By extracting information from noisy, aperiodic and intermittent signals, wavelet transforms are making an impact in medicine, astronomy, imaging and beyond

The little wave with the big future

Paul Addison

EUREKA moments are so rare in science that when you have one, you don't forget it in a hurry. My colleagues and I at Napier University in Edinburgh were fortunate enough to have one such moment a few years ago when we were studying the signals from a medical device called a pulse oximeter. Widely used in hospitals to measure the percentage of blood haemoglobin that is saturated with oxygen, the device also provides an accurate measure of a patient's heart rate.

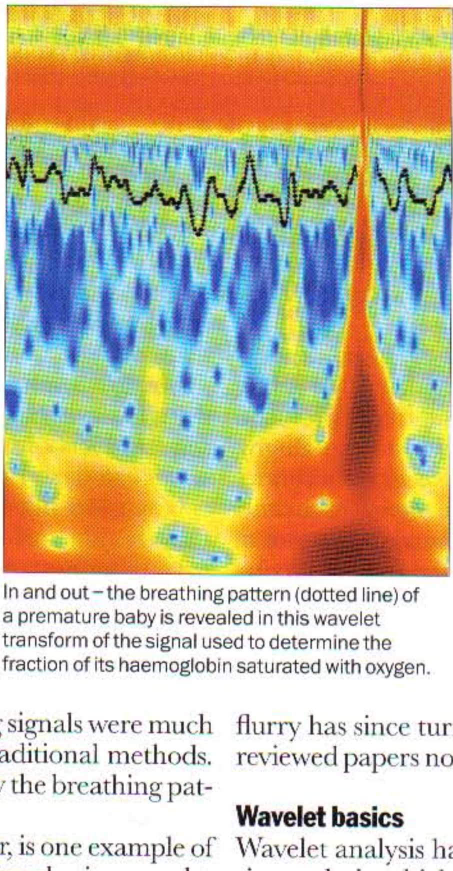

In and out – the breathing pattern (dotted line) of a premature baby is revealed in this wavelet transform of the signal used to determine the fraction of its haemoglobin saturated with oxygen.

We had been trying to use the oximeter to measure how ill patients were, based on subtle changes to these signals, which repeat regularly once every heartbeat. Using the relatively new technique of "wavelet transforms", we suddenly realized that some of the regular patterns that were appearing in our signal were not caused by the beating heart. They were, in fact, caused by the patient's breathing. Moreover, the breathing signals were much clearer than could be measured using traditional methods. We have since used this technique to study the breathing patterns of newborn babies.

This finding, which I shall return to later, is one example of the many ways in which data can be analysed using wavelet transforms. The technique is ideal for teasing out information from signals that are aperiodic, noisy, intermittent or transient. It has been used by many different researchers to study climate patterns and financial indices, to monitor heartbeats and rotating machinery, to de-noise seismic signals and astronomical images, to characterize cracks and turbulence, and to compress electronic and medical images.

Many of the ideas behind wavelet transforms have been around for a long time. Indeed, the first "wavelet" – a simple, square waveform – was developed by the mathematician Alfred Haar at the beginning of the last century. But it was not until the mid-1980s that true wavelet-transform analysis was developed by Jean Morlet and Alex Grossmann. Morlet, who was an engineer with the French oil firm Elf Aquitane, developed the technique to study seismic signals. He then teamed up with Grossmann, who worked at the CNRS Centre for Theoretical Physics in Marseille, to formalize the mathematics of the wavelet transform.

Despite their efforts, wavelet analysis initially remained confined to a small, mainly mathematical, community with only a handful of scientific papers being published each year. At the end of the 1980s, however, two further important mathematical advances were made by Ingrid Daubechies at the Courant Institute in New York and by Stéphane Mallat at the University of Pennsylvania. By the start of the 1990s the stage was set for the practical application of wavelet analysis in science and engineering. As more and more researchers spotted the potential of the technique, a flurry of papers began to appear. That flurry has since turned into a blizzard, with over 1000 peer-reviewed papers now appearing each year.

Wavelet basics

Wavelet analysis has some similarities with traditional Fourier analysis, which involves representing a signal as the sum of sines and cosines of various frequencies. A simple, square-wave signal with a magnitude $\pm a$, for example, is roughly equivalent to a series of cosines of increasing frequency: $a(\cos x) + a/3[\cos(x/3)] + a/5[\cos(x/5)] + \ldots$. Fourier analysis is particularly useful for gleaning the pertinent spectral characteristics in signals that do not change with time. It is also good at identifying the spectral components in relatively noisy signals that change in a regular fashion with time, such as the emission of radio waves from rotating neutron stars (pulsars). What Fourier analysis is not so good at, however, is investigating more complex, "non-stationary" signals, such as our heartbeat, where the component frequencies change with time. This is because it averages the key features of a

1 The secrets of the little wave

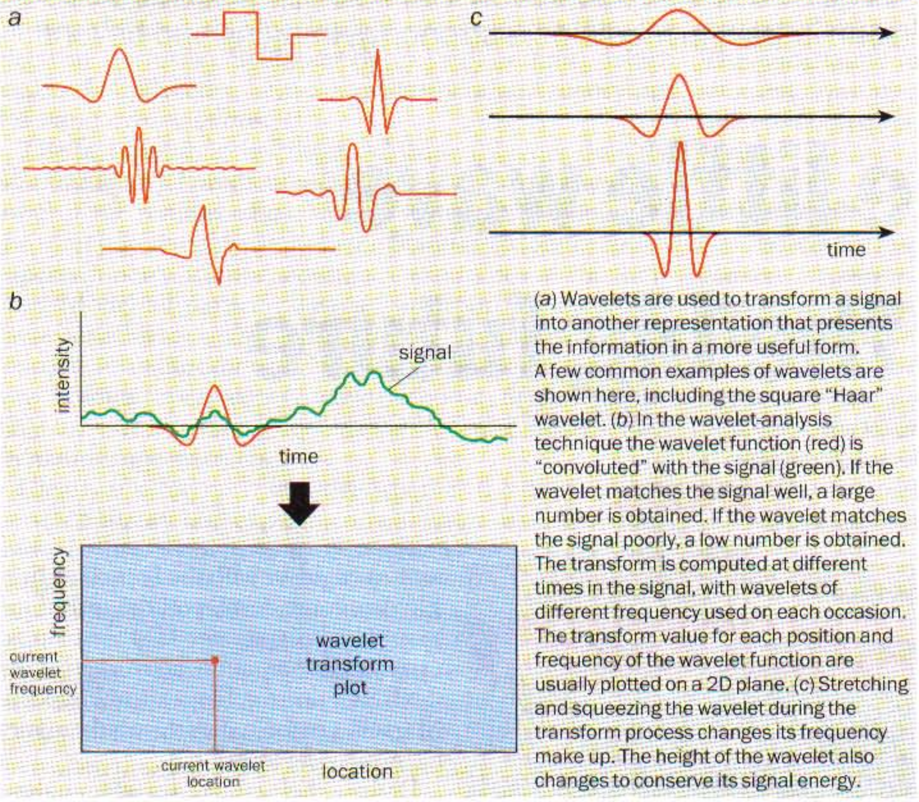

(*a*) Wavelets are used to transform a signal into another representation that presents the information in a more useful form. A few common examples of wavelets are shown here, including the square "Haar" wavelet. (*b*) In the wavelet-analysis technique the wavelet function (red) is "convoluted" with the signal (green). If the wavelet matches the signal well, a large number is obtained. If the wavelet matches the signal poorly, a low number is obtained. The transform is computed at different times in the signal, with wavelets of different frequency used on each occasion. The transform value for each position and frequency of the wavelet function are usually plotted on a 2D plane. (*c*) Stretching and squeezing the wavelet during the transform process changes its frequency make up. The height of the wavelet also changes to conserve its signal energy.

time-varying signal over the entire length of the signal, which means that fine detail is lost.

In 1946 the Hungarian-born physicist Dennis Gabor – who later won the 1971 Nobel Prize for Physics for inventing holography – found a way of getting round this problem. While working for the firm Thomson-Houston in Rugby, UK, Gabor realized that the answer was to perform a Fourier analysis of only a small segment of the signal at a time. Moving incrementally along the signal provides a series of spectral analyses, each of which is located at a particular point in time. But there is still a drawback with this "short-time Fourier transform" technique: it uses a window of fixed width. It therefore averages short-duration components and cannot capture those components that last longer than the window itself.

This is where wavelet analysis shines: it is able to take into account the scale of signal components. Wavelets are small, wave-like functions that come in all shapes and sizes (figure 1*a*). Some are smooth and look like a Mexican hat. Some are sharp functions with square edges. Others can be highly oscillatory, fractal or even complex. Wavelets are used to transform the signal under investigation into another representation that presents the signal information in a more useful form. Mathematically speaking, the wavelet transform is a convolution of the wavelet function with the signal (figure 1*b*). If the wavelet matches the shape of the signal well at a specific scale and location, then a large value of the transform is obtained. If, however, the wavelet and the signal do not correlate well then the transform value is small.

The choice of wavelet depends on the type of signal that is being investigated. Short-duration (high-frequency) features are best interrogated using narrow wavelets, while longer-lasting (low-frequency) features are more suited to wider wavelets. Changing the type of wavelet lets one zoom in on individual small-scale, high-frequency components or pan out to pick up larger-scale, low-frequency components (figure 1*c*). In practice the transform is computed at a series of points in time along the signal and for different sizes (i.e. frequencies) of wavelet at each point. The values of the transform at every time and frequency can then be plotted on a 2D "transform plane", with time down one axis and frequency along the other.

Wavelet transforms in practice

In my research, a three-minute signal will typically be transformed at 10 000 separate points in time using wavelets of 200 different sizes. The resulting two million transform points can be computed relatively quickly using a standard desktop PC with software such as Matlab or Mathcad. For real-time applications the wavelet transform is usually updated live. If complex wavelets – containing both a real and imaginary component – are used, both the phase and the modulus of the transform value can be used to study the signal. A time–frequency plot with energy density on the third axis – the wavelet scalogram – can then be obtained from the square of the modulus.

The exact number of transform points depends on how the transform has been calculated. If it is computed only at selected locations on the signal and for a restricted range of scales – so-called discrete wavelet transforms – the resolution is very poor. But if the transform is computed over a continuous range of wavelet scales and signal locations – so-called continuous wavelet transforms – the resolution can be significantly improved.

The advantage of the discrete approach is that it is faster because it uses a cunning piece of mathematical jiggery-pokery called "multi-resolution analysis". Formulated by Mallat while he was a graduate student at Pennsylvania in the late 1980s, it allows the wavelet transform to be calculated very quickly using a set of wavelets that are "orthogonal" to each other. In other words, the wavelet functions would have zero output if convolved with themselves.

The resulting discrete wavelet transform (DWT) is both "complete" and has "zero redundancy", which means that all the signal information is contained in the resulting transform and none is duplicated between transform coefficients. By converting the signal into its DWT coefficients and then removing all except those containing the most pertinent signal information, the resulting transform is much smaller in size, which provides a good way of compressing a signal. Performing an "inverse transform" on the remaining components recreates a signal that very nearly matches the original. This is the basis of image-compression algorithms such as those used in the new JPEG2000 format. Developed by the Joint Photographic Experts Group, it creates a file that is much smaller than – but looks similar to – the original.

Watermarks, water wakes and watercolours

Discrete wavelet transforms are ideal for analysing image files and other data sets that contain two or more dimensions because they can be rapidly calculated using multi-resolution analysis and because there are relatively few resulting transform coefficients. One application is in "digital watermarking", which protects the owners of digital multimedia objects against unauthorized copyright infringement. The technique involves adding information that is imperceptible to the human eye to a digital multimedia object. This added information can then be used to check if the object has been unlawfully copied at a later stage.

Many researchers – both in universities and industry – are now trying to incorporate wavelet transforms into the watermarking process by hiding secret information in the transform coefficients. When the inverse transform is calculated, the reconstructed image appears the same to the human eye, although not to a machine that can detect the watermark. Watermarking is still an evolving science; researchers are, for example, trying to find out if the watermarked image can retain its security features even if it has been blurred, rotated, added to or otherwise altered.

In another recent application of the discrete wavelet transform, Jin-Min Kuo, Kun-Shan Chen and colleagues at the Institute of Space Science at the National Central University of Taiwan have used the technique to monitor shipping traffic automatically using satellite images. They found that they could extract information about the ships' wakes with a higher degree of accuracy using wavelets than was possible with traditional methods. What the researchers did was to pass high-resolution radar images of the wakes through a computer algorithm that seeks strong correlations between the transform values obtained using wavelets of different scales. By doing so, they were able to accurately determine the angle between the two edges of the V-shaped wake in noisy images and so get a better estimate of the velocity of the vessel.

Mituo Kobayasi and Taizo Muroya at the University of Electro-communication in Tokyo have been using discrete wavelet transforms to analyse quite different types of images: fine-art paintings. They have recently proposed a way of measuring the intensities of different colours in these paintings using rectangular Haar wavelets. This has made it possible to quantify the various characteristics of paintings by van Gogh, Picasso, Monet, Klee and others, which are usually described qualitatively in terms of the "fine touch of brush" or the "broad arrangement of colours". The researchers found that their technique corresponds well with their intuitive impression of the paintings, and suggest that it could be used to automatically retrieve images from a database by applying wavelet analysis to the image file. Although not yet foolproof, their technique is the first attempt to study paintings in this way.

Continuous coverage

Discrete wavelet transforms have proved so popular because the technique is widely available within common software packages such as Matlab and Mathcad. It also fits nicely with existing digital signal-processing techniques. In practice, however, this has meant that some researchers have been using discrete wavelet transforms where the continuous wavelet transform would have been a better choice.

The continuous wavelet transform (CWT) is very different to its DWT sibling. It is highly redundant – i.e. there is a lot of repetition of information between transform coefficients – and it swallows up a lot of computer power. Another problem is that some signal energy (the integral of the square of the modulus) can be lost when the wavelet transform is calculated. Worse still, the inverse transform is not normally supplied in commercial signal-processing software toolboxes.

However, the CWT makes up for these negative traits by providing very high temporal and spectral resolution in the transform space. Subtle non-stationary or transient features – that might be missed using other methods – can therefore be observed, which makes the additional computational expense well worth the effort (figure 2). The computation can, however, be made to run quicker by expressing the continuous wavelet transform in Fourier space and performing the computation using a fast Fourier transform algorithm. Even redundancy – the big drawback with CWT – is becoming less of a problem because new approaches are being developed that can reduce the excess of information in CWT space down to a smaller subset containing just the essentials.

2 Forget Fourier

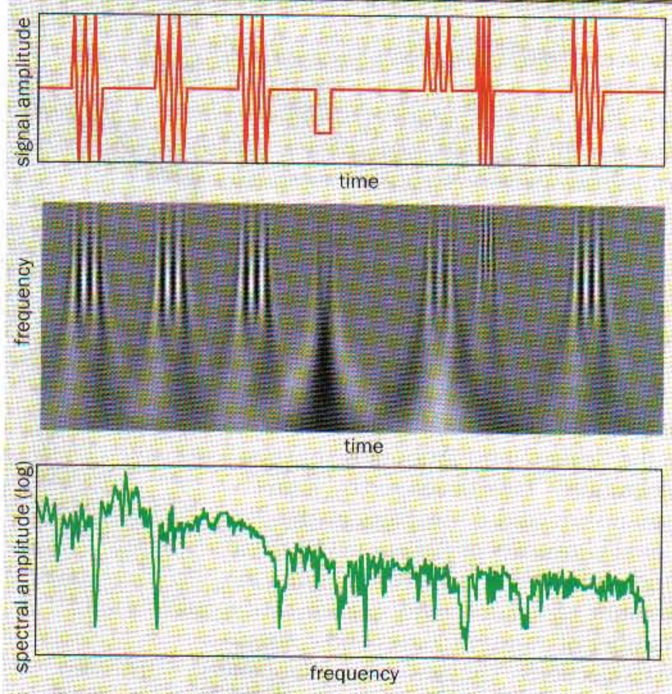

Wavelet transforms are good at picking out features in a signal that occur only intermittently. This advantage can be seen with a synthetic signal (top), which contains a number of isolated features. It has been transformed using a wavelet that is in the form of a "Mexican hat". The four identical wavegroups all have the same morphology in the wavelet-transform plot (middle), while the remaining features each have their own unique appearance. Although the correlation between features in the signal and those in the transform can be seen visually in this example, statistical techniques have to be used for the messier signals that are more likely to be encountered in real applications. In contrast, the conventional Fourier transform of the original signal (below) provides no useful information about the obvious features in the original signal.

Astronomical observations

One way of cutting the excess of information in CWT space is to use a technique called wavelet reassignment. It has been employed by Bruno Torressani and his team at the Université de Provence in Marseille, France, in an attempt to detect gravitational waves using data from interferometers. Gravitational waves – or ripples in space–time as they are poetically known – are produced whenever a large enough accelerating mass distorts space and time. Torressani's team thinks that it can recognize the signatures of the gravitational waves that are produced when binary stars coalesce by examining characteristic "ridges" in the wavelet transform. The wavelet-reassignment technique sharpens information within the time–frequency plane by shifting components nearer to their source location using phase information, which ought to allow faint signal components to be more easily identified.

Two other methods that also cut the excess of information in CWT space – the "ridge" and "modulus maxima" techniques – have been used by John Polygiannakis and co-workers at the National University of Athens in Greece to study sunspots. It is well known that the amount of solar activity – such as the number of sunspots, flares and coronal-mass ejections – changes in a regular pattern, with a dominant 11 year periodicity. However, there have been claims in the literature of a number of other, less dominant, periodicities.

To try to get to the bottom of the problem, Polygiannakis and co-workers have looked at records of the number of sunspots between 1831 and 2002. They have found that wavelet-based frequency spectra obtained from the maxima lines – the maxima with respect to time on the transform scalogram – give a clearer indication of these other periodicities. In particular, the researchers found that many of these components are harmonics of two basic periodicities – the 11 year oscillation and a two year oscillation. Their result thus supports the novel "double-dynamo" model of solar convection proposed by Elena Benevolenskaya of Stanford University. It suggests that solar activity is caused not only by tubes of magnetic field flowing radially outwards at the "convection zone" (the outer 200 000 km of the Sun) but also by a shorter latitudinal movement of the field at the photosphere, which is the shell that lies above the convection zone.

3 Finger on the pulse

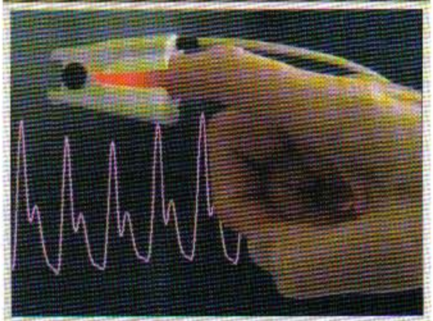

A pulse oximeter provides a continuous measure of the level of oxygen in haemoglobin in the blood by alternately shining red and infrared light into the finger. The light comes from a diode in the top arm of the clip and passes through the fingernail. The transmitted light is measured by a detector placed in the bottom arm of the clip. The different absorption of the red and infrared light allows the oxygen saturation of the blood to be determined. The measured infrared signal is shown in the background; it repeats itself every heartbeat.

Baby breathing

In my own application of wavelet transforms, my colleagues and I have been developing methods to detect the very subtle respiration signals obtained through a patient's fingertip using a pulse oximeter. Based on the different ways in which oxygenated and deoxygenated haemoglobin absorb light, the device provides an accurate measure of not only blood-oxygen levels but also heart rate. It is widely used to monitor patients who are at risk of hypoxia – blood with too little oxygen – and can be found on any hospital ward.

The device is basically a box of electronic gadgetry that includes a display connected to a light probe. The probe is attached to the patient's body, usually the finger. Red and infrared light is then shone sequentially through the body tissue via a fast-switching light-emitting diode circuit. The transmitted light is picked up by a suitably positioned photodiode. The resulting signal repeats its waveform every time the heart beats (figure 3).

We found that wavelet transforms – in the form of a complex sinusoidal function within a Gaussian window – were good at finding faint perturbations in the signal caused by the patient's breathing. In fact, we found that some of the regular patterns in the transform plots were too long to have been caused by the beating heart. It was then that the penny dropped. Our method was, in fact, picking up the rhythm of the patient's breathing – and was doing so with much more clarity than could be achieved using traditional methods.

We have recently used our technique in clinical trials with premature babies. By attaching the probe of the pulse oximeter to their feet (their fingers are just too small) we were able to track their respiration in real time and pick up individual breaths (figure 4). These little patients take between 60 and 90 breaths per minute and their pulse rate is even higher at about 150 beats per minute. Breathing rates had been measured before, but no technique is as reliable as our pulse oximeter. Indeed, premature babies push our analysis techniques to the limit because they breathe in a highly irregular fashion and wriggle a lot during data-taking, which means that the raw data signals are very noisy and non-stationary.

This work reveals the benefits of the high resolution that is available with the continuous wavelet transform. It lets us distinguish between the breathing and the pulse information – and allows us to track rapid changes of the breathing signal with time. The fact that we can measure respiration directly from the pulse-oximeter signal means that additional, more obtrusive and expensive medical equipment is not required for the task. More recently, we have found that by carrying out a second wavelet transform of the pulse-signal information in the original transform, it is even possible to pick out tiny modifications to the pulse that are caused by the baby's breathing.

Watch this wavelet space

I remember attending a Royal Society conference on wavelet transforms in 1999, when an eminent professor, whom I would rather not name, was asked if the technique would ever replace Fourier analysis. "I wouldn't like to say that wavelet analysis will supersede Fourier analysis," he replied, "but rather that Fourier methods will gradually become marginalized." Certainly the academic literature – and, perhaps more tellingly, the patent literature – is now full of wavelet-based time–frequency techniques, many of which supersede traditional Fourier spectral methods.

But if the professor's prediction is ever to come true, a widespread and deep-rooted Fourier mentality needs to be overcome. The challenge now is for researchers to appreciate the advantages of wavelet transforms over other time–frequency techniques. It would be great if wavelet transforms were taught to all physics and engineering undergraduates, but I am not holding my breath. Such a change will probably be driven by calls from industry once the new techniques have become widely used in the commercial world.

It is, however, going to take quite a few years yet for the dust to settle in wavelet space. When it does – and the merits of the

4 How premature babies breathe

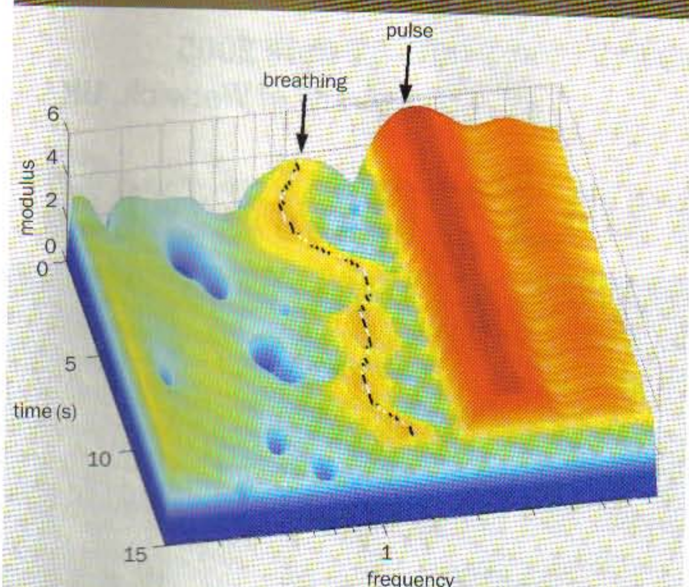

A wavelet transform of a pulse-oximeter signal from a premature baby. Data were taken over a 15 second period using a wavelet of different frequencies. The bright red band corresponds to the baby's regular pulse rate. The orange ridge corresponds to the baby's breathing rate, while the alternate black and white markings that lie along the peak of the ridge indicate the inspiration and expiration of breath, respectively. This signal is much more erratic than the pulse rate because premature babies wriggle a lot during data-taking and are prone to highly irregular breathing. Indeed, they can sometimes even stop breathing altogether for short times (not shown). This breathing ridge is determined automatically by an algorithm that searches the transform surface for maxima and decides which of these correspond to respiration. The decision criteria are based on a careful study of the link between wavelet features and patient respiration. A complex wavelet was used in the analysis, which means that the *z* axis is actually the square of the modulus of the transform.

new wavelet-based analysis methods become more obvious – I believe that the wavelet transform, in all its forms, will become the time-frequency analysis tool of choice. By allowing signal features and the frequency of their occurrence to be determined simultaneously, wavelet transforms have a lot to offer the scientific community. The little wave, I have no doubt, has a big future.

Further reading

P S Addison 2002 *The Illustrated Wavelet Transform Handbook* (Institute of Physics Publishing, Bristol)

W-Y Chen and C-H Chen 2003 Robust watermarking scheme for still images using frequency shift keying with high-variance block selection *Opt. Eng.* **42** 1826–1835.

M Kobayasi and T Muroya 2003 A spatial wavelength analysis of coarseness or fineness of color variation in painting arts *Pattern Recognition Lett.* **24** 1737–1749

J-M Kuo and K-S Chen 2003 The application of wavelets-correlator for ship wake detection in SAR images *IEEE Trans. Geoscience Remote Sensing* **41** 1506–1511

J Polygiannakis *et al.* 2003 On signal-noise decomposition of time-series using the continuous wavelet transform: application to sunspot index *Mon. Not. R. Astron. Soc.* **343** 725–734

wavelet digest newsgroup: www.wavelet.org

Paul Addison is currently on secondment to CardioDigital Ltd, Elvingston Science Centre, Edinburgh EH33 1EH, UK – a spin-off company that he co-founded to commercialize novel time-frequency methods for medical devices, e-mail p.addison@cardiodigital.com