

The Undecimated Wavelet Decomposition and its Reconstruction

Jean-Luc Starck, Jalal Fadili, and Fionn Murtagh

Abstract—This paper describes the undecimated wavelet transform and its reconstruction. In the first part, we show the relation between two well known undecimated wavelet transforms, the standard undecimated wavelet transform and the isotropic undecimated wavelet transform. Then we present new filter banks specially designed for undecimated wavelet decompositions which have some useful properties such as being robust to ringing artifacts which appear generally in wavelet-based denoising methods. A range of examples illustrates the results.

Index Terms—Denoising, multiresolution, multiscale decomposition, restoration, wavelet.

I. INTRODUCTION

MULTISCALE methods have become very popular in the last couple of decades, especially with the development of wavelets. Background texts on the wavelet transform include [1]–[5]. The most widely used wavelet transform (WT) algorithm is certainly the decimated bi-orthogonal wavelet transform (DWT) which is used in JPEG2000. While the bi-orthogonal wavelet transform has led to successful implementation in image compression, results were far from optimal for other applications such as filtering, deconvolution, detection, or more generally, analysis of data. This is mainly due to the loss of the translation-invariance property in the DWT, leading to a large number of artifacts when an image is reconstructed after modification of its wavelet coefficients.

For this reason, some physicists and astronomers have preferred to continue working with the continuous wavelet transform [6], [7], even if the price to pay is 1) a great amount of redundancy in the transformation (i.e., there are many more pixels in the transformed data than in the input image) and 2) perfect reconstruction is not possible (i.e., an image cannot be reconstructed from its coefficients). For some applications like fractal analysis, these drawbacks have no impact because there is no need to apply a reconstruction and computers can support the redundancy.

For other applications where a reconstruction is needed, some researchers have chosen an intermediate approach, which consists of keeping the filter bank construction with fast and dyadic

algorithms, but eliminating the decimation step in the orthogonal wavelet transform [8], [9]. In Starck *et al.* [10], it was shown that thresholding using an undecimated transform rather than a decimated one can improve the result by more than 2.5 dB in denoising applications. The undecimated decomposition is computed by using the same filter bank as in the standard decimated bi-orthogonal wavelet transform and it leads to a three-directional analysis (horizontal, vertical, diagonal). Each band has the same size as the original image.

Because astronomical images contain mostly isotropic sources (stars, galaxies, etc.), astronomers prefer generally to use another transform, the Isotropic Undecimated Wavelet Transform (IUWT) [11]. Such isotropic image content also typifies many classes of images in biology. The IUWT also uses a filter bank but its filters do not verify the dealiasing condition, and decimation cannot be applied.

These two undecimated multiscale methods (i.e., the general undecimated wavelet transform to be discussed in Section II-A below as the UWT, and the IUWT) are very powerful for image restoration but, as with any other wavelet decomposition, present the drawback of creating ringing artifacts around singularities or edges. This has motivated the recent development of iterative techniques combining at the same time a multiscale method and a penalization term such as the Total Variation (TV) [12]–[14] or the l_1 norm of coefficients in the wavelet or curvelet decomposition [15].

Even if the wavelet frame theory and oversampled filter banks are well understood [3], [16]–[18], relatively few studies have been dedicated to the development of oversampled filter banks [19], [16].

This Paper

We show that the redundancy of the decomposition can be used for designing new filter banks. As a consequence, we are able to build a filter bank such that the filters \tilde{h} and \tilde{g} used in the reconstruction are both positive, which makes the reconstruction very robust to the ringing artifact problem. The decomposition is done using wavelets and the reconstruction using only scaling functions. This aspect can be very important in some applications such as edge detection.

Section II introduces the undecimated wavelet transform and establishes the relation between the UWT and the IUWT. New filter banks are discussed Section III. In Section IV, we show that the reconstruction from the thresholded nonsubsampling coefficients is not straightforward and we discuss an iterative scheme which gives better reconstruction results than just a simple direct synthesis. A range of experiments illustrates the results in Section V.

Manuscript received June 20, 2005; revised July 17, 2006. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Srdjan Stankovic.

J.-L. Starck is with the CEA-Saclay, DAPNIA/SEDI-SAP, Service d'Astrophysique, F-91191 Gif sur Yvette, France.

J. Fadili is with the GREYC CNRS UMR 6072, Image Processing Group, ENSICAEN 14050, Caen Cedex, France.

F. Murtagh is with the Department of Computer Science, Royal Holloway, University of London, Egham TW20 0EX, U.K.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2006.887733

Notation

For a real discrete-time filter whose impulse response is $h[n]$, $\tilde{h}[n] = h[-n]$, $n \in \mathbb{Z}$ is its time-reversed version. The hat notation will be used for the Fourier transform of square-integrable signals. For a filter h , its z -transform is written $H(z)$. The convolution product of two signals in $\ell^2(\mathbb{Z})$ will be written $*$. For the octave band nonsubsampling wavelet representation, analysis (respectively, synthesis) filters are denoted h and g (respectively, \tilde{h} and \tilde{g}). The scaling and wavelet functions using for the analysis (respectively, synthesis) are denoted ϕ ($\phi(x/2) = \sum_k h[k]\phi(x-k)$, $x \in \mathbb{R}$ and $k \in \mathbb{Z}$) and ψ ($\psi(x/2) = \sum_k g[k]\phi(x-k)$, $x \in \mathbb{R}$ and $k \in \mathbb{Z}$) (respectively, $\tilde{\phi}$ and $\tilde{\psi}$). We also define the scaled dilated and translated version of ϕ at scale j and position k as $\phi_{j,k}(x) = 2^{-j}\phi(2^{-j}x - k)$, and similarly for ψ , $\tilde{\phi}$ and $\tilde{\psi}$.

II. UNDECIMATED WAVELET TRANSFORM

A. Two-Dimensional Standard Undecimated Wavelet Transform

The undecimated wavelet transform (UWT) \mathcal{W} using the filter bank (h, g) of a 1-D signal c_0 leads to a set $\mathcal{W} = \{w_1, \dots, w_J, c_J\}$ where w_j are the wavelet coefficients at scale j and c_J are the coefficients at the coarsest resolution. The passage from one resolution to the next one is obtained using the “à trous” algorithm [9], [20]

$$\begin{aligned} c_{j+1}[l] &= (\tilde{h}^{(j)} * c_j)[l] = \sum_k h[k]c_j[l + 2^j k] \\ w_{j+1}[l] &= (\tilde{g}^{(j)} * c_j)[l] = \sum_k g[k]c_j[l + 2^j k] \end{aligned} \quad (1)$$

where $h^{(j)}[l] = h[l]$ if $l/2^j$ is an integer and 0, otherwise. For example, we have

$$h^{(1)} = (\dots, h[-2], 0, h[-1], 0, h[0], 0, h[1], 0, h[2], \dots).$$

The reconstruction is obtained by

$$c_j[l] = \frac{1}{2} \left[(\tilde{h}^{(j)} * c_{j+1})[l] + (\tilde{g}^{(j)} * w_{j+1})[l] \right]. \quad (2)$$

The filter bank $(h, g, \tilde{h}, \tilde{g})$ needs only to verify the exact reconstruction condition

$$H(z^{-1})\tilde{H}(z) + G(z^{-1})\tilde{G}(z) = 1. \quad (3)$$

This provides us with a higher degree of freedom when designing the synthesis prototype filter bank.

The à trous algorithm can be extended to 2-D by

$$\begin{aligned} c_{j+1}[k, l] &= (\tilde{h}^{(j)}\tilde{h}^{(j)} * c_j)[k, l] \\ w_{j+1}^1[k, l] &= (\tilde{g}^{(j)}\tilde{h}^{(j)} * c_j)[k, l] \\ w_{j+1}^2[k, l] &= (\tilde{h}^{(j)}\tilde{g}^{(j)} * c_j)[k, l] \\ w_{j+1}^3[k, l] &= (\tilde{g}^{(j)}\tilde{g}^{(j)} * c_j)[k, l] \end{aligned} \quad (4)$$

where $hg * c$ is the convolution of c by the separable filter hg (i.e., convolution first along the columns by h and then convolution along the rows by g). At each scale, we have three wavelet images, w^1, w^2, w^3 , and each has the same size as the original image. The redundancy factor is, therefore, $3(J-1) + 1$ [3].

B. Two-Dimensional Isotropic Undecimated Wavelet Transform

The IUWT algorithm is well known in the astronomical domain, because it is well adapted to astronomical data where objects are more or less isotropic in most cases [11]. Requirements for a good analysis of such data are as follows.

- Filters must be symmetric ($\tilde{h}[k] = h[k]$, and $\tilde{g}[k] = g[k]$).
- In 2-D or a higher dimension, h, g, ψ, ϕ must be nearly isotropic.

Filters do not need to be orthogonal or bi-orthogonal and this lack of the need for orthogonality or bi-orthogonality is beneficial for design freedom. For computational reasons, we also prefer to have the separability; $h[k, l] = h[k]h[l]$. Separability is not a required condition, but it allows us to have a fast calculation, which is important for a large data set.

This has motivated the following choice for the analysis scaling and wavelet functions [11]:

$$\begin{aligned} \phi_1(x) &= \frac{1}{12} (|x-2|^3 - 4|x-1|^3 + 6|x|^3 \\ &\quad - 4|x+1|^3 + |x+2|^3) \\ \phi(x, y) &= \phi_1(x)\phi_1(y) \\ \frac{1}{4}\psi\left(\frac{x}{2}, \frac{y}{2}\right) &= \phi(x, y) - \frac{1}{4}\phi\left(\frac{x}{2}, \frac{y}{2}\right) \end{aligned} \quad (5)$$

where $\phi_1(x)$ is the spline of order 3, and the wavelet function is defined as the difference between two resolutions. The related filters h and g are defined by

$$\begin{aligned} h^{(1D)}[k] &= \frac{[1, 4, 6, 4, 1]}{16}, \quad k = -2, \dots, 2 \\ h[k, l] &= h^{(1D)}[k]h^{(1D)}[l] \\ g[k, l] &= \delta[k, l] - h[k, l] \end{aligned} \quad (6)$$

where δ is defined as $\delta[0, 0] = 1$ and $\delta[k, l] = 0$ for all (k, l) different from $(0, 0)$.

The following useful properties characterize any pair of even-symmetric analysis FIR (finite impulse response) filters $(h, g = \delta - h)$ such as those of (6).

1) *Property 1:* For any pair of even symmetric filters h and g such that $g = \delta - h$, the following holds.

- 1) This FIR filter bank implements a frame decomposition, and perfect reconstruction using FIR filters is possible.
- 2) The above filters do not implement a tight frame decomposition.

Proof: For the filter bank (6), the z -transforms $H(z)$ and $G(z)$ have no zeros in common. Thus, statement 1) follows directly from [16, Theorem 1 and Proposition 3]. The second part of 1) can be alternatively obtained by the Bezout theorem for polynomials. Statement 2) is due to the obvious fact that

even-symmetric h and $g = \delta - h$ cannot be power complementary. For an analysis filter bank h and g to be power-complementary, it must satisfy

$$|\hat{h}(\nu)|^2 + |\hat{g}(\nu)|^2 = c, \quad c > 0. \quad (7)$$

As h was set to have unit average and $g = \delta - h$, c is necessarily 1 for the above relation to hold at $\nu = 0$. Furthermore, for real even-symmetric filter h , its Fourier transform is also real and even symmetric. Thus, the left hand side of the above relation can be rewritten

$$\hat{h}^2(\nu) + (1 - \hat{h}(\nu))^2 = 1 + 2\hat{h}(\nu)(\hat{h}(\nu) - 1). \quad (8)$$

Hence, for (7) to hold, $h = \delta$ is the only trivial solution. Then, applying [16, Theorem 2], the result follows. ■

From the structure of g , it is easily seen that the wavelet coefficients are obtained just by taking the difference between two resolutions

$$w_{j+1}[k, l] = c_j[k, l] - c_{j+1}[k, l] \quad (9)$$

where $c_{j+1}[k, l] = (\bar{h}^{(j)}\bar{h}^{(j)} * c_j)[k, l]$. At each scale j , we obtain one set $\{w_j\}$ (and not three as in the undecimated WT, denoted UWT above) which has the same number of pixels as the input image.

The reconstruction is obtained by a simple co-addition of all wavelet scales and the final smoothed array, namely

$$c_0[k, l] = c_J[k, l] + \sum_{j=1}^J w_j[k, l]. \quad (10)$$

That is, the synthesis filters are $\tilde{h} = \delta$ and $\tilde{g} = \delta$, which are indeed FIR as expected from Property 1(i). This wavelet transformation is very well adapted to the analysis of images which contain isotropic objects such as in astronomy [11] or in biology [21]. This construction has a close relation to the Laplacian pyramidal construction introduced by Burt and Adelson [22] or the FFT-based pyramidal wavelet transform [4]

C. Relation Between the UWT and the IUWT

Equivalence between the UWT and Mallat's "à trous" algorithm has been previously reported by Shensa [20]. However, to the best of our knowledge, there is no work that has yet shed light on the relation between the UWT and the IUWT.

Since the dealiasing filter bank condition is not required anymore in the UWT decomposition, we can build the standard three-directional undecimated filter bank using the non-(bi-)orthogonal "Astro" filter bank ($h^{1D} = [1, 4, 6, 4, 1]/16$, $g^{1D} = \delta - h^{1D} = [-1, -4, 10, -4, -1]/16$ and $\tilde{h} = \tilde{g} = \delta$). In two dimensions, this filter bank leads to a wavelet decomposition with three orientations w_j^1, w_j^2, w_j^3 at each scale j , but with the same property as for the IUWT, i.e., the sum of all scales reproduces the original image

$$c_0[k, l] = c_J[k, l] + \sum_{j=1}^J \sum_{d=1}^3 w_j^d[k, l]. \quad (11)$$

Indeed, a straightforward calculation immediately shows that

$$w_j^1 + w_j^2 + w_j^3 = c_j - c_{j+1}. \quad (12)$$

Therefore, the sum of the three directions reproduces the IUWT detail band at scale j . Fig. 1 shows the UWT of the galaxy NGC2997. When we add the three directional wavelet bands at a given scale, we recover exactly the isotropic undecimated scale. When we add all bands, we recover exactly the original image. The relation between the two undecimated decompositions is clear.

III. DESIGNING NEW FILTER BANKS

A. A Surprising Result

Because the decomposition is nonsubsampling, there are many ways to reconstruct the original image from its wavelet transform. For a given filter bank (h, g) , any filter bank (\tilde{h}, \tilde{g}) which satisfies the reconstruction condition of (3) leads to exact reconstruction. For instance, for isotropic h , if we choose $\tilde{h} = h$ (the synthesis scaling function $\tilde{\phi} = \phi$) we obtain a filter \tilde{g} defined by

$$\tilde{g} = \delta + h.$$

Again, as expected from Property 1, the analysis filter bank $(h, g = \delta - h)$ implements a (nontight) frame decomposition for FIR symmetric h , where $\tilde{h} = h$ and $\tilde{g} = \delta + h$ are also FIR filters. For instance, if $h = [1, 4, 6, 4, 1]/16$, then $\tilde{g} = [1, 4, 22, 4, 1]/16$. \tilde{g} is **positive**. This means that \tilde{g} is no longer related to a wavelet function. The synthesis scaling function related to \tilde{g} is defined by

$$\frac{1}{2}\tilde{\psi}\left(\frac{x}{2}\right) = \phi(x) + \frac{1}{2}\phi\left(\frac{x}{2}\right). \quad (13)$$

Finally, note that choosing $\tilde{\phi} = \phi$, any synthesis function $\tilde{\psi}$ which satisfies

$$\hat{\tilde{\psi}}(2\nu)\hat{\psi}(2\nu) = \hat{\phi}^2(\nu) - \hat{\phi}^2(2\nu) \quad (14)$$

leads to an exact reconstruction [3] and $\tilde{\psi}(0)$ can take any value. The synthesis function $\tilde{\psi}$ does not need to verify the admissibility condition (i.e., to have a zero mean).

Fig. 2 shows the two scaling functions $\tilde{\phi}(x) (= \phi)$ and $\tilde{\psi}(x)$ used in the reconstruction in 1-D, corresponding to the synthesis filters $\tilde{h} = h$ and $\tilde{g} = \delta + h$. Fig. 3 shows the backprojection of a wavelet coefficient in 2-D (all wavelet coefficients are set to zero, except one), when the nonzero coefficient belongs to different bands. We can see that the reconstruction functions are positive.

Finally, we have an expansion of a 1-D signal s

$$s(x) = \sum_k c_J[k]\tilde{\phi}_{J,k}(x) + \sum_{j=1}^J \sum_k w_j[k]\tilde{\psi}_{j,k}(x) \quad (15)$$

where $\tilde{\phi}$ and $\tilde{\psi}$ are not wavelet functions (both of them have a nonzero mean and are positive), but the w_j are wavelet coefficients.




Fig. 1. UWT of the galaxy NGC2997 using the Astro filter bank. The addition of three bands at a given scale is exactly the band related to the isotropic wavelet transform. Addition of all bands reproduces exactly the original image.




Fig. 2. Left: $\bar{\phi}$ synthesis scaling function. Right: $\bar{\psi}$ detail synthesis function.

B. Reconstruction From the Haar Undecimated Coefficients

The Haar filters ($h = \tilde{h} = [1/2, 1/2]$, $g = \tilde{g} = [-1/2, 1/2]$) are not considered good filters because of their lack of regularity. They are, however, very useful in many situations such as denoising where their simplicity allows us to derive analytical or semi-analytical detection levels even when the noise does not follow a Gaussian distribution. To our knowledge, there is no real alternative to the Haar filters for Poisson noise, even if they are known to produce block artifacts in the reconstruction after thresholding. Recent papers using the Haar filters for Poisson noise are [23]–[29]. The Haar transform has also close relations with the Total Variation norm (TV) [30], and it has been shown that, for 1-D denoising, the undecimated Haar soft thresholding produces similar results to the TV norm regularization [30].

Adopting the same design approach as before, we can reconstruct a signal from its Haar wavelet coefficients choosing a smooth scaling function. For instance, if $\tilde{h} = [1, 4, 6, 4, 1]/16$, it is easy to see that the z transforms of these three filters are, respectively

$$\begin{aligned} H(z) &= \frac{1 + z^{-1}}{2}, & G(z) &= \frac{z^{-1} - 1}{2} \\ \tilde{H}(z) &= \frac{z^2 + 4z + 6 + 4z^{-1} + z^{-2}}{16}. \end{aligned} \quad (16)$$

From the exact reconstruction condition in (3), we obtain

$$\tilde{G}(z) = \frac{1 - \tilde{H}(z)H(z^{-1})}{G(z^{-1})}. \quad (17)$$




Fig. 3. Back projection: Each image corresponds to the backprojection of one wavelet coefficient. All of these reconstructed images are positive (no negative values). From left to right, the coefficient belongs to the vertical, horizontal, and diagonal direction. From top to bottom, the scale index increases.

In the case of the spline filter bank, this yields after some re-arrangement (where we used simple convolution properties of splines)

$$\begin{aligned} \tilde{G}(z) &= -2 \frac{1 - z^3 \left(\frac{1+z^{-1}}{2}\right)^5}{1 - z^{-1}} \\ &= z^3 \frac{1 + 6z^{-1} + 16z^{-2} - 6z^{-3} - z^{-4}}{16} \end{aligned} \quad (18)$$

which is the z -transform of the corresponding filter $\tilde{g} = [1, 6, 16, -6, -1]/16$.

The Haar analysis filters fulfill the following property.

1) *Property 2:* Haar analysis filters implement a tight frame expansion. Perfect reconstruction with FIR synthesis filters is possible.

Proof: Proof of the second statement is obviously the same as in Property 1.

Haar analysis filters have no zeros in common and are power complementary. Therefore, the first statement is a consequence of [16, Theorem 2]. An alternative way to prove this result is to show the existence of the tight frame bounds in the same way as in [31]. ■

Fig. 4, upper left and right, shows the coarsest scale and a wavelet scale of the Haar transform when the input signal con-

tains only zero values except one sample (Dirac). Fig. 4, bottom left, shows the backprojection of a Dirac at the coarsest scale (all coefficients are set to zero) and Fig. 4, bottom right, shows the backprojection of a Haar wavelet coefficient. Since the synthesis filters are regular, the backprojection of a Dirac does not produce any block artifact. Finally, we would like to point out that other alternatives exist. For example the filter bank $h = [1/2, -1/2]$, $g = [-1/4, 1/2, -1/4]$, $\tilde{h} = [1, 3, 3, 1]/8$ and $\tilde{g} = [1, 6, 1]/4$ leads also to an interesting solution where the synthesis filters are both positive.

C. Another Interesting Filter Bank

A particular case is obtained when $\hat{\phi} = \hat{\psi}$ and $\hat{\psi}(2\nu) = (\hat{\phi}^2(\nu) - \hat{\phi}^2(2\nu))/\hat{\phi}(\nu)$, which leads to a filter g equal to $\delta - h * h$. In this case, the synthesis function $\tilde{\psi}$ is defined by $(1/2)\tilde{\psi}(x/2) = \phi(x)$ and the filter $\tilde{g} = \delta$ is the solution to (3). We end up with a synthesis scheme where only the smooth part is convolved during the reconstruction. Furthermore, for a symmetric FIR filter h , it can be easily shown that this filter bank fulfills the statements of Property 1.

Deriving h from a spline scaling function, for instance B_1 ($h_1 = [1, 2, 1]/4$) or B_3 ($h_3 = [1, 4, 6, 4, 1]/16$) (note that $h_3 = h_1 * h_1$), since h is even-symmetric (i.e.,

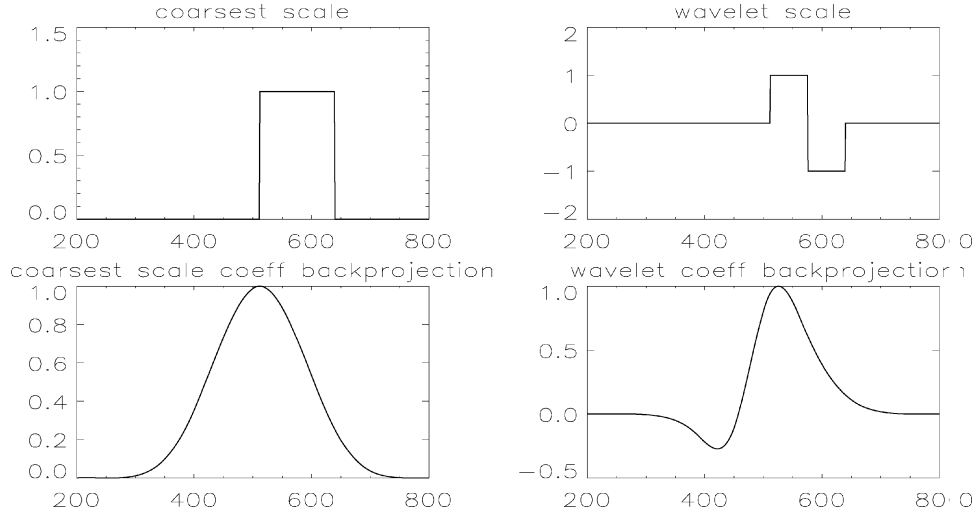


Fig. 4. Haar Undecimated Transform. Upper left: Coarsest scale when the signal contains only one sample (a Dirac). Upper right: One wavelet scale of the Dirac decomposition. Bottom left: Backprojection of a Dirac at the coarsest scale. Bottom right: Backprojection of a Haar wavelet coefficient.




Fig. 5. Left: $\tilde{\phi}$ analysis scaling function. Right: $\tilde{\psi}$ analysis wavelet function. The synthesis functions $\bar{\phi}$ and $\bar{\psi}$ are the same as those in Fig. 2.

$H(z) = H(z^{-1})$, the z -transform of g is (19), shown at the bottom of the page, which is the z -transform of the filter $g = [-1, -8, -28, -56, 186, -56, -28, -8, -1]/256$. We get the following filter bank:

$$h = h_3 = \tilde{h} = \frac{[1, 4, 6, 4, 1]}{16}$$

$$g = \delta - h * h = \frac{[-1, -8, -28, -56, 186, -56, -28, -8, -1]}{256} \quad (20)$$

$$\tilde{g} = \delta. \quad (21)$$

With this filter bank, there is a no convolution with the filter \tilde{g} during the reconstruction. Only the low-pass synthesis filter \tilde{h} is used. The reconstruction formula is

$$c_j[l] = (h^{(j)} * c_{j+1})[l] + w_{j+1}[l] \quad (22)$$

and denoting $L^j = h^{(0)} * \dots * h^{(j-1)}$ and $L^0 = \delta$, we have

$$c_0[l] = (L^J * c_J)[l] + \sum_{j=1}^J (L^{j-1} * w_j)[l]. \quad (23)$$

Each wavelet scale is convolved with a low-pass filter.

Fig. 5 shows the analysis scaling and wavelet functions. The synthesis functions $\tilde{\phi}$ and $\tilde{\psi}$ are the same as those in Fig. 2. We will see in the experimental section that such filters are much more robust to the ringing artifact than classical filters.

IV. ITERATIVE RECONSTRUCTION

Denoting \mathcal{W} the undecimated wavelet transform operator and \mathcal{R} the reconstruction operator, and thanks to the exact reconstruction formulae, we have the relation: $\alpha_S = \mathcal{W}\mathcal{R}\alpha_S$, where S is an image and α_S its wavelet coefficients (i.e., $\alpha_S = \mathcal{W}S$), but we lose one fundamental property of the (bi-)orthogonal

$$G(z) = 1 - H^2(z) = 1 - z^4 \left(\frac{1 + z^{-1}}{2} \right)^8 = \frac{(-z^4 - 8z^3 - 28z^2 - 56z + 186 - 56z^{-1} - 28z^{-2} - 8z^{-3} - z^{-4})}{256} \quad (19)$$

WT. Indeed, the relation $\alpha = \mathcal{WR}\alpha$ is not true for all α sets. For example, if we set all wavelet coefficients to zero except one at a coarse scale, there is no image such that its UWT would produce a Dirac at a coarse scale. Another way to understand this point is to consider the Fourier domain of a given undecimated scale. Indeed, wavelet coefficients α_j at scale j obtained using the wavelet transform operator will contain information only localised at a given frequency band, but any modification of the coefficients at this scale, such as a thresholding ($\alpha_T = \Delta_T(\alpha)$, where Δ_T is the thresholding operator with threshold T and α_T are the thresholded coefficients), will introduce some frequency components which should not exist at this scale j , and we have $\alpha_T \neq \mathcal{WR}\alpha_T$.

A. Reconstruction From a Subset of Coefficients

If only a subset of coefficients (for instance after thresholding) is different from zero, we would like to reconstruct an image \tilde{S} such that its wavelet transform reproduces the nonzero wavelet coefficients. This can be seen as an inverse problem. We want to solve the following optimization problem $\min_{\tilde{S}} \|M(\alpha_T - \mathcal{W}\tilde{S})\|_2^2$ where $M_{j,k}$ is the multiresolution support of α , i.e., $M_{j,k} = 1$ if the wavelet coefficient $\alpha_{j,k}$ at scale j and at position k is different from zero, and $M_{j,k} = 0$ otherwise. A solution can be obtained using the Landweber iterative scheme [4], [32]

$$\tilde{S}^{n+1} = \tilde{S}^n + \mathcal{R}M[\alpha_T - \mathcal{W}\tilde{S}^n]. \quad (24)$$

If the solution is known to be positive, the positivity constraint can be introduced using the following equation:

$$\tilde{S}^{n+1} = P_+ \left(\tilde{S}^n + \mathcal{R}M[\alpha_T - \mathcal{W}\tilde{S}^n] \right) \quad (25)$$

where P_+ is the projection on the cone of nonnegative images.

For denoising applications, additive constraints such as the TV or the l_1 norm can be added as well.

B. Equivalence With Alternating Projection

The alternating projection method [33] is a known technique for consistent reconstruction with frames, frame design and more generally for inverse eigenvalue problems (see survey in [34]). As stated above, in our setting, we are seeking a consistent reconstruction, that is the reconstructed image must satisfy some structural constraints (multiresolution support and positivity). The closure of the reconstruction set of $\tilde{\alpha}$ can be represented as the intersection of the following sets.

- W is the range of the wavelet transform: $W = \{\alpha : \alpha = \mathcal{WS} \text{ for some } S \in \ell_2(\mathbb{Z}^2)\}$.
- \mathcal{M} is determined by the requirement that the elements of $\tilde{\alpha}$ must reproduce the coefficients of interest (i.e., $\tilde{\alpha}_T = \alpha_T$ or $M\tilde{\alpha} = M\alpha$).
- \mathcal{C} is the set of all $\tilde{\alpha}$ such that the corresponding reconstructed image is positive valued.

The alternating projection algorithm starts from an initial estimate $\tilde{\alpha}_0$, and then alternately projects onto the sets W , \mathcal{M} and \mathcal{C} , and it repeats the process *ad infinitum*. W is a subspace of ℓ_2 , the sets \mathcal{M} and \mathcal{C} are obviously nonempty closed convex sets. The alternating projections in this case is equivalent to the

method of projection onto convex sets (POCS) [35]. By standard convergence results about cyclic projections, the algorithm is supposed to converge (here in a strong sense) to a point in the intersection of the above sets.

The projector onto the range of the wavelet transform is $\mathcal{P}_W = \mathcal{WR}$. The projection of some $\tilde{\alpha}$ onto \mathcal{M} is obtained by

$$\mathcal{P}_{\mathcal{M}}\tilde{\alpha}_{j,k} = \begin{cases} \alpha_{T,j,k}, & \text{if } M_{j,k} = 1 \\ \tilde{\alpha}_{j,k}, & \text{otherwise.} \end{cases} \quad (26)$$

The projection of some $\tilde{\alpha}$ in a Hilbert space onto \mathcal{C} is given by

$$\mathcal{P}_{\mathcal{C}}\tilde{\alpha} = \mathcal{W}P_+\mathcal{R}\tilde{\alpha}. \quad (27)$$

One can easily verify that these two operators are idempotent (projectors). Applying one iteration of alternating projections to some $\tilde{\alpha}^n = \mathcal{W}\tilde{S}^n$ at iteration n yields a solution at $n+1$

$$\tilde{\alpha}^{n+1} = \mathcal{P}_{\mathcal{C}} \circ \mathcal{P}_{\mathcal{M}} \circ \mathcal{P}_W \tilde{\alpha}^n = \mathcal{W}P_+\mathcal{R}[M\alpha_T + (I-M)\mathcal{WR}\tilde{\alpha}^n] \quad (28)$$

where I is the identity matrix of the same size as M . Finally, applying the synthesis operator \mathcal{R} to both sides, rearranging the terms inside the brackets and recalling that by definition of the (weak generalized) left inverse $\mathcal{R}\mathcal{WR} = \mathcal{R}$, one can easily see that this equation is exactly the same as the one of the iterative scheme in (25).

It is worth pointing out that convergence properties of this iterative reconstruction scheme are influenced by the choice of the analysis/reconstruction filter bank. Therefore, they should be designed cautiously. Authors in [19] have also observed this problem. However, this does not mean that the iterative algorithm will not converge with left inverses other than the Moore–Penrose pseudo-inverse (i.e., synthesis frame minimal dual of the analysis frame). The design of general synthesis filters which would guarantee convergence of the POCS-based reconstruction algorithm in the case where \mathcal{R} is not the Moore–Penrose inverse of \mathcal{W} is still an open question. Actually, with our experiments, we observed that the algorithm always converged and gave very good results.

V. EXPERIMENTS

A. Nonlinear Approximation With an Undecimated Transform

In order to compare how well an UWT is able to represent an image, we can plot the nonlinear approximation curve. Since we are considering here undecimated decompositions, we plot the error as a function of the threshold level rather than the number of coefficients.

1) *Lena Image*: Fig. 6 shows such curves for a threshold varying from 0 to 30 and with different filter banks on the Lena image. From top to bottom, we see:

- 1) the undecimated WT using the nonorthogonal filter banks of (21), with direct reconstruction;
- 2) the standard decimated bi-orthogonal wavelet transform with the 7/9 filters [36];
- 3) the UWT (7/9 filters) with direct reconstruction;
- 4) the UWT using the nonorthogonal filter banks and an iterative reconstruction;
- 5) the UWT (7/9 filters) with an iterative reconstruction.

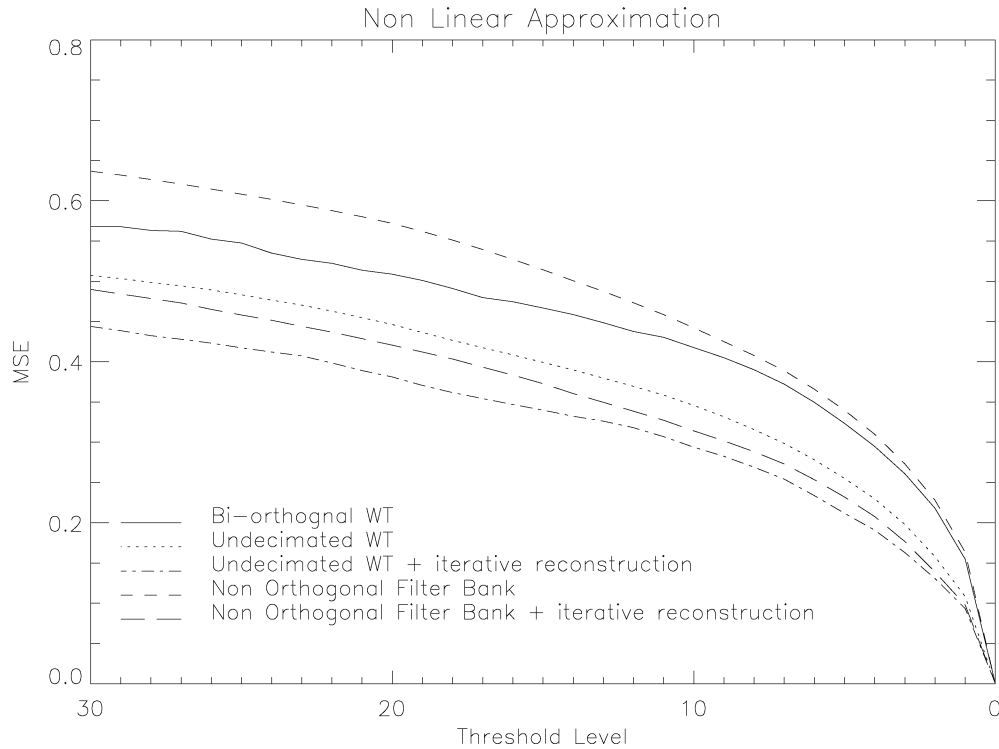


Fig. 6. Nonlinear approximation: Mean Square Error versus the threshold level with a bi-orthogonal DWT, the UWT (7/9 filters), the UWT using the nonorthogonal filter banks and an iterative reconstruction, and the UWT (7/9 filters) with an iterative reconstruction.




Fig. 7. Left: Truncated Gaussian. Right: One row of the left image.

From these curves, we conclude the following.

- Without iterative reconstruction, the nonorthogonal filter bank produces very poor results, even worse than the decimated bi-orthogonal WT.
- Using an iterative reconstruction, the nonorthogonal WT works much better, and the results are now better than the UWT.
- The UWT with the 7/9 filters can also be improved using an iterative reconstruction. This method yields the lowest approximation error among its competitors.

Therefore, these experiments confirm that when a redundant transform is used, there is real interest in using an iterative re-

construction technique when we want to reconstruct an image from a subset of its coefficients. This is also in agreement with recent findings in [37].

2) *Truncated Gaussian Image*: Fig. 7 left shows the truncated Gaussian image (a piecewise smooth function with a sharp transition). The Gaussian standard deviation is 25, and it is normalized to have a maximum intensity equal to 1. Fig. 7 right shows one row of the image. The nonlinear approximation curve is plotted in Fig. 8.

We see that this time the best results are not obtained using the 7/9 filter bank, but by using the nonorthogonal filter bank. This means that the best filter bank is data-dependent, but whatever

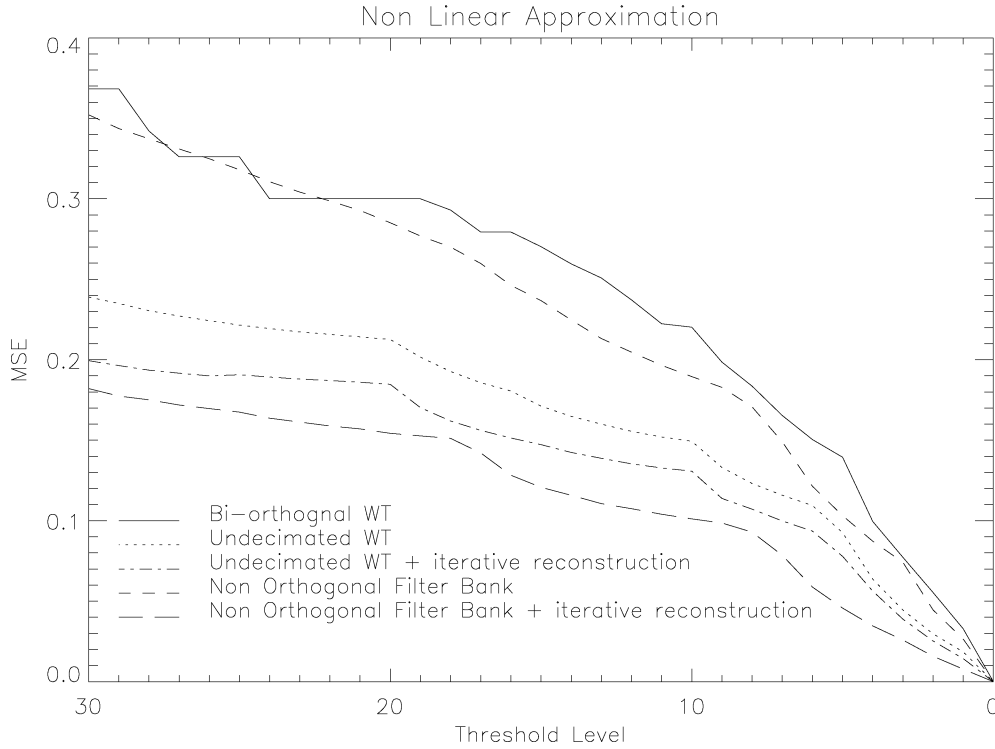


Fig. 8. Nonlinear approximation for the truncated Gaussian image.

the filter bank chosen, and especially if it is not (bi-)orthogonal, iterating improves clearly the nonlinear approximation curve.

B. Ringing Artifact

Fig. 9 shows one row of the reconstructed truncated Gaussian image after thresholding in the wavelet domain with a threshold equal to 2.5. Plots on the right correspond to reconstructions using a positivity constraint. We show from top to bottom, the undecimated decomposition with direct reconstruction (7/9 filters), the undecimated decomposition with iterative reconstruction (7/9 filters), the undecimated decomposition with iterative reconstruction (with third order polynomial spline Battle–Lemarié filters [38], [3]) and the nonorthogonal filter bank. Here, a FIR version of Battle–Lemarié filters was implemented. From this experiment, we conclude the following.

- If the positivity constraint can be used in a given application, it will help a lot for the reduction of the ringing effect along discontinuities.
- Iterating without positivity reduces the oscillations which are not close to the discontinuity, but amplify those very close to it.
- The nonorthogonal filter bank with positivity and an iterative reconstruction produces impressively good results for the reconstruction of the truncated Gaussian. This is certainly related to the fact that the scaling function is very close to a Gaussian, and explains why such a scaling function is so popular in the astronomical domain.

C. Edge Detection

Fig. 10, top, shows a simulated image containing a square (pixel value = 10) and some Gaussian noise

(standard deviation = 3) and the detected edges using a Canny detector (the standard deviation of the Gaussian kernel is 3). The noisy image has been filtered using the iterative denoising procedure described previously, using the UWT with both the 7/9 filters and the nonorthogonal filters described in Section III-A (i.e., $h^{1D} = \tilde{h}^{1D} = [1, 4, 6, 4, 1]/16$, $g^{1D} = \delta - h^{1D}$, $\tilde{g}^{1D} = \delta + h^{1D}$). A simple pixel-difference edge detector has been applied on both denoised images. We can see that the latter leads to less spurious detected edges than the 7/9-filter wavelet denoising.

D. MCA

The Morphological Component Analysis method (MCA) [10], [39], [40] is a method which allows us to decompose a single signal into two or more layers, each layer containing only one kind of feature in the input signal. The separation can be achieved when each kind of feature is well represented by a given transformation. For instance, line and Gaussian in a image can be separated using the ridgelet transform (which represents lines well) and the wavelet transform [10] (for the Gaussians), or the texture can be separated from the piecewise smooth content using the local DCT and the curvelet transform [10]. A full description of MCA is given in [10].

We have applied MCA on a 1-D signal containing a sine, three bumps and some Gaussian noise. The sine are well represented by the DCT and the bumps by the wavelet transform. Denoting respectively by \mathbf{W} and \mathbf{C} the matrices related to the wavelet and the DCT transforms, the MCA algorithm finds the solution of the following minimization problem:

$$\begin{aligned} \min \quad & \|\mathbf{W}X_b\|_1 + \|\mathbf{C}X_s\|_1 \\ \text{subject to} \quad & \|Y - X_b - X_s\|_2 < \sigma \end{aligned} \quad (29)$$




Fig. 9. One row of the reconstructed truncated Gaussian image after thresholding in the wavelet domain with a threshold equal to 2.5. From top to bottom, panels correspond respectively to the undecimated decomposition with direct reconstruction (7/9 filters), the undecimated decomposition with iterative reconstruction (7/9 filters), the undecimated decomposition with iterative reconstruction (Battle-Lemarié filters), and the nonorthogonal filter bank. Plots on the right correspond to reconstructions with the same decomposition but using a positivity constraint.

where Y is the noisy data, σ is the noise standard deviation and X_b and X_s are the two components to recover, one (i.e., bumps) being sparse in the wavelet representation (\mathbf{W} matrix) and the second (i.e., sine) in the DCT domain (\mathbf{C} matrix).

The MCA algorithm relies on an iterative alternate projection and thresholding scheme. At the k th iteration, we have two estimates $\tilde{X}_b^{(k)}$, $\tilde{X}_s^{(k)}$ of X_b and X_s ($\tilde{X}_b^{(0)} = 0$, $\tilde{X}_s^{(0)} = 0$), and $\tilde{X}_b^{(k+1)}$ (resp., $\tilde{X}_s^{(k+1)}$) is obtained by applying a thresholding of the residual using W (resp., C)

$$\begin{aligned}\tilde{X}_b^{(k+1)} &= \mathbf{F}\mathbf{W}_{,\lambda_k} \left(Y - \tilde{X}_s^{(k)} \right) \\ \tilde{X}_s^{(k+1)} &= \mathbf{F}\mathbf{C}_{,\lambda_k} \left(Y - \tilde{X}_b^{(k+1)} \right)\end{aligned}\quad (30)$$

where $\mathbf{F}\mathbf{T}_{,\lambda}(x)$ consists in decomposing x using the transform \mathbf{T} ($\alpha = \mathbf{T}x$), threshold the obtained coefficients α with the threshold λ ($\tilde{\alpha} = \Delta_\lambda(\alpha)$), and reconstruct \tilde{x} from $\tilde{\alpha}$. The thresholding operator Δ_λ can be either a hard or a soft thresholding. In practice, hard thresholding leads generally to better results. The threshold λ_k decreases linearly toward zero, starting from a first

threshold set to a large enough value. At each iteration the positivity of \tilde{X}_b can be enforced by replacing negative values with zero. More details can be found in [10], [39]. Fig. 11 shows the result. With or without positivity, we can see that the nonorthogonal filters produce a solution with less ringing.

VI. CONCLUSION

We have shown in this paper that reconstruction from undecimated wavelet transform coefficients can be addressed in a very different way compared to the usual one. The nondecimation gives us additional freedom for designing filters. As a result, we have seen that nonnegative reconstruction filters can be used or that regular reconstruction can be obtained from Haar wavelet coefficients. Finally, we have shown that the concept of reconstruction from partial information in the case of undecimated decompositions is different from (bi)orthogonal decomposition. Therefore, we confirm the results presented in recent papers that the multiscale denoising can be improved when an iterative approach is performed. Furthermore, an additional constraint such as the TV or l_1 norm of the wavelet coefficients can easily be incorporated within the iterative scheme.




Fig. 10. Top: Noisy image containing a square and some noise and the detected edges using the Canny detector. Bottom left: Pixel-difference detected edges on the UWT denoising image using the 7/9 filters. Bottom right: The same processing but using the nonorthogonal filters.




Fig. 11. Top: From left to right, input signal contains three bumps, a sine and Gaussian noise; the three bumps; and the bumps recovered by MCA without positivity constraint using the DCT and UWT (Battle-Lemarié filters). Bottom left: Bumps recovered by MCA without positivity constraint using the DCT and the nonorthogonal filters. Bottom middle: MCA recovered bumps using the DCT and UWT Battle-Lemarié transform plus the positivity constraint. Bottom right: MCA recovered bumps using the DCT and the nonorthogonal filters plus positivity constraint.

Our work opens up also new questions: Which properties should the analyzing filters and the synthesis filters verify for good image restoration, for a given application?

The iterative methods for the inversion can certainly also be improved by using an additional constraint similar to what is

used in inverse problem methods. This will be investigated in future work.

ACKNOWLEDGMENT

The authors would like to thank B. Zhang, E. Candès, and D. Donoho for useful discussions.

REFERENCES

- [1] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA: SIAM, 1992.
- [2] G. Strang and T. Nguyen, *Wavelet and Filter Banks*. Cambridge, U.K.: Wellesley-Cambridge, 1996.
- [3] S. Mallat, *A Wavelet Tour of Signal Processing*. New York: Academic, 1998.
- [4] J.-L. Starck, F. Murtagh, and A. Bijaoui, *Image Processing and Data Analysis: The Multiscale Approach*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [5] A. Cohen, *Numerical Analysis of Wavelet Methods*. New York: Elsevier, 2003.
- [6] J. P. Antoine and R. Murenzi, "Two dimensional wavelet analysis in image processing," *Phys. Mag.*, vol. 16, pp. 105–134, 1994.
- [7] A. Arneodo, F. Argoul, E. Bacry, J. Elezgaray, and J. F. Muzy, *Ondelettes, Multifractales et Turbulences*. Paris, France: Diderot, Arts et Sciences, 1995.
- [8] P. Dutilleul, "An implementation of the "algorithme à trous" to compute the wavelet transform," in *Wavelets: Time-Frequency Methods and Phase-Space*, J. M. Combes, A. Grossmann, and P. Tchamitchian, Eds. New York: Springer, 1989.
- [9] M. Holschneider, R. Kronland-Martinet, J. Morlet, and P. Tchamitchian, "A real-time algorithm for signal analysis with the help of the wavelet transform," in *Wavelets: Time-Frequency Methods and Phase-Space*. New York: Springer-Verlag, 1989, pp. 286–297.
- [10] J.-L. Starck, M. Elad, and D. L. Donoho, "Redundant multiscale transforms and their application for morphological component analysis," *Adv. Imag. Electron Phys.*, vol. 132, 2004.
- [11] J.-L. Starck and F. Murtagh, *Astronomical Image and Data Analysis*. New York: Springer-Verlag, 2002.
- [12] S. Durand and J. Froment, "Reconstruction of wavelet coefficients using total variation minimization," *SIAM J. Sci. Comput.*, vol. 24, no. 5, pp. 1754–1767, 2003.
- [13] F. Malgouyres, "Minimizing the total variation under a general convex constraint for image restoration," *IEEE Trans. Image Process.*, vol. 11, no. 12, pp. 1450–1456, Dec. 2002.
- [14] E. J. Candès and F. Guo, "New multiscale transforms, minimum total variation synthesis: Applications to edge-preserving image reconstruction," *Signal Process.*, vol. 82, no. 11, pp. 1519–1543, 2002.
- [15] J.-L. Starck, D. L. Donoho, and E. J. Candès, "Very high quality image restoration by combining wavelets and curvelets," *Proc. SPIE*, vol. 4478, pp. 9–19, 2001.
- [16] Z. Cvetkovic and M. Vetterli, "Oversampled filter banks," *IEEE Trans. Signal Process.*, vol. 46, no. 5, pp. 1245–1255, May 1998.
- [17] H. Bolcskei, F. Hlawatsch, and H. G. Feichtinger, "Frame-theoretic analysis of oversampled filter banks," *IEEE Trans. Signal Process.*, vol. 46, no. 12, pp. 3256–3268, Dec. 1998.
- [18] V. K. Goyal, M. Vetterli, and N. T. Thao, "Quantized overcomplete expansions in \mathbf{R}^n : Analysis, synthesis, and algorithms," *IEEE Trans. Inf. Theory*, vol. 44, no. 1, pp. 16–31, Jan. 1998.
- [19] Z. Cvetkovic and M. Vetterli, "Discrete-time wavelet extrema representation: Design and consistent reconstruction," *IEEE Trans. Image Process.*, vol. 43, no. 3, pp. 1245–1255, Mar. 1995.
- [20] M. J. Shensa, "Discrete wavelet transforms: Wedding the à trous and Mallat algorithms," *IEEE Trans. Signal Process.*, vol. 40, no. 10, pp. 2464–2482, Oct. 1992.
- [21] A. Genovesio and J.-C. Olivo-Marin, "Tracking fluorescent spots in biological video microscopy," in *Three-Dimensional and Multidimensional Microscopy: Image Acquisition and Processing X*, J.-A. Conchello, C. J. Cogswell, and T. Wilson, Eds. Bellingham, WA: SPIE, 2003, vol. 4964, pp. 98–105.
- [22] P. J. Burt and A. E. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.*, vol. COM-31, no. 4, pp. 532–540, Apr. 1983.
- [23] P. Fryzlewicz and G. P. Nason, "A Haar-Fisz algorithm for Poisson intensity estimation," *J. Comput. Graph. Statist.*, vol. 13, pp. 621–638, 2004.
- [24] E. D. Kolaczyk, "Nonparametric estimation of intensity maps using Haar wavelets and Poisson noise characteristics," *Astrophys. J.*, vol. 534, pp. 490–505, 2000.
- [25] —, "Wavelet shrinkage estimation of certain Poisson intensity signals using corrected thresholds," *Statist. Sinica*, vol. 9, pp. 119–135, 1999.
- [26] R. D. Nowak and R. G. Baraniuk, "Wavelet-domain filtering for photon imaging systems," *IEEE Trans. Image Process.*, vol. 8, no. 5, pp. 666–678, May 1999.
- [27] A. Antoniadis and T. Sapatinas, "Wavelet shrinkage for natural exponential families with quadratic variance functions," *Biometrika*, vol. 88, pp. 805–820, 2001.
- [28] K. E. Timmermann and R. D. Nowak, "Multiscale modeling and estimation of Poisson processes with application to photon-limited imaging," *IEEE Trans. Inf. Theory*, vol. 45, no. 3, pp. 846–862, Apr. 1999.
- [29] G. Jammal and A. Bijaoui, "Dequant: A flexible multiresolution restoration framework," *Signal Process.*, vol. 84, no. 7, pp. 1049–1069, 2004.
- [30] G. Steidl, J. Weickert, T. Brox, P. Mrázek, and M. Welk, "On the Equivalence of Soft Wavelet Shrinkage, Total Variation Diffusion, Total Variation Regularization, and Sides," Tech. Rep. 26, Dept. Math., Univ. Bremen, Bremen, Germany, 2003.
- [31] M. Unser, "Texture classification and segmentation using wavelet frames," *IEEE Trans. Image Process.*, vol. 4, no. 11, pp. 1549–1560, Nov. 1995.
- [32] J.-L. Starck, A. Bijaoui, and F. Murtagh, "Multiresolution support applied to image filtering and deconvolution," *CVGIP: Graph. Models Image Process.*, vol. 57, pp. 420–431, 1995.
- [33] I. Halperin, "The product of projection operators," *Acta Sci. Math.*, vol. 23, pp. 96–99, Mar. 1962.
- [34] J. A. Tropp, I. Dhillon, R. W. Heath, and T. Strohmer, "Designing structured tight frames via an alternating projection method," *IEEE Trans. Inf. Theory*, vol. 51, no. 1, pp. 188–209, Jan. 2005.
- [35] D. C. Youla, "Mathematical theory of image restoration by the method of convex projections," in *Image Recovery: Theory and Application*, H. Starck, Ed. New York: Academic, 1987.
- [36] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image coding using wavelet transform," *IEEE Trans. Image Process.*, vol. 1, no. 2, pp. 205–220, Apr. 1992.
- [37] M. Elad, "Why Simple Shrinkage is Still Relevant for Redundant Representations," to be published, 2006.
- [38] G. Battle, "A block spin construction of ondelettes. Part I: Lemarié functions," *Comm. Math. Phys.*, vol. 110, pp. 601–615, 1987.
- [39] J.-L. Starck, M. Elad, and D. Donoho, "Image decomposition via the combination of sparse representation and a variational approach," *IEEE Trans. Image Process.*, vol. 14, no. 10, pp. 1570–1582, Oct. 2005.
- [40] M. Elad, J.-L. Starck, D. Donoho, and P. Querre, "Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA)," *J. Appl. Comput. Harmon. Anal.*, vol. 19, pp. 340–358, 2006.



Jean-Luc Starck received the Ph.D. degree from the University Nice-Sophia Antipolis, France, and the Habilitation degree from the University Paris XI, Paris, France.

He was a Visitor at the European Southern Observatory (ESO) in 1993; at the University of California, Los Angeles, in 2004; and at the Statistics Department, Stanford University, Stanford, CA, in 2000 and 2005. He has been a Researcher at CEA, Gif sur Yvette, France, since 1994. His research interests include image processing, statistical methods in astrophysics, and cosmology. He is an expert in multiscale methods such as wavelets and curvelets. He is the Leader of the project Multiresolution at CEA, and he is a core team member of the PLANCK ESA project. He has published more than 200 papers in different areas in scientific journals and conference proceedings. He is also the author of two books entitled *Image Processing and Data Analysis: The Multiscale Approach* (Cambridge Univ. Press, 1998) and *Astronomical Image and Data Analysis* (2nd Ed., Springer, 2006).



Jalal Fadili received the degree from the Ecole Nationale Supérieure d'Ingénieurs (ENSI) de Caen, Caen, France, and the M.Sc. and Ph.D. degrees in signal and image processing from the University of Caen.

He was a Research Associate at the University of Cambridge, Cambridge, U.K., in 2000. He has been an Assistant Professor of signal and image processing at ENSI since September 2001. His research interests include statistical estimation and detection and multiscale methods in signal and image processing. His areas of application include medical imaging.



Fionn Murtagh received the B.A. and BAI degrees in mathematics and engineering science and the M.Sc. in computer science from Trinity College, Dublin, Ireland, the Ph.D. degree in mathematical statistics from the Université P. & M. Curie Paris VI, Paris, France, and the Habilitation degree from the Université L. Pasteur, Strasbourg, France.

His previous posts have included Senior Scientist with the Space Science Department of the European Space Agency; Professor of computer science at Queen's University, Belfast, U.K.; and visiting appointments with the European Commission's Joint Research Centre, and the Department of Statistics, University of Washington, Seattle. He is currently a Professor of computer science and Head of the Computer Science Department, University of London, London, U.K.

Dr. Murtagh is the Editor-in-Chief of *The Computer Journal* and a Fellow of the British Computer Society.