The Continuous Wavelet Transform and Variable Resolution Time–Frequency Analysis

Amir-Homayoon Najmi and John Sadowsky

Wavelet transforms have recently emerged as a mathematical tool for multiresolution decomposition of signals. They have potential applications in many areas of signal processing that require variable time–frequency localization. The continuous wavelet transform is presented here, and its frequency resolution is derived analytically and shown to depend exclusively on one parameter that should be carefully selected in constructing a variable resolution time–frequency distribution for a given signal. Several examples of application to synthetic and real data are shown.
(Keywords: Continuous wavelets, Time–frequency analysis, Signal processing.)

TIME–FREQUENCY DECOMPOSITION OF SIGNALS AND IMPLEMENTATION OF MORLET WAVELET TRANSFORM

In most signal processing applications, we are interested in constructing a transformation that represents signal features simultaneously in time t and frequency f. Standard Fourier analysis decomposes signals into frequency components but does not provide a time history of when the frequencies actually occur. When the frequency content of a signal is time-varying, the Fourier transform $S(f)$ of a signal $s(t)$,

$$S(f) = \int_{-\infty}^{\infty} s(t) \exp(-2\pi i f t)\, dt\,, \qquad (1)$$

is incapable of capturing any local time variations and so would not be suitable for the analysis of nonstationary signals. A partial solution to this problem was provided by Gabor,[1] who in 1946 described the short-time Fourier transform in which a fixed-duration window over the time function extracts all the frequency content in that time interval. Denoting the window function by $w(t)$ and its midpoint position by τ, the Gabor transform is given by

$$S(f,\tau) = \int_{-\infty}^{\infty} s(t)\, w(t-\tau) \exp(-2\pi i f t)\, dt\,. \qquad (2)$$

Gabor actually used a Gaussian window, but in general any window function can be used. The transformation can be thought of as an expansion in terms of basis functions, which are generated by modulation, and by translation of the window $w(t)$, where f and τ are the modulation and translation parameters, respectively.

The main problem with the Gabor transform is that the fixed-duration window function is accompanied by a fixed frequency resolution. Thus, this transform allows only a fixed time–frequency resolution. Furthermore, let us define the time width and the frequency width of the window function by σ_t and σ_f, respectively:

$$\sigma_t^2 = \frac{\int_{-\infty}^{\infty} t^2 w^2(t)\, dt}{\int_{-\infty}^{\infty} w^2(t)\, dt}\,, \quad \sigma_f^2 = \frac{\int_{-\infty}^{\infty} f^2 W^2(f)\, df}{\int_{-\infty}^{\infty} W^2(f)\, df}\,, \qquad (3)$$

where $W(f)$ is the Fourier transform of the window function. It is well known (the uncertainty principle) that $\sigma_t\sigma_f \geq (2\pi)^{-1}$, with the minimum achieved for a Gaussian window function. Increasing τ simply amounts to translating the window (also known as the mother function) in time while its spread is kept fixed. Similarly, as the modulation parameter f increases, the transform translates in frequency, retaining a constant width. Thus, the resolution cells in the time–frequency plane have dimensions σ_t and σ_f, which are fixed for all τ and f. Figure 1 shows the constant-resolution cells for the short-time Fourier transform in which a sliding time window is centered at integral multiples of τ and the transforms are evaluated at bin frequencies centered at integral multiples of f_0.

The wavelet transform, on the other hand, is based on a set of basis functions formed by dilation (as opposed to modulation) and translation of a prototype mother function $\psi(t)$. The dilation of the mother function produces short-duration, high-frequency and long-duration, low-frequency functions. These basis functions are clearly better suited for representing short bursts of high-frequency or long-duration slowly varying signals. Mathematically, if $\Psi(f)$ is the Fourier transform of the mother function $\psi(t)$, then the dilated (and normalized) function $\left(1/\sqrt{a}\right)\psi\left(t/a\right)$ will have $\sqrt{a}\,\Psi(af)$ as its Fourier transform, where a is the scale parameter. Thus, a contraction in time results in an expansion in frequency and vice versa. This procedure of dilating and translating is analogous to constant Q filters in which the ratio of the root-mean-square bandwidth to center frequency of all dilated functions is a constant; i.e., each dilated wavelet will have a spread in the frequency domain equal to (σ_f/a) and a center frequency of (f_0/a), which will have a constant ratio of (σ_f/f_0) for all the dilated functions.

The continuous wavelet transform of a signal $s(t)$ is then defined by

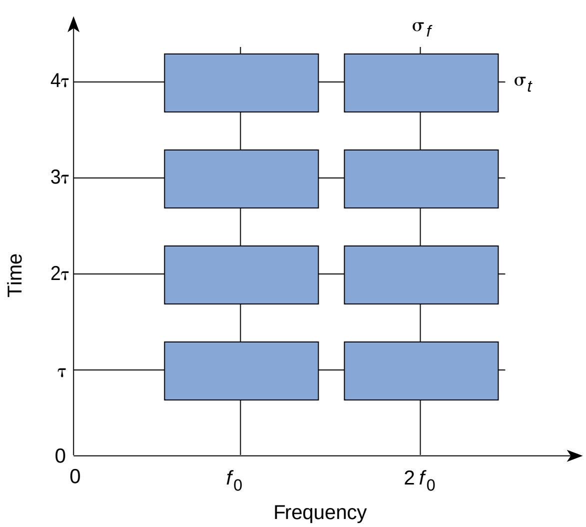

Figure 1. Resolution cells for the fixed-duration, time-window, short-time Fourier transform.

$$W(a,b) = \frac{1}{\sqrt{a}} \int_{-\infty}^{\infty} s(t)\psi\left(\frac{t-b}{a}\right) dt \; , \tag{4}$$

where $a > 0$. Mother wavelet functions of interest are bandpass filters that are oscillatory in the time domain. Thus, for large values of a, the basis function becomes a stretched version of the mother wavelet, i.e., a low-frequency function, whereas for small values of a, the basis function is a contracted version of the mother wavelet, which is a short-duration, high-frequency function. Parameter b defines a translation of the wavelet and provides for time localization.

The transformation is invertible[2] if and only if the following admissibility condition holds:

$$c_\psi \equiv \int_0^{\infty} \frac{\left|\Psi(f)\right|^2}{\left|f\right|} df < \infty \; , \tag{5}$$

which implies that the DC component $\Psi(0)$ must vanish. Thus, $\psi(t)$ is a bandpass signal that should decay sufficiently fast to provide good time resolution. The Parseval relation for the wavelet transform is

$$\int_{a=0}^{\infty} \int_{b=-\infty}^{\infty} \left|W(a,b)\right|^2 \frac{db\,da}{a^2} = c_\psi \int_{-\infty}^{\infty} \left|s(t)\right|^2 dt \; . \tag{6}$$

The orthonormal wavelet transform preserves energy between the different scales, which are parametrized by a, in the sense that

$$\int_{-\infty}^{\infty} \left|\psi(t)\right|^2 dt = \int_{-\infty}^{\infty} \left|\frac{1}{|a|}\psi\left(\frac{t-b}{a}\right)\right|^2 dt \; . \tag{7}$$

The Morlet wavelet[3] is a good example of a mother function for the construction of the continuous wavelet transform.[4] It is defined by

$$\psi(t) = \left[\exp(-2i\pi f_0 t) - \exp\left(-2\pi^2 f_0^2 \sigma^2\right)\right] \exp\left(-\frac{t^2}{2\sigma^2}\right) . \tag{8}$$

Its Fourier transform is

$$\Psi(f) = \sqrt{2\pi\sigma^2}\left\{\exp\left[-2\pi^2\sigma^2(f-f_0)^2\right] - \exp\left(-2\pi^2\sigma^2 f^2\right)\exp\left(-2\pi^2\sigma^2 f_0^2\right)\right\} , \tag{9}$$

which satisfies the admissibility condition $\Psi(0) = 0$. With this choice of mother function, the continuous wavelet transform upon time discretization $t = n\Delta T$, where ΔT is the sampling time in seconds, becomes

$$W(a,b) = \frac{\Delta T}{\sigma\sqrt{2\pi a}} \sum_{n=-\infty}^{\infty} s(n) \exp\left[-\frac{(n\Delta T - b)^2}{2a^2\sigma^2}\right] \times \exp\left(-2i\pi f_0 \frac{n\Delta T - b}{a}\right). \quad (10)$$

We choose a number of octaves and voices within each octave. Denoting the product of the number of octaves and the number of voices by M, we use $a_k = 2^{k/V}$, where V is the number of voices and $1 \le k \le M$. We further discretize the translation parameter b and finally obtain

$$W(k,l) = \frac{\Delta T}{\sigma\sqrt{2\pi 2^{k/M}}} \sum_{n=-\infty}^{\infty} s(n) \exp\left[-\frac{(n-l)^2}{2^{1+(2k/M)} f_s^2 \sigma^2}\right] \times \exp\left[-2i\pi \frac{f_0}{f_s} 2^{-k/M}(n-l)\right], \quad (11)$$

where $1 \le k \le M$, $-\infty \le l \le \infty$, and f_s is the sampling frequency in hertz. Although this equation represents a discretization of the dilation and the translation parameters, it is not the discrete wavelet transform.[5] The latter is defined by a fixed set of coefficients that represent the bandpass signal at all scales (see the boxed insert on this page).

The dilation operation for the wavelet transform also divides the time–frequency plane into resolution cells, analogous to the division of the time–frequency plane by the operation of modulation for the short-time Fourier. The difference in the two sets of operations is that whereas the resolution cells have fixed values for the short-time Fourier transform (fixed duration of the time window), the resolution cells for the wavelet transform have variable lengths, depending on the scale parameter a. The relationship between a and f is obtained by calculating an exact expression for the response to a sinusoid. Letting $s(t) = \exp(2i\pi ft)$ be the signal, and then using Eqs. 4 and 8,

$$W(a,b) = \sqrt{a}\exp(2i\pi bf)\exp\left[-2\pi^2 a^2 \sigma^2 \left(f - \frac{f_0}{a}\right)^2\right]. \quad (12)$$

This response is peaked at $a = f_0/f$, and so frequency and scale are related by $f = f_0/a$. The frequency resolution cells are, therefore, centered at f_0/a for each scale value a. Now, if the original mother wavelet has a frequency width and a time width denoted by σ_f and σ_t, respectively, the dilated wavelet at scale a has a frequency width and time width equal to σ_f/a and $a\sigma_t$, respectively. Thus, the wavelet transform provides a variable resolution in the time–frequency plane, as shown in Fig. 2.

Furthermore, the preceding response indicates that the spread in the frequency domain for the dilated

DISCRETE WAVELET TRANSFORM VS. THE CONTINUOUS WAVELET TRANSFORM

The signal transform computed in the article is the continuous wavelet transform (CWT), even though a discrete-time formulation is used. This formulation differs, however, from the discrete wavelet transform (DWT). In the article, time has been discretized to correspond with the sampling of the physical signals, and the summation form follows directly from the Riemann sum approximation to the integral in the CWT definition. The CWT displays the distribution of signal amplitude and phase in two variables, time and scale. The DWT, on the other hand, plays a role similar to Fourier series coefficients. One can create a set of basis functions from a mother wavelet by rescaling the wavelet over octaves (powers of 2) and translating the wavelet over discrete, scale-dependent time steps. The DWT computes the coefficients of a signal with respect to this basis. Thus, in effect, it acts as a rectangular array of filters centered about specific times and scales (frequencies). Although these coefficients can provide an approximation to the CWT, in a way similar to the approximation of a Fourier transform using a discrete Fourier series, they should more properly be multiplied by the associated basis functions to obtain an approximation to the CWT. The DWT is an important structure to study in its own right, both in the theory of wavelets and in applications such as signal and image compression, subband coding, and pattern recognition. Many features of the DWT make it an attractive tool for signal processing, such as the existence of Mallat's linear-time algorithm to compute the DWT, which is more efficient than the $N\log(N)$ fast Fourier transform algorithm, and a generalization of representations in bases to representations in overdetermined sets of functions, called frames, in which uniqueness of representation is sacrificed in favor of more localized influence of coefficients.

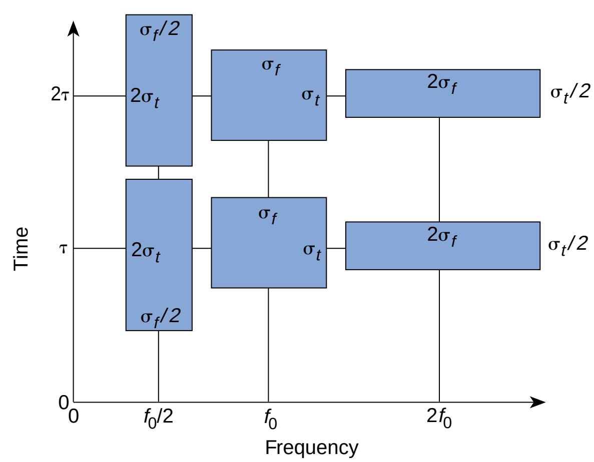

Figure 2. Variable length resolution cells for the continuous wavelet transform.

Morlet wavelet is equal to $1/\pi a\sigma$, and so it is inversely proportional to σ, where σ is the time width of the undilated (scale 1) Morlet mother wavelet. Thus, for the higher frequencies (smaller scales), the frequency resolution can be improved by choosing larger values for σ, and for lower frequencies (larger scales), frequency resolution is improved by choosing smaller values for σ. The frequency f_0 must be chosen so as to ensure the bandpass property of the mother function, which in practice can be achieved by choosing a value equal to one-half the sampling frequency of the input discrete time data. Equation 11 has a simple and fast convolutional implementation using the fast Fourier transform. Our implementation on an HP-735 workstation computes and displays a 6400 by 120 pixel image for a 6400-point signal with 10 octaves and 12 voices per octave, without any subsampling, in approximately 20 s.

WAVELET TIME–FREQUENCY DISTRIBUTIONS AND COMPARISON WITH WIGNER DISTRIBUTION

The first example data set is shown in Fig. 3. It includes a linearly chirped frequency ranging between 2.5 and 15 Hz, added to two short-duration sinusoids with frequencies of 25 and 52 Hz, respectively. Figure 4 shows the magnitude of the continuous wavelet time–frequency transform for $\sigma = 1$(measured in units of the sampling time). The y axis shows time in seconds, and the x axis shows the octave numbers. There are five octaves separated by vertical white lines, and 12 voices per octave. The minimum and maximum frequencies within each octave are shown on the top horizontal axis. Clearly, there is excellent time resolution but very poor frequency resolution for the chirp and the sinusoids. Figure 5 shows the transform computed for $\sigma = 7.5$, which produces excellent frequency resolution for all the signal components while also providing very good time resolution. Frequency resolution improves with decreasing scale (increasing frequency) in agreement with our frequency spread calculation for the Morlet wavelet. Although this value of σ appears to be an excellent compromise for most of the time–frequency plane of this specific signal, it does not allow the determination of the lower end of the chirp's frequency range in the fourth and fifth octaves. This problem can be remedied simply by varying σ by some prescribed rule or in a data adaptive manner. Figure 6 shows the output when an exponentially decreasing set of values

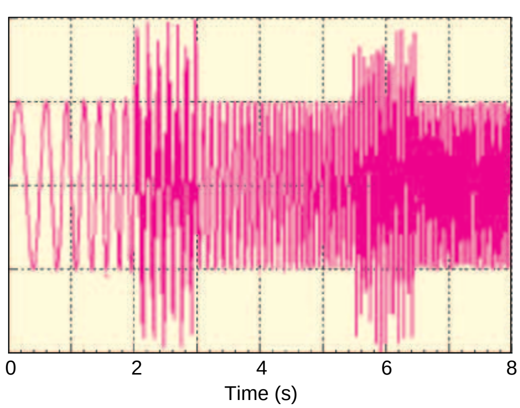

Figure 3. Time series of two short single-frequency signals together with a linearly chirped signal.

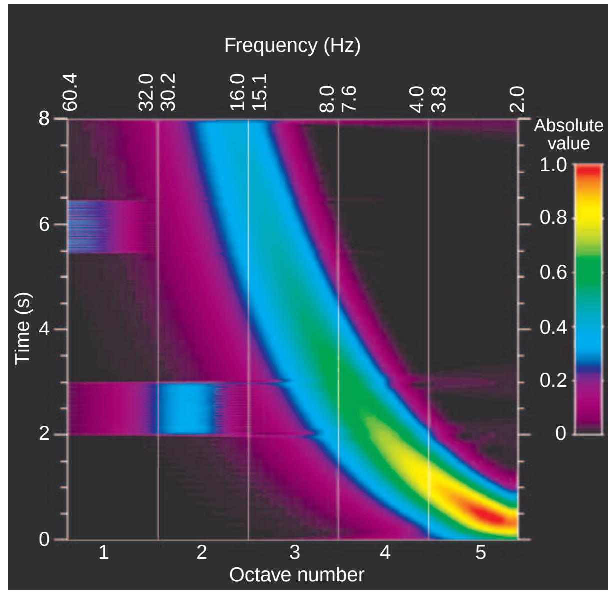

Figure 4. The continuous wavelet transform of the example in Fig. 3 for a very narrow mother wavelet.

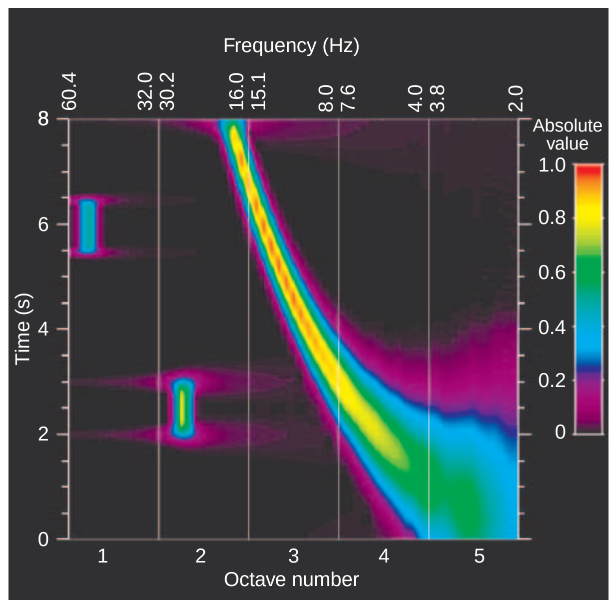

Figure 5. The continuous wavelet transform of the example in Fig. 3 for a wider mother wavelet.

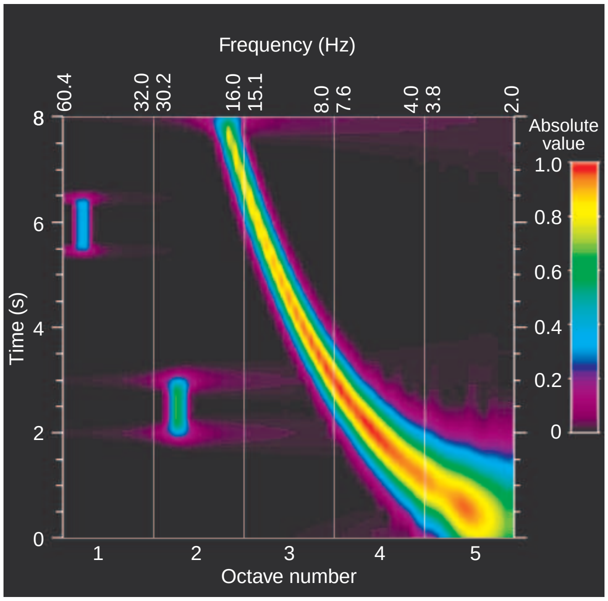

Figure 6. The continuous wavelet transform of the example in Fig. 3 for an exponentially varying set of the width parameter.

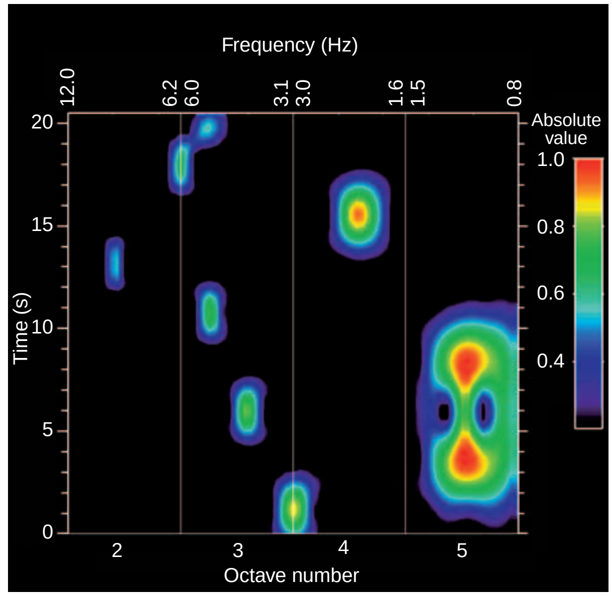

Figure 8. The continuous wavelet transform of the example in Fig. 7.

for σ between 3.5 and 10.0 is used. The minimum frequency of the chirp, in the fifth octave, is now quite apparent and measurable.

The second example data set is that of a hopped frequency pattern. The data and its spectrum are shown in Fig. 7. The wideband spectrum of these communication signals makes the task of detecting the individual frequencies impossible without performing a joint time–frequency analysis. Figure 8 shows the continuous wavelet transform of this data with an exponentially decreasing set of values for σ between 3.0 and 10.0. The frequencies are well resolved in time and can be easily detected to have the values 3, 1, 4, 5, 9, 2, 6, and 5 Hz.

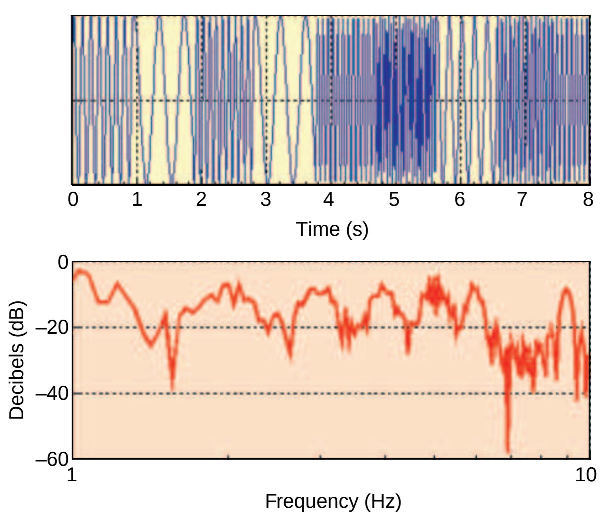

Figure 7. The time series of a hopped-frequency signal (top) and its Fourier spectrum (bottom).

Figure 9 shows a bowhead whale sound. The data set has been analyzed widely using the Wigner distribution (see the boxed insert on the next page) in which the complicated chirp signals are well resolved both in time and frequency. The main difficulty with this nonlinear procedure is to reduce the interference terms between the various components of the signal. The continuous wavelet transform of this same signal over two octaves spanning the frequencies between 78 and 300 Hz is shown in Fig. 10. The values of σ ranged between 12.0 and 15.0 in this case. The continuous wavelet transform seems to have resolved the components quite well both in time and frequency. Although the overall resolution is only slightly worse than the Wigner distribution, there are no cross terms to be suppressed here, and the computation time is reduced by at least a factor of 4.

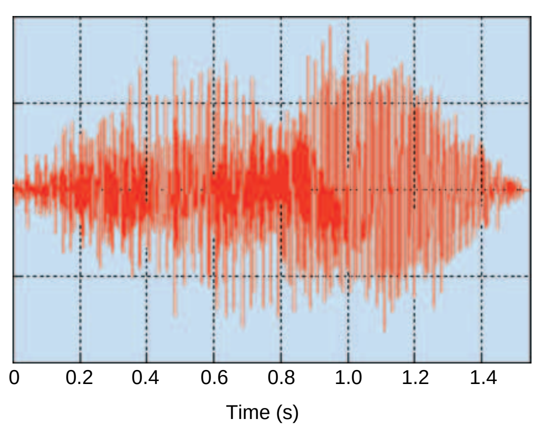

Figure 9. The time series of a bowhead whale sound.

THE WIGNER DISTRIBUTION

The Wigner distribution for a real waveform $s(t)$ is defined by

$$P_{wigner}(t,f) = \int_{-\infty}^{\infty} \tilde{s}\left(t+\frac{\tau}{2}\right)\tilde{s}^*\left(t+\frac{\tau}{2}\right)\exp(-2i\pi f\tau)\,d\tau\ ,$$

where $\tilde{s}(t)$ is the analytic signal whose real part is the original waveform and whose imaginary part is the Hilbert transform of the real part,[6] t is time, and f is frequency. An example of this distribution on the bowhead whale sound data is shown in the figure in which nonlinear effects appear in between the well-resolved chirps.

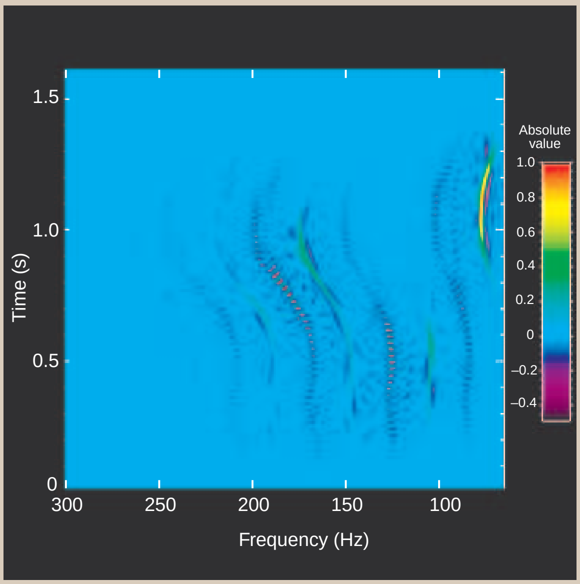

Wigner distribution of the bowhead whale sound data.

CONCLUSION

The continuous wavelet transform, using a Morlet mother wavelet with frequency parameter f_0, and bandwidth parameter σ, is a good signal analysis tool for displaying and estimating features in wideband signals with time-dependent frequency and scale characteristics. With judicious selection of the wavelet parameters, the continuous wavelet transform performs well compared with other methods used for time–frequency representation, such as the Wigner distribution. The continuous wavelet transform is, however, less susceptible to false features due to cross product terms because it is a linear function of the signal. As a processing option in a signal analysis toolbox, it is proving to be quite valuable for qualitative and quantitative investigations of signal features, for example, as a preliminary step in the design of feature detection algorithms.

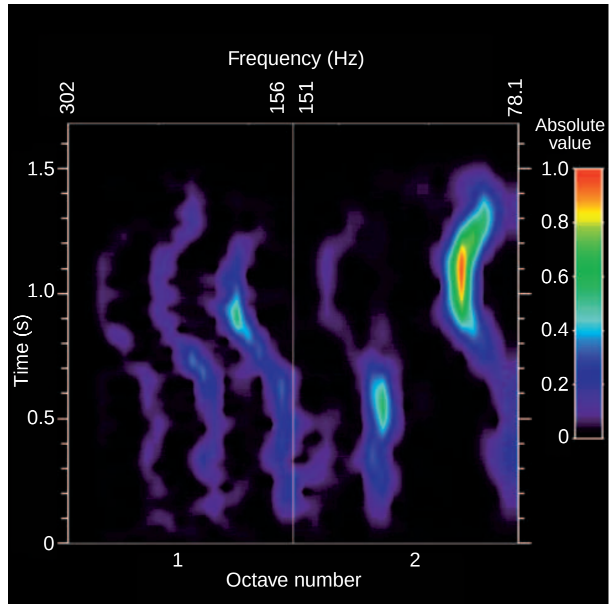

Figure 10. The continuous wavelet transform of the bowhead whale sound.

In this article, we derived analytic expressions relating the scale at which features occur in the continuous wavelet transform to the associated, time-localized frequency and the frequency parameter of the Morlet wavelet. We also derived an analytic expression for the spread in scale and, hence, frequency in the transform as a function of the bandwidth parameter. Thus, one can derive quantitative information about temporal–local behavior of the signal from features in the surface and, in an iterative process, refine the mother wavelet to present the features of interest optimally.

REFERENCES

[1] Gabor, D., "Theory of Communication," *J. IEE* **93**, 429–457 (1946).

[2] Daubechies, I., *Ten Lectures on Wavelets*, SIAM Press, Philadelphia (1992).

[3] Meyer, Y., *Ondelettes et Operateurs I: Ondelettes*, Hermann Editeurs des Sciences et des Arts, Paris (1990).

[4] Sadowsky, J., "The Continuous Wavelet Transform: A Tool for Signal Analysis and Understanding," *Johns Hopkins APL Tech. Dig.* **15**(4), 306–318 (1994).

[5] Mallat, S., "Multiresolution Approximations and Wavelet Orthogonal Bases of $L^2(R)$," *Trans. Am. Math. Soc.* **315**(1), 69–87 (1989).

[6] Najmi, A.-H., "The Wigner Distribution: A Time-Frequency Analysis Tool," *Johns Hopkins APL Tech. Dig.* **15**(4), 298–305 (1994).

THE AUTHORS

AMIR-HOMAYOON NAJMI received B.A. and M.A. degrees in mathematics from the University of Cambridge (1977) where he completed the Mathematical Tripos, and the D.Phil. degree in theoretical physics from the Astrophysics Department of the University of Oxford (1982). He was a Fulbright scholar at the Relativity Center of the University of Texas at Austin (1977–1978), Research Associate and Instructor at the University of Utah (1982–1985), and Research Physicist at the Shell Oil Company Bellaire Geophysical Research Center in Houston (1985–1990). Dr. Najmi joined APL in December 1990. His e-mail address is Najmi@jhuapl.edu.

JOHN SADOWSKY is the assistant supervisor of the Advanced Signal and Information Processing Group of the M. S. Eisenhower Research and Technology Development Center at APL. He earned a B.A. in mathematics from The Johns Hopkins University in 1971 and a Ph.D. in mathematics from the University of Maryland in 1980. His work includes advanced transforms and algorithms for signal processing and, in collaboration with researchers in The Johns Hopkins University Schools of Engineering, Arts and Sciences, and Medicine, he coordinates the sensory engineering program. Before joining APL in August 1989, he was a principal scientist and supervisor at the Systems Engineering and Development Corporation (now a division of the Essex Corporation), where he led programs to develop advanced transforms, algorithms, and architectures for processing systems for electronic and communication signals intelligence. His e-mail address is John.Sadowsky@jhuapl.edu.