients at different scales.
- Poor directional selectivity for diagonal features, because the wavelet filters are separable and real.

A well-known way of providing shift invariance is to use the undecimated form of the dyadic filter tree, but this suffers from increased computation requirements and high redundancy in the output information, making subsequent processing expensive too.

We introduce the Dual-Tree Complex Wavelet Transform (DT CWT) with the following properties:

- Approximate shift invariance;
- Good selectivity and directionality in 2-dimensions (2-D) with Gabor-like filters (also for higher dimensionality);

- Perfect reconstruction (PR) using short linear-phase filters;
- Limited redundancy, independent of the number of scales, = 2 : 1 for 1-D, $2^m : 1$ for m -D;
- Efficient order- N computation – only 2^m times the simple DWT for m -D.

We propose the DT CWT as a useful front-end for many multi-dimensional signal analysis and reconstruction tasks, and demonstrate this with simple examples of edge enhancement and denoising.

2 THE DUAL FILTER TREE FOR ONE-DIMENSIONAL SIGNALS

Our work with complex wavelets for motion estimation [2] showed that complex wavelets could provide approximate shift invariance and good directionality. Unfortunately we were unable to obtain PR and good frequency characteristics using short support complex FIR filters in a single tree (eg. fig. 1 Tree *a*). This is because the complex filters, in order to be useful, should be designed to emphasise positive frequencies and reject negative frequencies (or vice-versa); and it is then not possible for the 2-band reconstruction block to have a flat overall frequency response, as required if $y \equiv x$.

However we can achieve approximate shift invariance with a *real* DWT by doubling the sampling rate at each level of the tree. For this to work, the samples must be evenly spaced. We can double all the sampling rates in Tree *a* of fig. 1 by eliminating the down-sampling by 2 after the level 1 filters, H_{0a} and H_{1a} . This is equivalent to two parallel fully-decimated trees, *a* and *b*, provided that the delays of H_{0b} and H_{1b} are one sample offset from H_{0a} and H_{1a} . We then find that, to get uniform intervals between samples from the two trees below level 1, the filters in one tree must provide delays that are half a sample different (at the filter input rate) from those in the other tree. For linear phase, this requires *odd-length* filters in one tree and *even-length* filters in the other. This is probably the most novel aspect of the dual-tree transform. Greater symmetry between the two trees




Figure 1: Dual tree of filters for the complex wavelet transform.

occurs if each tree uses odd and even filters alternately from level to level. For example in fig. 2, we show the positions of the output samples when the filters are odd and even as in fig. 1.




Figure 2: Effective sampling points of odd and even filters in fig. 1, assuming zero phase responses.

If the filters are from linear-phase PR biorthogonal sets, the odd-length highpass filters have even symmetry about their mid point, while the even-length highpass filters have odd symmetry. The impulse responses of these then look like the real and imaginary parts of a

complex wavelet! So this is how we choose to use them.

For our dual filter tree, we selected two linear-phase PR biorthogonal filter sets with odd and even lengths respectively and the additional desirable conditions that the filters should be nearly orthogonal and have good smoothness and rational coefficients. For the odd-length set, we chose (13,19)-tap filters [3], designed using the transformation of variables method. A (12,16)-tap even-length set was then designed such that the impulse responses to x_{00a} and x_{00b} were as similar as possible (in a minimum mean-squared error sense).




Figure 3: Impulse responses at level 4 of complex scaling function, $g_{0000b} + j g_{0000a}$, and wavelet, $g_{0001b} + j g_{0001a}$.




Figure 4: Frequency responses of complex wavelets at levels 1 to 4 and of the level 4 scaling function.

These filters may be implemented efficiently using ladder structures, the odd filter pair requiring 4 multiplies and 6 additions per input sample, and the even pair 7.5 multiplies and 7 additions. Reconstruction filters are formed from the inverse ladder structures and figs. 3 and 4 show their impulse and frequency responses. The analysis filters have similar responses since both filter sets are almost orthogonal. Observe the approximately gaussian shape of the response envelopes in fig. 3 and the low sidelobe levels and the low gain at negative frequencies in fig. 4. Note that much simpler PR filter sets can be used in the dual tree, but usually this is at the expense of wavelet smoothness or other of the desirable conditions.

In fig. 1 the input signal may be reconstructed exactly from the filter output samples from either tree, but it is preferable to average the two reconstruction tree outputs so that the system becomes approximately shift invariant. Fig. 5a shows the output waveforms reconstructed from just the x_{0000a} and x_{0000b} coefficients, when the input is 16 shifts of a unit step function; and fig. 5b shows the same from just the x_{0001a} and x_{0001b} coefficients. Figs. 5c and 5d show the equivalent responses for a conventional real DWT, employing the same odd-length filters. The CWT responses are clearly much more consistent with shift (shift invariant).

3 EXTENSION TO TWO DIMENSIONS

Extension to 2-D is achieved by separable filtering along columns and then rows. However, if column and row filters both suppress negative frequencies, then only the first quadrant of the 2-D signal spectrum is retained. Two adjacent quadrants of the spectrum are required to represent fully a real 2-D signal, so we also filter with complex conjugates of the row filters. This gives 4 : 1 redundancy in the transformed 2-D signal.

A normal 2-D DWT produces three bandpass subimages at each level, corresponding to lo-hi, hi-hi and hi-lo




Figure 5: 16 shifted step responses of the complex scaling functions (a) and wavelets (b) at level 4, compared with the same for the real DWT, (c) and (d).

filtering. Our 2-D CWT produces three subimages in each of spectral quadrants 1 and 2, giving six bandpass subimages of complex coefficients at each level, which are strongly oriented at angles of $\pm 15^\circ, \pm 45^\circ, \pm 75^\circ$ as shown from their Gabor-like impulse responses in fig. 6. The strong orientation occurs because complex filters can separate positive from negative frequencies vertically and horizontally. For comparison, the equivalent 3 bandpass responses for a real DWT are shown below the CWT responses, and the absence of directional selectivity in the DWT at 45° is clear.




Figure 6: 2-D impulse responses of the complex wavelets at level 4 (6 bands at angles from -75° to $+75^\circ$), and equivalent responses for a real wavelet transform (3 bands).




Figure 7: Mesh plot of a gaussian blurred edge, enhanced using the 2-D dual-tree CWT (a), and using the 2-D real DWT (b).

4 IMAGE PROCESSING APPLICATIONS

The shift invariance and directionality of the CWT may be applied to advantage in many areas of image processing, for example: denoising, restoration, texture modelling, steerable filtering, and registration / motion processing. We have space for only two examples here.

Fig. 7a shows a mesh plot of a bandlimited image edge, which has been contrast enhanced by scaling up the CWT coefficients by $1.3^{(5-level)}$ over levels 1 to 4 of the transform. The original edge was oriented at an angle of 10° to the vertical to test the shift invariance, and was bandlimited by a gaussian filter of standard deviation 1.5 pixels. The overshoot of the enhanced edge in fig. 7a is relatively constant with shift, varying from 11.2% to 14.3%. If an equivalent real DWT is used instead, as shown in fig. 7b, the overshoot varies from 7.4% to 17.1%, showing poor shift invariance.

In fig. 8 we show an example of denoising. Image (d) is the result of denoising image (a) using the DT CWT and a simple soft thresholding method which suppresses all complex wavelet coefficients of low amplitude with a raised cosine gain law: $g(x) = \frac{1}{2}(1 - \cos\{|x|/(\pi T)\})$ for $|x| < T$, and $g(x) = 1$ elsewhere. For comparison we show images (b) and (c) which were obtained using the same soft thresholding method with the real DWT in

its decimated and undecimated forms respectively. (b) shows significantly worse artifacts than (d), while (c) is very similar to (d) but requires about five times as much computation. In all cases the thresholds T were selected so as to get minimum mean-squared error from the original (clean) image. In practice, more complicated thresholding methods may be used, such as in [4] which uses Markov random fields in conjunction with an undecimated WT. It is likely that, by replacing the undecimated WT with the CWT, the effectiveness of the MRFs at coarser wavelet levels can be improved, owing to the more appropriate sampling rates of the CWT.



Figure 8: 128×128 pel portions of Lena image: (a) with white gaussian noise added to give SNR = 3.0 dB; (b) denoised with real DWT, SNR = 11.67 dB; (c) denoised with undecimated WT, SNR = 12.82 dB; (d) denoised with dual-tree CWT, SNR = 12.99 dB.

References

- [1] S G Mallat, "A theory for multiresolution signal decomposition: The wavelet representation", *IEEE Trans. PAMI*, 11(7), pp 674-693, 1989.
- [2] J F A Magarey and N G Kingsbury, "An improved motion estimation algorithm using complex wavelets", *Proc. ICIP-96*, Lausanne, pp 969-972, Sept. 1996.
- [3] D B H Tay and N G Kingsbury, "Flexible design of multidimensional perfect reconstruction FIR 2-band filters using transformations of variables", *IEEE Trans. Image Proc.*, 2(4), pp 466-480, 1993.
- [4] M Malfait and D Roose, "Wavelet-based image denoising using a Markov random field a priori model", *IEEE Trans. Image Proc.*, 6(4), pp 549-565, 1997.