WAVELET ANALYSIS OF SIGNALS AND IMAGES, A GRAND TOUR

J.-P. Antoine, Institut de Physique Théorique, Université Catholique de Louvain
B - 1348 Louvain-la-Neuve, Belgium

ABSTRACT
We review the general properties of the wavelet transform, both in its continuous and its discrete versions, in one or two dimensions, and we describe some of its applications in signal and image processing. We also consider its extension to higher dimensions and to the space-time context, for the analysis of moving objects.

1. MOTIVATION: WHAT IS WAVELET ANALYSIS?

Wavelet analysis is a particular time- or space-scale representation of signals which has found a wide range of applications in physics, signal processing and applied mathematics in the last few years. In order to get a feeling for it and to understand its success, let us consider first the case of one-dimensional signals.

It is a fact that most real life signals are nonstationary. They often contain transient components, sometimes very significant physically, and mostly cover a wide range of frequencies. In addition, there is frequently a direct correlation between the characteristic frequency of a given segment of the signal and the time duration of that segment. Low frequency pieces tend to last a long interval, whereas high frequencies occur in general for a short moment only. Human speech signals are typical in this respect. Vowels have a relatively low mean frequency and last quite long, whereas consonants contain a wide spectrum, up to very high frequencies, especially in the attack, but they are very short.

Clearly standard Fourier analysis is inadequate for treating such signals, since it looses all information about the time localization of a given frequency component. In addition, it is very uneconomical. If a segment of the signal is almost flat, i.e., uninteresting, one still has to sum an infinite series for reproducing it. Worse yet, Fourier analysis is highly unstable with respect to perturbation, because of its global character. For instance, if one adds an extra term, with a very small amplitude, to a linear superposition of sine waves, the signal will barely be modified, but the Fourier spectrum will be completely perturbed. This does not happen if the signal is represented in terms of *localized* components.

Therefore, signal analysts turn to *time-frequency* (TF) representations. The idea is that one needs *two* parameters. One, called a, characterizes the frequency, the other one, b, indicates the position in the signal. This concept of a TF representation is in fact quite old and familiar. The most obvious example is simply a musical score!

If one requires, in addition, the transform to be *linear*, a general TF transform will take the form:

$$s(x) \mapsto S(b,a) = \int_{-\infty}^{\infty} \overline{\psi_{ba}(x)}\, s(x)\, dx, \tag{1.1}$$

where s is the signal and ψ_{ba} the analyzing function (we denote the time variable by x, in view of the extension to higher dimensions). Within this class, two TF transforms stand out as particularly simple and efficient, the windowed or short time Fourier transform (STFT) and the wavelet transform (WT). For both of them, the analyzing function ψ_{ba} is obtained by a group action on a basic (or mother) function ψ, only the group differs. The essential difference between the two is in the way the frequency parameter

a is introduced. For the WT, one takes:

$$\psi_{ba}(x) = \frac{1}{\sqrt{a}}\,\psi\left(\frac{x-b}{a}\right). \tag{1.2}$$

The action of a on the function ψ is a dilation ($a > 1$) or a contraction ($a < 1$): The shape of the function is unchanged, it is simply spread out or squeezed. As for b, it is simply a translation. By contrast, the STFT takes for ψ_{ba} the function $\psi_{ba}(x) = e^{ix/a}\,\psi(x-b)$. This means that the a-dependence is a modulation ($1/a \sim$ frequency); the window has constant width, but the lower a, the larger the number of oscillations in the window ψ.

Actually one should distinguish between *two* radically different versions of the wavelet transform, the continuous WT (CWT) and the discrete WT (DWT). The CWT plays the same rôle as the Fourier transform and is mostly used for analysis and feature detection in signals, whereas the DWT is the analogue of the Discrete Fourier Transform (see for instance (Rioul, 91)) and is more appropriate for data compression and signal reconstruction. Somewhat schematically, one may say that the CWT is more natural to the physicist, while the DWT is more congenial to the signal analyst and the numericist. In these lectures, we will review the CWT, both from the theoretical side and in its practical implementation. We will proceed in two steps. First the one-dimensional case, then the extension to higher dimensions, including space and time-dependent wavelets appropriate for motion analysis.

Let us begin by a few introductory remarks, to indicate where we are aiming. Both for the continuous and the discrete versions, the WT is given by the basic transformation formula, which reads, according to (1.1) and (1.2):

$$S(b,a) = a^{-1/2}\int_{-\infty}^{\infty} \overline{\psi\left(a^{-1}(x-b)\right)}\,s(x)\,dx, \tag{1.3}$$

where $a > 0$ is a scale parameter and $b \in \mathbb{R}$ a translation parameter. In the relation (1.3), s is a finite energy signal and the function ψ, the analyzing wavelet, is assumed to be well localized *both* in the time domain and in the frequency domain. In addition ψ must satisfy an admissibility condition, which guarantees the invertibility of the WT, and in most cases, may be reduced to the requirement that ψ has zero mean. Therefore, since ψ_{ba} acts like a filter (convolution), the WT $s \mapsto S$ provides a *local (i.e., bandpass) filtering*, both in space (b) and in scale (a). The transform $S(b,a)$ is nonnegligible only when the wavelet ψ_{ba} matches the signal, that is, the WT selects the part of the signal, if any, that lives around the time b and the scale a. In addition, combining this feature with the localization properties of $\psi(t)$ and its Fourier transform $\widehat{\psi}(\omega)$, we see that the WT works at constant relative bandwidth, $\Delta\omega/\omega =$ constant. Thus it is more efficient at high frequency, i.e., small scales, in particular for the detection of *singularities* in the signal.

In addition, the transformation $W_\psi : s(x) \mapsto S(b,a)$ may be inverted exactly, which yields a reconstruction formula :

$$s(x) \simeq \int_{-\infty}^{\infty} db \int_0^{\infty} \frac{da}{a^2}\;\psi_{ba}(x)\,S(b,a), \tag{1.4}$$

This means that the WT provides a decomposition of the signal as a linear superposition of the wavelets ψ_{ba} with coefficients $S(b,a)$ — the analogy with Fourier integrals or series is clear.

We will discuss the mathematical properties of the CWT and its actual implementation for applications in Sections 2 and 3, respectively.

All this concerns the continuous WT. But, in practice, for the actual computations, the transform must be *discretized*, by restricting the parameters a and b in (1.3) to the points of a lattice, typically a dyadic one:

$$c_{j,k} = 2^{-j/2}\int_{-\infty}^{\infty} \overline{\psi(2^{-j}x-k)}\,s(x)\,dx,\ j,k \in \mathbb{Z}. \tag{1.5}$$

Then the reconstruction formula (1.4) becomes simply

$$s(x) = \sum_{j,k \in \mathbb{Z}} c_{j,k} \, \widetilde{\psi_{j,k}}(x), \tag{1.6}$$

where the function $\widetilde{\psi_{j,k}}$ may be explicitly constructed from $\psi_{j,k}$. In this way, one arrives at the theory of *frames* or nonorthogonal expansions (Daubechies, 86) (Daubechies, 92), which offer a good substitute to orthonormal bases. Very general functions ψ satisfying the admissibility condition described above will yield a good frame, but not an orthonormal basis, since the functions $\{\psi_{j,k}(x) \equiv 2^{j/2}\psi(2^j x - k), j, k \in \mathbb{Z}\}$ are in general not orthogonal to each other! This problem of discretization will be evoked in Section 4.

Yet orthonormal bases of wavelets can be constructed, but by a totally different approach, based on the concept of *multiresolution analysis*. We emphasize that the discretized version of the CWT just described is totally different in spirit and method from the genuine DWT. Although this is not the main topic of these lectures, we will provide a rapid glimpse of the DWT in Section 5. The full story may be found in (Daubechies, 92), for instance.

As we will see below, the CWT is extremely efficient and has a number of nice mathematical properties. Where does all that come from? The answer, as so often in physics, lies in group theory. In a nutshell, wavelets are 'coherent states' associated to the natural representation in the space of finite energy signals of the affine group of the line (the $ax + b$ group, consisting of dilations and translations). Now a general theory is available for constructing similar families of vectors from representations of various types of groups (Ali, 95) (Ali, 99). Thus a general pattern emerges (see Section 6), that will yield wavelets in two or more dimensions, wavelets on the sphere, or space-time wavelets designed for the analysis of moving objects. To give an example, in 2-D, the analog of (1.3) is

$$S(\boldsymbol{b}, a, \theta) \;=\; a^{-1} \int_{\mathbb{R}^2} \overline{\psi(a^{-1} r_{-\theta}(\boldsymbol{x} - \boldsymbol{b}))} s(\boldsymbol{x}) \, d^2\boldsymbol{x}. \tag{1.7}$$

where, as before, $s(\boldsymbol{x})$ is the 2-D signal (the image to be analyzed) and $\psi(\boldsymbol{x})$ is the analyzing wavelet, which is translated by the vector $\boldsymbol{b} \in \mathbb{R}^2$, dilated by $a > 0$ and rotated by an angle θ ($r_{-\theta}$ is the rotation operator). Since the wavelet ψ is required to have zero mean, we have again a bandpass filtering effect, i.e., the analysis is *local* in all four parameters $a, \theta, \boldsymbol{b}$ and, here too, it is particularly efficient at detecting discontinuities or specific sharp features in images. We will discuss the 2-D case in detail in Section 7 and the other generalizations in Section 8.

Finally a word about references. The literature on wavelet analysis is growing exponentially, so that some guidance may be helpful. As a first contact, the introductory articles (Rioul, 91) and (Heil, 89) may be a good suggestion, followed by the popular book (Hubbard, 98) or the elementary book (Meyer, 93). For a survey of the various applications, and a good glimpse of the chronological evolution, there is still no better place to look than the proceedings of the three large wavelet conferences, Marseille 1987 (Combes, 90), Marseille 1989 (Meyer, 91) and Toulouse 1992 (Meyer, 93). Finally a systematic study requires a textbook. Among the increasing number of books and special issues of journals appearing on the market, we note in particular the volumes (Daubechies, 92), (Chui, 92), (Kaiser, 94) and (Mallat, 99), the collection of review articles in (Ruskai, 92) and two special issues of IEEE journals (IEEE, 92) (IEEE, 96).

2. THE ONE-DIMENSIONAL CWT

2.1. Basic definitions

As is clear from (1.3), the CWT is a projection of the signal, in the L^2 sense, onto the family $\{\psi_{ba}, a > 0, b \in \mathbb{R}\}$ generated from the single function ψ by translation and dilation:

$$
\begin{aligned}
S(b,a) &= \langle \psi_{ba} | s \rangle \\
&= a^{-1/2} \int_{-\infty}^{\infty} \overline{\psi\left(a^{-1}(x-b)\right)}\, s(x)\, dx, \qquad (2.1) \\
&= a^{1/2} \int_{-\infty}^{\infty} \overline{\widehat{\psi}(a\omega)} \widehat{s}(\omega) e^{ib\omega}\, d\omega, \qquad (2.2)
\end{aligned}
$$

where the hat denotes a Fourier transform. Thus the transform $S(b,a)$ lives in the half-plane $\mathbb{R}^2_+ = \{a > 0, b \in \mathbb{R}\}$. The analyzing wavelet ψ satisfies a number of conditions.

(i) For the formalism to make sense, $\psi(x)$, hence also $\widehat{\psi}(\omega)$, should be square integrable: $\psi \in L^2(\mathbb{R})$.

(ii) ψ must be *admissible*, that is, the following integral must converge:

$$
c_\psi \equiv 2\pi \int_{-\infty}^{\infty} |\widehat{\psi}(\omega)|^2 \, \frac{d\omega}{|\omega|} < \infty. \qquad (2.3)
$$

This condition implies (and for ψ regular enough, is equivalent to)

$$
\widehat{\psi}(0) = 0, \qquad (2.4)
$$

which in turn is equivalent to the *zero mean* condition (thus ψ must be oscillating):

$$
\int_{-\infty}^{\infty} \psi(x)\, dx = 0. \qquad (2.5)
$$

(iii) In order to get an efficient transform (good bandpass filtering, both in space *and* in frequency), $\psi(x)$ and $\widehat{\psi}(\omega)$ should be both well localized (it suffices to require that ψ be also integrable ($\psi \in L^1 \cap L^2$), but in practice a better localization will be useful).

(iv) In addition to (ii), ψ *may* be required to have a certain number of *vanishing moments:*

$$
\int_{-\infty}^{\infty} x^n\, \psi(x)\, dx = 0, \ n = 0, 1, \ldots N \qquad (2.6)
$$

(this property improves the efficiency of ψ at detecting singularities in the signal, since it is blind to polynomials up to order N).

(v) Finally, ψ is often required to be *progressive*, that is, $\widehat{\psi}(\omega)$ is real and $\widehat{\psi}(\omega) = 0$ for $\omega \leq 0$ (such a ψ is also called an analytic signal or a Hardy function).

We want to stress the difference between these requirements: whereas (i) and (ii) are essential for having a CWT at all, (iii) ensures a clear interpretation, and is thus practically necessary, finally (iv) and (v) make life easier, but are not compulsory.

As we shall see in the following sections, a wavelet ψ that satisfies these requirements generates by (2.1) a transform $W_\psi : s(x) \mapsto S(b,a)$ that yields a good analysis of the signal, and allows an efficient reconstruction $S(b,a) \mapsto s(x)$ of the signal from its transform.

2.2. Two common wavelets

The two analyzing wavelets most used in practice are the Mexican hat and the Morlet wavelet.

The *Mexican hat* is simply the second derivative of the Gaussian,

$$
\psi_H(x) = (1 - x^2)\, e^{-x^2/2}, \ \widehat{\psi}_H(\omega) = \omega^2\, e^{-\omega^2}. \qquad (2.7)
$$

This is an admissible, real wavelet, with two vanishing moments ($n = 0, 1$).

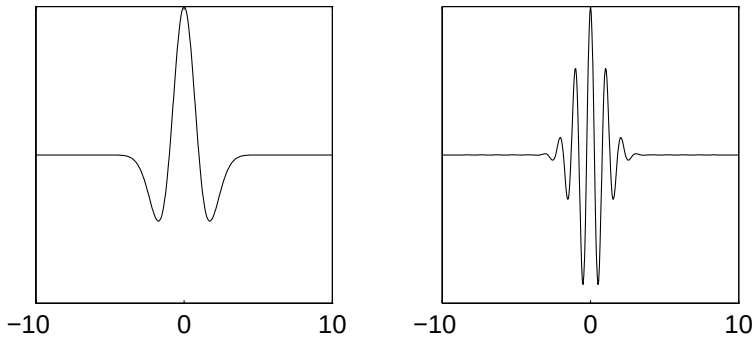

Figure 1: Two usual one-dimensional wavelets: (left) The Mexican hat or Marr wavelet; (right) The real part of the 1-D Morlet wavelet, for $k_o = 5.6$.

The *Morlet wavelet* is just a modulated Gaussian, given by

$$\psi_M(x) = e^{i\omega_o x}\, e^{-x^2/2\sigma_o^2} + \text{ corr.}, \qquad \widehat{\psi}_M(\omega) = \sigma_o e^{-[(\omega-\omega_o)\sigma_o]^2/2} + \text{ corr.} \tag{2.8}$$

In fact, the first term alone does *not* satisfy the admissibility conditions (2.4), (2.5), hence the necessity of a correction. However, for ω_o large enough (typically $\omega_o\sigma_o \geq 5.5$), this correction term (of Gaussian type) is numerically negligible. Notice that, without the correction term, (2.8) is just a Gabor function, the most common function used for STFT.

The Morlet wavelet is complex, hence the corresponding transform $S(b, a)$ is also complex, and one may treat separately its phase and its modulus. It turns out that the phase of the transform is a crucial ingredient for the algorithm of singularity detection in a signal, for instance, the localization of spectral lines (Section 3).

Notice that none of these wavelets is progressive.

2.3. Localization properties and interpretation

The main virtues of the CWT follow from the support properties of ψ. Assume ψ and $\widehat{\psi}$ to be as well localized as possible (compatible with the Fourier uncertainty principle). More specifically, assume that ψ has an 'essential' support of width L, centered around 0, while $\widehat{\psi}$ has an essential support of width Ω, centered around ω_o. Then the transformed wavelets ψ_{ba} and $\widehat{\psi}_{ba}$ have, respectively, an essential support of width aL around b and an essential support of width Ω/a around ω_o/a. This behavior is illustrated in Figure 2, which shows the Morlet wavelet in the time and frequency domains, for three successive scales $a = 0.5, 1$ and 2, from left to right.

Notice that the product of the two widths is constant. We know it has to be bounded below by a fixed constant, by the (Fourier) uncertainty principle. We illustrate this vital fact in Figure 3, which is a time-frequency representation. Remember that $1/a$ behaves like a frequency. Therefore:

• if $a \gg 1, \psi_{ba}$ is a wide window, whereas $\widehat{\psi_{ba}}$ is very peaked around a small frequency ω_o/a: this transform is most sensitive to *low frequencies*.

• if $a \ll 1, \psi_{ba}$ is a narrow window and $\widehat{\psi_{ba}}$ is wide and centered around a high frequency ω_o/a: this wavelet has a good localization capability in the space domain and is mostly sensitive to *high frequencies*.

Thus we have obtained a tool that reproduces the correlation between duration and average frequency discussed in the introduction. Low frequency portions of the signal tend to be long, whereas high frequencies occur briefly in general.

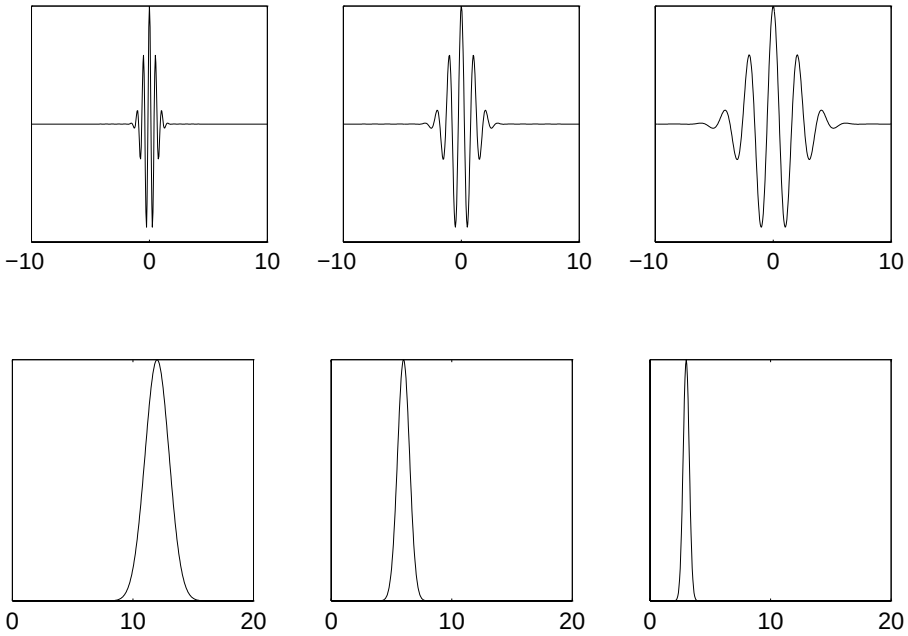

Figure 2: Support properties of the Morlet wavelet ψ_M: for $a = 0.5, 1, 2$ (left to right), ψ_{ba} has width 3, 6, 12, respectively (top), while $\widehat{\psi_{ba}}$ has width 3, 1.5, 0.75, and peaks at 12, 6, 3 (bottom).

Combining now these localization properties with the zero mean condition and the fact that ψ_{ba} acts like a filter (convolution), we see that the CWT performs a *local filtering*, both in time and in scale. The wavelet transform $S(b, a)$ is nonnegligible only when the wavelet ψ_{ba} matches the signal $s(x)$, that is, it filters the part of the signal, if any, that lives around the time b and the scale a.

Taking all these properties together, one is naturally led to the interpretation of the CWT as a *mathematical microscope*, with optics ψ, position b and global magnification $1/a$ (Arnéodo, 91). In addition, the analysis works at constant relative bandwidth ($\Delta\omega/\omega =$ constant), so that it has a better resolution at high frequency, i.e., small scales. This property makes it an ideal tool for detecting *singularities* (for instance, discontinuities in the signal or one of its derivatives), and also scale dependent features, in particular, for analyzing *fractals* (Holschneider, 88) (Arnéodo, 91).

2.4. More on the mathematical side

Given an admissible wavelet ψ, i.e., such that $c_\psi < \infty$ [see (2.3)], the corresponding CWT $W_\psi : s(x) \mapsto S(b, a)$ is a linear map, with the following properties:

(1) W_ψ is *covariant* under translation and under dilation (scale change):

$$W_\psi \;:\; s(x - x_o) \mapsto S(a, b - x_o); \tag{2.9}$$

$$W_\psi \;:\; \frac{1}{\sqrt{a_o}} s(\frac{x}{a_o}) \mapsto S(\frac{a}{a_o}, \frac{b}{a_o}); \tag{2.10}$$

(2) W_ψ conserves the energy of the signal:

$$\int_{-\infty}^{\infty} |s(x)|^2 \, dx = c_\psi^{-1} \iint_{\mathbb{R}_+^2} |S(b, a)|^2 \, \frac{da \, db}{a^2}. \tag{2.11}$$

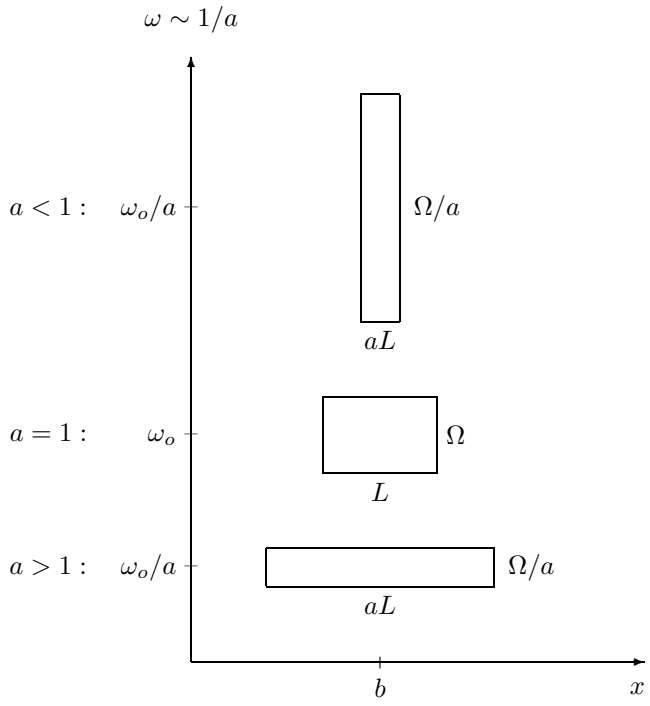

Figure 3: Support properties of ψ_{ba} and $\widehat{\psi_{ba}}$.

From this condition, one sees that $|S(b,a)|^2$ may be interpreted as an energy density in the (b,a)-half-plane, and that the natural geometry of the latter is not the usual Euclidean one. Indeed, the measure $da\,db/a^2$ is invariant under time translation and dilation.

The relation (2.11) means that the map W_ψ is an isometry from the space of signals $L^2(\mathbb{R})$ onto a closed subspace of $L^2(\mathbb{R}^2_+, da\,db/a^2)$. An equivalent statement is that the wavelet ψ generates a *resolution of the identity:*

$$c_\psi^{-1} \iint_{\mathbb{R}^2_+} |\psi_{ba}\rangle \langle\psi_{ba}| \frac{da\,db}{a^2} = I. \tag{2.12}$$

(3) As a consequence, the map W_ψ is invertible on its range, and the inverse transformation is simply the adjoint of W_ψ. Thus the signal $s(x)$ may be reconstructed from its wavelet transform by the formula:

$$s(x) = c_\psi^{-1} \iint_{\mathbb{R}^2_+} \psi_{ba}(x)\, S(b,a) \frac{da\,db}{a^2} \ . \tag{2.13}$$

This means that the WT provides a decomposition of the signal as a linear superposition of the wavelets ψ_{ba} with coefficients $S(b,a)$ — exactly as for the Fourier transform.

(4) The projection from $L^2(\mathbb{R}^2_+, da\,db/a^2)$ onto the range of W_ψ, that is, the space of wavelet transforms, is an integral operator, whose kernel,

$$K(a',b';a,b) = c_\psi^{-1} \langle \psi_{a'b'} | \psi_{ba} \rangle \ , \tag{2.14}$$

is the autocorrelation function of ψ. It is also called a *reproducing kernel*, because the statement above means precisely that a function $f \in L^2(\mathbb{R}^2_+, da\,db/a^2)$ is the WT of a certain signal iff it satisfies the reproduction property:

$$f(a',b') = \iint_{\mathbb{R}^2_+} K(a',b';a,b)\, f(b,a) \frac{da\,db}{a^2} \ . \tag{2.15}$$

Remark: The relation (2.13) means that the *same* wavelet ψ has been used for the analysis *and* for the reconstruction. This is an unnecessary restriction, however. Indeed one may reconstruct the signal by using a wavelet χ different from the analyzing wavelet ψ (Daubechies, 92):

$$s(x) = c_{\chi\psi}^{-1} \iint_{\mathbb{R}_+^2} \chi_{ba}(x)\,(W_\psi s)(b,a)\frac{da\,db}{a^2}\ , \tag{2.16}$$

provided χ and ψ satisfy the compatibility condition

$$0 < \left| \int_{-\infty}^{\infty} \widehat{\psi}(\omega)\,\widehat{\chi}(\omega)\,\frac{d\omega}{|\omega|} \right| < \infty. \tag{2.17}$$

In this way one may obtain simpler formulas, in particular for reconstruction (as shown originally by Morlet). The same idea, when transposed to the DWT, leads to the so-called biorthogonal wavelet bases (see Section 5).

3. IMPLEMENTATION OF THE CWT

3.1. Academic signals

Faced with this new tool, one must begin by learning the rules of the trade, that is, one must learn how to read and understand a CWT (Grossmann, 90). The simplest way is to get some practice on very simple academic signals, such as a simple discontinuity in time or a monochromatic signal (pure sinusoid). Two remarks are in order here. First, the L^2 normalization used in (2.1) was chosen mainly for mathematical reasons. It is the one for which the wavelet transform $s(x) \mapsto S(b,a)$ is a unitary map. However, in practice, this may be changed. For instance, it is common to use instead the so-called L^1 normalization, which amounts to replace the factor $a^{-1/2}$ in (2.1) by a factor a^{-1}. The effect is to enhance the small scales, that is, to make more conspicuous the discontinuities in the signal, which is precisely one of the aims of the WT. Second, it is natural to use a logarithmic scale for the scale parameter a. The visual effect is that the lines, $b/a =$ constant, are not straight lines, but hyperbolic curves; at the same time, the horizon $a = 0$ recedes to (minus) infinity (see Figure 4 below).

We will now analyze the two academic signals mentioned above. The analyzing wavelet ψ is supposed to be complex, so that we may treat separately the modulus and the phase of the transform. The scale axis, in units of $\log a$, points downward, so that high frequencies (small a) correspond to the top of the plots, and low frequencies (large a) to the bottom. The results are presented by coding the height of the function by density of points (12 levels of grays, from white to black). The phase is 2π-periodic. When it reaches 2π, it is wrapped around to the value 0. Thus the lines of constant phase with value $2k\pi$ are lines of discontinuity, where the density of points drops abruptly from 1 (black) to 0 (white). In addition, the functions plotted are thresholded at 1% of the maximum value of the modulus of $S(b,a)$.

3.1.1 A simple discontinuity

The simplest signal is a simple discontinuity in time, at $x = x_o$, modeled by $s(x) = \delta(x - x_o)$. The WT is obtained immediately and reads

$$S(b,a) = a^{-1/2}\,\overline{\psi\,(a^{-1}(x_o - b))}. \tag{3.1}$$

The following features may be read off (3.1):

• The phase of $S(b,a)$ is constant on the lines $b/a =$ constant, originating from the the point $b = x_o$ on the horizon. These lines point towards the position of the singularity, like a finger.

• On the same lines of constant phase, the modulus of $S(b,a)$ increases as $a^{-1/2}$ when $a \to 0$, so that the singularity is enhanced. The effect is even more pronounced if one uses the L^1 normalization.

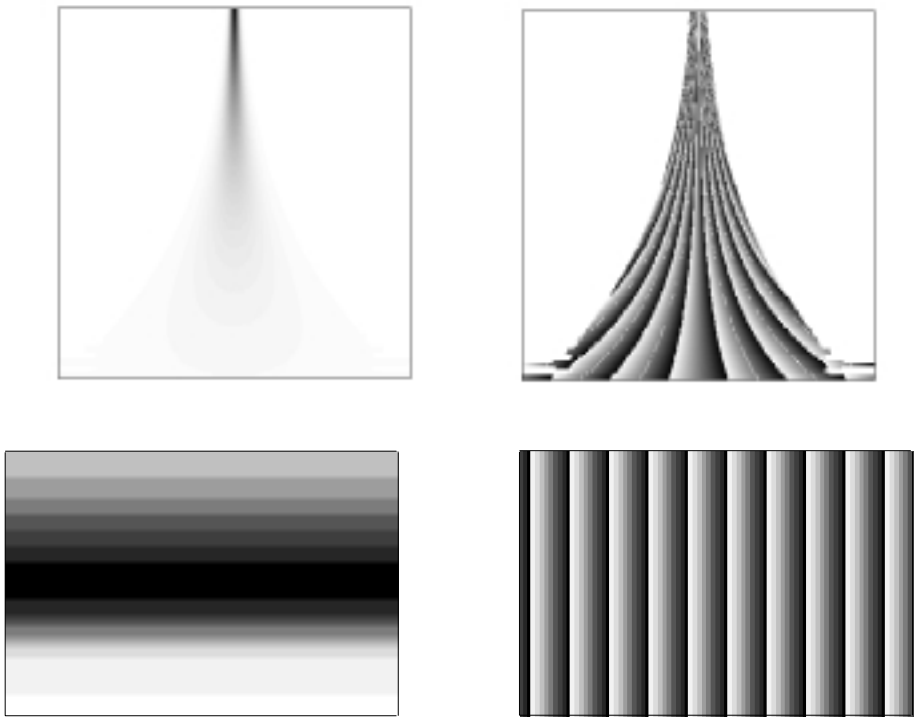

Figure 4: WT of the two academic signals with the Morlet wavelet (modulus on the left, phase on the right): (top) A δ function; (bottom) A single sinusoid.

This is illustrated on Figure 4 (top), which presents the modulus and phase of the WT of a δ function, using a standard Morlet wavelet (but the result is independent of the choice of ψ).

These properties explain why the CWT is primarily a tool for detecting singularities in the signal (discontinuity in $s(x)$ or in one of its derivatives). Indeed they follow only from the fact that the δ function is homogeneous of order -1:

$$\delta(\lambda(x - x_o)) = \frac{1}{\lambda}\,\delta(x - x_o), \quad \lambda > 0, \tag{3.2}$$

and they remain true for any signal which is homogeneous of order α around a given point x_o:

$$s(\lambda(x - x_o)) = \lambda^{\alpha} s(x - x_o)) \;\text{ for }\; x \sim x_o, \lambda > 0, \tag{3.3}$$

implies (using the covariance property of the CWT under dilation)

$$S(a, b - x_o) = a^{\alpha+\frac{1}{2}}\; S\left(1, a^{-1}(b - x_o)\right). \tag{3.4}$$

Such a homogeneous behavior, with different values of α on the left and on the right of x_o, models almost any type of discontinuity in x_o (Holschneider, 95). In fact, the CWT allows also to estimate the strength of singularities. Indeed, for such a homogeneous function (3.3), the exponent α characterizing

the singularity on the righthand side of x_o may be read off a plot of $\log|S(b,a)|$ vs. $\log a$, as b tends to x_o from the right (this is in fact a particular case of a general analysis of fractal signals through their WT (Holschneider, 88) (Arnéodo, 91)), and similarly from the left.

The interesting point is that this behavior is extremely robust. For instance, the 'finger' pointing to a δ-singularity remains clearly visible when the latter is superposed on a continuous signal (even if the amplitude of the δ function is too small to be visible on the signal itself), or even in the presence of substantial background noise. Similarly, the discontinuity corresponding to the abrupt onset of a signal is readily identified with the CWT (a situation common in seismology, for example). We refer to (Grossmann, 90) for several spectacular examples.

3.1.2 A single monochromatic wave

Equally simple is a single harmonic signal (monochromatic wave):

$$s(x) = e^{i\omega_s x} \;\Leftrightarrow\; \widehat{s}(\omega) = \frac{1}{\sqrt{2\pi}}\,\delta(\omega-\omega_s), \tag{3.5}$$

which gives

$$S(b,a) = \sqrt{\frac{a}{2\pi}}\;\widehat{\psi}(a\omega_s)\,e^{i\omega_s b} = S(a,0)\;e^{i\omega_s b}. \tag{3.6}$$

The same relations remain true for a real monochromatic signal, $s(x)=\sin\omega_s x$ or $s(x)=\cos\omega_s x$, if the wavelet ψ is progressive ($\widehat{\psi}(\omega)=0$ for $\omega\le 0$).

Again two important properties may be read off immediately from (3.6):

• The modulus of $S(b,a)$ is independent of b. Hence, the graph of $|S(b,a)|$ consists of horizontal bands and the profile for a fixed time b essentially reproduces the profile of $\widehat{\psi}$.

• The phase of $S(b,a)$ is linear in b. Since the phase is 2π-periodic, the graph of $\Phi(b,a)\equiv\arg S(b,a)$ is a linear sawtooth function:

$$\Phi(b,a) = \omega_s b \;(\mathrm{mod}\;2\pi). \tag{3.7}$$

These properties are illustrated on Figure 4 (bottom) for a single sine function analyzed with a Morlet wavelet.

Both the modulus and the phase allow to determine the frequency ω_s of the signal. If the modulus of the wavelet $\widehat{\psi}(\omega)$ has a single maximum for $\omega=\omega_o$, (3.6) gives immediately $\omega_s=\omega_o/a_M$, where a_M is the scale corresponding to the maximum in the profile of $|S(b,a)|$ for fixed b. For instance, the (truncated) Morlet wavelet $\psi(x)=\exp(i\omega_o x)\,\exp(-x^2/2)$ yields:

$$S(b,a) = \sqrt{a}\;e^{-\frac{1}{2}(a\omega_s-\omega_o)^2}\;e^{i\omega_s b}, \tag{3.8}$$

and the result is obvious. As for the phase, (3.7) gives, at least locally:

$$\frac{\partial\Phi(b,a)}{\partial b} = \omega_s = \frac{\omega_o}{a_M}\;. \tag{3.9}$$

For an asymptotic, locally monochromatic signal [see (4.5)]

$$s(x) = A(x)\exp(i\phi(x)), \tag{3.10}$$

the result given in (3.9) may be recovered if one introduces the notion of *instantaneous frequency:*

$$\omega_{\mathrm{inst}} = \frac{\partial\Phi(b,a)}{\partial b}\,. \tag{3.11}$$

Then it can be shown that (3.9) generalizes to:

$$\frac{\partial \Phi(b,a)}{\partial b} = \frac{\omega_o}{a}. \tag{3.12}$$

The solution of (3.12) corresponds to the points where the instantaneous frequency of the signal coincides with the frequency of the wavelet, i.e., the values of a corresponding to the peaks of the signal (center of the spectral lines). This relation may then be further extended to a general asymptotic signal analyzed by an asymptotic wavelet, and this leads precisely to the notion of ridge, sketched in Section 4.2 (Delprat, 92).

3.2. Physical applications

The CWT has found a wide variety of applications in various branches of physics and/or signal processing. We will list here a representative selection; most of them may be found in the proceedings volumes (Combes, 90), (Meyer, 91) and (Meyer, 93). In all cases, the CWT is primarily used for analyzing transient phenomena, detecting abrupt changes in a signal or comparing it with a given pattern.

- *Sound and acoustics:*

The first applications of the CWT were in the field of acoustics. A few examples are musical synthesis, speech analysis and modeling of the sonar system of bats and dolphins. Other examples include various problems in underwater acoustics, such as the disentangling of the different components of an underwater refracted wave (Saracco, 90) and the identification of an obstacle (a submarine is a good example!).

- *Geophysics:*

This is the origin of the method, which was designed in an empirical fashion by J. Morlet for analyzing the recordings of microseisms used in oil prospection. More recently, the CWT has been applied to the analysis of various long time series of geophysical origin, e.g., in gravimetry (fluctuations of the local gravitational field), in geomagnetism (fluctuation of the Earth's magnetic field (Alexandrescu, 95)) or in astronomy (fluctuations of the length of the day, variations of solar activity, measured by the sunspots, etc.).

- *Fractals, turbulence:*

As mentioned above, the CWT is an ideal tool for studying fractals or, more generally, phenomena with particular properties under scale changes. Thus it is quite natural that the CWT has found many applications in the analysis of (1-D and 2-D) fractals, artificial (diffusion limited aggregates) or natural (arborescent growth phenomena) (Arnéodo, 91). Related to these is the use of the CWT in the analysis of developed turbulence (identification of coherent structures, uncovering of hierarchical structure) (Farge, 92).

- *Spectroscopy:*

This was one of the earliest and most successful applications, in particular for NMR spectroscopy, where the method proves extremely efficient in subtracting unwanted spectral lines or filtering out background noise (Guillemain, 91) (Delprat, 92). We will discuss this example in more detail below.

- *Medical applications:*

The CWT has been used for analyzing or monitoring various electrical or mechanical phenomena in the brain (EEG) (Thonet, 98) or the heart (ECG). See (Aldroubi, 96) for a comprehensive review.

- *Industrial applications:*

Here again the important aspect is monitoring, for instance in detecting anomalies in the functioning of nuclear, electrical or mechanical installations, or analyzing the behavior of materials under impact (Vandergheynst, 99).

- *Shape recognition:*

A standard problem in artificial vision is to identify an object by its shape. An interesting method

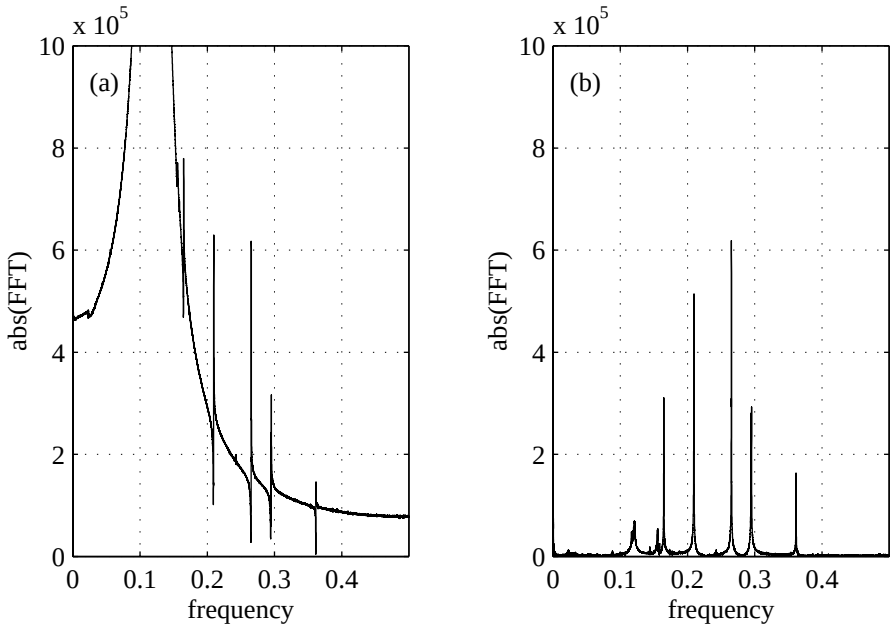

Figure 5: Application of the CWT in NMR spectroscopy: Subtraction of an unwanted peak. (a) The original spectrum; (b) The spectrum reconstructed after subtraction of the water peak.

consists in treating the contour of the object as a curve in the (complex) plane and analyzing its real and imaginary components with a 1-D CWT (Antoine, 97).

Before concluding this section, we illustrate the use of the CWT by the example of NMR spectroscopy. The physical phenomenon may be described as follows. When a sample is placed in a static magnetic field, nuclei with a magnetic moment align along this applied field, resulting in a net magnetization. In its equilibrium state, the magnetization is static and does not induce a signal in the receiver antenna. In order to obtain information, one must first excite the nuclei with a radio frequency pulse. After such a pulse, the magnetizations precess around the static field at angular frequencies characteristic of their chemical environnement and relax to their equilibrium state. This precession induces a signal in the receiver antenna. The signal to be analyzed is the Fourier transform (spectrum) of the damped response curve of the protons. It contains a large number of narrow peaks, the spectral lines, but many among them are useless, coming for instance from the protons of the solvent. These peaks, which may be quite big, must be subtracted, and the position of the relevant ones measured with precision. In addition, the spectra are often quite noisy, and must be 'cleaned' before any useful measurement can be performed.

A typical analysis of NMR spectra with help of wavelets is given in Figures 5 and 6 (Guillemain, 91) (Delprat, 92) (Barache, 97). The first one is an example of peak subtraction. The original spectrum (left) exhibits a huge parasite peak, due to the protons of the solvent (water), that masks to a large extent the interesting structures. The analysis consists in isolating this peak in the CWT, subtracting it from the spectrum and reconstructing the remaining part. The result (right) is a spectrum where all the fine details are now clearly visible, and have not been perturbed by the removal of the large peak. Indeed, the prominent structures appear at exactly the same place on the frequency (horizontal) axis in both pictures. The reason for the remarkable efficiency of the method in this case is that the huge line and the rest of the spectrum live at different scales, hence they are decoupled in the CWT ('unfolding') and

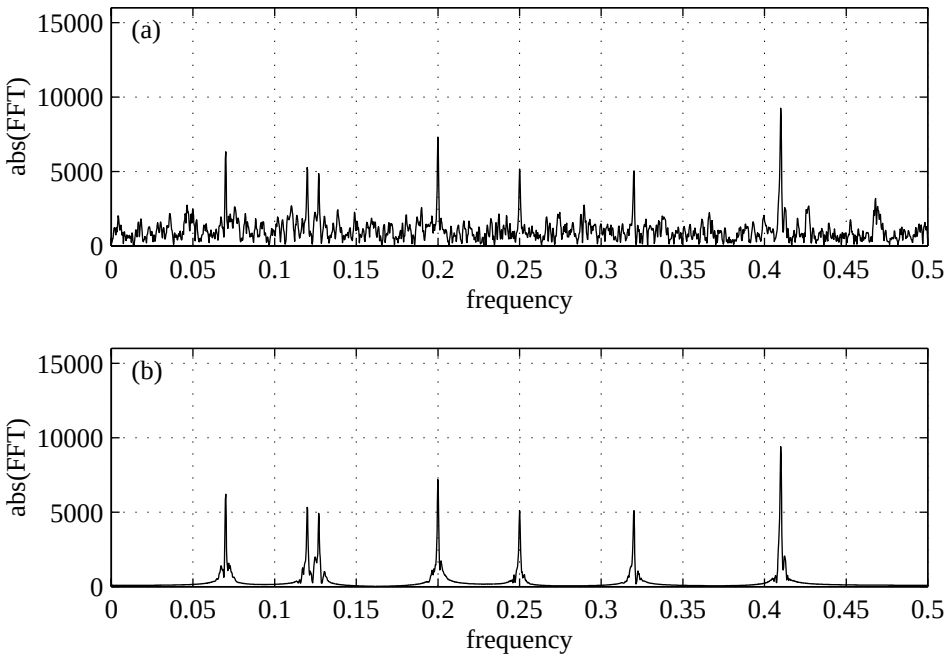

Figure 6: Noise filtering with the CWT: (a) The original NMR spectrum; (b) The spectrum reconstructed after noise removal.

can be readily separated with very little distortion.

Figure 6 is an example of noise filtering. The original signal (top) consists of a number of damped sinusoids enbedded in noise. The dominant peaks are localized with help of the CWT ridge algorithm, the remnant of the spectrum is subtracted and the filtered spectrum is reconstructed. The result (bottom) compares well with standard methods of noise suppression (signal filtering).

4. DISCRETIZATION OF THE 1-D CWT

The reproduction property (2.15) implies that the information content of the CWT $S(b, a)$ is highly redundant. In fact the signal has been unfolded from one to two dimensions, and this explains the practical efficiency of the CWT for disentangling parts of the signal that live at the same time, but on different scales. This redundancy may be exploited in several ways. The principal one is that it must be possible to obtain the full information about the signal from a small subset of the values of the transform $S(b, a)$. Two applications at least have been described in the literature, the theory of frames (Daubechies, 86) and the ridge or skeleton representation (Delprat, 92).

4.1. Frames

Let $\Gamma = \{a_j, b_k, j, k \in \mathbb{Z}\}$ be a discrete lattice in the (b, a)-half-plane. We say that Γ yields a good discretization for the CWT if an arbitrary signal $s(x)$ may be represented as a superposition

$$s(x) = \sum_{j,k \in \mathbb{Z}} \langle \psi_{jk} | s \rangle \, \widetilde{\psi_{jk}}(x), \tag{4.1}$$

where $\psi_{jk} \equiv \psi_{b_k a_j}$ and $\widetilde{\psi_{jk}}$ may be explicitly constructed from ψ_{jk}. We emphasize that (4.1) must be an *exact* representation, i.e., there is no loss of information as compared to the continuous reconstruction (2.13). Actually (4.1) means that the signal $s(x)$ may be replaced by the set $\{\langle\psi_{jk}|s\rangle\}$ of its wavelet coefficients. In addition, one wants the reconstruction of $s(x)$ from its coefficients to be numerically stable (i.e., a small error in the coefficients implies a small error in the reconstructed signal). As shown in (Daubechies, 86), (Daubechies, 90) and (Daubechies, 92), this is achieved provided the following condition holds true, for some constants $A > 0, B < \infty$:

$$A \|s\|^2 \;\leq\; \sum_{j,k\,\in\,\mathbb{Z}} |\langle\psi_{jk}|s\rangle|^2 \;\leq\; B \|s\|^2 \tag{4.2}$$

(the lower bound guarantees the numerical stability). In that case, one says that the set $\{\psi_{jk}\}$ constitutes a *frame*, with *frame bounds* A and B (this notion was originally introduced in the 1950s, in the context of nonharmonic analysis). If $A = B > 1$, the frame is said to be *tight*. Of course, if $A = B = 1$, and $\|\psi\| = 1$, the set $\{\psi_{jk}\}$ is simply an orthonormal basis. The important point here is that, for all practical purposes, a good frame is almost as good as an orthonormal basis. By 'good frame', we mean that the expansion (4.1) converges sufficiently fast. The detailed analysis of Daubechies shows this to be the case if $|B/A - 1| \ll 1$, thus in particular if the frame is tight.

Precisely at this point arises the basic difference between the discretized CWT and the DWT. In the former case, the wavelet ψ is chosen *a priori* (with very few constraints, see Section 2), and the question is whether one can find a lattice Γ such that $\{\psi_{jk}\}$ is a frame with decent frame bounds A, B. In the other approach, one imposes that the set $\{\psi_{jk}\}$ be an orthonormal basis and proves the existence of a function ψ to that effect. The construction is rather indirect and the resulting function is usually very complicated (often it has a fractal behavior). We will discuss briefly the discrete-time WT in the next section.

Clearly, for a given a wavelet ψ, the determination of a lattice Γ that might lead to a suitable frame must take into account the (non-Euclidean) geometry of the time-scale half-plane. Usually, this lattice is chosen to be invariant under discrete dilations and translations:

- for the scales, one chooses naturally $a_j = a_o^j, j \in \mathbb{Z}$, for some $a_o > 1$;
- for the times, one takes $b_k = k\, b_o\, a_o^j, j, k \in \mathbb{Z}$. Thus

$$\psi_{jk}(x) = a_o^{-j/2}\, \psi(a_o^{-j} x - k), \quad j, k \in \mathbb{Z}. \tag{4.3}$$

The most common choice is $a_o = 2$ (octaves!) and $b_o = 1$, which results in

$$\psi_{jk}(x) = 2^{-j/2}\, \psi(2^{-j} x - k), \quad j, k \in \mathbb{Z}. \tag{4.4}$$

It is worth noticing that this so-called *dyadic* lattice is exactly the same that indexes the DWT (see Section 5 below), although the two approaches are totally different.

For this choice indeed, the relevant theorems may be proven (Daubechies, 86) (Daubechies, 92). Both the Mexican hat and the Morlet wavelet yield good frames. Actually the same question may be asked for the STFT. There the geometry is Euclidean. Typically Γ is a square lattice and the frame theorem simply says that a frame is obtained if the density of Γ is larger than a critical value. This is the well-known result of von Neumann concerning the density of canonical coherent states (Klauder, 85), and it is closely related to the standard theorem known in different circles under the names of Fourier, Heisenberg (uncertainty relations), Nyquist or Shannon.

In conclusion, the two standard wavelets yield very good frames (nontight, however). This explains their efficiency in signal analysis, and thus their popularity for applications. We refer to (Daubechies, 90) (Daubechies, 92) for a detailed analytical and numerical discussion of this problem.

4.2. Ridges and skeletons

Real signals are frequently very entangled and noisy, and their WT is difficult to understand. However, a clever exploitation of the intrinsic redundancy of the CWT is often able to bypass the difficulty and thus to improve the efficiency and the range of applicability of wavelet analysis.

The idea is that many signals are well approximated by a superposition of spectral lines:

$$s(x) = \sum_n A_n(x) \exp(i\phi_n(x)), \tag{4.5}$$

where the amplitude $A_n(x)$ varies slowly with respect to the phase $\phi_n(x)$. Such signals are sometimes called *asymptotic*. Typical examples are spectra in NMR spectroscopy (Guillemain, 91). For a signal of this kind, the CWT (2.1) in the time domain is a rapidly oscillating integral, the essential contribution to which is given by the stationary points of the phase of the integrand. Assume for simplicity there is only one such point, $x_s = x_s(b,a)$. Then the *ridge* of the WT is defined as the set of points (b,a) for which $x_s(b,a) = b$. These constitute a curve in the (b,a)-half-plane and a detailed analysis shows that, on this curve, the WT $S(b,a)$ coincides, up to a small correction, with the analytic signal $Z(b)$ associated to $s(x)$ (Delprat, 92). It follows that the restriction of the WT $S(b,a)$ to its ridge, the so-called *skeleton*, contains the whole information. In particular, the so-called frequency modulation law $x^{-1}\arg\{s(x)\}$ of $s(x)$ is easily recovered from the skeleton. Thus, it is not necessary to compute the whole CWT, but only its skeleton. This is of course much less costly computationally, because there are fast algorithms available. Spectacular applications of this method may be found, for instance, in spectroscopy (Delprat, 92), geomagnetism (Alexandrescu, 95) or shape determination (Antoine, 97).

5. THE DISCRETE WT: ORTHONORMAL WAVELET BASES

One of the successes of the WT was the discovery that it is possible to construct functions ψ for which $\{\psi_{jk}, j,k \in \mathbb{Z}\}$ is indeed an orthonormal basis of $L^2(\mathbb{R})$. In addition, such a basis still has the good properties of wavelets, including space *and* frequency localization. In addition, it yields fast algorithms, and this is the key to the usefulness of wavelets in many applications.

The construction is based on two facts: first, almost all examples of orthonormal bases of wavelets can be derived from a multiresolution analysis, and then the whole construction may be transcripted into the language of Quadrature Mirror Filters (QMF), familiar in the signal processing literature.

A *multiresolution analysis* of $L^2(\mathbb{R})$ is an increasing sequence of closed subspaces

$$\ldots \subset V_{-2} \subset V_{-1} \subset V_0 \subset V_1 \subset V_2 \subset \ldots, \tag{5.1}$$

with $\bigcup_{j\in\mathbb{Z}} V_j$ dense in $L^2(\mathbb{R})$ and $\bigcap_{j\in\mathbb{Z}} V_j = \{0\}$, and such that

(1) $f(x) \in V_j \Leftrightarrow f(2x) \in V_{j+1}$

(2) There exists a function $\phi \in V_0$, called a *scaling* function, such that $\{\phi(x-k), k \in \mathbb{Z}\}$ is anorthonormal basis of V_0.

Combining (1) and (2), one gets an orthonormal basis of V_j, namely $\{\phi_{j,k}(x) \equiv 2^{j/2}\phi(2^j x - k),\ k \in \mathbb{Z}\}$.

Each V_j can be interpreted as an approximation space: the approximation of $f \in L^2(\mathbb{R})$ at the resolution 2^j is defined by its projection onto V_j, and no scale is privileged, by (1). The additional details needed for increasing the resolution from 2^j to 2^{j+1} are given by the projection of f onto the orthogonal complement W_j of V_j in V_{j+1}:

$$V_j \oplus W_j = V_{j+1}, \tag{5.2}$$

and we have:

$$\bigoplus_{j\in\mathbb{Z}} W_j = L^2(\mathbb{R}). \tag{5.3}$$

Then the theory asserts the existence of a function ψ, called the *mother* wavelet, explicitly computable from ϕ, such that $\{\psi_{j,k}(x) \equiv 2^{j/2}\psi(2^j x - k), k \in \mathbb{Z}\}$ constitutes an orthonormal basis of W_j. Thus $\{\psi_{j,k}, j, k \in \mathbb{Z}\}$ is an orthonormal basis of $L^2(\mathbb{R})$, and these are the *orthonormal wavelets*.

Various additional conditions may be imposed on the function ψ (hence on the basis wavelets): arbitrary regularity, several vanishing moments (in any case, ψ has always zero mean), fast decrease at infinity, even compact support. The technique consists in translating the multiresolution structure into the language of QMF filters, and putting suitable constraints on the filter coefficients. For instance, ψ has compact support if only finitely coefficients differ from zero (in technical terms, one has an FIR filter).

The simplest example of this construction is the Haar basis, which comes from the scaling function $\phi(x) = 1$ for $0 \leq x < 1$ and 0 otherwise. Similarly, various spline bases may be obtained along the same line. Other explicit examples may be found in (Chui, 92) or (Daubechies, 92).

Although appropriate filters generate orthonormal wavelet bases, the resulting scheme turns out to be too rigid for many applications and various generalizations have been proposed.

(i) Biorthogonal wavelet bases:
As we mentioned at the end of Section 2.4, the wavelet used in the CWT for reconstruction need not be the same as that used for decomposition, the two have only to satisfy a cross-compatibility condition. The same idea in the discrete case leads to biorthogonal bases, i.e., one has two hierarchies of approximation spaces, V_j and $\hat{V}_j$, with cross-orthogonality relations. This gives a better control, for instance, on the regularity or decrease properties of the wavelets (Cohen, 92).

(ii) Wavelet packets:
The construction of orthonormal wavelet bases leads to a special subband coding scheme, rather asymmetrical. Each approximation space V_j gets further decomposed into V_{j-1} and W_{j-1}, whereas the detail space W_j is left unmodified. More flexible subband schemes have been considered, called *wavelet packets*; they provide rich libraries of orthonormal bases, and also strategies for determining the optimal basis in a given situation (Coifman, 93).

(iii) Second generation wavelets:
One can go further and abandon the regular dyadic scheme and the Fourier transform altogether. Using the 'lifting scheme', one obtains the so-called *second-generation wavelets* (Sweldens, 96). The same scheme applies in 2-D as well. For instance, Schröder and Sweldens (Schröder, 95) have applied it to the design of wavelets on the sphere, with a very convincing application to the reproduction of coastlines on a terrestrial globe.

(iv) Integer wavelet transforms:
In their standard numerical implementation, the classical (discrete) WT converts floating point numbers into floating point numbers. However, in many applications (data transmission from satellites, multimedia), the input data consists of integer values only and one cannot afford to lose information: only lossless compression scheme are allowed. Recent developments have produced new methods that allow one to perform all calculations in integer arithmetic (Calderbank, 98).

Of course, the literature on the DWT and its applications is growing fast. Besides the classical book (Daubechies, 92), it may be useful to mention several books, which are specifically designed for signal processing, for instance, (Akansu, 95), (Mallat, 99), (Strang, 96), (Vetterli, 95) or (Wickerhauser, 94).

6. HOW TO GENERALIZE THE 1-D CWT?

As we have seen in Section 2, the natural geometry of the (b, a)-half-plane $\mathbb{R}^2_+$ is not the usual Euclidean one. Indeed the measure $da\,db/a^2$ is invariant not only under time translation, but also under dilation. The reason behind these facts and the nice properties described above is to be found in group representation theory. The natural operations on a 1-D signal are precisely time translations and dilations, and these together constitute the $ax + b$ group, that is, the connected affine group G_{aff} of the line. Then the

relation

$$(U(b,a)f)(x) \equiv f_{ba}(x) = a^{-1/2}\, f(a^{-1}(x-b)), \quad a>0,\ b\in\mathbb{R}, \tag{6.1}$$

defines a unitary representation of $G_{\rm aff}$ in the Hilbert space $L^2(\mathbb{R}, dx)$ of finite energy signals. This means that, for every $g \equiv (b,a) \in G_{\rm aff}$, $U(g)$ is a unitary operator and one has, for any $g, g' \in G_{\rm aff}$, $U(g)U(g') = U(gg')$, $U(g^{-1}) = U(g)^\dagger$, and thus $U(e) = I$, where $e = (1,0)$ denotes the unit element of $G_{\rm aff}$.

The representation U becomes irreducible when restricted to the subspace of progressive signals $\mathcal{H}_+ = \{f \in L^2(\mathbb{R}, dx), \widehat{f}(\omega) = 0 \text{ for } \omega \le 0\}$, that is, $\mathcal{H}_+$ contains no subspace invariant under U, except the trivial one $\{0\}$. Furthermore, and this is the crucial feature, the restricted representation U_+ is *square integrable*, that is, there exists at least one (and in fact a dense set of) admissible vectors, i.e., vectors ψ such that the matrix element $\langle U_+(b,a)\psi|\psi\rangle$ is square integrable over the group, with respect to the (left) invariant measure, namely $da\, db/a^2$. Then a straightforward calculation shows that the squared norm of that matrix element is proportional to the constant c_ψ given in (2.3). Thus the two notions of admissibility we have introduced indeed coincide.

From this fact follow all the mathematical properties described in Section 2.4, covariance, energy conservation (2.11), inversion formula (2.13), reproducing kernel (2.15). This is of course no accident! It simply reflects the fact that the 1-D CWT is a particular case of the general theory of coherent states associated to group representations (Ali, 95) (Ali, 99). This observation is of central importance, for it is this approach that allows a natural and easy extension of the 1-D CWT to higher dimensions.

How does this come about? Where does the appropriate group come from? The answer lies in the notion of *symmetry*, which plays such a prominent part in the whole of physics. Suppose indeed that the signal possesses certain symmetry properties. It is natural to build these into the wavelet transform itself, and this clearly requires the use of the continuous approach. What emerges here is a general pattern, that we now describe.

Consider the class of finite energy signals living on a manifold Y, i.e., $s \in L^2(Y, d\mu) \equiv \mathcal{H}$. For instance, Y could be space $\mathbb{R}^n$, the 2-sphere S^2, space-time $\mathbb{R}\times\mathbb{R}$ or $\mathbb{R}^2\times\mathbb{R}$, etc. Suppose there is a group G of transformations acting (transitively) on Y : $y \mapsto g[y]$, with $g[g'[y]] = gg'[y]$, $e[y] = y$, and for any pair $y, y' \in Y$, there is at least one $g \in G$ such that $g[y] = y'$. Then we may let the group G act linearly on signals, $s \mapsto U(g)s$, and consistency requires that U should be a unitary representation of G in the space $\mathcal{H}$ of signals:

$$\langle U(g)s|U(g)s'\rangle = \langle s|s'\rangle,\ \forall\, g\in G,\ s, s' \in \mathcal{H}. \tag{6.2}$$

In order to get a wavelet analysis on Y, adapted to the symmetry group G, three conditions must be met: (1) G contains *dilations* of some kind; (2) U is irreducible; and (3) U is *square integrable*, i.e., there exists at least one nonzero vector $\psi \in \mathcal{H}$, called *admissible*, such that the matrix element $\langle U(g)\psi|\psi\rangle$ is square integrable as a function on G.

Under these three conditions, a G-adapted wavelet analysis on Y may be constructed, following the general construction of coherent states on Y associated to G (Ali, 95) (Ali, 99).

Choose a fixed admissible vector $\psi \in \mathcal{H}$ as analyzing wavelet (normalized to $c_\psi = 1$ for convenience). Then the wavelets are the vectors $\psi_g = U(g)\psi \in \mathcal{H}\,(g \in G)$, and the corresponding continuous wavelet transform (CWT) is defined as the linear map $W_\psi : \mathcal{H} \to L^2(G, dg)$ given by $(W_\psi s)(g) \equiv S_\psi(g) = \langle \psi_g|s\rangle$ (dg denotes the left invariant measure on G). The CWT has the following properties:

(1) *Energy conservation:*

$$\int_G |S_\psi(g)|^2\, dg = \int_Y |s(y)|^2\, d\mu(y), \tag{6.3}$$

i.e., W_ψ is an isometry (in other words, it generates a resolution of the identity); hence its range, the space of wavelet transforms, is a *closed* subspace $\mathcal{H}_\psi$ of $L^2(G, dg)$.

(2) By (1), W_ψ may be inverted on its range by the transposed map, which gives the *reconstruction formula:*

$$s(y) = \int_G S_\psi(g)\psi_g(y)\, dg. \tag{6.4}$$

(3) The projection from $L^2(G, dg)$ onto $\mathcal{H}_\psi$ is an integral operator with kernel $K(g, g') = \langle \psi_g | \psi_{g'} \rangle$, that is, the auto-correlation function of ψ, also called a *reproducing kernel*; in other words, a function $f \in L^2(G, dg)$ is a WT iff it satisfies the reproducing relation:

$$f(g) = \int_G \langle \psi_g | \psi_{g'} \rangle f(g')\, dg'. \tag{6.5}$$

(4) The CWT is *covariant* under the action of the group G:

$$W_\psi[U(g)s](g_o) = (W_\psi s)(g^{-1}g_o),\ \forall\, g \in G. \tag{6.6}$$

Moreover, if the analyzing wavelet ψ has a nontrivial stability subgroup H_ψ (even up to a phase), the whole construction may be performed under a slightly less restrictive condition (the representation U need only to be square integrable on the coset space $X = G/H_\psi$), and yields wavelets indexed by the points of X. We will encounter this situation both in the 2-D and in the 3-D case. One may even allow the parameter space to be a general coset space $X = G/H$, but then technical difficulties arise. We refer the reader to (Ali, 95) (Ali, 99) for a thorough discussion of this extended theory.

This formalism is general enough to design a symmetry-adapted CWT in all cases of physical interest, while, of course, reproducing the familiar 1-D CWT. First, one should notice that the Weyl-Heisenberg group, which consists of phase space translations (translations and modulations), yields the STFT (the corresponding wavelets are called gaborettes in the wavelet community, while quantum physicists call them canonical coherent states (Klauder, 85) (Ali, 95) (Ali, 99)). In the following sections, we will apply this technique and obtain the extension of the CWT to 2 and 3 space dimensions, to space-time (time-dependent signals or images, such as TV or video sequences), including relativistic effects (using wavelets associated to the affine Galilei or Poincaré group) and the 2-sphere (a tool most wanted by geophysicists).

7. THE CWT IN TWO DIMENSIONS

According to the discussion above, the first thing to do for extending the CWT to 2-D case, is to identify the relevant transformation group and a suitable representation in the space of signals. The discussion below follows (Antoine, 99a), where further details may be found.

7.1. Mathematical properties

By an image, we mean a 2-D signal of finite energy, represented by a complex-valued, square integrable function $s \in L^2(\mathbb{R}^2, d^2\boldsymbol{x})$. This condition may be relaxed, to allow, for instance, a plane wave or a δ function. In practice, a black and white image will be represented by a bounded non-negative function: $0 \leq s(\boldsymbol{x}) \leq M$, $\forall\, \boldsymbol{x} \in \mathbb{R}^2$ $(M > 0)$, the discrete values of $s(\boldsymbol{x})$ corresponding to the level of gray of each pixel.

All the operations we will apply to a signal s are obtained by combining three elementary transformations of the plane, namely, translations, dilations and rotations. These transformations are represented by the following unitary operators in the space $L^2(\mathbb{R}^2, d^2\boldsymbol{x})$ of finite energy signals:

(1) translation : $(T^{\boldsymbol{b}}s)(\boldsymbol{x}) = s(\boldsymbol{x} - \boldsymbol{b}),\ \boldsymbol{b} \in \mathbb{R}^2$,
(2) dilation : $(D^a s)(\boldsymbol{x}) = a^{-1}s(a^{-1}\boldsymbol{x}),\ a > 0$,
(3) rotation : $(R^\theta s)(\boldsymbol{x}) = s(r_{-\theta}(\boldsymbol{x})),\ \theta \in [0, 2\pi)$,

where $\boldsymbol{b} \in \mathbb{R}^2$ is the translation parameter, $a > 0$ the dilation parameter, θ the rotation angle and the rotation $r_\theta \in SO(2)$ acts on $\boldsymbol{x} = (x, y)$ as usual :

$$r_\theta(\boldsymbol{x}) = (x\cos\theta - y\sin\theta, x\sin\theta + y\cos\theta),\ 0 \le \theta < 2\pi. \tag{7.1}$$

Combining now the three operators, we define the unitary operator $U(\boldsymbol{b}, a, \theta) = T^{\boldsymbol{b}} D^a R^\theta$, which acts on a given function s as:

$$(U(\boldsymbol{b}, a, \theta)s)(\boldsymbol{x}) \equiv s_{\boldsymbol{b},a,\theta}(\boldsymbol{x}) = a^{-1}\, s(a^{-1}\, r_{-\theta}(\boldsymbol{x} - \boldsymbol{b})), \tag{7.2}$$

or, equivalently, in the space of Fourier transforms:

$$\widehat{s_{\boldsymbol{b},a,\theta}}(\boldsymbol{k}) = a\, e^{-i\boldsymbol{b}\cdot\boldsymbol{k}}\, \widehat{s}(a r_{-\theta}(\boldsymbol{k})). \tag{7.3}$$

Clearly we are following the general pattern described in Section 6. The three operations of translation, (global) dilation and plane rotation constitute the similitude group of the plane, $G \equiv SIM(2) = \mathbb{R}^2 \rtimes (\mathbb{R}_*^+ \times SO(2))$ ($\rtimes$ denotes a semidirect product), U is the natural representation of $SIM(2)$ in the Hilbert space $L^2(\mathbb{R}^2)$, and it is unitary, irreducible and square integrable. The admissibility condition for a wavelet $\psi \in L^2(\mathbb{R}^2)$ reads:

$$c_\psi \equiv (2\pi)^2 \int |\widehat{\psi}(\boldsymbol{k})|^2\, \frac{d^2\boldsymbol{k}}{|\boldsymbol{k}|^2} < \infty. \tag{7.4}$$

If ψ is regular enough ($\psi \in L^1(\mathbb{R}^2) \cap L^2(\mathbb{R}^2)$ suffices), the admissibility condition simply means that the wavelet has zero mean:

$$\widehat{\psi}(\boldsymbol{0}) = 0 \iff \int \psi(\boldsymbol{x})\, d^2\boldsymbol{x} = 0. \tag{7.5}$$

Clearly the three unitary operators $T^{\boldsymbol{b}}, D^a, R^\theta$ preserve the admissibility condition, and so does therefore $U(\boldsymbol{b}, a, \theta)$. Hence, any function $\psi_{\boldsymbol{b},a,\theta} = U(\boldsymbol{b}, a, \theta)\psi$ obtained from a wavelet ψ by translation, rotation or dilation is again a wavelet. Thus the given wavelet ψ generates the whole family $\{\psi_{\boldsymbol{b},a,\theta}\}$, indexed by the elements $(\boldsymbol{b}, a, \theta) \in SIM(2)$, and the linear span of this family is dense in $L^2(\mathbb{R}^2)$.

Now we have only to follow the general scheme of Section 6. Given an image $s \in L^2(\mathbb{R}^2)$, its WT with respect to the wavelet ψ is:

$$\begin{aligned} S(\boldsymbol{b}, a, \theta) &= \langle \psi_{\boldsymbol{b},a,\theta} | s \rangle \\ &= a^{-1} \int \overline{\psi(a^{-1} r_{-\theta}(\boldsymbol{x} - \boldsymbol{b}))}\, s(\boldsymbol{x})\, d^2\boldsymbol{x}. \qquad (7.6) \\ &= a \int e^{i\boldsymbol{b}\cdot\boldsymbol{k}}\, \overline{\widehat{\psi}(a r_{-\theta}(\boldsymbol{k}))}\, \widehat{s}(\boldsymbol{k})\, d^2\boldsymbol{k}. \qquad (7.7) \end{aligned}$$

The wavelet $\psi \in L^2(\mathbb{R}^2)$ and its Fourier transform $\widehat{\psi}$ are both supposed to be well localized. In addition, ψ is often required to have a certain number of vanishing moments.

Then we may summarize as follows the main properties of the wavelet transform $W_\psi : s \mapsto S$:

- W_ψ conserves the energy of the signal:

$$c_\psi^{-1} \iiint |S(\boldsymbol{b}, a, \theta)|^2\, \frac{da}{a^3} d\theta d^2\boldsymbol{b} = \int |s(\boldsymbol{x})|^2\, d^2\boldsymbol{x}, \tag{7.8}$$

where $dg \equiv a^{-3} da d\theta d^2\boldsymbol{b}$ is the natural invariant measure on $SIM(2)$.

- Reconstruction formula:

$$s(\boldsymbol{x}) = c_\psi^{-1} \iiint \psi_{\boldsymbol{b},a,\theta}(\boldsymbol{t})\, S(\boldsymbol{b}, a, \theta)\, \frac{da}{a^3}\, d\theta\, d^2\boldsymbol{b}\ . \tag{7.9}$$

In other words, the 2-D CWT provides a decomposition of the signal in terms of the analyzing wavelets $\psi_{\boldsymbol{b},a,\theta}$, with coefficients $S(\boldsymbol{b},a,\theta)$. As in 1-D, one can also reconstruct the signal by resumming only over scales:

$$s(\boldsymbol{x}) \sim \int S(a,\boldsymbol{x})\ \frac{da}{a^2}. \tag{7.10}$$

- Reproducing kernel:

$$K(\boldsymbol{b}',a',\theta'|\boldsymbol{b},a,\theta) = c_\psi^{-1}\ \langle \psi_{\boldsymbol{b}',a',\theta'}|\psi_{\boldsymbol{b},a,\theta}\rangle, \tag{7.11}$$

and corresponding reproduction property:

$$S(\boldsymbol{b}',a',\theta') = \iiint K(\boldsymbol{b}',a',\theta'|\boldsymbol{b},a,\theta)\ S(\boldsymbol{b},a,\theta)\ \frac{da}{a^3}\, d\theta\, d^2\boldsymbol{b}. \tag{7.12}$$

- W_ψ is *covariant* under translations, dilations and rotations (Murenzi, 90a) (Murenzi, 90b).

7.2. The various representations

The first problem one faces in practice is one of visualization. Indeed, $S(\boldsymbol{b},a,\theta)$ is a function of four variables, two position variables $\boldsymbol{b} \in \mathbb{R}^2$, and the pair $(a,\theta) \in \mathbb{R}_*^+ \times [0,2\pi) \simeq \mathbb{R}_*^2$. It turns out that this splitting has an intrinsic geometrical meaning. Indeed, it can be shown (Antoine, 93) (Antoine, 96) that the pair (a^{-1},θ) plays the role of spatial frequency, expressed in polar coordinates, and so the full four-dimensional parameter space of the 2-D WT may be interpreted as a phase space, in the sense of classical mechanics.

Now, to compute and visualize the full CWT in all four variables is hardly possible. Therefore, in order to obtain a manageable tool, some of the variables, a,θ,b_x,b_y must be frozen. In other words, one must restrict oneself to a *section* of the parameter space. The geometrical considerations made above indicate that two of them are more natural: Either (a,θ) or (b_x,b_y) are fixed, and the WT is treated as a function of the two remaining variables. The corresponding representations are called *position representation* and the *scale-angle representation*, respectively. Whereas the former is useful for the general purpose of image processing, the latter is particularly interesting whenever scaling behavior (as in fractals) or angular selection is important, for instance, when directional wavelets are used. In fact, both representations are needed for a full understanding of the properties of the CWT in all four variables, and the reproducing kernel K should be studied in both (see Section 7.4 below and (Antoine, 93) (Antoine, 96)).

For the numerical evaluation, discretization of the WT in any of these representations, and systematic use of the FFT algorithm will lead to a numerical complexity of $3N_1N_2\log(N_1N_2)$, where N_1,N_2 denote the number of sampling points in the two remaining variables.

Whichever representation we use, we end up with a function of two variables, which may be real or complex. In the latter case, it will be often represented through its modulus and phase. It turns out that the phase is particularly instructive, as was already the case in 1-D. We refer to (Antoine, 93) for a more detailed discussion.

7.3. Choice of the analyzing wavelet

The next step is to choose an analyzing wavelet ψ. As this point, there are two possibilities, depending on the problem at hand.

(i) *Isotropic wavelets:*

If one wants to perform a pointwise analysis, that is, when no oriented features are present or relevant in the signal, one may choose an analyzing wavelet ψ which is invariant under rotation. Then the θ dependence drops out, for instance, in the reconstruction formula (7.9) (in the mathematical formulation of Section 6, the isotropy subgroup of ψ is $H_\psi = SO(2)$ and one gets $X = \mathbb{R}^2 \times \mathbb{R}_*^+$). A typical example is the isotropic 2-D Mexican hat wavelet.

(ii) *Anisotropic wavelets:*
When the aim is to detect directional features in an image, for instance, to perform directional filtering, one has to use a wavelet which is *not* rotation invariant (H_ψ trivial). The best angular selectivity will be obtained if ψ is *directional*, which means that its (essential) support in spatial frequency space is contained in a convex cone with apex at the origin. Typical directional wavelets are the 2-D Morlet wavelet or the Cauchy wavelets (Antoine, 96) (Antoine, 99c).

Let us in more detail examine some examples of wavelets of each kind.

7.3.1 Isotropic wavelets

• *The 2-D Mexican hat or Marr wavelet:*
In its isotropic version, this is simply the Laplacian of a Gaussian:

$$\psi_H(\boldsymbol{x}) = (2 - |\boldsymbol{x}|^2)\ \exp(-\tfrac{1}{2}|\boldsymbol{x}|^2). \tag{7.13}$$

This is a real, rotation invariant wavelet, originally introduced by Marr (Marr, 82). There exists also an anisotropic version, obtained by replacing in (7.13) $\boldsymbol{x}$ by $A\boldsymbol{x}$, where $A = \mathrm{diag}[\epsilon^{-1/2}, 1], \epsilon \geq 1$, is a 2×2 anisotropy matrix. However, this wavelet is of little use in practice, because it is still *not* a directional wavelet, in the technical sense defined below. Hence the Mexican hat will be efficient for a fine pointwise analysis, but not for detecting directions.

• *Difference wavelets:*
An interesting class consists of wavelets obtained as the difference of two positive functions, for instance a single function h and a contracted version of the latter. If h is a smooth nonnegative function, integrable and square integrable, with all moments of order one vanishing at the origin, then the function ψ given by the relation :

$$\psi(\boldsymbol{x}) = \alpha^{-2}\ h(\alpha^{-1}\boldsymbol{x}) - h(\boldsymbol{x})\ \ (0 < \alpha < 1) \tag{7.14}$$

is easily seen to be a wavelet satisfying the admissibility condition (7.5). A typical example is the Difference-of-Gaussians or DOG wavelet, obtained by taking for h a Gaussian

$$\psi_D(\boldsymbol{x})\ =\ \alpha^{-2}\ e^{-|\boldsymbol{x}|^2/2\alpha^2} - e^{-|\boldsymbol{x}|^2/2},\ \ (0 < \alpha < 1). \tag{7.15}$$

The DOG filter is a good substitute for the Mexican hat (for $\alpha^{-1} = 1.6$, their shapes are extremely similar), frequently used in psychophysics works (De Valois, 88) (Duval-Destin, 91). Notice that h, and thus also ψ, need not be isotropic.

7.3.2 Directional wavelets

If one wants to detect oriented features (segments, edges, vector field,...), one needs a wavelet which is directionally selective. To be precise, we will say that a given wavelet ψ is *directional* if the effective support of its Fourier transform $\widehat{\psi}$ is contained in a convex cone in spatial frequency space $\{\boldsymbol{k}\}$, with apex at the origin.

According to this definition, the anisotropic Mexican hat is not directional, since the support of $\widehat{\psi}_H$ is centered at the origin, no matter how big its anisotropy is, as can be seen on Figure 7. Indeed, detailed tests confirm its poor performances in selecting directions (Antoine, 93).

• *The 2-D Morlet wavelet*
This is the prototype of a directional wavelet:

$$\psi_M(\boldsymbol{x}) = \exp(i\boldsymbol{k}_o \cdot \boldsymbol{x})\ \exp(-\tfrac{1}{2}|A\boldsymbol{x}|^2), \quad \widehat{\psi}_M(\boldsymbol{k}) = \sqrt{\epsilon}\ \exp\left(-\tfrac{1}{2}[\epsilon k_x^2 + (k_y - k_o)^2]\right). \tag{7.16}$$

The parameter $\boldsymbol{k}_o$ is the wave vector, and A the anisotropy matrix as above. As in 1-D, we should add a correction term to (7.16) to enforce the admissibility condition $\widehat{\psi}_M(\boldsymbol{0}) = 0$. However, since it is

numerically negligible for $|\boldsymbol{k}_o| \geq 5.6$, we have dropped it altogether. The modulus of the (truncated) wavelet ψ_M is a Gaussian, elongated in the x direction if $\epsilon > 1$, and its phase is constant along the direction orthogonal to $\boldsymbol{k}_o$. Thus the wavelet ψ_M smoothes the signal in all directions, but detects the sharp transitions in the direction perpendicular to $\boldsymbol{k}_o$. The angular selectivity increases with $|\boldsymbol{k}_o|$ and with the anisotropy ϵ. The best selectivity will be obtained by combining the two effects, i.e., by taking $\boldsymbol{k}_o = (0, k_o)$. The function $\widehat{\psi}_M$, with $\epsilon = 5$ and rotated by $\theta = 45^o$, is shown (in level curves) in Figure 7 . Its effective support is centered at $\boldsymbol{k}_o$ and is contained in a convex cone, that becomes narrower as ϵ increases.

• *Conical or Cauchy wavelets:*

In order to achieve a genuinely oriented wavelet, it suffices to consider a smooth function $\widehat{\psi}_S(\boldsymbol{k})$ with support in a strictly convex cone S in spatial frequency space and behaving inside S as $P(k_x, k_y)e^{-(\boldsymbol{\alpha}\cdot\boldsymbol{k})}$, with $\boldsymbol{\alpha} \in S$, or $P(k_x, k_y)e^{-|\boldsymbol{k}|^2}$, where $P(.)$ denotes a polynomial in two variables. Examples of the first type are given in (Antoine, 96) (Antoine, 99c), under the name of Cauchy wavelets.

7.4. Performance evaluation of the 2-D CWT

Given a wavelet, what are its angular and scale selectivity (resolving power)? What is the minimal sampling grid for the reconstruction formula (7.9) that guarantees that no information is lost? The answer to both questions resides in a *quantitative* knowledge of the properties of the wavelet, that is, the tool must be *calibrated*.

To that effect, one takes the WT of particular, standard signals. Three such tests are useful, and in each case the outcome may be viewed either at fixed (a, θ) (position representation) or at fixed $\boldsymbol{b}$ (scale-angle representation).

• *Point signal:* for a snapshot at the wavelet itself, one takes as the signal a δ function, i.e., one evaluates the impulse response of the filter:

$$\langle \psi_{\boldsymbol{b},a,\theta} | \delta \rangle = a^{-1}\, \overline{\psi(a^{-1} r_{-\theta}(-\boldsymbol{b}))}. \tag{7.17}$$

• *Benchmark signals:* for testing particular properties of the wavelet, such as its ability to detect a discontinuity or its angular selectivity in detecting a particular direction, one may use appropriate 'benchmark' signals.

• *Reproducing kernel:* taking as the signal the wavelet ψ itself, one obtains the reproducing kernel K, which measures the *correlation length* in each variable $\boldsymbol{b}, a, \theta$:

$$c_\psi\, K(\boldsymbol{b}, a, \theta | 1, 0, \boldsymbol{0}) = \langle \psi_{\boldsymbol{b},a,\theta} | \psi \rangle = a^{-1} \int \overline{\psi(a^{-1} r_{-\theta}(\boldsymbol{x} - \boldsymbol{b}))}\, \psi(\boldsymbol{x})\, d^2\boldsymbol{x}. \tag{7.18}$$

A detailed analysis of K yields a definition of the resolving power of the wavelet ψ in each variable, and the result fits perfectly with empirical definitions (Antoine, 93).

7.4.1 Calibration of a wavelet with benchmark signals

As we said above, particular properties of the wavelet may be tested on appropriate benchmark signals. For instance, its capacity at detecting a discontinuity may be measured on the signal consisting of an infinite rod (see (Antoine, 93) for the full discussion). The result is that both the Mexican hat and the Morlet wavelet are efficient in this respect.

For testing the angular selectivity of a wavelet, one computes the WT of a semi-infinite rod, as a function of the difference in orientation, $\Delta\theta$, between the wavelet and the rod. First a Morlet wavelet ψ_M with excentricity $\epsilon = 5$ detects the orientation of a rod with a precision of the order of 5°. That is, the WT of the rod is a "wall" of constant height as long as $\Delta\theta$ is smaller than 5°. For increasing $\Delta\theta$, the

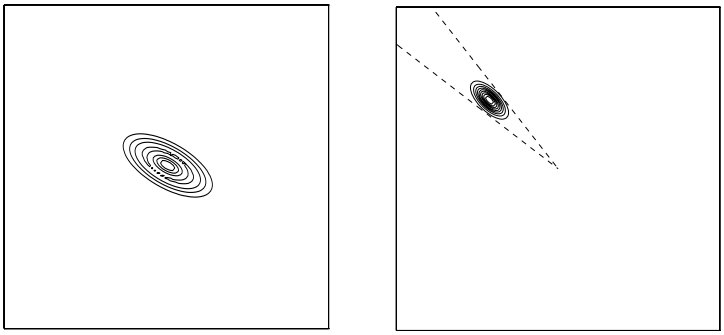

Figure 7: Frequency space realization of two basic wavelets ($\epsilon = 5, \theta = 45^o$): (left) The anisotropic Mexican hat $\widehat{\psi}_H$; (right) The Morlet wavelet $\widehat{\psi}_M$.

wall gradually collapses, and essentially disappears for $\Delta\theta > 15°$. Only the tip of the rod remains visible, and for large $\Delta\theta$, it gives a sharp peak (it is a point singularity). Essentially the same result is obtained with a Cauchy wavelet supported in a cone of opening angle 20°, whereas the same test performed with an anisotropic Mexican hat gives a result almost independent of $\Delta\theta$. The conclusion is that the Morlet and the Cauchy wavelets are highly sensitive to orientation, but the Mexican hat is not.

Let us illustrate these features by two academic examples, both of which have direct physical applications.

(a) Detection of contours with the Mexican hat:
Exactly as in the 1-D case, the WT is especially useful to detect *discontinuities* in images, for instance the *contour* or the *edges* of an object (Murenzi, 90b) (Antoine, 93). For that purpose, an isotropic wavelet may be chosen, such as the radial Mexican hat ψ_H given in (7.13). In that case the effect of the WT consists in smoothing the signal with a Gaussian and taking the Laplacian of the result. Thus large values of the amplitude will appear at the location of the discontinuities, in particular the contour of objects (which is a discontinuity in luminosity).

In order to test this property, we may compute the WT of a simple object, such as a square or a set with the shape of a thick letter, represented by its characteristic function, for various values of the scale parameter a. Then, for large values of a, the WT sees only the object as a whole, thus allowing the determination of its position in the plane. When a decreases, increasingly finer details appear. In this simple case, the WT vanishes both inside and outside the contour, since the signal is constant there, and thus only the contour remains. Of course, if one takes values of a that are too small, numerical artifacts (aliasing) appear and spoil the result. It should be noted that the corners of the figure are highlighted in the WT by sharp peaks. The amplitude is larger at these points, since the signal is singular there in *two* directions, as opposed to the edges. In addition the WT detects the *convexity* of each corner. Convex corners give rise to positive peaks, whereas concave ones yield a negative peak. Here we see the advantage of using a real wavelet and plotting the WT itself, *not* its modulus, which is a frequent practice.

(b) Directional filtering with a Morlet or a Cauchy wavelet:

As a consequence of their good directional selectivity, the Morlet and the Cauchy wavelets are quite efficient for directional filtering. In order to illustrate the point, we analyze in Figure 8 a pattern made of rods in many different directions (a). Applying the CWT with a fixed direction, here horizontal, selects all those rods with roughly the same direction (b), wheras the other ones, which are misaligned, yield only a faint signal corresponding to their tips, in agreement with the behavior discussed above. Since this is in fact noise, one performs a thresholding to remove it, thus getting a clear picture (c). The same

Figure 8: Directional filtering with the CWT: (a) The pattern; (b) CWT with a Cauchy wavelet supported in $\mathcal{C}(-10°, 10°)$ (c) The same after thresholding.

two operations are then repeated with various successive orientations of the wavelet. In this way, one can count the number of objects that lie in any particular direction.

7.5. Physical applications of the 2-D CWT

The 2-D CWT has been used by a number of authors, in a wide variety of physical problems (Combes, 90) (Meyer, 91) (Meyer, 93). In all cases, its main use is for the *analysis* of images. It can be used for the detection of specific features, such as a hierarchical structure, edges, filaments, contours, boundaries between areas of different luminosity, etc. Of course, the type of wavelet chosen depends on the precise aim. An isotropic wavelet (Mexican hat) suffices for pointwise analysis, but an oriented wavelet (Morlet, Cauchy) is more efficient for the detection of oriented features in the signal, that is, regions where the amplitude is regular along one direction and has a sharp variation along the perpendicular direction.

Among the most successful applications, we may select the following ones.

• *Analysis of astronomical images:*
The CWT has been used for several purposes: noise filtering (background sky), with a technique known is astronomy as 'unsharp masking', unraveling of the hierarchical structure of a galactic nebula, or that of the universe itself (galaxy counts, detection of galaxy clusters or voids).

• *Analysis of 2-D fractals:*
Fractals, either artificial (numerical snowflakes, diffusion limited aggregates) or natural (electrodeposition clusters, various arborescent phenomena, clouds) have been thoroughly investigated with the 2-D CWT (Argoul, 90) (Arnéodo, 91); here the (a, θ) representation is useful, since it presents the signal at all scales at once. Particular applications include the measurement of the fractal dimensions and the unraveling of universal laws (mean angle between branches, azimuthal Cantor structures, etc.).

• *Turbulence in fluid dynamics:*
Analysis of 2-D developed turbulence in fluids, especially localization of small scales in the distribution of energy or enstrophy (Farge, 92). Other applications in fluid dynamics include the visualization and measurement of a velocity field with help of an oriented wavelet (see the discussion below).

• *Medical physics and psychophysics:*
Modelling of human vision, e.g., definition of local contrast in images (Duval-Destin, 91), medical imaging, in particular 2-D NMR imaging and tomography.

• *Determination of local regularity of a signal:*
The tool here is the estimation of local Lipschitz exponents (Mallat, 92).

Among the applications of 2-D wavelets in fluid dynamics, we would like to mention two, which both rely on the possibility of directional filtering with directional wavelets. The first one is a straightforward application of the method described in Section 7.4.1 (Wisnoe, 93a) (Wisnoe, 93b). The aim is to measure the velocity field of a 2-D turbulent flow around an obstacle. Velocity vectors are materialized by small segments (tiny plastic balls are injected into the stream and photographed with a fast CCD camera). Then the WT with a Morlet wavelet is computed twice. First the WT selects those vectors that are closely aligned with the wavelet. Then the analysis is repeated with a wavelet oriented in the orthogonal direction, thus completely misoriented with respect to the selected vectors. Now the WT sees only the tips of the vectors and their length may be easily measured. Using appropriate thresholdings, as in Figure 8, and repeating the operation a certain number of times, with the wavelet rotated each time by a fixed angle (typically 10°), the complete velocity field may thus be obtained, in a totally automated fashion, with an efficiency sensibly better than with more traditional methods.

A second example concerns the disentangling of a wave train, represented by a linear superposition of damped plane waves. The problem originates from underwater acoustics. When a point source emits a sound wave above the surface of water, the wave hitting the surface splits into several components of very different characteristics (called respectively direct, lateral and transient). The goal is to measure the parameters of all components. This phenomenon has been analyzed successfully with the WT both in 1-D (Saracco, 90) and in 2-D (Antoine, 96). In the latter case, the signal representing the underwater wave train is taken as a linear superposition of damped plane waves:

$$f(\boldsymbol{x}) = \sum_{n=1}^{N} c_n \, e^{i\boldsymbol{k}_n \cdot \boldsymbol{x}} \, e^{-\boldsymbol{l}_n \cdot \boldsymbol{x}}, \tag{7.19}$$

where, for each component, $\boldsymbol{k}_n$ is the wave vector, $\boldsymbol{l}_n$ is the damping vector, and c_n a complex amplitude. Then, using successively the scale-angle and the position representations, one is able to measure all the $6N$ parameters of this signal with remarkable ease and precision. And the extension to a 3-D version is straightforward.

7.6. The 2-D discrete WT

In fact, the first extension of wavelet analysis to 2-D signals was proposed by Mallat (Mallat, 89a) (Mallat, 89b), who developed systematically the 2-D *discrete* WT. This generalization is indeed a very natural one, if one realizes that the whole idea of multiresolution analysis lies at the heart of human vision. In fact, most of the concepts are indeed already present in the pioneering work of Marr (Marr, 82) on vision modelling. As in 1-D, this discrete WT has a close relationship with numerical filters (QMFs and their generalizations) and related techniques of signal analysis, such as subband coding. It has been applied successfully to several standard problems of image processing. As a matter of fact, all the approaches that we have mentioned above in the 1-D case have been extended to 2-D: orthonormal bases, biorthogonal bases, wavelet packets, lifting scheme.

However, the 2-D DWT suffers from intrisic limitations, because it is essentially linked to a Cartesian geometry (except for the lifting scheme). Indeed, a 2-D multiresolution is simply the tensor product of two 1-D schemes, one for the horizontal direction and one for the vertical direction (in technical terms, one uses only separable filters). And one needs three different analyzing wavelets, corresponding to horizontal, vertical and oblique details, respectively. More isotropic schemes have been designed, based on more complicated dilation operations, such as the so-called 'quincunx' lattice (see for instance (Kovačević, 92)). Yet the CWT remains more flexible, especially for detecting oriented features in an image, for instance, in the classical problem of *edge detection*. Of course, many applications require only a pointwise analysis (fractals are typical), hence the 2-D WT is often used as a 'mathematical microscope', like in 1-D, thus ignoring directions. But if directions are important, the full CWT, including the orientation parameter θ, should be used, and it is a very efficient tool in this respect, provided one uses a wavelet which has itself an intrinsic orientation, that is, a directional wavelet.

8. THE CWT IN HIGHER DIMENSIONS

8.1 The 3-D case

Some physical phenomena are intrinsically multiscale and three-dimensional. Typical examples may be found in fluid dynamics, for instance the appearance of coherent structures in turbulent flows, or the disentangling of a wave train in (mostly underwater) acoustics, as discussed above. In such cases, a 3-D wavelet analysis is clearly more adequate and likely to yield a deeper understanding (Astruc, 93). Hence we will also describe briefly the 3-D CWT, following again the pattern of Section 6.

Given a 3-D signal $s \in L^2(\mathbb{R}^3, d^3\boldsymbol{x})$,with finite energy, one may act on it by translation, dilation and rotation:

$$s_{a,\gamma,\boldsymbol{b}}(\boldsymbol{x}) \equiv [U(a, r(\gamma), \boldsymbol{b})s](\boldsymbol{x}) = a^{-\frac{3}{2}} s(a^{-1}r(\gamma)^{-1}(\boldsymbol{x} - \boldsymbol{b})), \tag{8.1}$$

where $a > 0,\ \gamma \in SO(3),\ \boldsymbol{b} \in \mathbb{R}^3$ and $r(\gamma) \in SO(3)$ is a 3×3 rotation matrix. The element $\gamma \in SO(3)$ may be parametrized, for instance, in terms of three Euler angles. These three operations generate the 3-D Euclidean group with dilations, that is, the similitude group of $\mathbb{R}^3$, $SIM(3) = \mathbb{R}^3 \rtimes (\mathbb{R}^+_* \times SO(3))$. Then (8.1) is a unitary representation of $SIM(3)$ in $L^2(\mathbb{R}^3, d^3\boldsymbol{x})$, which is irreducible and square integrable. Hence, it generates a CWT exactly as before.

Wavelets are taken in $L^2(\mathbb{R}^3, d^3\boldsymbol{x})$ and the admissibility condition is now

$$\int |\widehat{\psi}(\boldsymbol{k})|^2 \, \frac{d^3\boldsymbol{k}}{|\boldsymbol{k}|^3} < \infty. \tag{8.2}$$

Also the two familiar wavelets have a 3-D realization.

• *The 3-D Mexican hat* is given by

$$\psi_{H}(\boldsymbol{x}) = (3 - |A\boldsymbol{x}|^2)\, \exp(-\tfrac{1}{2}|A\boldsymbol{x}|^2). \tag{8.3}$$

where $A = \mathrm{diag}[{\epsilon_1}^{-1/2}, {\epsilon_2}^{-1/2}, 1],\ \epsilon_1 \geq 1, \epsilon_2 \geq 1$, is a 3×3 anisotropy matrix. We distinguish three cases:

(1) If $\epsilon_1 \neq \epsilon_2 \neq 1$, one has the fully anisotropic 3-D Mexican hat. Here, H_ψ is trivial.

(2) If $\epsilon_1 = \epsilon_2 = 1$, one has the isotropic, $SO(3)$-invariant, 3-D Mexican hat. Now, $H_\psi = SO(3)$.

(3) If $\epsilon_1 = \epsilon_2 \equiv \epsilon \neq 1$, the wavelet is *axisymmetric*, i.e., $SO(2)$-invariant, but not isotropic, so that $H_\psi = SO(2)$.

• *The 3-D Morlet wavelet* is given by

$$\psi(\boldsymbol{x}) = \exp(i\boldsymbol{k}_o \cdot \boldsymbol{x})\, \exp(-\tfrac{1}{2}|A\boldsymbol{x}|^2), \tag{8.4}$$

where A is the same 3×3 anisotropy matrix as in the first example. Here again, for $\epsilon_1 = \epsilon_2 \equiv \epsilon \neq 1$ and $\boldsymbol{k}_o$ along the z-axis, the wavelet ψ is invariant under $SO(2)$.

Then, given a signal $s \in L^2(\mathbb{R}^3)$, its CWT with respect to the admissible wavelet ψ is given as

$$S(\boldsymbol{b}, a, \gamma) = a^{-3/2} \int \overline{\psi(a^{-1}r(\gamma)^{-1}(\boldsymbol{x} - \boldsymbol{b}))}\, s(\boldsymbol{x})\, d^3\boldsymbol{x}. \tag{8.5}$$

As compared with (7.6), the only differences are in the normalization factors and the rotation matrices. Since the structure of the formulas is the same as before, so are the interpretation and the consequences (local filtering, reproducing kernel, reconstruction formula, etc.). Thus the CWT (8.5) may be interpreted as a mathematical *camera* with *magnification* $1/a$, *position* $\boldsymbol{b}$ and *directional selectivity*, given, in the axisymmetric case, by the rotation parameters $\varpi \equiv (\theta, \varphi)$. As for the visualization, the full CWT $S(a, \gamma, \boldsymbol{b})$ is a function of 7 variables. However, if the wavelet ψ is chosen *axisymmetric*, i.e., $SO(2)$-invariant, S depends on 6 variables only, $a > 0,\ \varpi \in S^2 \simeq SO(3)/SO(2)$, the unit sphere in $\mathbb{R}^3$, and $\boldsymbol{b} \in \mathbb{R}^3$. In this case again, (a^{-1}, ϖ) may be interpreted as polar coordinates in spatial frequency space

(this is in fact true in any number of dimensions). It follows that, here too, there are two natural representations for the visualization of the WT, the position representation (a, ϖ fixed) and the scale-orientation (or spatial frequency) representation ($\boldsymbol{b}$ fixed). Of course, there are many other posssible representations that may be useful.

In conclusion, let us discuss briefly a simple example, the detection of 3-D objects in a cluttered medium. We consider a scene with 3-D objects (targets) immersed in a cluttered medium, modeled by the signal:

$$s(\boldsymbol{x}) = \sum_{m=1}^{N} s_m(\boldsymbol{x}) + n(\boldsymbol{x}), \tag{8.6}$$

where $s_m(\boldsymbol{x})$ denotes the density of the target m, and $n(\boldsymbol{x})$ the density of the medium.

Since the density of the targets is very different from that of the medium, there will be a high density gradient at the boundary between the objects and the medium. In this situation, the wavelet transform $S(\boldsymbol{b}, a, \theta, \varphi)$ may be used to extract the 3-D objects and determine their characteristics, position (range and orientation) and spatial frequency. Further details may be found in (Antoine, 95b), where a detailed strategy is explained for the 2-D version of the same problem.

8.2 Wavelets on the sphere S^2

A recurring problem for applications (in geophysics, for instance) is the extension of wavelet analysis to the sphere (or more general manifolds). On the discrete side, an efficient solution may be obtained with the so-called lifting scheme (Schröder, 95), but this obviously misses the particular symmetry of the sphere. It turns out that the general formalism developed in Section 6 yields an elegant solution to the problem.

As usual, finite energy signals are taken in $L^2(S^2, d\varpi)$, where $d\varpi = \sin\theta\, d\theta\, d\phi$ denotes the standard measure on the sphere. The natural operations on such signals are translations (on the sphere) and local dilations. The former are given by rotations from $SO(3)$. Dilations around the North Pole are obtained by considering ordinary dilations in the tangent plane and lifting them to S^2 by stereographic projection from the South Pole. Thus, a dilation by a becomes a nonlinear map $\theta \mapsto \theta_a$ acting on the azimuthal angle:

$$\tan\frac{\theta_a}{2} = a\, \tan\frac{\theta}{2}\ . \tag{8.7}$$

As for dilations around any other point $\varpi \in S^2$, it suffices to bring ϖ to the North Pole by a rotation, perform the dilation and go back by the inverse rotation. In this way one defines the parameter space of the putative CWT on the sphere as $X_S = SO(3) \cdot \mathbb{R}^+_*$ (Holschneider, 96). The problem is that this set is not a group. Obviously translations and dilations do not commute. However, the only group that can be obtained by combining only $SO(3)$ and the dilation group $\mathbb{R}^+_*$ is their direct product, which cannot be the similitude group of the sphere.

A way out of this difficulty (Antoine, 99b) is to embed the two groups into the Lorentz group $SO_o(3,1)$, using the *Iwasawa decomposition*, $SO_o(3,1) = SO(3) \cdot \mathbb{R}^+_* \cdot N$, where $N \sim \mathbb{C}$ is two-dimensional and abelian. (The Lorentz group is the conformal group of the sphere S^2.) The isotropy subgroup of the North Pole is the so-called minimal parabolic subgroup of $SO_o(3,1)$, namely $P = SO(2) \cdot \mathbb{R}^+_* \cdot N$, where $SO(2)$ consists of rotations around the z-axis. Hence $SO_o(3,1)/P \simeq SO(3)/SO(2) \simeq S^2$ and the group $SO_o(3,1)$ acts transitively on S^2 (this action is obtained easily from the Iwasawa decomposition). The manifold $X_S = SO_o(3,1)/N$ may then be lifted to $SO_o(3,1)$ by the map (section) $\sigma : X_S \to SO_o(3,1)$, given as

$$\sigma(\gamma, a) = \gamma\, a,\ \gamma \in SO(3),\ a \in \mathbb{R}^+_*. \tag{8.8}$$

Clearly, one is in a more general situation than the one described in Section 6 and the general machinery developed in (Ali, 95) (Ali, 99) must be used.

There is a natural UIR of $SO_o(3,1)$ in the space of signals $L^2(S^2, d\varpi)$ given by

$$(U(g)f)(\varpi) = \lambda(g,\varpi)^{1/2}\, f(g^{-1}\varpi). \tag{8.9}$$

In this relation, $\lambda(g,\varpi)$ is the correcting factor (Radon-Nikodym derivative), that takes into account the fact that the measure $d\varpi$ on S^2 is not invariant under dilations. Then the construction of the CWT starts by considering the restriction of this representation to X_s, by putting $U_s(\gamma, a) = U(\sigma(\gamma, a))$, in which case one gets simply

$$\lambda(a,\varpi) = 4a^2\left[(a^2-1)\cos\theta + (a^2+1)\right]^{-2}.$$

This representation U_s is infinite dimensional and its restriction to $SO(3)$ decomposes into the direct sum of all the familiar $(2\ell+1)$-dimensional representations, $\ell = 0, 1, \ldots$. It is also square integrable on X_s, but the admissibility condition is somewhat complicated. A necessary condition for the admissibility of a wavelet ψ is that

$$\iint_{S^2} \frac{\psi(\theta,\phi)}{1+\cos\theta}\, \sin\theta\, d\theta d\phi = 0 \tag{8.10}$$

(hence ψ must vanish sufficiently fast when one approaches the South Pole ($\theta = \pi$), which corresponds to the point at infinity under the stereographic projection). This is a zero mean condition, so that we have the filtering effect, as usual. Thus, a genuine CWT on the sphere has been obtained.

An additional bonus is that this CWT has the expected Euclidean limit. By this we mean the following. Consider instead of the unit sphere S^2 a sphere S^2_R of radius R and let $R \to \infty$. Then S^2_R becomes the plane $\mathbb{R}^2$, the group $SO(3)$ becomes the Euclidean group of $\mathbb{R}^2$, so that $\sigma(X_s) \sim SO(3)\cdot\mathbb{R}^+_*$ becomes $SIM(2)$, with corresponding transitive action, and the representation U_s becomes the natural representation (7.2) of $SIM(2)$ (in mathematical terminology, this limiting process is a group contraction). Furthermore, admissible vectors tend exactly to admissible vectors. Therefore, the wavelet analysis on S^2 goes into the usual wavelet analysis in the plane, as developed in Section 7.

9. MOTION ANALYSIS WITH CWT

9.1. Kinematical wavelets

An important aspect of signal and image processing is the analysis of time-dependent or moving signals, e.g., in television, and the CWT may be extended to this case too (Duval-Destin, 93). We consider first motion on the line. Finite energy signals are taken as functions $s(x,t) \in L^2(\mathbb{R}\times\mathbb{R}, dx\,dt)$ The natural transformations on such a signal are translations and dilations in space and time independently, $(x,t) \mapsto (a_1 x + b_1, a_0 t + b_0)$. However it is more convenient to replace the two independent dilations a_1, a_0 by a global dilation a and a so-called speed-tuning transformation c, defined as:

$$\begin{array}{rcl} s(x,t) & \mapsto & a^{-1}s(a^{-1}x, a^{-1}t), \quad a>0; \\ s(x,t) & \mapsto & s(c^{1/2}x, c^{-1/2}t), \quad c>0. \end{array} \tag{9.1}$$

This transformation comes from the physiological characteristics of motion perception by our visual system. In order to be visible, fast moving objects must be wide, and narrow objects must move slowly (for a typical example, think of the inscriptions on a departing train car).

Combining the transformation (9.1) with space and time translations, we obtain the affine group of space-time. This group has a natural unitary irreducible representation in $L^2(\mathbb{R}^2, dx\,dt)$:

$$[U(b_0,b_1,a,c,\epsilon)s](x,t) = \frac{1}{a}\, s\left(\frac{\sqrt{c}}{a}(x-b_1), \frac{\epsilon}{a\sqrt{c}}(t-b_0)\right), \tag{9.2}$$

where (b, τ) denote space-time translations and $\epsilon = \pm 1$ corresponds to time-reflection (this additional operation is needed for irreducibility). In addition, the representation U is square integrable. A wavelet ψ is admissible iff it satisfies the condition

$$\iint \frac{|\widehat{\psi}(k,\omega)|^2}{|k||\omega|}\, dk\, d\omega < \infty. \tag{9.3}$$

From here on, everything follows exactly the general pattern. Thanks to the filtering property in a and c, the resulting CWT (called *kinematical*) is efficient in detecting moving objects. The dilation parameter a catches the size of the target, while the new parameter c adjusts the speed of the wavelet to that of the target. Thus the spatio-temporal CWT is a tool for motion tracking. Clearly there are plenty of applications in which such a technique might be used.

The extension of these considerations to higher dimensions is straightforward. First, in n dimensions, the dilation and speed tuning operations (9.1) become:

$$x \mapsto a^{-1}c^{1/n+1}x, \quad t \mapsto a^{-1}c^{-n/n+1}t. \tag{9.4}$$

Then one has to add rotations, as usual, and follow the general pattern of Section 6.

9.2. Relativistic wavelets

The kinematical wavelets just described may not always be sufficient, depending on the type of signal to be analyzed. One may wish to consider a specific form of movement, i.e., choose a particular relativity group. Three examples may be of interest (we begin again with one space dimension).

(i) *Galilean wavelets:* here we add to the transformations discussed above the Galilei boosts, thus getting $(x, t) \mapsto (a_1 x + a_0 v t + b_1, a_0 t + b_0)$. The resulting group G_{aff}, called the affine Galilei group, is quite complicated. It has a natural unitary representation in the space of finite energy signals, which splits into the direct sum of four irreducible ones. And each of these is square integrable, so that wavelets may be constructed in the usual way. In addition, more restricted wavelets may be constructed by taking as parameter space various quotient spaces G_{aff}/H, where H is *not* the stability subgroup of the basic wavelet. See (Antoine, 99d) for details.

(ii) *Schrödinger wavelets:* one obtains an interesting subclass of the previous one by imposing the relation $a_0 = a_1^2$, so that the transformations leave invariant the Schrödinger (or the heat) equation. Then, the unitary irreducible representation U_G splits into the direct sum of two square integrable UIRs. Thus, once again a CWT is at hand, which may prove useful for describing, for instance, the motion of quantum particles on the line.

(iii) *Poincaré wavelets:* in order to get a CWT in the relativistic regime, it suffices to replace Galilei transformations by Poincaré ones, while of course imposing the relation $a_0 = a_1$ to space and time dilations. The resulting affine Poincaré group has a square integrable unitary irreducible representation, defined on the solid future light cone (Bohnké, 91). The Poincaré wavelets might be useful, for instance, in the presence of electromagnetic fields.

Of course, this analysis extends in a straightforward way to higher dimensions, just by adding rotations.

10. OUTCOME

As a general conclusion, it is fair to say that the wavelet techniques have become an established tool in signal and image processing, both in their CWT and DWT incarnations and their generalizations. We want to emphasize here that the CWT and the DWT have almost opposite properties, hence their ranges of application differ widely too. The CWT is very efficient at detecting specific features in signals

or images, such as in pattern recognition or directional filtering. On the other hand, the DWT and its generalizations are extremely fast and economical. For instance, they yield impressive data compression rates, which is especially useful in image processing, where huge amount of data, mostly redundant, have to be stored and transmitted.

Both are powerful tools, and very flexible ones, thanks to their adaptive character. And both have become a significant element in the standard toolbox of signal processing, which finds its way into an increasing number of reference books and software codes. As a consequence, they have found applications in many branches of physics, such as acoustics, spectroscopy, geophysics, astrophysics, fluid mechanics (turbulence), medical imagery, atomic physics (laser-atom interaction), solid state physics (structure calculations), ... Clearly wavelets are here to stay, and one should expect this trend to continue, with an increasingly diverse spectrum of physical applications.

References

AKANSU, A.N.; M. SMITH, eds. (1995): Subband and Wavelet Transforms: Design and Applications, Kluwer, Dordrecht.

ALDROUBI, A.; M. UNSER, eds. (1996): Wavelets in Medicine and Biology, CRC Press, Boca Raton, FL.

ALEXANDRESCU, M.; D. GIBERT; G. HULOT; J-L. LE MOUËL; G. SARACCO (1995): Detection of Geomagnetic Jerks Using Wavelet Analysis, J. Geophys. Res. 100, p. 12,557–12,572.

ALI, S.T.; J-P. ANTOINE; J-P. GAZEAU; U.A. MUELLER (1995): Coherent States and Their Generalizations: A Mathematical Overview, Reviews Math. Phys. 7, p. 1013–1104.

ALI, S.T.; J-P. ANTOINE; J-P. GAZEAU (1999): Coherent States, Wavelets and Their Generalizations, Springer-Verlag, New York.

ANTOINE, J-P.; P. CARRETTE; R. MURENZI; B. PIETTE (1993): Image Analysis with 2D Continous Wavelet Transform, Signal Process. 31, p. 241–272.

ANTOINE, J-P.; R. MURENZI (1995a): The Continuous Wavelet Transform, from 1 to 3 Dimensions, in (Akansu, 95), p. 149–187.

ANTOINE, J-P.; P. VANDERGHEYNST; K. BOUYOUCEF; R. MURENZI (1995b): Target Detection and Recognition Using Two-Dimensional Isotropic and Anisotropic Wavelets, Automatic Object Recognition V, SPIE Proc., 2485, p. 20–31.

ANTOINE, J-P.; R. MURENZI (1996): Two-Dimensional Directional Wavelets and the Scale-Angle Representation, Signal Process. 53, p. 259–281.

ANTOINE, J-P.; D. BARACHE; R.M. CESAR Jr.; L. da F. COSTA (1997): Shape Characterization with the Wavelet Transform, Signal Process. 62, p. 265–290.

ANTOINE, J-P. (1999a): The 2-D Wavelet Transform, Physical Applications and Generalizations, in (van den Berg, 99), p. 23–75.

ANTOINE, J-P.; P. VANDERGHEYNST (1999b): Wavelets on the 2-Sphere: A Group-Theoretical Approach, Appl. Comput. Harmon. Anal. 7, p. 1–30.

ANTOINE, J-P.; R. MURENZI; P. VANDERGHEYNST (1999c): Directional Wavelets Revisited: Cauchy Wavelets and Symmetry Detection in Patterns, Appl. Comput. Harmon. Anal. 6, p. 314–345.

ANTOINE, J-P.; I. MAHARA (1999d): Galilean Wavelets: Coherent States for the Affine Galilei Group, J. Math. Phys. 40, p. 5956–5971.

ARGOUL, F.; A. ARNÉODO; J. ELEZGARAY; G. GRASSEAU; R. MURENZI (1990): Wavelet Analysis of the Self-Similarity of Diffusion-Limited Aggregates and Electrodeposition Clusters, Phys. Rev. A 41, p. 5537–5560.

ARNÉODO, A.; F. ARGOUL; E. BACRY; J. ELEZGARAY; E. FREYSZ; G. GRASSEAU; J.F. MUZY; B. POULIGNY (1991): Wavelet Transform of Fractals, in (Meyer, 91), p. 286–352.

ASTRUC, D.; L. PLANTIÉ; R. MURENZI; Y. LEBRET; D. VANDROMME (1993): On the Use of the 3D Wavelet Transform for the Analysis of Computational Fluid Dynamics Results, in (Meyer, 93), p. 463–470.

BARACHE, D.; J-P. ANTOINE; J-M. DEREPPE (1997): The Continuous Wavelet Transform, a Tool for NMR Spectroscopy, J. Magn. Reson. A 128, p. 1–11.

van den BERG J.C., ed. (1999): Wavelets in Physics, Cambridge Univ. Press, Cambridge.

BOHNKÉ, G. (1991): Treillis d'Ondelettes Associés aux Groupes de Lorentz, Ann. Inst. H. Poincaré 54, p. 245–259 .

CALDERBANK, A.R.; I. DAUBECHIES; W. SWELDENS; B. L. YEO (1998): Wavelets that Map Integers to Integers, Appl. Comput. Harmon. Anal., 5, p. 332–369.

CHUI, C.K. (1992): An Introduction to Wavelets, Academic Press, New York, London.

COHEN, A.; I. DAUBECHIES; J-C. FEAUVEAU (1992): Biorthogonal Bases of Compactly Supported Wavelets, Commun. Pure Appl. Math. 45, p. 485–560.

COIFMAN, R.R.; Y. MEYER; S. QUAKE; M.V. WICKERHAUSER (1993): Signal Processing and Compression with Wavelet Packets, in (Meyer, 93), p. 77–93.

COMBES, J-M.; A. GROSSMANN; Ph. TCHAMITCHIAN, eds. (1990): Wavelets, Time-Frequency Methods and Phase Space (Proc. Marseille 1987), 2d ed., Springer-Verlag, Berlin.

DAUBECHIES, I.; A. GROSSMANN; Y.MEYER (1986): Painless Nonorthogonal Expansions, J. Math. Phys. 27, p. 1271–1283.

DAUBECHIES, I. (1990): The Wavelet Transform, Time-Frequency Localization and Signal Analysis, IEEE Trans. Inform. Theory 36, p. 961–1005.

DAUBECHIES, I. (1992): Ten Lectures on Wavelets, SIAM, Philadelphia, PA.

DELPRAT, N.; B. ESCUDIÉ; P. GUILLEMAIN; R. KRONLAND-MARTINET; Ph. TCHAMITCHIAN; B. TORRÉSANI (1992): Asymptotic Wavelet and Gabor Analysis: Extraction of Instantaneous Frequencies, IEEE Trans. Inform. Theory 38, p. 644–664.

DE VALOIS R.; K. DE VALOIS (1988): Spatial Vision, Oxford Univ. Press, New York.

DUVAL-DESTIN, M. (1991): Analyse Spatiale et Spatio-Temporelle de la Stimulation Visuelle à l'Aide de la Transformée en Ondelettes, Thèse de Doctorat, Université d'Aix-Marseille II.

DUVAL-DESTIN, M.; R. MURENZI (1993): Spatio-Temporal Wavelets: Application to the Analysis of Moving Patterns, in (Meyer, 93), p. 399–408.

FARGE, M. (1992): Wavelet Transforms and their Applications to Turbulence, Annu. Rev. Fluid Mech. 24, p. 395–457.

GROSSMANN, A.; R. KRONLAND-MARTINET; J. MORLET (1990): Reading and Understanding Continuous Wavelet Transforms, in (Combes, 90), p.2–20.

GUILLEMAIN, P.; R. KRONLAND-MARTINET; B. MARTENS (1991): Estimation of Spectral Lines with the Help of the Wavelet Transform. Applications in N.M.R. Spectroscopy, in (Meyer, 91), p. 38–60.

HEIL, C.; D. WALNUT (1989): Continuous and Discrete Wavelet Transforms, SIAM Review 31, p. 628–666.

HOLSCHNEIDER, M. (1988): On the Wavelet Transformation of Fractal Objects, J. Stat. Phys. 50, p. 963–993.

HOLSCHNEIDER, M. (1995): Wavelets, An Analysis Tool, Oxford U. Press, Oxford.

HOLSCHNEIDER, M. (1996): Continuous Wavelet Transforms on the Sphere, J. Math. Phys., 37, p. 4156–4165.

HUBBARD, B.B. (1998): The World According to Wavelets, 2nd ed., A.K. Peters, Wellesley, MA.

IEEE Trans. Inform. Theory, Special Issue on Wavelet Transforms and Multiresolution Signal Analysis, 38, No.2, March 1992.

Proc. IEEE, Special Issue on Wavelets, 84, No.4, April 1996.

KAISER, G. (1994): A Friendly Guide to Wavelets, Birkhäuser, Boston-Basel-Berlin.

KLAUDER, J.R.; B.S. SKAGERSTAM (1985): Coherent States – Applications in Physics and Mathematical Physics, World Scientific, Singapore.

KOVAČEVIĆ J.; M. VETTERLI (1992): Nonseparable Multidimensional Perfect Reconstruction Filter Banks and Wavelet Bases for $\mathbb{R}^n$, IEEE Trans. Inform. Theory. 38, p. 533–555.

MALLAT, S.G.(1989a): Multifrequency Channel Decompositions of Images and Wavelet Models, IEEE Trans. Acoust., Speech, Signal Proc. 37, p. 2091–2110.

MALLAT, S.G. (1989b): A Theory for Multiresolution Signal Decomposition: The Wavelet Representation, IEEE Trans. Pattern Anal. Machine Intell. 11, p. 674–693.

MALLAT, S.; W.L. HWANG (1992): Singularity Detection and Processing with Wavelets, IEEE Trans. Inform. Theory 38,p. 617–643.

MALLAT, S.G. (1999): A Wavelet Tour of Signal Processing, 2d ed., Academic Press, San Diego.

MARR, D. (1982): Vision, Freeman, San Francisco.

MEYER, Y. ed. (1991): Wavelets and Applications (Proc. Marseille 1989), Springer-Verlag, Berlin, and Masson, Paris.

MEYER, Y. (1993): Wavelets: Algorithms and Applications, SIAM, Philadelphia, PA.

MEYER, Y.; S. ROQUES, eds. (1993): Progress in Wavelet Analysis and Applications (Proc. Toulouse 1992), Ed. Frontières, Gif-sur-Yvette.

MURENZI, R. (1990a) Wavelet Transforms Associated to the n-Dimensional Euclidean Group with Dilations: Signals in More than one Dimension, in (Combes, 90), p. 239–246.

MURENZI, R. (1990b): Ondelettes Multidimensionnelles et Applications à l'Analyse d'Images, Thèse de Doctorat, Univ. Cath. Louvain, Louvain-la-Neuve.

RIOUL, O.; M. VETTERLI (1991): Wavelets and Signal Processing, IEEE SP Magazine, October 1991, 14–38.

RUSKAI, M.B.; G. BEYLKIN; R. COIFMAN; I. DAUBECHIES; S. MALLAT; Y. MEYER; L. RAPHAEL, eds. (1992): Wavelets and Their Applications, Jones and Bartlett, Boston, MA.

SARACCO, G.; A. GROSSMANN; Ph. TCHAMITCHIAN (1990): Use of Wavelet Transforms in the Study of Propagation of Transient Acoustic Signals Across a Plane Interface Between two Homogeneous Media, in (Combes, 90), p.139–146.

SCHRÖDER, P.; W. SWELDENS (1995): Spherical Wavelets: Efficiently Representing Functions on the Sphere, Computer Graphics Proc. (SIGGRAPH95), ACM Siggraph, p.161–175.

STRANG, G.; T. NGUYEN (1996): Wavelets and Filter Banks, Wellesley-Cambridge Press, Wellesley, MA.

SWELDENS, W. (1996): The Lifting Scheme: A Custom-Design Construction of Biorthogonal Wavelets, Applied Comput. Harm. Anal. 3, p. 1186–200.

THONET, G.; O. BLANC; P. VANDERGHEYNST; E. PRUVOT; J-M. VESIN; J-P. ANTOINE (1998): Wavelet-Based Detection of Ventricular Ectopic Beats in Heart Rate Signals, Applied Sign. Process. 5, p. 170–181.

VANDERGHEYNST, P.; E. VAN VYVE ; A. GOLDBERG; J-P. ANTOINE; I. DOGRI (1999): Modelling and Simulation of an Impact Test Using Wavelets, Analytical Solutions and Finite Elements, preprint UCL-IPT-99-05, Louvain-la-Neuve (submitted).

VETTERLI, M.; J. KOVAČEVIĆ (1995): Wavelets and Subband Coding, Prentice Hall, Englewood Cliffs, NJ.

WICKERHAUSER, M.V. (1994): Adapted Wavelet Analysis from Theory to Software, A.K. Peters, Wellesley, MA.

WISNOE,W. (1993a): Utilisation de la Méthode de Transformée en Ondelettes 2D pour l'Analyse de Visualisation d'Écoulements, Thèse de Doctorat ENSAE, Toulouse.

WISNOE, W.; P. GAJAN; A. STRZELECKI; C. LEMPEREUR; J-M. MATHÉ (1993b): The Use of the Two-Dimensional Wavelet Transform in Flow Visualization Processing, in (Meyer, 93), p. 455–458.