COMPACT IMAGE REPRESENTATION FROM MULTISCALE EDGES

Sifen Zhong and Stephane Mallat

Courant Institute of Mathematical Sciences
New York University
251 Mercer Street
New York, NY 10012

Abstract

The sharper variation points of a signal, measured at different scales, can be detected from the local maxima of its wavelet transform. We describe an algorithm that reconstructs one-dimensional signals and images from their sharper variation points at dyadic scales. This algorithm reconstructs exactly images from their multiscale edges. We also prove that the evolution across scales of the wavelet maxima characterizes the local shape of the signal sharp variations. We can thus not only detect edges but also classify them. The wavelet maxima representation is a new reorganization of the image information that enables us to develop algorithms uniquely based on edges for solving image processing and computer vision problems. We describe an "intelligent" compact coding algorithm based on this representation, that takes into account the information content of the image.

numerically that the whole signal information can be decomposed in multiscale edges.

An important issue in multiscale analysis is to relate the local properties of the signal to the evolution of the transform values when the scale varies. The wavelet theory proves that this evolution across scales characterizes the local Lipschitz regularity of the signal. The local maxima of the wavelet transform thus not only detect the signal sharper variations but also characterize their "local shape". This complement of information is very important to distinguish different type of edges and discriminate textures. We study the application of this multiscale edge representation to compact image coding. From the behavior of the edges across scales, we show that one can select the important information that is required to reconstruct good quality images while removing a large portion of the edges. This opens a new approach on compact image coding based on the analysis of the image information from the multiscale edges. We describe a coding procedure that leads to a compression by a factor of 40. The application of this representation to pattern recognition is also discussed.

2. Dyadic Wavelet Transform in One Dimension

For a complete presentation of the wavelet theory the reader is referred to a general review [8] or an advanced functional analysis book of Meyer [12]. A wavelet is a function $\psi(x) \in L^2(\mathbf{R})$ whose integral is zero. Let us denote by $\psi_s(x)$ the dilation of $\psi(x)$ by a scale factor s:

$$\psi_s(x) = \frac{1}{s}\,\psi\left(\frac{x}{s}\right) . \tag{1}$$

The wavelet transform of a function $f(x)$ at the scale s and position x is given by the convolution product:

$$W_s f(x) = f * \psi_s(x) . \tag{2}$$

Let us concentrate on scales s that varies along the dyadic sequence $\left(2^j\right)_{j\in\mathbf{Z}}$. We call *dyadic wavelet transform* the sequence of functions

$$\left(W_{2^j} f(x)\right)_{j\in\mathbf{Z}} . \tag{3}$$

The Fourier transform of $W_{2^j} f(x)$ is

$$\hat{W} f_{2^j}(\omega) = \hat{f}(\omega)\,\bar{\hat{\psi}}(2^j\omega) . \tag{4}$$

By imposing that

$$\sum_{j=-\infty}^{+\infty} |\hat{\psi}(2^j\omega)|^2 = 1 , \tag{5}$$

we insure that the whole frequency axis is covered by a dilation of $\hat{\psi}(\omega)$ by the scales factors $\left[2^j\right]_{j\in Z}$. One can then derive [9] that $f(x)$ is reconstructed from its dyadic wavelet transform by

$$f(x) = \sum_{j=-\infty}^{+\infty} W_{2^j} f * \bar{\psi}_{2^j}(x) \ .$$

It is important to observe that any sequence of functions $\left[g_j(x)\right]_{j\in Z}$ is not a priori the dyadic wavelet transform of some function $f(x)$. One can prove that the sequence $\left[g_j(x)\right]_{j\in Z}$ is the dyadic wavelet transform of some function $f(x)$ if and only it satisfies the following reproducing kernel equation:

$$\forall j \in Z \quad \sum_{l=-\infty}^{+\infty} g_l * K_{l,j}(x) = g_j(x) \ , \text{ with} \tag{6}$$

$$K_{l,j}(x) = \bar{\psi}_{2^l} * \psi_{2^j}(x) \ .$$

This reproducing kernel equation expresses the inner redundancy within a wavelet transform at different scales. This condition shows that the vector space $\mathbf{V}$ of all dyadic wavelet transform is smaller than the space of all sequences of $\mathbf{L}^2(\mathbf{R})$ functions. One can prove that the reproducing kernel equation (6) defines an orthogonal projector $\mathbf{P_V}$ on the space $\mathbf{V}$. This means that even if the sequence of functions $\left[g_j(x)\right]_{j\in Z}$ is not the wavelet transform of some function in $\mathbf{L}^2(\mathbf{R})$, if $\left[h_j(x)\right]_{j\in Z}$ is defined by:

$$h_j(x) = \sum_{l=-\infty}^{+\infty} g_l * K_{l,j}(x) \ ,$$

then $\left[h_j(x)\right]_{j\in Z} \in \mathbf{V}$. The operator $\mathbf{P_V}$ is particularly important in the rest of this paper.

In practice, the input signal is measured at a finite resolution so we can not compute the wavelet transform at an arbitrary fine scale. Also, we must limit our computations to a finite larger scale. Let us suppose for normalization purpose that the signal was measure at the resolution 1 and thus that the finner scale is 1. Let 2^J be the larger scale. The wavelet transform at all the scales larger than 2^J, $\left[W_{2^j} f(x)\right]_{1\le j<+\infty}$ carries the low frequency information of $f(x)$. One can prove that these low frequencies can also be obtained with the convolution of $f(x)$ with a low-pass filter $\phi_{2^J}(x)$:

$$S_{2^J} f(x) = f * \phi_{2^J}(x) \ . \tag{7}$$

The Fourier transform $\hat{\phi}_{2^J}(\omega)$ of $\phi_{2^J}(x)$ can be expressed from the Fourier transform of the

wavelet $\psi(x)$:

$$|\hat{\phi}_{2^J}(\omega)|^2 = \sum_{j=J+1}^{+\infty} |\hat{\psi}(2^j\omega)|^2 \quad .$$

We call *finite dyadic wavelet transform* of $f(x)$ between the scales 1 and 2^J, the set of functions

$$\left\{ S_{2^J} f(x) \, , \left[W_{2^j} f(x) \right]_{1 \le j \le J} \right\} \quad . \tag{8}$$

The discretization of this model is carefully studied in [9]. If the original discrete signal is given by N samples, one can compute a uniform discretization of a finite scale dyadic wavelet transform with an algorithm of complexity $O(N \log(N))$. This algorithm is based on a cascade of discrete convolutions with a low-pass and a band-pass filter. The reconstruction of the original signal from its dyadic wavelet transform is exact and also requires $O(N \log(N))$ computations. Fig. 1 shows the dyadic wavelet transform of a signal, computed between the scales 1 and 2^5.

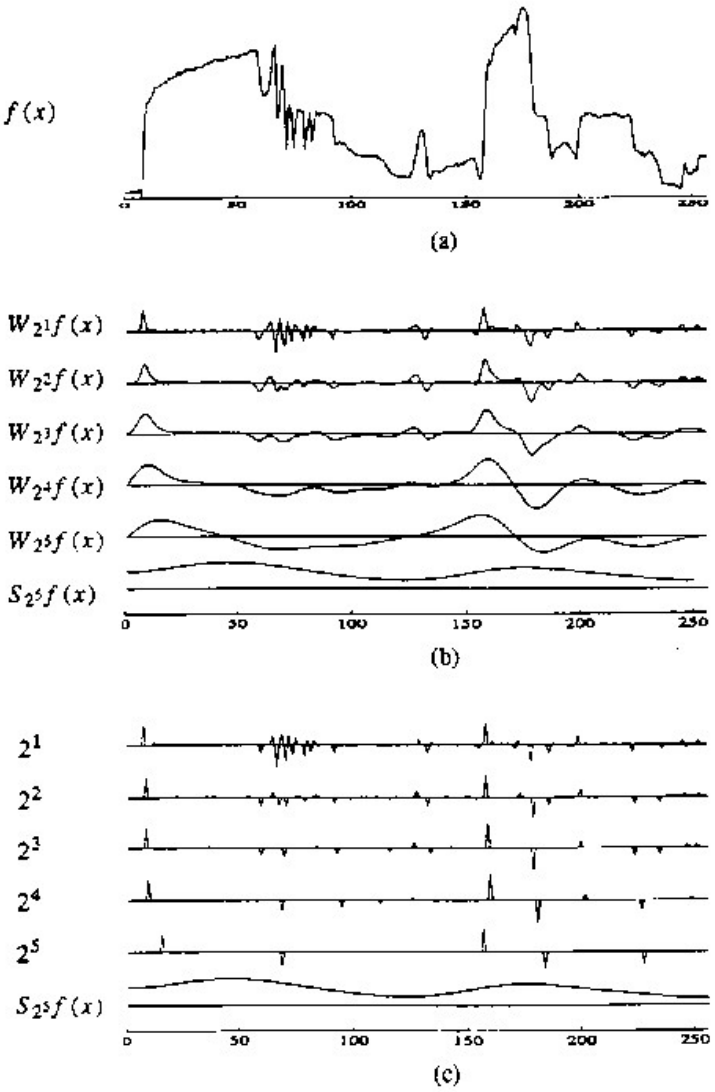

Fig. 1: *(a): Image scan-line of 256 samples. (b): Dyadic wavelet transform of signal 1(a), on 5 scales. Since we decompose up to a finite coarser scale (2^5), we keep the remaining low-frequencies $S_{2^5}f(x)$ to have a complete representation. (c): Maxima representation of the dyadic wavelet transform shown in (b). Each dirac indicates the position and amplitude of a local maxima of the at the corresponding scale of the dyadic wavelet transform.*

3. Wavelet Transform Maxima

The concept of multiscale is particularly important for edge detection. In order to separate the fine structures from the larger one, many researchers have studied the properties of a multiscale edge detection [13]. Let us show that for some particular wavelet functions $\psi(x)$, the local maxima of the wavelet transform indicate the position of the multiscale edges.

Let $\theta(x)$ be a smoothing function. A classical example often used in computer vision is the Gaussian. Let $\psi(x)$ be the first order derivative of $\theta(x)$:

$$\psi(x) = \frac{d\theta(x)}{dx} . \tag{9}$$

Let us denote $\theta_{2^j}(x) = \frac{1}{2^j}\,\theta(\frac{x}{2^j})$. The wavelet transform at the scale 2^j is given by:

$$W_{2^j} f(x) = f * \psi_{2^j}(x) = f * (2^j \frac{d\theta_{2^j}}{dx})(x) = 2^j \frac{d}{dx}(f * \theta_{2^j})(x) . \tag{10}$$

The wavelet transforms $W_{2^j} f(x)$ is proportional to the first derivative of $f(x)$ smoothed by $\theta_{2^j}(x)$. The maxima of $|W_{2^j} f(x)|$ are thus the maxima of the modulus of the derivative of $f * \theta_{2^j}(x)$. It corresponds to the sharper variation points of the signal smoothed at the scale 2^j. This is illustrated in Fig. 2. These sharper variation points are also called edge points. The maxima detection of such a wavelet transform is essentially equivalent to a Canny edge detection [1]. In fig. 1(c) the maxima of the wavelet transform are indicated by diracs. As expected, these maxima indicate the position of the sharper variation points of the signal smoothed at different scales. As we shall see in section 6, the evolution of the amplitude of these maxima across scales, characterizes the local shape of the sharp variations.

Let us observe that if the wavelet $\psi(x)$ was the second derivative of $\theta(x)$, the sharp variation points of the signal smoothed at the scale 2^j would be detected from the zero-crossings of the wavelet transform. This can be viewed as a Marr Hildreth [11] edge detector. The inconvenience of this approach is that we can not distinguish maximum variation points from minimum variation points (see fig. 2).

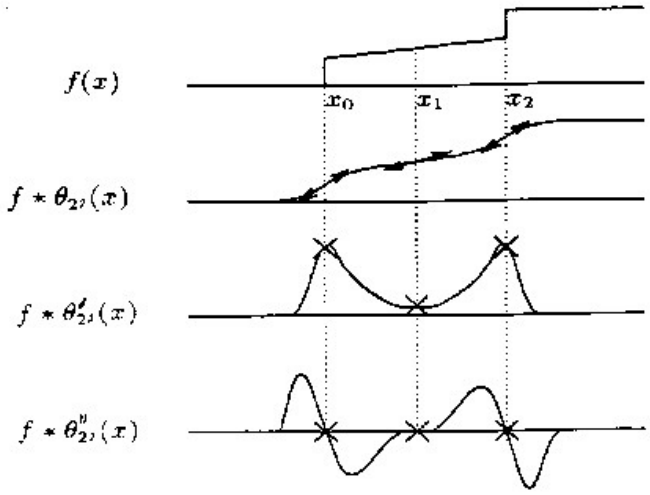

Fig. 2: *The inflection points of $f * \theta_{2^j}(x)$ correspond to the extrema of $W_{2^j}f(x)$ if $\psi(x)$ is the first derivative of $\theta(x)$ and to the zero crossings of $W_{2^j}f(x)$ if $\psi(x)$ is the second derivative of $\theta(x)$. For the first derivative case, we only record the points of abscissa x_0 and x_2 where $|W_{2^j}f(x)|$ is maximum because they locate the sharper variation points of $f(x)$ smoothed at the scale 2^j.*

4. Completeness and Stability

4.1. Previous Results

A fundamental issue is to understand whether the multiscale local maxima or the zero-crossings define a complete and stable representation of the original signal. Most previous works have been done with zero-crossings but these results can easily be extended in the maxima framework. The most classical result concerning the characterization of a signal from its zero-crossings is due to Logan [7]. Let $g(x) \in L^2(\mathbf{R})$ and let us suppose that its Fourier transform has a support included in one octave intervals. Logan theorem [7] proves that if $g(x)$ does not share any zero-crossings with its Hilbert transform, then it is uniquely characterized by its zero-crossings. Let us give an intuitive justification of this result. We know that there exists ω_0 such that the Fourier of $g(x)$ has a support included in the intervals $[-2\omega_0, -\omega_0] \cup [\omega_0, 2\omega_0]$. The Nyquist theorem proves that such a signal is characterized by a uniform sampling at the rate $\frac{\omega_0}{\pi}$. One can also prove that this signal changes sign approximatively as frequently as the function $\sin(\omega_0 x)$. The number of zero-crossings is therefore of the same order than the number of values needed to characterize the signal with a uniform sampling. Of course, the zero-crossing

problem is different since zero-crossings are not uniformly distributed but one can see that qualitatively the same amount of information is available. The zero-crossing characterization as explained by Logan is not stable: "the problem of actually recovering (the signal) from its sign changes appears to be very difficult and impractical". This unstability means that a small error on the zero-crossing position might create an arbitrary large error on the reconstructed function. The Logan theorem can not be applied in our framework because the functions $W_{2^j} f(x)$ are not one-octave band-limited. As a consequence, zero-crossings occur selectively at sharp variation points of $f(x)$ and not at regular intervals like for one-octave band-limited functions.

The Logan result has been extended by several researchers by supposing only that $g(x)$ is band-limited. The proofs on the completeness of the zero-crossings are based on an analytical extension of these signals in the complex plane. All these proofs do not provide any stability result since they are based on non-stable characterization of analytical functions [2, 14, 18]. The reader is referred to a review by Hummel and Moniot for more details [4].

Many studies have also described the properties of zero-crossings of functions convolved with the Laplacian of a Gaussian. This convolution is equivalent to a wavelet transform built from a wavelet $\psi(x)$ equal to the Laplacian of a Gaussian. Such a wavelet transform can be interpreted as the result of a heat diffusion process [5]:

$$\frac{\partial W_s f(x)}{\partial s} = \frac{\partial^2 W_s f(x)}{\partial x^2} \quad . \tag{11}$$

Using the properties of the heat diffusion differential equations several authors have proved interesting properties of the propagation of zero-crossings across scales [3, 5, 17]. Hummel & Moniot as well as Yuille & Poggio have also proved that the position of the zero-crossings of $W_s f(x)$ give a complete characterization of the function $f(x)$ [3]. These proofs are based on deconvolution arguments which are highly instable and they do not take a full advantage of the information given by the zero-crossings at all scales. They can therefore not be used for the reconstruction of $f(x)$ from the zero-crossings of its wavelet transform. The differential equation (11) gives the evolutionary properties of $W_s f(x)$ when the scale s and the abscissa x vary. It expresses the redundancy of the functions $W_s f(x)$ at different scales. In all applications, the scale parameter varies on sparse discrete sequences such as the dyadic sequence $\left(2^j\right)_{j \in Z}$. One therefore cannot use the heat partial differential equation to express the redundancy specifically at these scales.

The wavelet transform reproducing kernel is an alternative approach to formalize the redundancy of the functions $W_s f(x)$. In the next section we reformalize the completeness problem by using the wavelet transform reproducing kernel equation.

4.2. One-Dimensional Signal Reconstruction from the Wavelet Maxima

In this section, we study the reconstruction of a function from the maxima of its wavelet transform. We formalize the completeness problem within the wavelet framework and then derive an algorithm to perform the reconstruction. Let $f(x) \in \mathbf{L}^2(\mathbf{R})$ and $\left\{W_{2^j} f(x)\right\}_{j \in \mathbf{Z}}$ be its dyadic wavelet transform. Since $f(x)$ can be recovered from its dyadic wavelet transform, we first try to reconstruct $\left\{W_{2^j} f(x)\right\}_{j \in \mathbf{Z}}$ given the local maxima of each function $W_{2^j} f(x)$, $j \in \mathbf{Z}$. Clearly, for any scale 2^j, there exists an infinite number of functions $g_j(x)$ which have the same local maxima as $W_{2^j} f(x)$. However, any such sequence of functions $\left\{g_j(x)\right\}_{j \in \mathbf{Z}}$ is not necessarily the dyadic wavelet transform of some function in $\mathbf{L}^2(\mathbf{R})$. Indeed, we saw in section 2 that for being a dyadic wavelet transform, it must satisfy the reproducing kernel conditions (6). Let us recall from section 2 that the space of all dyadic wavelet transforms is denoted $\mathbf{V}$. In order to express the conditions given by the maxima of the wavelet transform of $f(x)$, we define the set Γ of all sequences $\left\{g_j(x)\right\}_{j \in \mathbf{Z}}$ such that for all scales 2^j, $g_j(x)$ has the same maxima than $W_{2^j} f(x)$. The local maxima representation is complete if and only there exists no dyadic wavelet transform different from $\left\{W_{2^j} f(x)\right\}_{j \in \mathbf{Z}}$ that has the same local maxima. In other words, the intersection of Γ with $\mathbf{V}$ must be reduced to one element:

$$\Gamma \cap \mathbf{V} = \left\{ \left\{W_{2^j} f(x)\right\}_{j \in \mathbf{Z}} \right\} .$$

In order to verify numerically this assertion, we describe an algorithm that reconstructs the intersection of Γ with $\mathbf{V}$.

The set Γ is almost convex [9]. A classical technic for recovering the intersection of a convex set with a linear space is to iterate on alternative projections on the convex and the linear space [16]. For any $\left\{g_j(x)\right\}_{j \in \mathbf{Z}}$ in this Hilbert space, we can define [9] a projection $\mathbf{P}_\Gamma$ on Γ that transforms $\left\{g_j(x)\right\}_{j \in \mathbf{Z}}$ into the sequence of functions $\left\{h_j(x)\right\}_{j \in \mathbf{Z}} \in \Gamma$ that is the closest to $\left\{g_j(x)\right\}_{j \in \mathbf{Z}}$. The functions $h_j(x)$ are smoothed deformations of $g_j(x)$ in order to match the maxima constraint. The deformation is minimum when measured with an $\mathbf{H}^1(\mathbf{R})$ norm. Let $\mathbf{P}_\mathbf{V}$ be the orthogonal projection on the space $\mathbf{V}$ and $\mathbf{P} = \mathbf{P}_\Gamma \circ \mathbf{P}_\mathbf{V}$ be the composition of $\mathbf{P}_\Gamma$ and $\mathbf{P}_\mathbf{V}$. Clearly any element at the intersection of Γ and $\mathbf{V}$ is a fixed point of $\mathbf{P}$. To compute such a fixed point, we iterate on the operator $\mathbf{P}$. The algorithm is illustrated in fig. 3. Let $\mathbf{P}^{(n)}$ be the composition n times of the operator $\mathbf{P}$. The convergence of $\mathbf{P}^{(n)}\left\{g_j(x)\right\}_{j \in \mathbf{Z}}$ to an element of $\Gamma \cap \mathbf{V}$ is guaranteed if one keeps all the maxima and minima of the wavelet transform [9].

When only the maxima information is kept, we have not proved that the convergence is guaranteed.

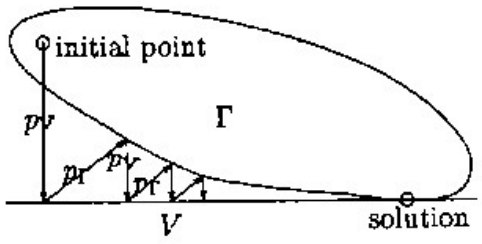

Fig. 3: *The reconstruction of the wavelet transform of $f(x)$ is done with an alternating projection of the set Γ and the space $\mathbf{V}$ of all dyadic wavelet transform.*

5. Numerical Results in One Dimension

Within a discrete framework, the operators $\mathbf{P}_V$ and $\mathbf{P}_\Gamma$ can be implemented with a complexity of $O(N \log(N))$. Each iteration on the operator $\mathbf{P}$ thus requires $O(N \log(N))$ computations. After reconstruction the dyadic wavelet transform by iterating on $\mathbf{P}$, we can then reconstruct the corresponding signal. The error is defined as the difference between this reconstructed signal and the original one. We compute the error to signal ratio with a mean-square measure. Fig. 4 is an example of reconstructed signal after 15 iterations. The error to signal ratio is $3\,10^{-2}$. Fig. 5 plots the error to signal ratio as a function of the number of iterations on the operator $\mathbf{P}$. The decay of the error is fast during the first 20 iterations. Then it is slow but approximatively constant. The remaining errors are then concentrated in the highest frequencies, which are slower to reconstruct exactly. This is due to the structure of the reproducing kernel at the finner scale [9]. The reproducing kernel expresses the redundancy between the functions $W_{2^j} f(x)$ at different scales 2^j. This redundancy is maximum for consecutive scales: 2^j with 2^{j+1} and 2^{j-1}. At the scale 2^1, there is no information available at the finner scale 2^0 so the correlation constraint is weaker. After 4000 iterations, the error to signal ration is $2\,10^{-5}$. In these computations, the floating point numbers precision is 10^{-7} and the error does not further decrease because of this limit.

All the numerical results obtained so far, seem to indicate that the maxima of a dyadic wavelet transform do provide a complete and stable representation. This algorithm has been tested on a large class of signals including all types of classical functions. We also used different wavelets $\psi(x)$ with similar convergence results [9]. We therefore conjecture that for a large class of wavelets, the maxima of $\left\{W_{2^j} f(x)\right\}_{j \in \mathbf{Z}}$ provide a complete and stable representation of $f(x)$. The class of wavelet for which this is true remains to be defined. The reconstruction

algorithm previously described is however sufficient for the applications.

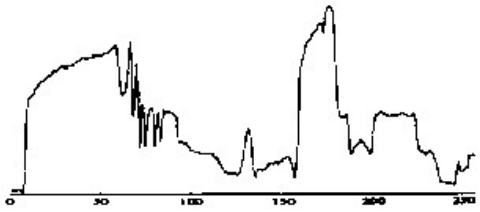

Fig. 4: *Reconstruction with 15 iterations from the maxima representation given in fig. 1(c). The error to signal ratio of the reconstruction* $2\ 10^{-3}$.

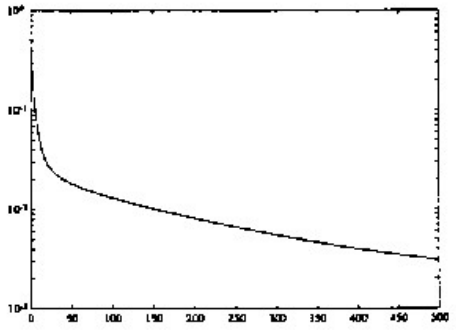

Fig. 5: *Evolution of error to signal ratio of the reconstructed signal from the maxima representation of fig. 1(c), depending upon the number of iterations on the operator* P.

6. Behavior of Maxima Amplitude Across Scales

The previous section showed experimentally that we can reconstruct the original signal from the maxima of the wavelet transform and thus that this representation is complete and stable. Let us now explain how to use directly the maxima representation for characterizing the different type of signal variations.

The singularity of a function at a point x_0, can be characterized with a Lipschitz exponent α. A function $f(x)$ is Lipschitz α regular in x_0 $(0 \leq \alpha \leq 1)$, if and only if for all x in a neighborhood of x_0, we have

$$|f(x) - f(x_0)| = O(|x - x_0|^{\alpha}) . \tag{12}$$

A step edge for example is characterized by $\alpha = 0$. The following theorem relates the Lipschitz

regularity of a function to its wavelet transform amplitude.

Theorem 1

Let $f(x) \in L^2(R)$, $f(x)$ is Lipschitz α in all points of an open interval if and only if for all x in this interval the wavelet transform satisfies

$$|W_s f(x)| = O(s^\alpha) . \tag{13}$$

This theorem remains valid if the scale s is restricted to a dyadic sequence $\left[2^j\right]_{j \in Z}$. The proof can be found in [12]. Theorem 1 proves that we obtain the Lipschitz regularity of a signal from the behavior of its wavelet transform amplitude when the scale decreases. By applying theorem 1, we can thus recover the Lipschitz regularity of a sharp variation points from the evolution of the maxima across scales.

Often the signal singularities are blurred due to some diffusion process. We thus get smoothed "singularities" and it is important to estimate this smoothing factor. For example, the shadows in an image do not produce sharp discontinuities of the image intensity but relatively smooth variations because of the diffraction effect. In this case, the image intensity is locally equal to the convolution of a step edge with a smoothing kernel. Let us suppose that our signal $f(x)$ is locally equal to a singularity $h(x)$ that has been smoothed by a Gaussian of variance σ:

$$f(x) = h * g_\sigma(x) \quad \text{where} \quad g_\sigma(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp(-\frac{x^2}{2\sigma^2}) .$$

We saw that the wavelet transform of $f(x)$ can be written

$$W_{2^j} f(x) = 2^j \frac{d}{dx} (f * \theta_{2^j})(x) = 2^j \frac{d}{dx} (h * g_\sigma * \theta_{2^j})(x) . \tag{14}$$

Let us suppose that the function $\theta(x)$ is close to a Gaussian function. In this case we have:

$$\theta_{2^j} * g_\sigma(x) = \theta_{s_0}(x) \quad \text{with} \quad s_0 = \sqrt{2^{2j} + \sigma^2} .$$

Equation (14) can thus be rewritten:

$$W_{2^j} f(x) = 2^j \frac{d}{dx} (h * \theta_{s_0})(x) = \frac{2^j}{s_0} W_{s_0} h(x) . \tag{15}$$

In other words, the wavelet transform of a singularity smoothed by a Gaussian of variance σ, at the scale 2^j, is equal to the wavelet transform of the non-smoothed singularity at the scale $s_0 = \sqrt{2^{2j} + \sigma^2}$. Theorem 1 relates the decay of the maxima of $|W_s h(x)|$ to the Lipschitz exponent of the singularity of $h(x)$. We can thus derive from equation (15) the decay of the maxima generated by a smoothed singularity.

Fig. 6 gives several examples of singularities smoothed by Gaussians of different variances. The decay of the maxima are clearly affected by the different Lipschitz exponent as well as the variance of the Gaussian smoothing. Let us now explain how to compute numerically the Lipschitz exponent α and the smoothing scale σ. After renormalization, we can only compute the wavelet transform at scales larger than 1. Let us suppose that a singularity in a point x_0 creates maxima of amplitude a_j at each scale 2^j between the scale 2^1 and some larger scale 2^{j_0}. If there is no smoothing factor ($\sigma = 0$), theorem 1 implies that a_j decays like $2^{\alpha j}$. In order to estimate the Lipschitz exponent α associated to this singularity, we compute the constants K and α that minimize

$$E(\alpha,K) = \sum_{j=1}^{j_0} (a_j - K2^{\alpha j})^2 \ . \tag{16}$$

For isolated singularities, this simple procedure provides an estimation of α with a few percent error [9]. If we want also to estimate the smoothing scale σ, then we estimate the three constants K, α and σ that minimize

$$E(\sigma,\alpha,K) = \sum_{j=1}^{j_0} (a_j - K(\sqrt{\sigma^2+2^{2j}})^{\alpha})^2 \ . \tag{17}$$

The error of estimation of the Lipschitz exponent α depends upon the size of the smoothing scale. The larger σ the larger the error on the estimation of α. More numerical results are described in [9].

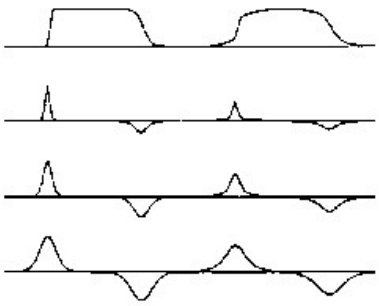

Fig. 6: *(a): The signal $S^q f$ is built from a function $f(x)$ which has four singularities respectively characterized by $(\alpha = 0\ ,\ \sigma = 0)$, $(\alpha = 0\ ,\ \sigma = 5.8)$, $(\alpha = 0.25\ ,\ \sigma = 0)$ and $(\alpha = 0.25\ ,\ \sigma = 5.8)$.*
(b): The behavior of the maxima of $W_{2^1}^q f$, $W_{2^2}^q f$ and $W_{2^3}^q f$ depend upon the Lipschitz regularity α and the smoothing scale σ.

7. Wavelet Maxima Representation of Images

The wavelet transform can be extended in $L^2(R^2)$ by introducing several wavelets $\Psi^k(x,y) \in L^2(R^2)$. We study more particularly a two-dimensional wavelet decomposition which has only two orientation selectivities: a horizontal and a vertical one. This wavelet transform is defined with two wavelets $\Psi^1(x,y)$ and $\Psi^2(x,y)$ given by:

$$\Psi^1(x,y) = \frac{\partial \Theta^1(x,y)}{\partial x} \quad \text{and} \quad \Psi^2(x,y) = \frac{\partial \Theta^2(x,y)}{\partial y} \ . \tag{18}$$

$\Theta^1(x,y)$ and $\Theta^2(x,y)$ are two smoothing functions. Let us denote $\Psi^1_{2^j}(x,y) = \frac{1}{2^{2j}} \Psi^1(\frac{x}{2^j}, \frac{y}{2^j})$ and $\Psi^2_{2^j}(x,y) = \frac{1}{2^{2j}} \Psi^2(\frac{x}{2^j}, \frac{y}{2^j})$. The wavelet transform of a function $f(x,y) \in L^2(R^2)$ at the scale 2^j, along the horizontal and vertical orientations is given by:

$$W^1_{2^j} f(x,y) = f * \Psi^1_{2^j}(x,y) \quad , \quad W^2_{2^j} f(x,y) = f * \Psi^2_{2^j}(x,y) \ . \tag{19}$$

We call *two-dimensional wavelet transform* of $f(x,y)$ the set of functions

$$\left(W^1_{2^j} f(x,y) \ , \ W^2_{2^j} f(x,y) \right)_{j \in Z} \ . \tag{20}$$

With a particular choice of $\Theta^1(x,y)$ and $\Theta^2(x,y)$ we can have:

$$\sum_{j=-\infty}^{+\infty} \left(|\hat{\Psi}^1(2^j\omega_x, 2^j\omega_y)|^2 + |\hat{\Psi}^2(2^j\omega_x, 2^j\omega_y)|^2 \right) = 1 \ . \tag{21}$$

Given this constraint, the two-dimensional dyadic wavelet transform has the same properties than in one dimension. The function $f(x,y)$ can be reconstructed from its wavelet transform. Similarly to the one-dimensional case, any sequence of two dimensional functions $\left(g^1_j(x,y) \ , \ g^2_j(x,y) \right)_{j \in Z}$ is not a priori the dyadic wavelet transform of some two-dimensional function. In order to be a dyadic wavelet transform, such a sequence must satisfy a reproducing kernel equation similar to equation (6). Like in one dimension, the set of all sequence of functions which are the dyadic wavelet transform of some function in $L^2(R^2)$ is a vector space that we denote V. This vector space is smaller than the space of all sequences of $L^2(R^2)$ functions. We can build an orthogonal projector on V from the two-dimensional reproducing kernel equation.

Images are measured at a finite resolution so we can not compute the wavelet transform at a scale below the limit set by this resolution. Like in one dimension, we must also limit the larger scale to a finite value. Let us suppose for normalization purpose that the finner scale is 1 and the larger one is 2^J. One can prove that the low-frequency information carried by the wavelet transform $\left(W^1_{2^j} f(x,y) \ , \ W^2_{2^j} f(x,y) \right)_{J < j < +\infty}$ at the scales larger than 2^J, is given by the convolu-

tion of $f(x,y)$ with a smoothing function $\Phi_{2^j}(x,y)$:

$$S_{2^j} f(x,y) = f * \Phi_{2^j}(x,y) \quad . \tag{22}$$

The Fourier transform of $\Phi_{2^j}(x,y)$ is given by:

$$|\hat{\Phi}_{2^j}(\omega_x,\omega_y)|^2 = \sum_{j=J+1}^{+\infty} \left(|\hat{\Psi}^1(2^j\omega_x,2^j\omega_y)|^2 + |\hat{\Psi}^2(2^j\omega_x,2^j\omega_y)|^2 \right) . \tag{23}$$

The discretization of this model has been studied in [9]. If the original signal $f(x,y)$ is uniformly sampled, one can compute a uniform sampling of the two-dimensional dyadic wavelet transform. If the image has N pixels than we need $O(N \log(N))$ computations to compute the corresponding dyadic wavelet transform. The reconstruction of the original image from its wavelet transform also requires $O(N \log(N))$ computations.

Let us denote by $\Theta^1_{2^j}(x,y) = \frac{1}{2^{2j}} \Theta^1(\frac{x}{2^j},\frac{y}{2^j})$ and $\Theta^2_{2^j}(x,y) = \frac{1}{2^{2j}} \Theta^2(\frac{x}{2^j},\frac{y}{2^j})$. Equations (18) and (19) yield:

$$W^1_{2^j} f(x,y) = 2^j \frac{\partial}{\partial x} (f * \Theta^1_{2^j})(x,y) \quad \text{and} \tag{24}$$

$$W^2_{2^j} f(x,y) = 2^j \frac{\partial}{\partial y} (f * \Theta^2_{2^j})(x,y) \quad . \tag{25}$$

The dyadic wavelet transform $W^1_{2^j} f(x,y)$ and $W^2_{2^j} f(x,y)$ are respectively the partial derivative along the horizontal and vertical directions of $f(x,y)$ smoothed at the scale 2^j. In practice, the two functions $\Theta^1(x,y)$ and $\Theta^2(x,y)$ are very similar so that the wavelet transforms $W^1_{2^j} f(x,y)$ and $W^2_{2^j} f(x,y)$ can be viewed as the two components of the gradient vector of $f(x,y)$ smoothed at the scale 2^j. Fig. 7 shows the dyadic wavelet transform of a circle image decomposed between the scales 1 and 2^4.

Let us now define a modulus and angle image at each scale 2^j:

$$M_{2^j} f(x,y) = \sqrt{|W^1_{2^j} f(x,y)|^2 + |W^2_{2^j} f(x,y)|^2} \quad , \tag{26}$$

$$A_{2^j} f(x,y) = argtan\left(\frac{W^2_{2^j} f(x,y)}{W^1_{2^j} f(x,y)} \right) \quad . \tag{27}$$

$M_{2^j} f(x,y)$ can be interpreted as the modulus of the gradient vector of $f(x,y)$ smoothed at the scale 2^j whereas $A_{2^j} f(x,y)$ gives the angle of the gradient vector orientation. The modulus and angle images of the circle are also shown in fig. 7.

The sharper variation points of $f(x,y)$ smoothed at the scale 2^j correspond to the points where $M_{2^j} f(x,y)$ has a local maxima along the gradient direction. For each of these maxima, we record the corresponding amplitude of $M_{2^j} f(x,y)$ and the angle $A_{2^j} f(x,y)$. This maxima detection is essentially equivalent to Canny's non-maxima suppression [1]. In fig. 7 these

maxima are located at the border of the circle. The first column of fig. 8 gives an example of modulus images $M_{2^j} f(x,y)$ at four scale levels for a lady image. The second column gives the corresponding position of the maxima. The third and fourth columns display the maxima whose amplitude are respectively larger than 4 and 8. The higher amplitude maxima correspond to the most important edges of the image.

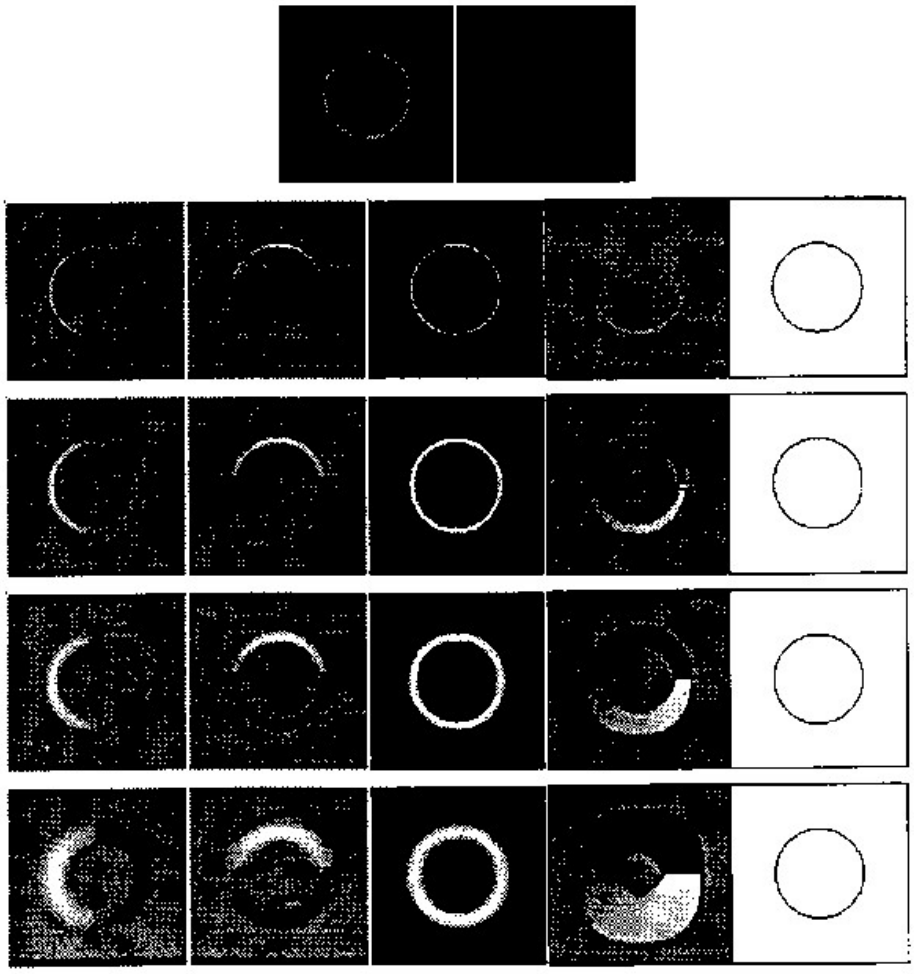

Fig. 7: *The original image is at the top left and next to it the low-frequency image* $S_{2^4}f(x,y)$. *The first column from the left gives the images* $\left[W^1_{2^j}f(x,y)\right]_{1\le j\le 4}$ *and the scale increases from top to bottom. The second columns is* $\left[W^2_{2^j}f(x,y)\right]_{1\le j\le 4}$. *Black, grey and white pixels indicate respectively negative, zero and positive sample values. The third column displays the modulus images* $\left[M_{2^j}f(x,y)\right]_{1\le j\le 4}$, *black pixels indicate zero values whereas white one correspond to the highest value. The fourth column gives the angle images* $\left[A_{2^j}f(x,y)\right]_{1\le j\le 4}$, *and we can see that the angle value turns from* π *to* $-\pi$ *along the circle contour. The fifth column displays in black the position of the maxima of* $\left[M_{2^j}f(x,y)\right]_{1\le j\le 4}$.

8. Reconstruction of Two-Dimensional Signals from the Wavelet Maxima

The reconstruction of the original image from its wavelet maxima is done with a direct extension of the one-dimensional algorithm. Let $f(x,y) \in L^2(\mathbf{R}^2)$ and $\left\{ W_{2^j}^1 f(x,y) , W_{2^j}^2 f(x,y) \right\}_{j \in \mathbf{Z}}$ be its dyadic wavelet transform. For each scale 2^j, we detect the maxima of $M_{2^j} f(x,y)$ along the direction of the gradient given by the angle image $A_{2^j} f(x,y)$. We record the position of each maxima as well as the value of $M_{2^j} f(x,y)$ and $A_{2^j} f(x,y)$ at the corresponding location. It is possible to reconstruct the dyadic wavelet transform of $f(x,y)$ given these maxima if and only if there is no other wavelet transform having the same maxima (position, amplitude and angle).

Let us denote by Γ the set of all sequences of functions $\left\{ g_j^1(x,y) , g_j^2(x,y) \right\}_{j \in \mathbf{Z}}$ such that for all $j \in \mathbf{Z}$ the functions $\sqrt{|g_j^1(x,y)|^2 + |g_j^2(x,y)|^2}$ and $\sqrt{|W_{2^j}^1 f(x,y)|^2 + |W_{2^j}^2 f(x,y)|^2}$ have the same local maxima and the same angle value at the maxima locations:

$$\frac{g_j^1(x,y)}{g_j^2(x,y)} = \frac{W_{2^j}^1 f(x,y)}{W_{2^j}^2 f(x,y)} .$$

Let us recall that $\mathbf{V}$ is the space of all possible dyadic wavelet transform. The reconstruction is possible if and only if the intersection of $\mathbf{V}$ with Γ is reduced to the wavelet transform of $f(x,y)$:

$$\mathbf{V} \cap \Gamma = \left\{ \left\{ W_{2^j}^1 f(x,y) , W_{2^j}^2 f(x,y) \right\}_{j \in \mathbf{Z}} \right\} . \tag{28}$$

As in the one-dimensional case, we take a numerical approach to this problem and develop an algorithm for reconstructing the intersection of Γ and $\mathbf{V}$. Like in one dimension, the algorithm iterates on a projector on $\mathbf{V}$ and a projector on Γ. We saw in the previous section that we can define a two-dimensional reproducing kernel equation that defines an orthogonal projection $\mathbf{P_V}$ on $\mathbf{V}$. We also describe in [9] a projector $\mathbf{P_\Gamma}$ that transforms any sequence $\left\{ g_j^1(x,y) , g_j^2(x,y) \right\}_{j \in \mathbf{Z}}$ of two-dimensional function into a new sequence $\left\{ h_j^1(x,y) , h_j^2(x,y) \right\}_{j \in \mathbf{Z}}$ which is in Γ. The projector $\mathbf{P_\Gamma}$ makes a non-linear deformation of $(g_j^1(x,y) , g_j^2(x,y))$ so that the resulting functions $(h_j^1(x,y) , h_j^2(x,y))$ have the same maxima (position, modulus, angle) than $(W_{2^j}^1 f(x,y) , W_{2^j}^2 f(x,y))$. The algorithm iterates on the projectors $\mathbf{P_V}$ and $\mathbf{P_\Gamma}$ to reach the intersection of $\mathbf{V}$ and Γ.

Fig. 8: *The first column gives the modulus images $M_{2^j} f(x,y)$ for $1 \leq j \leq 4$ of the lady image shown at the top left of fig. 9. The second column displays the position of the maxima of $M_{2^j} f(x,y)$. The third and fourth columns display the position of the local maxima whose amplitude are respectively larger than 4 and 8. The maxima that have been removed correspond essentially to the noise and the light texture irregularities.*

Fig. 9: *The upper left is the original lady image. The upper right image is a reconstruction from the maxima representation shown in the second column of fig. 8. This reconstruction is performed with 8 iterations and the noise to signal ratio is* $6.6\ 10^{-2}$. *The lower left and lower right images have been reconstructed from the maxima representation shown respectively in the third and fourth column of fig. 8 (thresholding by the factors 4 and 8). The light textures have disappeared but the strong edges and textures remain unchanged.*

9. Numerical Reconstruction of Images from the Wavelet Maxima

The upper left image of fig. 9 is the original image whereas the upper right is the reconstructed image from the maxima representation with 8 iterations. These two images are visually identical on a good quality image display. If the original image has N pixels, the implementation of both $\mathbf{P_V}$ and $\mathbf{P_\Gamma}$ requires $O(N \log(N))$ computations. Each iterations thus has a complexity of $O(N \log(N))$. The error to signal ratio of this reconstruction is $6.6\ 10^{-2}$. Fig. 10 gives the error to signal ratio when reconstructing the lady image from the wavelet maxima obtained with these wavelets. After 10 iterations, the error to signal ratio is approximatively $3\ 10^{-2}$. In this case, the error to signal ratio decreases steadily and reaches $2.5\ 10^{-5}$ after 5000 iterations. Like in one dimension, the reconstruction errors are concentrated in the high frequencies and the decay rate of the signal to error ratio reflects the decay rate of the high frequency errors. After 300 iterations, the error made on the value of any pixel in the reconstructed image is always smaller than 0.5. Since the pixel values of the original image are coded with integers between 0 and 255, we can recover exactly this image with a round-off operation. The reconstruction algorithm has been tested for a large collection of images including special two-dimensional functions such as Diracs, sinusoidal waves, step edges, Brownian noises... For all these experiments, the error to signal ratio behaves similarly to fig. 10. Let us emphasize that for image processing applications, we need at most 10 iterations for reconstructing an image with no perceivable distortions. The computations required by the algorithms are cascade of convolutions and can thus be implemented in real time on a pipe-line hardware architecture.

The stability of the convergence enables us to slightly perturbate the local maxima representation and reconstruct a close image. The lower left and lower right images in fig. 9 are reconstructed from the maxima whose amplitude are respectively larger than 4 and 8. Here, we removed the maxima produced by the noise and the light textures. As expected, the fine textures disappear in the reconstructed images but the sharp image variations are not affected. The higher the threshold, the more textures disappear. In section 11 we study more carefully how to select the maxima that can be removed from the representation without affecting the quality of the reconstructed image.

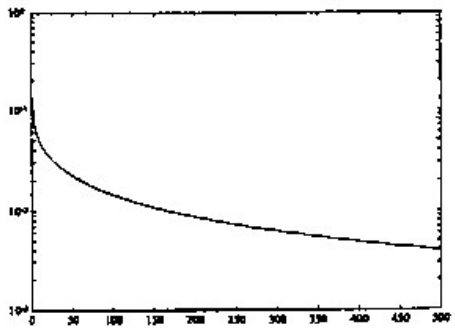

Fig. 10: *Evolution of the error to signal ration for the reconstruction of the lady image from its maxima representation, depending upon the number of iterations on the operator* **P**.

10. Maxima Chain Representation

The meaningful features of the maxima representation are not the local maxima by themselves but rather the chains of maxima points that correspond to the borders of the different image structures. In fig. 7, at each scale, these chains correspond to the border of the circle. For a given scale 2^j, we must chain together the maxima points of $M_{2^j}f(x,y)$ in order to recover these boundaries. Let us study more carefully the chaining procedure. At each maxima point, we know the value of $A_{2^j}f(x,y)$ which is the angle of the gradient vector of $f(x,y)$ smoothed at the scale 2^j. At a given maxima point, clearly the gradient vector is perpendicular to the boundary which goes through this point. We therefore chain together maxima whose respective position is perpendicular to the direction indicated by $A_{2^j}f(x,y)$. Along a given chain, we also impose that the value of $M_{2^j}f(x,y)$ varies smoothly which indicates that the gray level profile of the boundary varies smoothly. We thus chain together two maxima only if their modulus value $M_{2^j}f(x,y)$ is close enough. With this chaining procedure, we build a representation that is a set of chains of maxima for each scale 2^j.

As we just mentioned, the angle $A_{2^j}f(x,y)$ at a maxima point, is perpendicular to the tangent of the maxima chain that goes through this maxima. Once we have chained together the maxima point, we do not need to keep any more this angle information since it is encoded into the chain orientation. This is however only an approximation since the two smoothing functions $\Theta^1(x,y)$ and $\Theta^2(x,y)$ are not strictly equal so the angle $A_{2^j}f(x,y)$ is not exactly perpendicular to the chain tangent. In practice, the error of estimation on the angle $A_{2^j}f(x,y)$ is negligible. The image in fig. 11 is reconstructed after chaining together the maxima at each scale and suppressing the angle information. Of course, one can compute the the tangent to the maxima

chain only when the maxima chain has at least 2 maxima points so we suppress all the chains of length one. A noise removal operation, described in the next section, is performed before the chaining operation. There are almost no visual difference between this image and the original one. The errors introduced by the estimation of the angle are too small to be visible.

We have now reorganized the image information into a set of chains which correspond to the boundary of the image structures. Along each chain we only record the value of the modulus $M_{2^j}f(x,y)$. The reconstruction algorithm shows that this representation carries all the necessary information have a stable reconstruction with no visible distortion. In the next section, we study the application of this image representation to compact image coding.

11. Compact Image Coding

As explained in the introduction, our goal is not to obtain an image coding that reproduces exactly the image with a small mean-square error but an image that is visually of good quality. By adapting the coding to the properties of the human visual perception and to the information content of the image, we want to obtain higher compression rates than with classical image processing methods. There are two separate steps in such a coding procedure. We first select the information that can be removed from the maxima chain representation without affecting the visual quality of the reconstructed image. This selection clearly depends upon the a priori information that we have on the information content of the image. For example, we might we able to remove the light textures of the image without affecting much the overall quality of the image and its information content. We might also be able to remove some image patterns after recognizing them based on the shape of their contours. The second step is to code efficiently the remaining maxima chains. This is a purely signal processing problem that we can solve independently.

The removal of information from the maxima representation requires to understand how these chains relate to the information in the original image. In section 6, we saw that the evolution of the maxima amplitude across scales characterize the type of irregularities it corresponds too in the original signal. In order to compute the type of these singularities, we must relate across scales the maxima that correspond to the same sharp variation in the image. This is done before chaining. We shall consider that two maxima appearing respectively at the scales 2^{j-1} and 2^j do correspond to the same sharp variation of the image if their positions and angle values $A_{2^j}f(x,y)$ are close enough. We insure that this propagation is a one to one relationship across scale. When the scale increases, the number of maxima decreases so clearly not all the maxima that appear at the finner scale 2^1 do propagate at the scales 2^2 and 2^3.

The maxima that appear at the scale 2^1 but do not propagate at the scale 2^2 correspond to the small high frequency image variations that are generally due to the image noise. In the lady

image, there are 54% such maxima at the scale 2^1. Fig. 11 is a reconstruction of the lady image after removing these maxima from the maxima representation. This removal introduces almost no visible distortion on the reconstructed image. One can also remove all the maxima at the scale 2^1 and 2^2 that do not propagate up to the scale 2^3. By doing this, we remove 78% of the maxima at the scale 2^1 and 59% at the scale 2^2. Fig. 12 shows the image reconstructed after this removal. Although the reconstructed image is still of good quality, here we can see that the lighter textures disappeared from the image. All the main boundaries remain unchanged. In the following, we keep only the maxima that do propagate up to the scale 2^3 and build the maxima chain representation from these maxima.

Among all the chains of the maxima representation, we want to suppress the one that correspond to non-important features in the image. Such a removal can be quite sophisticated if one takes into account the information content of the image. For example, the eyes of the lady are very important for the image visualization and must therefore not be affected. Here, we are removing some of the maxima chain with a simple thresholding. Let us call chain amplitude the sum of the maxima modulus along a particular chain. This chain amplitude increases with the length of the chain and the value of the modulus of each maxima in the chain. In the right of fig. 13, we show all the chains at the scale 2^2 whose total amplitude is larger than 32. These chains carry the information of the most important structures in the image. We also suppress at the scales 2^1 and 2^3, all the chains which propagate at the scale 2^2 to a chain that has been removed. Any chain at a scale 2^2 now corresponds to a chain at a scale 2^1 and to another one at the scale 2^3. A priori, the position of these chains might differ slightly due to the blurring effect. In order to save bits for the coding, we are going to suppose that all these positions are the same and equal to the position at the scale 2^2. This means that we modify the position of the chains at the scale 2^1 and 2^3 in order to match exactly the position at the scale 2^2. The complement of information provided at the scale 2^1 and 2^3 is therefore reduced to the value of the modulus at these scales. This characterizes the type of the edges as we explained in section 6. The left image of fig. 13 was reconstructed from this perturbated maxima chain representation. Although the representation was seriously modified, the reconstructed image is of good quality.

From the bandwidth of the image at each scale, we can derive that about $\frac{7}{8}$ of the total image information is concentrated at the three finner scales 2^1, 2^2 and 2^3. We shall therefore concentrate on the coding of the maxima chains After the modification of the maxima chain representation explained in the previous paragraph, the remaining information to code is the position of the chains at the scale 2^2 and the value of the modulus $M_{2^j}f(x,y)$ along these chains, at the scales 2^1, 2^2 and 2^3. We record the position of the first point of each chain and the geometry of the chain is coded with cubic splines. Along each chain, we use a predictive coding technic to record the value of $M_{2^j}f(x,y)$ at the scales 2^1, 2^2 and 2^3. In order to obtain a

compact coding, we quantize the prediction errors on few bits. Errors on the value of the modulus along the chain do not produce large visual distortions on the original image. The image in fig. 14 required less than $\frac{1}{5}$ bit per pixel in order to code the information at the scales 2^1, 2^2 and 2^3. In general, the number of bits required to code the image depends upon the complexity of the image. For the circle image of fig. 7, only 4 10^{-3} bits per pixels are necessary for the first three scale levels.

The coding procedure that we described can be certainly enhanced. The thresholding procedure to remove some of the maxima chain is clearly too brutal. In particular, it creates visible distortions around the eyes of the lady. The information content of the image must be taken into account more carefully. Also, the predictive coding technic that we used to record the modulus amplitude is quite simple and one can probably develop a more efficient coding scheme.

Fig. 11: *The right image shows the position of the maxima at the finner scale 2^1, after removing the maxima that do not propagate at the scale 2^2. The left image is a reconstruction after this maxima removal. The angle information $A_{2^1}f(x,y)$ was not kept but estimated from the tangent of the maxima chains. These modifications introduce very few distortion.*

Fig. 12: *The right image shows the position of the maxima at the finner scale 2^1, after removing the maxima that do not propagate at the scale 2^2 and 2^3. At the scale 2^2 was also removed the maxima that do not propagate at the scale 2^3. The left image is a reconstruction from the maxima chain representation built after this maxima removal. Some light textures disappeared in this reconstructed image but the strong edges and textures are not affected.*

Fig. 13: *We removed all the maxima at the scale 2^1 and 2^2 that do not propagate up to the scale 2^3. We then thresholded the maxima chain representation and removed all the chains at the scale 2^2, whose total amplitude is larger than 32. The right image shows the remaining chains at the scale 2^2. Only 200 chains remain with a total of 5000 maxima points. At the scale 2^1 and 2^3, we kept only the chains which correspond to a chain at the scale 2^2 which has not been removed. We also moved the position of the chain at the scales 2^1 and 2^3 so that their position is the same than at the scale 2^2. The left image is the image reconstructed from the corresponding maxima representation. Although the maxima chain representation was heavily modified, the reconstructed image is of good quality.*

Fig. 14: *This image was reconstructed after making a compact coding of the chains that were kept in fig. 13. The coding introduces some further distortions but it requires less than $\frac{1}{5}$ bit per pixel (with no entropy coding) to code the maxima information at the scales $2^1, 2^2$ and 2^3.*

12. Application to image understanding

The maxima chain representation provides a reorganization of the whole image information into multiscale edges. There is no need to further motivate the importance of edges for image pattern recognition since many existing vision algorithms are edge based. An important contribution of the wavelet theory is the understanding of the behavior of edges through scale. This enable us to characterize the types of edges through their Lipschitz exponent and their blurring factor σ. Instead of viewing edge detection as a binary process that classifies the image pixels into edge or non-edges, we can obtain a complement information about the local shape of the edge. We indeed believe that edge detection should not only separate smooth regions from discontinuities but rather characterize the different type of discontinuities. Even if the image intensity is extremely irregular almost everywhere, often there are some singularities that one might want to distinguish from others. For example, in the lady image of fig. 9, the irregularities of the fur texture of the hat are very different from the edges corresponding to the border of the shoulders. A better understanding of the classifications of the image singularities is necessary to make a stable distinction between edges and textures and for texture discrimination.

13. Conclusion

We saw that multiscale edges can be obtained from the local maxima of a wavelet transform and we explained how to reconstruct the original signal from these local maxima. The reconstruction algorithm can restore exactly the image and thus gives a verification of Marr conjecture on the completeness of multiscale edges. The mathematical study of the completeness problem remains on open question. We also proved that one can characterize the type of the edge points from the decay of the wavelet maxima amplitude across scales. This is particularly important to discriminate the different image sharp variations. The wavelet maxima representation is a new reorganization of the image information in order to develop solutions to image processing problems, from the properties of the image edges. In particular, we showed that it enables to build a compact image coding that takes into account the information content of the image.

References

1. Canny, J., "A Computational approach to edge detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 8, pp. 679-698, 1986.

2. Curtis, S., Shitz, S., and Oppenheim, V., "Reconstruction of Non-periodic two-dimensional signals from zero-crossings," *IEEE Trans. Acoustic Speech and Signal Process.*, vol. 35, pp. 890-893, 1987.

3. Hummel, R. and Moniot, R., "A network approach to reconstruction from zero-crossings," *Proc. of IEEE Workshop on computer vision*, Dec. 1987.

4. Hummel, R. and Moniot, R., "Reconstruction from Zero-Crossings in Scale-Space," *IEEE Trans. on Acoustic Speech and Signal Processing*, Dec. 1989.

5. Koenderink, J., "The structure of images," *Biological Cybernetics*, Springer Verlag, 1984.

6. Kunt, M., Ikonomopoulos, A., and Kocher, M., "Second generation image coding technics," *Proceed. of the IEEE*, vol. 74, pp. 549-575, April 1985.

7. Logan, B., "Information in the zero-crossings of band pass signals," *Bell Systems Tech. Journ.*, vol. 56, p. 510, 1977.

8. Mallat, S., "Review of multifrequency channel decompositions of images and wavelet models," *IEEE Trans. on Acoustic Speech and Signal Processing*, Dec. 1989.

9. Mallat, S. and Zhong, S., "Complete signal representation with multiscale edges," *NYU, Computer Science Tech. Report 483*, Dec. 1989.

10. Marr, D., in *Vision*, W.H.Freeman and Company, 1982.

11. Marr, D. and Hildreth, E., "Theory of edge detection," *Proc. of the Royal Society of London*, vol. 207, pp. 187-217, 1980.

12. Meyer, Y., in *Ondelettes et Operateurs*, Hermann, 1988.

13. Rosenfeld, A. and Thurston, M., "Edge and curve detection for visual scene analysis," *IEEE Trans. on Computers*, 1971 vol C-20.

14. Sanz, J. and Huang, T., "Theorem and experiments on image reconstruction from zero-crossings," *Research report RJ5460*, IBM.

15. Witkin, A., "Scale space filtering," *Proc. Int. Joint Conf. Artificial Intell.*, 1983.

16. Youla, D. and Webb, H., "Image restoration by the method of convex projections," *IEEE Trans. Medical Imaging*, vol. 1, pp. 81-101, Oct. 1982.

17. Yuille, A. and Poggio, T., "Scaling theorems for zero crossings," *IEEE trans on PAMI*, vol. 8, Jan 1986.

18. Zeevi, Y. and Rotem, D., "Image reconstruction from zero-crossings," *IEEE Acoustic Speech and Signal Proc.*, vol. 34, pp. 1269-1277, 1986.