

The Pairing of a Wavelet Basis With a Mildly Redundant Analysis via Subband Regression

Michael Unser, *Fellow, IEEE*, and Dimitri Van De Ville, *Member, IEEE*

Abstract—A distinction is usually made between wavelet bases and wavelet frames. The former are associated with a one-to-one representation of signals, which is somewhat constrained but most efficient computationally. The latter are over-complete, but they offer advantages in terms of flexibility (shape of the basis functions) and shift-invariance. In this paper, we propose a framework for improved wavelet analysis based on an appropriate pairing of a wavelet basis with a mildly redundant version of itself (frame). The processing is accomplished in four steps: 1) redundant wavelet analysis, 2) wavelet-domain processing, 3) projection of the results onto the wavelet basis, and 4) reconstruction of the signal from its nonredundant wavelet expansion. The wavelet analysis is pyramid-like and is obtained by simple modification of Mallat’s filterbank algorithm (e.g., suppression of the down-sampling in the wavelet channels only). The key component of the method is the subband regression filter (Step 3) which computes a wavelet expansion that is maximally consistent in the least squares sense with the redundant wavelet analysis. We demonstrate that this approach significantly improves the performance of soft-threshold wavelet denoising with a moderate increase in computational cost. We also show that the analysis filters in the proposed framework can be adjusted for improved feature detection; in particular, a new quincunx Mexican-hat-like wavelet transform that is fully reversible and essentially behaves the $(\gamma/2)$ th Laplacian of a Gaussian.

Index Terms—Denoising, feature detection, fractals, frames, isotropy, Mexican-hat filter, pyramid, wavelets.

I. INTRODUCTION

DURING the past 15 years, wavelets have had a profound impact on signal and image processing. They have led to the development of a variety of new algorithms for signal compression (JPEG2000) [1], [2], denoising (soft-thresholding and its multiple variants) [3]–[6], enhancement, and restoration (e.g., deconvolution) [7], [8], as well as signal analysis (e.g., fractals, singularities, texture) [9], [10]. From the practitioner point of view, there are two categories of wavelet transforms: nonredundant and redundant ones.

The first type (one-to-one mapping) is necessarily associated with a wavelet basis that can be defined either in the continuous domain—i.e., for the space $L_2(\mathbb{R}^d)$ of d -dimensional finite energy signals [11]—, or in the discrete domain by considering the

space of square summable sequences $\ell_2(\mathbb{Z}^d)$ [12]. Wavelet bases are elegant mathematically, as well as most efficient computationally, thanks to Mallat’s fast filterbank algorithm. They can be designed to be orthogonal, which is advantageous for optimal bit allocation (thanks to Parseval’s identity), and also for statistical purposes because white noise remains white in the transformed domain. A key property that justifies their success is that they yield sparse representations of signals that are piecewise smooth. These features are beneficial for data compression, and for designing simple, efficient wavelet-domain algorithms such as soft-threshold denoising. The only downside of wavelet bases is their lack of shift-invariance because the underlying sampling is critical.

The second type of redundant transformation is associated with wavelet frames. The simplest versions are redundant variants of a decomposition into wavelet bases; they go under the name of cycle spinning (where one iterates the transform on a sequence of translated input signals) [13], the “à trous” wavelets which correspond to the undecimated version of Mallat’s algorithm [14]–[16], or discrete wavelet frames [17]. There are also other less redundant transforms such as Kingsbury’s complex dual tree decomposition which combines two wavelet bases that are in quadrature [18], [19], and Simoncelli’s steerable pyramid [20]. Presently, the consensus among researchers is that redundant wavelets are better suited for signal and image analysis, and that they also yield better results for tasks such as denoising. While it is straightforward to obtain a redundant decomposition from a basis (using cycle spinning or the undecimated version of Mallat’s algorithm), this approach tends to be very heavy in term of memory usage and computation because the resolution of the wavelet subband remains the same irrespective of the scale of the analysis. This can be a serious limitation, especially in two or higher dimensions. If one analyzes the situation in term of sampling density, there is an obvious waste in conciseness as one moves to coarser scales because the bandwidth of wavelet coefficients gets narrower as a result of the convolution with the dilated analysis wavelet.

Our purpose in this paper is to present an effective way of getting rid of the less informative part of this redundancy without compromising perfect reconstruction. Specifically, we are interested in maintaining some controlled level of redundancy within the wavelet subbands to improve shift-invariance, while retaining the sub-sampling (pyramid structure) as we move to coarser resolution to avoid explosion in data size. Yang and Jing have recently proposed a practical solution for getting closer to that goal by simply replacing the original sampling matrix by a more progressive one (e.g., quincunx instead of dyadic) which can be viewed as an incomplete version of cycle spinning [21]. The present approach is based on a different principle which does not require any factorization of the dilation matrix. It is also

Manuscript received January 9, 2008; revised July 4, 2008. Current version published October 10, 2008. This work was supported in part by the Center for Biomedical Imaging (CIBM) of the Geneva and Lausanne Universities and the EPFL, in part by the foundations Leenaards and Louis-Jeantet, and in part by the Swiss National Science Foundation under Grant 200020–109415. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Srdjan Stankovic.

The authors are with the Biomedical Imaging Group (BIG), École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland (e-mail: michael.unser@epfl.ch; dimitri.vandeville@epfl.ch).

Digital Object Identifier 10.1109/TIP.2008.2004607

more flexible from a wavelet designer point of view. Our proposal is a two-step strategy that tries to make the best use of both wavelet worlds: 1) frames for a (mildly) redundant analysis and for data processing, and 2) a wavelet basis for the representation and efficient reconstruction of the output signal. The conversion between the two representations is performed by a regression module that orthogonally projects the redundant wavelet coefficients onto the corresponding wavelet subspace at a given resolution.

The paper is organized as follows. In Section II, we first develop the concept in 1-D, which leads to a modification of the standard two-channel filterbank algorithm where the high-pass analysis branch (wavelet coefficients) is undecimated. We then generalize the system to multiple bands in Section III for the processing of higher dimensional signals. We also describe an extension of the method that allows using modified analysis wavelets while preserving the perfect reconstruction property. Finally, in Section IV, we validate the method by showing that it leads to some performance improvement for image denoising. We also illustrate the reshaping capabilities of the technique by presenting a new family of Mexican-hat-like wavelet transforms with improved isotropy properties.

II. ONE-DIMENSIONAL CASE

Following a brief review of the constituents of a nonredundant wavelet decomposition (basis), we show how the standard filterbank decomposition algorithm can be modified for a redundant analysis in the high-pass channel only.

A. Wavelet Basis

The cornerstone of Mallat’s multiresolution theory of the wavelet transform is the existence of a scaling function $\varphi(x) \in L_2(\mathbb{R})$. This function is often specified indirectly through the scaling relation

$$\varphi(x/2) = \sqrt{2} \sum_{k \in \mathbb{Z}} h[k] \varphi(x - k). \quad (1)$$

where h is a given (refinement) filter. The associated approximation spaces are

$$\mathcal{V}_i = \text{span}_{k \in \mathbb{Z}} \{ \sqrt{2^i} \varphi(2^i x - k) \}. \quad (2)$$

These spaces are embedded and form a valid multiresolution analysis of $L_2(\mathbb{R})$ under suitable technical conditions (i.e., existence of lower and upper Riesz bounds, and fulfillment of the partition of unity). This guarantees the existence of some corresponding wavelet basis of $L_2(\mathbb{R})$ —not necessarily unique.

In the so-called semi-orthogonal design, the wavelet is constructed such that it generates a Riesz basis of the orthogonal complement $\mathcal{W}_0 = \mathcal{V}_0 \ominus \mathcal{V}_{-1}$. In particular, this implies that $\langle \psi(\cdot - k) \varphi(\cdot) \rangle = 0$ and that the wavelet can be written as a linear combination of functions from the finer approximation space

$$\psi(x/2) = \sqrt{2} \sum_{k \in \mathbb{Z}} g[k] \varphi(x - k) \quad (3)$$

where g is the wavelet filter. The corresponding wavelet spaces are

$$\mathcal{W}_i = \text{span}_{k \in \mathbb{Z}} \{ \sqrt{2^i} \psi(2^i x - k) \}. \quad (4)$$




Fig. 1. Analysis-synthesis filterbank for the 1-D nonredundant wavelet decomposition.

The scaling function φ and wavelet ψ completely determine the synthesis part of the algorithm (wavelet expansion). They corresponds to a unique pair of dual (or biorthogonal) functions $\tilde{\varphi}$ and $\tilde{\psi}$ with associated scaling and wavelet filters \tilde{h} and \tilde{g} , respectively. The wavelet expansion of a function $f \in L_2(\mathbb{R})$ down to resolution J is then given by

$$f(x) = \sum_{k \in \mathbb{Z}} \overbrace{\langle f, \tilde{\varphi}_{J,k} \rangle}^{c_J[k]} \varphi_{J,k} + \sum_{i=J}^{\infty} \sum_{k \in \mathbb{Z}} \overbrace{\langle f, \tilde{\psi}_{i,k} \rangle}^{d_i[k]} \psi_{i,k}$$

with the short-hand notation $\psi_{i,k} = 2^{i/2} \psi(2^i x - k)$. In practice, the scaling and wavelet coefficients— $c_i[k]$, and $d_i[k]$, respectively—are computed iteratively by repeated application of Mallat’s algorithm. The corresponding perfect reconstruction filterbank is shown in Fig. 1. Note that the sampling is critical: decimation by a factor of two in both lowpass and wavelet channels.

B. Pyramid-Like Wavelet Decomposition

The idea here is not to decimate the wavelet subband in order to obtain a finer signal analysis (redundancy factor of two). The motivation is to improve shift-invariance which is advantageous for signal processing purposes (better detection and/or suppression of artifacts).

Fig. 2 shows the modified filterbank. The nondecimated wavelet coefficients are denoted as d'_{i-1} with $d'_{i-1}[k] = \langle f, \sqrt{2^i} \tilde{\psi}(2^i x - \frac{1}{2}k) \rangle$. At the synthesis side, the module “R” contains the algorithm that should estimate the coefficients \tilde{d}_{i-1} that are used for a nonredundant reconstruction.

Clearly, the reconstruction algorithm in this new configuration is no longer unique. For instance, we could simply pick $\tilde{d}_{i-1}[k] = d'_{i-1}[2k]$ which brings us back to the nonredundant case. This is fine in the absence of processing noise or quantization artifacts, but does not present any advantage otherwise.

In a practical system, the wavelet coefficients d'_{i-1} are often processed (module “P”) to yield some modified or “filtered” coefficients d''_{i-1} . We would, therefore, like to estimate the wavelet coefficients \tilde{d}_{i-1} to be injected into the nonredundant reconstruction algorithm such that the resulting signal is most consistent with those measurements.

To that end, we propose a “subband regression” that minimizes the quadratic error

$$\arg \min_{\tilde{d}_{i-1}} \left\| \sum_{n \in \mathbb{Z}} \left[\sum_{m \in \mathbb{Z}} \tilde{d}_{i-1}[m] g[n - 2m] + \sum_{m \in \mathbb{Z}} c_{i-1}[m] h[n - 2m] \right] \tilde{g}[n - \cdot] - d''_{i-1}[\cdot] \right\|_{l_2}^2$$




Fig. 2. Analysis-synthesis filterbank for the 1-D pyramid wavelet decomposition. Processing in the wavelet domain is done in the module “P.” The module “R” contains the subband regression algorithm.

which measures the difference between the (redundant) wavelet analysis of the reconstructed signal and the current (processed) wavelet coefficients d''_{i-1} .

To determine the solution, we first rewrite the criterion as

$$\left\| \sum_{m \in \mathbb{Z}} \tilde{d}_{i-1}[m] g_e[k - 2m] - \left(d''_{i-1}[\cdot] - \sum_{n \in \mathbb{Z}} c_{i-1}[n] h_e[\cdot - 2n] \right) \right\|_{l_2}^2$$

where we have introduced the composite filters

$$g_e = g * \tilde{g}^\triangleleft \quad (5)$$

$$h_e = h * \tilde{g}^\triangleleft \quad (6)$$

with the notation \tilde{g}^\triangleleft referring to the time-reversed filter $\tilde{g}[-k]$ whose z -transform is $\tilde{G}(z^{-1})$. Next, we identify the auxiliary residual signal

$$e[k] = d''_{i-1}[k] - \sum_{n \in \mathbb{Z}} c_{i-1}[n] h_e[k - 2n] \quad (7)$$

that can be readily computed from the known subband coefficients. The optimization problem then reads

$$\arg \min_{\tilde{d}_{i-1}} \left\| \sum_{m \in \mathbb{Z}} \tilde{d}_{i-1}[m] g_e[\cdot - 2m] - e[\cdot] \right\|_{l_2}^2 \quad (8)$$

which boils down to finding the least-squares approximation of $e[k]$ in $\text{span}_{m \in \mathbb{Z}} \{g_e[k - 2m]\}$. A standard argument (cf. [22]) then yields

$$\tilde{d}_{i-1}[m] = (g_{LS} * e)[2m] \quad (9)$$

where g_{LS} is the l_2 -dual of g_e

$$G_{LS}(z) = \frac{G_e(z^{-1})}{\frac{1}{2}(G_e(z)G_e(z^{-1}) + G_e(-z)G_e(-z^{-1}))}. \quad (10)$$

Note that this result can also be obtained as a special case of the multichannel solution which is derived in the next section. The solution is equivalent to digital filtering of the auxiliary signal e followed by down-sampling by a factor two. The complete subband regression algorithm is summarized in Fig. 3. This system will recover the original (nonredundant) wavelet coefficients of the signal exactly if the (redundant) analysis coefficients are fed into the algorithm unaltered; e.g., $d''_{i-1} = d_{i-1} \Rightarrow \tilde{d}_{i-1} = d_{i-1}$. Otherwise, it will lead to the reconstruction of the signal that is maximally consistent with the desired wavelet coefficient values.




Fig. 3. Overview of the 1-D subband regression algorithm (module “R” of Fig. 2), followed by subsampling.

III. MULTIPLE DIMENSIONS AND CHANNELS

We will now extend the above regression scheme to the case of multichannel wavelets. In doing so, we will adopt a vector formulation that is also valid for multidimensional signals.

A. Subband Regression: General Case

In the general d -dimensional case, the redundancy factor depends upon the subsampling ratio—i.e., $R = |\det(\mathbf{M})|$ where \mathbf{M} is the $d \times d$ subsampling matrix—and the number N of non-subsampled channels of the filterbank with the constraint that $N \leq R - 1$. It is given by

$$\frac{N}{1 - \frac{1}{R}} = \frac{RN}{R - 1}. \quad (11)$$

We achieve a redundancy factor of R by selecting $N = R - 1$ which is the case that will be treated here.

Let us denote the lowpass filter by $h[\mathbf{k}]$, $\mathbf{k} \in \mathbb{Z}^d$, and the N highpass filters on the analysis and synthesis side by

$$\tilde{\mathbf{g}}[\mathbf{k}] = \begin{pmatrix} \tilde{g}^{(1)}[\mathbf{k}] \\ \vdots \\ \tilde{g}^{(N)}[\mathbf{k}] \end{pmatrix}, \quad \mathbf{g}[\mathbf{k}] = \begin{pmatrix} g^{(1)}[\mathbf{k}] \\ \vdots \\ g^{(N)}[\mathbf{k}] \end{pmatrix} \quad (12)$$

respectively. Together with $\tilde{h}[\mathbf{k}]$, these specify a R -channel multidimensional perfect reconstruction filterbank analogous to the one displayed in Fig. 1.

We now consider the corresponding modified perfect reconstruction system whose schematic representation is shown in Fig. 4. The pooled effect of the wavelet filters is represented by the two $N \times N$ multidimensional transfer matrices $\tilde{\mathbf{G}}(\mathbf{z}^{-1})$ (analysis) and $\mathbf{G}(\mathbf{z})$ (synthesis), while the undecimated, processed, and decimated wavelet coefficients are collected in the vector signals $\mathbf{d}'_{i-1}[\mathbf{k}]$, $\mathbf{d}''_{i-1}[\mathbf{k}]$ and $\tilde{\mathbf{d}}_{i-1}[\mathbf{k}]$, respectively. Our task is to specify the estimation module “R” such that the critically-sampled coefficient sequence $\tilde{\mathbf{d}}_{i-1}[\mathbf{k}]$ provides the minimum-error approximation of the undecimated, and possibly altered, coefficients $\mathbf{d}''_{i-1}[\mathbf{k}]$.




Fig. 4. Analysis-synthesis filterbank for the R -channel pyramid wavelet decomposition. Processing in the wavelet domain is done in the module “P.” The module “R” contains the subband regression algorithm.

Our consistency measure for this general multidimensional setup is

$$\left\| \sum_{\mathbf{n} \in \mathbb{Z}^d} \left[\sum_{\mathbf{m} \in \mathbb{Z}^d} \mathbf{g}^T[\mathbf{n} - \mathbf{M}\mathbf{m}] \tilde{\mathbf{d}}_{i-1}[\mathbf{m}] + \sum_{\mathbf{m} \in \mathbb{Z}^d} c_{i-1}[\mathbf{m}] h[\mathbf{n} - \mathbf{M}\mathbf{m}] \right] \tilde{\mathbf{g}}[\mathbf{n} - \mathbf{k}] - \mathbf{d}''_{i-1}[\mathbf{k}] \right\|_{\mathbb{R}^N}^2 \Big\|_{l_2}^2$$

where the inner squared vector-norm sums the error contributions across the wavelet channels, and the outer $l_2(\mathbb{Z}^d)$ -norm performs the summation over the multidimensional signal domain (index \mathbf{k}). To simplify this expression, we introduce the $(N \times N)$ -matrix and N -vector filters \mathbf{g}_e and \mathbf{h}_e , respectively, whose entries are given by

$$\mathbf{g}_e^{(n_1, n_2)} = g^{(n_1)} * \tilde{g}^{(n_2) \triangleleft} \quad (13)$$

$$\mathbf{h}_e^{(n)} = h * \tilde{g}^{(n) \triangleleft} \quad (14)$$

with $n, n_1, n_2 = 1, \dots, N$ and $\tilde{g}^{(n) \triangleleft}[\mathbf{k}] \triangleq \tilde{g}^{(n)}[-\mathbf{k}]$. We then restate the error criterion as

$$\left\| \sum_{\mathbf{m} \in \mathbb{Z}^d} \mathbf{g}_e^T[\mathbf{k} - \mathbf{M}\mathbf{m}] \tilde{\mathbf{d}}_{i-1}[\mathbf{m}] - \mathbf{e}[\mathbf{k}] \right\|_{\mathbb{R}^N}^2 \Big\|_{l_2}^2 \quad (15)$$

with

$$\mathbf{e}[\mathbf{k}] = \mathbf{d}''_{i-1}[\mathbf{k}] - \sum_{\mathbf{n} \in \mathbb{Z}^d} \mathbf{h}_e[\mathbf{k} - \mathbf{M}\mathbf{n}] c_{i-1}[\mathbf{n}]. \quad (16)$$

As shown in the appendix, the minimization of this criterion leads to the set of normal equations

$$\sum_{\mathbf{m}' \in \mathbb{Z}^d} (\mathbf{g}_e^{\triangleleft} * \mathbf{g}_e^T)[\mathbf{M}(\mathbf{m} - \mathbf{m}')] \tilde{\mathbf{d}}_{i-1}[\mathbf{m}'] = (\mathbf{g}_e^{\triangleleft} * \mathbf{e})[\mathbf{M}\mathbf{m}]$$

for any $\mathbf{m} \in \mathbb{Z}^d$. Note that the left-hand side can be interpreted as the convolution between $\tilde{\mathbf{d}}_{i-1}[\mathbf{m}]$ and the \mathbf{M} -subsampled version of the “auto-correlation” matrix filter $(\mathbf{g}_e^{\triangleleft} * \mathbf{g}_e^T)[\mathbf{k}]$. This system of equations can be solved by inverse filtering,

$$\tilde{\mathbf{d}}_{i-1}[\mathbf{m}] = (\mathbf{g}_{\text{LS}} * \mathbf{e})[\mathbf{M}\mathbf{m}] \quad (17)$$




Fig. 5. Overview of the R -channel subband regression algorithm (module “R” of Fig. 4), followed by subsampling.

which is the direct matrix counterpart of (9); it involves a convolution with a $N \times N$ matrix filter \mathbf{g}_{LS} followed by a downsampling by \mathbf{M} . The individual channel components are, therefore, given by

$$\tilde{d}_{i-1}^{(n)}[\mathbf{m}] = \sum_{n'=1}^N (g_{\text{LS}}^{(n, n')} * e^{(n')})[\mathbf{M}\mathbf{m}] \quad (18)$$

where the convolutions are now scalar. The matrix filter \mathbf{g}_{LS} is specified by its $N \times N$ transfer function

$$\mathbf{G}_{\text{LS}}(\mathbf{z}) = R \left(\sum_{c=1}^R \mathbf{G}_e \left(\mathbf{E}_M^{(c)} \mathbf{z}^{-1} \right) \mathbf{G}_e^T \left(\mathbf{E}_M^{(c)} \mathbf{z} \right) \right)^{-1} \times \mathbf{G}_e(\mathbf{z}^{-1}) \quad (19)$$

where $\mathbf{z} = (z_1, \dots, z_d)$ is the d -dimensional complex z -transform-domain variable with the notational convention $\mathbf{z}^{-1} = (z_1^{-1}, \dots, z_d^{-1})$; the $\mathbf{E}_M^{(c)}$'s are $d \times d$ diagonal coset matrices which are given by $\mathbf{E}_M^{(c)} = \text{diag}(e^{j2\pi \mathbf{M}^{-T} \mathbf{l}^{(c)}})$ where $\mathbf{l}^{(c)}$ is the c -th coset generated by \mathbf{M} [23]. For example, in the case of the dyadic subsampling matrix, we have

$$\mathbf{l}^{(0)} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \mathbf{l}^{(1)} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \mathbf{l}^{(2)} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \mathbf{l}^{(3)} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

In effect, the term in parenthesis corresponds to the \mathbf{M} -downsampling-upsampling of the “auto-correlation” matrix filter $(\mathbf{g}_e^{\triangleleft} * \mathbf{g}_e^T)[\mathbf{k}]$ so that this operator matches the finer rate of the input data, while the matrix-inverse symbol “ $(\cdot)^{-1}$ ” implements the inverse filtering operation that is required to solve the normal equations.

The complete subband regression algorithm can thus be implemented using a block diagram that is the vector/matrix counterpart of the previous one, as shown Figs. 4 and 5. Note that this formulation lends itself well to an implementation in the Fourier domain: the (matrix) frequency response of \mathbf{g}_{LS} can be computed directly from (19), while the filtering is implemented by simple matrix multiplication with the N -vector Fourier transform of $\mathbf{e}[\mathbf{k}]$. If one includes the processing of the lowpass




Fig. 6. Wavelet decomposition of the test image “cameraman” for the orthogonal polyharmonic B-spline wavelet of order $\gamma = 4$ and quincunx subsampling matrix. (a) Original test image. (b) Nonredundant transform (compressed display). (c) Nonredundant transform where the wavelet subbands of odd decomposition levels are rotated by 45 degrees to maintain the geometrical organization of the quincunx lattice. (d) Pyramid transform as proposed in this paper. The redundancy factor with respect to (b) and (c) is $R = 2$.

branch, the subband regression can also be viewed as an $R = N + 1$ channel filterbank. Globally, this means that the cost of signal reconstruction is of the same order as that of the analysis part of the algorithm. Thus, one should expect an overall increase in computation over the nonredundant scheme that is roughly equivalent to the redundancy factor R .

B. General Two-Channel Case

The advantage of the two channel case ($R = |\det(\mathbf{M})| = 2$) is that there is only one wavelet, as opposed to $2^d - 1$ for the more conventional separable schemes which result from the application of a 1-D algorithm along the various dimensions of the data [24]. Specific nonseparable wavelet designs have been proposed for the quincunx subsampling matrix (2-D) and the FCC subsampling matrix (3-D) [25]. The most frequently used quincunx subsampling matrix is

$$\mathbf{M} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (20)$$

with the property $\mathbf{M}^2 = 2\mathbf{I}$.

For $R = 2$, we necessarily have $N = 1$ in our pyramid scheme, so that the filters \mathbf{g}_e and \mathbf{h}_e are scalar. The two corresponding coset matrices are $\mathbf{E}_M^{(1)} = \mathbf{I}$ and $\mathbf{E}_M^{(2)} = -\mathbf{I}$ where \mathbf{I} is the $d \times d$ identity matrix. The least-squares filter (19) then simplifies to

$$G_{LS}(\mathbf{z}) = \frac{G_e(\mathbf{z}^{-1})}{\frac{1}{2}(G_e(\mathbf{z})G_e(\mathbf{z}^{-1}) + G_e(-\mathbf{z})G_e(-\mathbf{z}^{-1}))} \quad (21)$$

which is essentially the same formula as in the 1-D case.

An example of such a wavelet image analysis is illustrated in Fig. 6, for the quincunx wavelet transform that is associated with

an orthogonal polyharmonic spline of order 4. The nonredundant transform can be depicted in two ways. Either the wavelet subbands of the odd decomposition levels can be squeezed together to fit into an image of the same size as the original one, as shown in Fig. 6(a), or those subbands can be displayed using a pixel arrangement that preserves relative spatial relationships (rotation by 45 degrees for odd decomposition levels), as shown in Fig. 6(b). The pyramid counterpart with a redundancy of two (no subsampling in the wavelet channels) is shown in Fig. 6(c). This corresponds to the analysis part of the algorithm. When the subband regression procedure (module “ R ”) is applied to the unaltered wavelet pyramid, it yields back the original nonredundant wavelet transform (consistent reconstruction). However, if the wavelet coefficients in the pyramid are modified through processing (e.g., pointwise nonlinearity), then the output of the regression module is different from the original nonredundant wavelet transform. The final image reconstruction is obtained by (nonredundant) inverse wavelet transformation and it is such that it is maximally consistent with the redundant wavelet processing.

C. Optional Wavelet Reshaping

The basic idea behind wavelet reshaping is that we can take further advantage of the frame redundancy and the generic structure of the analysis-synthesis algorithm in Fig. 4 to modify the analysis functions such that they are better suited for the application at hand. For instance, it may be desirable to optimize the wavelets to improve their space-domain localization or to make them more isotropic. Alternatively, one may also be interested in adapting their shape for better feature extraction.




Fig. 7. Analysis-synthesis filterbank for a wavelet decomposition in the 2-channel quincunx case. The shaping filter $Q'(\mathbf{z})$ allows to take advantage of the redundancy to better match a desired template; for instance, a Mexican-hat wavelet, as discussed in Section IV-B.

We now briefly describe the concept for the general two-channel case (cf. Subsection III.B) which is simpler to formulate because there is a single wavelet involved. In that setting, the original analysis wavelet space at resolution i is

$$\tilde{\mathcal{W}}_i = \text{span}_{\mathbf{k} \in \mathbb{Z}^d} \{ |\det(\mathbf{M})|^{i/d} \tilde{\psi}(\mathbf{M}^i \mathbf{x} - \mathbf{k}) \} \quad (22)$$

where $\tilde{\psi}$ is the biorthogonal analysis wavelet associated with the given wavelet basis. The analysis wavelet is usually expressed in terms of the analysis scaling function

$$\tilde{\psi}(\mathbf{M}^{-1} \mathbf{x}) = |\det(\mathbf{M})|^{1/d} \sum_{\mathbf{k} \in \mathbb{Z}^d} g[\mathbf{k}] \tilde{\varphi}(\mathbf{x} - \mathbf{k}) \quad (23)$$

where $\tilde{\varphi}$ is biorthogonal to φ and where $g[\mathbf{k}]$ corresponds to the wavelet analysis filter in the perfect reconstruct filterbank (cf. lower branch in Fig. 1).

Similarly, the modified wavelet system described by the block diagram in Fig. 4 yields a redundant signal description that is associated with the ‘‘augmented’’ analysis space

$$\tilde{\mathcal{W}}_i^+ = \text{span}_{\mathbf{k} \in \mathbb{Z}^d} \{ |\det(\mathbf{M})|^{i/d} \tilde{\psi}(\mathbf{M}^i \mathbf{x} - \mathbf{M}^{-1} \mathbf{k}) \} \quad (24)$$

because of the suppression of the down-sampling module in the highpass wavelet channel. Clearly, we have that $\tilde{\mathcal{W}}_i^+ \supset \tilde{\mathcal{W}}_i$, but one loses the orthogonality with the scaling space \mathcal{V}_i because of the inclusion of the additional ‘‘in-between’’ wavelets. On the other hand, it still holds that $\tilde{\mathcal{W}}_i^+ \subset \tilde{\mathcal{V}}_{i+1}$ where the latter is the approximation space (analysis) at the next finer resolution.

When working with a wavelet basis of $L_2(\mathbb{R}^d)$, the choice of the filter $g[\mathbf{k}]$ in (23) is quite constrained and often fixed *a priori*. If we now consider the enlarged space $\tilde{\mathcal{W}}_i^+$, we have many more degrees of freedom which makes it also possible to choose some other equivalent wavelet of the form

$$\tilde{\psi}'(\mathbf{M}^{-1} \mathbf{x}) = |\det(\mathbf{M})|^{1/d} \sum_{\mathbf{k} \in \mathbb{Z}^d} g'[\mathbf{k}] \tilde{\varphi}(\mathbf{x} - \mathbf{k}). \quad (25)$$

The only constraint is that $G'(e^{j\boldsymbol{\omega}}) = S(e^{j\boldsymbol{\omega}})G(e^{j\boldsymbol{\omega}})$ where $S(e^{j\boldsymbol{\omega}})$ is the frequency response of some reversible digital filter—i.e., $a_0 \leq |S(e^{j\boldsymbol{\omega}})| \leq A_0$ where a_0 and A_0 are some positive constants—so that ψ' is an equivalent generator of $\tilde{\mathcal{W}}_i^+$. The wavelet decomposition and subband regression algorithms can then be readily adapted by simply replacing $G(\mathbf{z})$ by $G'(\mathbf{z})$ in all relevant filter equations (cf. Fig. 7). The reconstruction will remain perfect because we are still spanning the same analysis and synthesis spaces.

While this modification of the wavelet filter seems rather innocuous, we will see that it can have a dramatic effect on the shape of the wavelet and on its feature extraction capabilities. In particular, we will illustrate the concept with the design of wavelets that essentially replicate a Mexican-hat-like analysis. This wavelet reshaping method is also transposable to the multichannel case where the degrees of freedom are even more numerous.

IV. APPLICATION EXAMPLES

The motivating idea for performing processing in a redundant wavelet domain is to ensure better shift-invariance. In practice, this tends to result in a finer detection of relevant image features, a better preservation of details, and a reduction of reconstruction artifacts. We will now investigate this effect for two basic image processing tasks for which wavelets have been found to be particularly successful: image denoising and multiscale feature extraction. Ideally, we would like to improve the performance of a given nonredundant wavelet-domain algorithm by only making use of a minimum of redundancy such as to keep the computational cost and memory requirement at a reasonable level.

A. Denoising

To compare the performance of the various transforms, we added white Gaussian noise to a series of test images and applied a standard denoising algorithm which consists in applying a soft-thresholding in the wavelet domain [3]. In order to have a fair comparison, we used a fixed soft-threshold across all subband that was optimized in each case for maximum signal-to-noise ratio (Oracle solution). We kept the denoising procedure simple on purpose to minimize the number of parameters and to focus on the comparison of transforms with various levels of redundancy: the standard nonredundant orthogonal wavelet transform, its pyramid extension obtained using the proposed subband regression, and a fully redundant scheme without any subsampling in the wavelet channels. We tested a combination of separable orthogonal dyadic transforms (Haar and cubic-spline Battle–Lemarié [11]) together with their nonseparable quincunx counterparts (orthogonal polyharmonic splines of order $\gamma = 2, 4$) [26]. We applied a comparable number of decomposition levels in both instances: $J = 4$ for the standard dyadic configuration and $J = 8$ for the quincunx case where the scale increases as a power of $\sqrt{2}$. The results are summarized in Table I

TABLE I
PSNR IMPROVEMENTS FOR DENOISING USING SOFT-THRESHOLDING WITH THE ORACLE SOLUTION.
DECOMPOSITION LEVELS: $J = 4$ (DYADIC) AND $J = 8$ (QUINCUNX)

		dyadic Haar wavelets			quincunx polyharmonic wavelets ($\gamma = 2$)		
	σ	non-redundant	pyramid	fully redundant	non-redundant	pyramid	fully redundant
cameraman	15.00	6.51	8.08	8.05	6.94	7.45	7.83
barbara	15.00	6.71	8.10	8.09	7.19	7.65	7.93
lena	15.00	6.79	8.60	8.52	7.47	8.11	8.45
einstein	15.00	9.21	10.51	10.62	9.67	10.15	10.53
peppers	15.00	8.84	11.40	11.23	10.04	10.80	11.25
cameraman	25.00	2.98	4.45	4.40	3.41	3.85	4.18
barbara	25.00	2.53	3.62	3.60	2.79	3.24	3.41
lena	25.00	2.81	4.32	4.20	3.32	3.84	4.14
einstein	25.00	4.06	5.43	5.41	4.34	4.81	5.17
peppers	25.00	4.26	6.63	6.40	5.38	6.07	6.47
redundancy		1	4	13	1	2	9
		dyadic cubic-spline wavelets			quincunx polyharmonic wavelets ($\gamma = 4$)		
cameraman	15.00	7.52	8.10	8.40	7.08	7.39	7.91
barbara	15.00	7.93	8.48	8.68	7.47	7.81	8.15
lena	15.00	8.11	8.73	8.96	7.78	8.11	8.56
einstein	15.00	10.04	10.57	10.87	9.92	10.17	10.65
peppers	15.00	10.74	11.29	11.65	10.32	10.61	11.22
cameraman	25.00	3.81	4.41	4.64	3.37	3.68	4.12
barbara	25.00	3.93	4.57	4.74	3.17	3.56	3.84
lena	25.00	3.87	4.46	4.67	3.45	3.76	4.18
einstein	25.00	4.78	5.41	5.66	4.46	4.82	5.27
peppers	25.00	5.87	6.53	6.83	5.50	5.84	6.42
redundancy		1	4	13	1	2	9

using the peak signal-to-noise ratio (in dB) as the measure of denoising quality.

According to expectations, the higher order most-redundant decomposition (cubic splines, fully redundant) gives the best results. The pyramid decomposition is competitive in most cases; in particular, it is always significantly better than the nonredundant decomposition (orthogonal wavelet basis). In the case of a quincunx transform where the redundancy is the least, it yields results that are approximately half-way in between the nonredundant ($R = 1$) and fully-redundant cases ($R = 9$) on the dB scale. The improvement is more striking in the dyadic case where the performance is very close to the fully redundant transform, despite the substantial reduction in computational and storage cost ($R = 4$ as compared to $R = 13$). The most remarkable results are those obtained in the Haar case where the denoising performance of the pyramid transform is as good—if not better—as the fully redundant system. Some representative examples of restoration are shown in Fig. 8. The results obtained using the nonredundant transform [cf. Fig. 8(b)] exhibit characteristic artifacts—these are the footprints of the underlying basis functions which are not shift-invariant. These artifacts are much attenuated in the more redundant schemes, which also yield better results in terms of the PSNR—the visual quality in both cases is quite comparable.

A finding that comes as a surprise is that the quincunx scheme (including the fully-redundant version) is not all that good; it is outperformed by both Haar and cubic spline pyramids in all but one experiment (“Einstein” with $\sigma = 15$). The observed findings are quite robust and reproducible; they hold across all noise conditions and for all test images to which we have applied the algorithm.

Note that the basic soft-thresholding algorithm that we have applied here is somewhat primitive when compared to state-of-the-art in wavelet denoising [6], [27]. Yet, it should still serve as a valid benchmark for the comparison of wavelet decompositions since the running conditions are the same in all cases and the value of the threshold is optimized for best performance (oracle solution). The main point that we want to make is that working with a fully redundant wavelet system is probably overkill, and that one should be able to obtain similar performance with a less redundant pyramid-like system.

While one could object that the oracle solution is not accessible in a practical situation where the noise-free signal is unknown, it turns out that the thresholding level can still be accurately optimized from the noisy data when the signal is corrupted by white Gaussian noise of known variance σ^2 . The statistical tool that makes this possible is Stein’s unbiased risk estimate (SURE) whose use for wavelet denoising was pioneered by Donoho in the case of an orthogonal transform [28]. Recently, it was shown how to extend the SURE technique to more general denoising schemes, including the ones considered here [29]. Concretely, this means that we could have obtained virtually identical denoising results using a purely data-driven approach which in the present case would require a few Monte-Carlo testing iterations (i.e., running the whole denoising algorithm several times on simulated data).

Based on these experiments, one may, therefore, conclude that the proposed pyramid algorithm stands as an attractive alternative to the fully redundant scheme (“à trous” or undecimated wavelet transform) because of its greater computational efficiency; in the present tests, it was about three to four times more efficient. While the speed-up increases linearly with the number




Fig. 8. Wavelet denoising by optimized soft-thresholding with various degrees of redundancy. (a) Input: 256×256 “cameraman” image corrupted by additive white Gaussian noise ($\sigma = 25$). (b) Denoising using the nonredundant Haar transform ($R = 1$) with $J = 4$ levels of decomposition; PSNR improvement 2.98 dB. (c) Denoising using the proposed Haar pyramid ($R = 4$); PSNR improvement 4.45 dB. (d) Denoising using the fully-redundant Haar transform ($R = 13$); PSNR improvement 4.40 dB.

of levels, it becomes even more significant when moving to higher dimensions (e.g., 2-D + time, or 3-D) which is often the rule in biomedical imaging applications. In fact, a fully-redundant denoising is often simply not feasible in higher dimensions because of the intrinsic limit on memory and computational resources—just think of the task of denoising a $512 \times 512 \times 512$ fluorescence micrograph which is a typical size for modern confocal microscopes, not to mention time-lapse imaging which would add another dimension [30].

B. Laplacian-Like Wavelet Analysis With Improved Isotropy

Several authors have applied wavelets to the detection of spots or blob-like objects in images [31], [32]. Specific ex-

amples of biomedical applications include the detection of microcalcifications in mammograms [33], [34] and the localization/tracking of GFP-labelled proteins in fluorescence micrographs [35], [36].

In essence, a wavelet transform performs a correlation analysis so that we can expect its output to be maximum when the input signal most resembles the analysis template $\tilde{\psi}(\frac{\mathbf{x}-\mathbf{x}_0}{a})$ at scale a and location \mathbf{x}_0 . Interestingly, we can show that the optimal detector for the type of biomedical signals mentioned above is a Laplacian-like wavelet. To that end, we consider the measurement model: $f(\mathbf{x}) = \varphi_a(\mathbf{x} - \mathbf{x}_0) + n(\mathbf{x})$ where $\varphi_a(\mathbf{x}) = \varphi(\mathbf{x}/a)$ is a known deterministic signal at scale a , \mathbf{x}_0 is the unknown location of the target, and $n(\mathbf{x})$ is additive Gaussian noise. It is well known from detection

theory that the optimal detector (prewhitened matched filter) corresponds to a correlation with the template $L^*L\varphi_a(\mathbf{x})$ where L is the whitening operator of the noise and where L^* denotes its adjoint. We can further show that this detector has the structure of a wavelet transform if and only if the whitening operator L commutes with dilations, which narrows down the possibilities to some fractional iterate of the Laplacian $L^*L = (-\Delta)^{\gamma/2} \xrightarrow{\mathcal{F}} \|\boldsymbol{\omega}\|^\gamma$ [37]. The stochastic processes that are whitened by $(-\Delta)^{\gamma/2}$ are the fractional Brownian fields—also commonly referred to as $1/\omega$ noises in the engineering literature [38], [39]. These processes are isotropic, but they have the particularity of not being stationary (their power spectrum is not defined in the conventional sense). Yet, they are frequently used to model natural images which tend to exhibit $1/\omega^\gamma$ -like spectral decay. Even more relevant to the present task is the fact such inverse-power-law spectral behavior (which is characteristic of fractals) has been observed in a variety of biomedical modalities including mammography [40] and fluorescence microscopy [41]. This points to an “optimal” wavelet detector of the form

$$\tilde{\psi}_\gamma(\mathbf{x}) = (-\Delta)^{\gamma/2}\varphi(\mathbf{x}) \quad (26)$$

where as $\varphi(\mathbf{x})$ is an isotropic smoothing function (ideally, a Gaussian for best space-frequency localization). In practice, the shape of the smoothing function should be matched to the type of object to detect (e.g., Gaussian spot or microcalcification) and the fractional order γ adjusted such that it essentially whitens the background signal in the image. Note that one recovers the classical Mexican hat filter for $\gamma = 2$.

To our knowledge, the only wavelet bases that exhibit a true Laplacian-like behavior are the polyharmonic spline wavelets which we have fully characterized for the quincunx subsampling matrix in our earlier work [26]. Note that quincunx subsampling is especially attractive for the present detection scenario because it provides a progressive scale transition—i.e., a reduction by a factor of $\sqrt{2}$ as opposed to 2 for a more traditional dyadic analysis. While the polyharmonic spline wavelets are compatible with (26) where $\gamma > 1$ is even allowed to be fractional, their downside is that the corresponding smoothing function is not truly isotropic nor Gaussian. In fact, as γ increases, the smoothing kernel for the orthogonal case converges to a sinc function which is clearly nonisotropic.

We will now show how to reshape these wavelets to make them perform as desired. The relevant functions for our purpose are the “improved” polyharmonic B-splines of order γ whose Fourier-domain definition is (cf. [26])

$$\hat{\beta}_\gamma(\boldsymbol{\omega}) = \frac{V(e^{j\boldsymbol{\omega}})^{\gamma/2}}{\|\boldsymbol{\omega}\|^\gamma} \quad (27)$$

with

$$V(e^{j\boldsymbol{\omega}}) = \left(\frac{8}{3}(\sin^2(\omega_1/2) + \sin^2(\omega_2/2)) + \frac{2}{3} \left(\sin^2\left(\frac{\omega_1 + \omega_2}{2}\right) + \sin^2\left(\frac{\omega_1 - \omega_2}{2}\right) \right) \right).$$

Note that $V(e^{j\boldsymbol{\omega}})$ is the transfer function of a discrete “most isotropic” version of the 2-D Laplace operator. Its Laplace-like behavior near the origin (e.g., $V(e^{j\boldsymbol{\omega}}) = \|\boldsymbol{\omega}\|^2 + \mathcal{O}(\|\boldsymbol{\omega}\|^4)$)

is crucial to neutralize the singularity of the denominator; it is the key to obtaining a B-spline function that is well localized in space (in particular, $\beta_\gamma(\mathbf{x}) \in L_2(\mathbb{R}^2)$) and essentially isotropic [26].

The polyharmonic B-spline is a valid scaling function for any real $\gamma > 1$; its refinement and autocorrelation filters are denoted by $H_\gamma(e^{j\boldsymbol{\omega}})$ and $A_\gamma(e^{j\boldsymbol{\omega}}) = \sum_{\mathbf{n} \in \mathbb{Z}^d} |\hat{\beta}_\gamma(\boldsymbol{\omega} + 2\pi\mathbf{n})|^2$, respectively. It is, therefore, possible to build multidimensional wavelet bases by suitable linear combination, as in (23) with $\varphi = \beta_\gamma$. Specifically, the corresponding generic Fourier-domain formula for quincunx polyharmonic spline wavelets is (cf. [26])

$$\begin{aligned} \hat{\psi}_\gamma(\mathbf{M}^T\boldsymbol{\omega}) &= 2^{-\frac{1}{2}} \underbrace{e^{j\omega_1} H_\gamma(-e^{j\boldsymbol{\omega}}) A_\gamma(-e^{j\boldsymbol{\omega}}) P(e^{j\mathbf{M}^T\boldsymbol{\omega}})}_{\propto Q(e^{j\boldsymbol{\omega}})} \hat{\beta}_\gamma(\boldsymbol{\omega}) \end{aligned} \quad (28)$$

where β_γ is the scaling function and where P is a filter (bounded from above and away from zero) that contains the degrees of freedom for a semi-orthogonal basis design. In the orthogonal case, $P(e^{j\boldsymbol{\omega}})$ is uniquely tied to the other quantities (refinement and autocorrelation filters H_γ and A_γ).

In the case of a redundant analysis, we are less constrained and we can replace $Q(e^{j\boldsymbol{\omega}})$ by any filter of our choice, as justified in Subsection III-C. In the present case, we select the more convenient wavelet filter

$$Q'(e^{j\boldsymbol{\omega}}) = V(e^{j\boldsymbol{\omega}})^{\gamma/2} \quad (29)$$

where $V(e^{j\boldsymbol{\omega}})^{\gamma/2}$ is the localization filter in the numerator of (27). Note that this corresponds to the reshaping factor

$$S(e^{j\boldsymbol{\omega}}) = \frac{V(e^{j\boldsymbol{\omega}})^{\gamma/2}}{e^{j\omega_1} H_\gamma(-e^{j\boldsymbol{\omega}}) A_\gamma(-e^{j\boldsymbol{\omega}}) P(e^{j\mathbf{M}^T\boldsymbol{\omega}})} \quad (30)$$

which has the required boundness property. The key observation here is that $V(e^{j\boldsymbol{\omega}})^{\gamma/2}$ and $H_\gamma(-e^{j\boldsymbol{\omega}})$ have the exact same $\mathcal{O}(\|\boldsymbol{\omega}\|^\gamma)$ behavior at the origin so that they essentially cancel each other, while all other terms are bounded from above and away from zero. The Fourier transform of the new analysis wavelet ψ' then simplifies to

$$\begin{aligned} \hat{\psi}'_\gamma(\mathbf{M}^T\boldsymbol{\omega}) &= 2^{-\frac{1}{2}} Q'(e^{j\boldsymbol{\omega}}) \hat{\beta}_\gamma(\boldsymbol{\omega}) \\ &= 2^{-\frac{1}{2}} V(e^{j\boldsymbol{\omega}})^{\gamma/2} \hat{\beta}_\gamma(\boldsymbol{\omega}) \\ &= 2^{-\frac{1}{2}} \|\boldsymbol{\omega}\|^\gamma \hat{\beta}_{2\gamma}(\boldsymbol{\omega}) \end{aligned}$$

which, in the space domain, is equivalent to

$$\psi'_\gamma(\mathbf{M}^{-1}\mathbf{x}) \propto (-\Delta)^{\gamma/2} \beta_{2\gamma}(\mathbf{x}). \quad (31)$$

This last equation is now fully compatible with (26). Specifically, the new wavelet corresponds to the $(\gamma/2)$ th Laplacian of a polyharmonic B-spline of order 2γ , where γ is the order parameter that is adjustable by the user in a continuous fashion. The smoothing function is now Gaussian-like and very nearly isotropic. In fact, we even have the guarantee that it converge to a 2-D isotropic Gaussian with standard deviation $\sqrt{(\gamma/6)}$ as γ increases [26]. Three examples of polyharmonic spline wavelets are shown in Fig. 9. The first two generate wavelet




Fig. 9. Polyharmonic spline wavelets for the quincunx subsampling matrix for order $\gamma = 6$. (a) Orthogonal polyharmonic spline wavelet. (b) Semi-orthogonal B-spline wavelet. (c) Reshaped Mexican-hat analysis wavelet. (d) Gaussian-like smoothing kernel for the wavelet in (c).

bases of $L^2(\mathbb{R}^2)$: wavelet (a) is orthogonal, while wavelet (b) is Gabor-like (modulated Gaussian); they are obviously both less isotropic than the reshaped Laplacian-like wavelet which is shown in Fig. 9(c). The corresponding smoothing filter is shown in Fig. 9(d) and is very much Gaussian-like. In principle, one may improve the isotropy of the third wavelet even further by matching higher-order Laplacian terms in the numerator of the Fourier-domain definition of the polyharmonic B-spline (data not shown, but available from the authors).

A concrete example of a Mexican-hat wavelet analysis of a fluorescence micrograph is displayed in Fig. 10(b). Note how the detection of the fluorescent spots is improved when using the reshaped wavelets instead of the orthogonal ones [cf. Fig. 10(c)].

V. CONCLUSION

We have introduced a general mechanism for improving the invariance properties and the feature extraction capabilities of any multidimensional wavelet basis. We believe that such a fix is essential if one wants to make this type of representation more attractive for analysis and processing purposes. Wavelet bases are elegant mathematically, but they are inherently constrained due to the requirement for a nonredundant representation. In other words, there is not much hope for coming up with an “ideal” wavelet basis that would simultaneously fulfill the following desirable properties for image processing: orthogonality (one-to-one representation), symmetry, good spatial localization, shift-invariance, scale-invariance, rotation-invariance, smoothness, etc. [42]. In fact, designing wavelet bases is essentially an art in finding the best compromise among the above conflicting goals. Yet, the invariance properties of the transform can be improved straightforwardly by adding basis functions, which then yields a wavelet frame. However, this needs to be accomplished parsimoniously, and this is where the technique proposed in this paper comes in very handy. Indeed, one can easily control the redundancy and maintain it within a reasonable range (2 to 4); this is much more difficult to

achieve with alternative schemes such as cycle spinning which are much more demanding computationally and storage-wise. The pyramid-like reduction that is provided by the present approach is advantageous from that perspective and also quite compatible with our present knowledge of the organization of primary visual cortex.

A final aspect that we have explored briefly is the possibility of changing the wavelet filters to improve the quality of the analysis. The possibilities there are almost limitless because there is essentially no constraint on the choice of the reshaping filter $Q'(e^{j\omega})$. We have illustrated the concept with the design of Mexican-hat-like wavelets whose distinguishing property is to be very nearly isotropic. Also, the fact that these wavelets behave like multiscale γ th order fractional derivative operators makes them prime candidates for the analysis of fractal processes. For instance, one can easily estimate the local (isotropic) Hurst exponent (or, equivalently, an average Hölder exponent) of a process based on the decay of the wavelet energy across scale (fitting of a straight line in a log-log plot). A further advantage of the present type of analysis is that it offers a progressive scale transition (a reduction by $\sqrt{2}$ as opposed to the more traditional factor of 2).

APPENDIX

DERIVATION THE SUBBAND REGRESSION EQUATIONS

Let $\mathbf{c}[\mathbf{k}]$ and $\mathbf{d}[\mathbf{k}]$ be two N -vector signals of dimension d . We define the corresponding vector ℓ_2 -inner product

$$\begin{aligned} \langle \mathbf{c}, \mathbf{d} \rangle_{\mathbb{R}^N \times \ell_2} &= \sum_{\mathbf{k} \in \mathbb{Z}^d} \mathbf{c}^T[\mathbf{k}] \mathbf{d}[\mathbf{k}] \\ &= \sum_{n=1}^N \sum_{\mathbf{k} \in \mathbb{Z}^d} c^{(n)}[\mathbf{k}] d^{(n)}[\mathbf{k}] \end{aligned}$$




Fig. 10. (a) Original fluorescence micrograph. (b) Its pyramid decomposition using the mexican-hat like wavelet of order $\gamma = 2$. (c) Its pyramid decomposition using the orthogonal wavelet of order $\gamma = 2$.

where $c^{(n)}[\mathbf{k}]$ denotes the n th component of $\mathbf{c}[\mathbf{k}]$. A basic signal processing operation for such signals is the matrix convolution, which is defined as

$$(\mathbf{H} * \mathbf{c})[\mathbf{k}] = \sum_{\mathbf{k}' \in \mathbb{Z}^d} \mathbf{H}[\mathbf{k} - \mathbf{k}'] \mathbf{d}[\mathbf{k}'].$$

The corresponding adjoint relation is

$$\langle \mathbf{d}, \mathbf{H} * \mathbf{c} \rangle_{\mathbb{R}^N \times \ell_2} = \langle \mathbf{H}^{\triangleleft \mathbf{T}} * \mathbf{d}, \mathbf{c} \rangle_{\mathbb{R}^N \times \ell_2}$$

where $\mathbf{H}^{\triangleleft \mathbf{T}}[\mathbf{k}]$ with $\mathbf{H}^{\triangleleft}[\mathbf{k}] \triangleq \mathbf{H}[-\mathbf{k}]$ is the transposed, space-reversed version of the matrix impulse response \mathbf{H} .

This allows us to write the error criterion (15) as

$$\left\| \sum_{\mathbf{m}' \in \mathbb{Z}^d} \mathbf{g}_e^T[\cdot - \mathbf{M}\mathbf{m}'] \tilde{\mathbf{d}}_{i-1}[\mathbf{m}'] - \mathbf{e}[\cdot] \right\|_{\mathbb{R}^N \times \ell_2}^2.$$

The least-squares approximation corresponds to the orthogonal projection of \mathbf{e} onto the subspace spanned by the vector signals $\{\mathbf{g}_e^{(n)}[\mathbf{k} - \mathbf{M}\mathbf{m}]\}_{\mathbf{m} \in \mathbb{Z}^d}$ with $n = 1, \dots, N$, where $\mathbf{g}_e^{(n)}$ denotes the n th row vector of \mathbf{g}_e . The solution is found by imposing the orthogonality of the approximation error

$$\left\langle \sum_{\mathbf{m}' \in \mathbb{Z}^d} \mathbf{g}_e^T[\cdot - \mathbf{M}\mathbf{m}'] \tilde{\mathbf{d}}_{i-1}[\mathbf{m}'] - \mathbf{e}[\cdot], \mathbf{g}_e^{(n)}[\cdot - \mathbf{M}\mathbf{m}] \right\rangle_{\mathbb{R}^N \times \ell_2} = 0$$

for $\mathbf{m} \in \mathbb{Z}^d$. By expanding and combining these equations for $n = 1, \dots, N$, we get

$$\begin{aligned} \sum_{\mathbf{k} \in \mathbb{Z}^d} \sum_{\mathbf{m}' \in \mathbb{Z}^d} \mathbf{g}_e^{\triangleleft}[\mathbf{M}\mathbf{m} - \mathbf{k}] \mathbf{g}_e^T[\mathbf{k} - \mathbf{M}\mathbf{m}'] \tilde{\mathbf{d}}_{i-1}[\mathbf{m}'] \\ = \sum_{\mathbf{k} \in \mathbb{Z}^d} \mathbf{g}_e^{\triangleleft}[\mathbf{M}\mathbf{m} - \mathbf{k}] \mathbf{e}[\mathbf{k}] \end{aligned}$$

where we have introduced the space-reversed version of the basis function matrix $\mathbf{g}_e^{\triangleleft}[\mathbf{k}] = \mathbf{g}_e^T[-\mathbf{k}]$. We then make the change of variable $\mathbf{k}' = \mathbf{k} - \mathbf{M}\mathbf{m}'$

$$\begin{aligned} \sum_{\mathbf{m}' \in \mathbb{Z}^d} \sum_{\mathbf{k}' \in \mathbb{Z}^d} \mathbf{g}_e^{\triangleleft}[\mathbf{M}(\mathbf{m} - \mathbf{m}') - \mathbf{k}'] \mathbf{g}_e^T[\mathbf{k}'] \tilde{\mathbf{d}}_{i-1}[\mathbf{m}'] \\ = \sum_{\mathbf{k} \in \mathbb{Z}^d} \mathbf{g}_e^{\triangleleft}[\mathbf{M}\mathbf{m} - \mathbf{k}] \mathbf{e}[\mathbf{k}]. \end{aligned}$$

Finally, we identify the inner sums as convolution operations, which yields the canonical form of the normal equations

$$\sum_{\mathbf{m}' \in \mathbb{Z}^d} (\mathbf{g}_e^{\triangleleft} * \mathbf{g}_e^T)[\mathbf{M}(\mathbf{m} - \mathbf{m}')] \tilde{\mathbf{d}}_{i-1}[\mathbf{m}'] = (\mathbf{g}_e^{\triangleleft} * \mathbf{e})[\mathbf{M}\mathbf{m}]$$

for any $\mathbf{m} \in \mathbb{Z}^d$.

ACKNOWLEDGMENT

The authors would like to thank Dr. P. Gönczy for kindly providing the image that was used in Fig. 10.

REFERENCES

- [1] 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.
- [2] C. Christopoulos, A. Skodras, and T. Ebrahimi, "The JPEG2000 still image coding system: An overview," *IEEE Trans. Consum. Electron.*, vol. 16, no. 4, pp. 1103–1127, Apr. 2000.
- [3] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation via wavelet shrinkage," *Biometrika*, vol. 81, pp. 425–455, 1994.
- [4] S. G. Chang, B. Yu, and M. Vetterli, "Spatially adaptive wavelet thresholding with context modeling for image denoising," *IEEE Trans. Image Process.*, vol. 9, no. 9, pp. 1522–1531, Sep. 2000.
- [5] L. Sendur and I. W. Selesnick, "Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency," *IEEE Trans. Signal Process.*, vol. 50, no. 11, pp. 2744–2756, Nov. 2002.
- [6] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of gaussians in the wavelet domain," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.
- [7] M. A. T. Figueiredo and R. D. Nowak, "An EM algorithm for wavelet based image restoration," *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 906–916, Aug. 2003.
- [8] I. Daubechies, M. DeFrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 57, pp. 1413–1457, 2004.
- [9] S. Mallat and W. Hwang, "Singularity detection and processing with wavelets," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 617–643, Feb. 1992.
- [10] G. W. Wornell, "Wavelet-based representations for the $1/f$ family of fractal processes," *Proc. IEEE*, vol. 81, no. 10, pp. 1428–1450, Oct. 1993.
- [11] S. G. Mallat, "A theory of multiresolution signal decomposition: The wavelet representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 7, pp. 674–693, Jul. 1989.
- [12] O. Rioul, "A discrete-time multiresolution theory," *IEEE Trans. Signal Process.*, vol. 41, no. 8, pp. 2591–2606, Aug. 1993.
- [13] R. R. Coifman and D. L. Donoho, *Wavelets and Statistics*. Berlin, Germany: Springer-Verlag, 1995, ch. Translation-Invariant De-noising, pp. 125–150.
- [14] P. Dutilleul, *Wavelets: Time-Frequency Methods and Phase Space*. New York: Springer-Verlag, 1989, ch. An Implementation of the Algorithm à trous to Compute the Wavelet Transform, pp. 298–304.
- [15] M. J. Shensa, "The discrete wavelet transform: Wedding the à trous and Mallat algorithms," *IEEE Trans. Signal Process.*, vol. 40, no. 10, pp. 2464–2482, Oct. 1992.
- [16] M. Lang, H. Guo, J. E. Odegard, C. S. Burrus, and R. O. Wells Jr., "Noise reduction using an undecimated discrete wavelet transform," *IEEE Signal Process. Lett.*, vol. 3, no. 1, pp. 10–12, Jan. 1996.
- [17] M. Unser, "Texture classification and segmentation using wavelet frames," *IEEE Trans. Image Process.*, vol. 4, no. 11, pp. 1549–1560, Nov. 1995.
- [18] N. G. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *J. Appl. Comput. Harmon. Anal.*, vol. 10, no. 3, pp. 234–253, May 2001.
- [19] I. W. Selesnick, R. G. Baraniuk, and N. C. Kingsbury, "The dual-tree complex wavelet transform," *IEEE Signal Process. Mag.*, vol. 22, no. 6, pp. 123–151, Jun. 2005.
- [20] E. P. Simoncelli and W. T. Freeman, "The steerable pyramid: A flexible architecture for multi-scale derivative computation," in *Proc. Int. Conf. Image Processing*, 1995, vol. 3, pp. 444–447.
- [21] B. Yang and Z. Jing, "A simple method to build oversampled filter banks and tight frames," *IEEE Trans. Image Process.*, vol. 16, no. 11, pp. 2682–2687, Nov. 2007.
- [22] M. Unser, "An improved least squares Laplacian pyramid for image compression," *Signal Process.*, vol. 27, no. 2, pp. 187–203, May 1992.
- [23] E. Viscito and J. P. Allebach, "The analysis and design of multidimensional fir perfect reconstruction filter banks for arbitrary sampling lattices," *IEEE Trans. Circuits Syst.*, vol. 38, no. 1, pp. 29–42, Jan. 1991.
- [24] S. Mallat, *A Wavelet Tour of Signal Processing*. San Diego, CA: Academic, 1998.
- [25] J. Kovačević and M. Vetterli, "Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for \mathbb{R}^n ," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 533–555, Mar. 1992.
- [26] D. Van De Ville, T. Blu, and M. Unser, "Isotropic polyharmonic b-splines: Scaling functions and wavelets," *IEEE Trans. Image Process.*, vol. 14, no. 11, pp. 1798–1813, Nov. 2005.
- [27] F. Luisier, T. Blu, and M. Unser, "A new SURE approach to image denoising: Interscale orthonormal wavelet thresholding," *IEEE Trans. Image Process.*, vol. 16, no. 3, pp. 593–606, Mar. 2007.
- [28] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *J. Amer. Statist. Assoc.*, vol. 90, no. 432, pp. 1200–1224, Dec. 1995.
- [29] S. Ramani, T. Blu, and M. Unser, "Monte-Carlo SURE: A black-box optimization of regularization parameters for general denoising algorithm," *IEEE Trans. Image Process.*, vol. 17, no. 9, pp. 1540–1554, Sep. 2008.
- [30] C. Vonesch, F. Aguet, J.-L. Vonesch, and M. Unser, "The colored revolution of bioimaging," *IEEE Signal Process. Mag.*, vol. 23, no. 3, pp. 20–31, May 2006.
- [31] M. Unser and A. Aldroubi, "A review of wavelets in biomedical applications," *Proc. IEEE*, vol. 84, no. 4, pp. 626–638, Apr. 1996.

- [32] R. Strickland and H. I. Hahn, "Wavelet transform methods for object detection and recovery," *IEEE Trans. Image Process.*, vol. 6, no. 5, pp. 724–735, May 1997.
- [33] A. Laine, S. Schuler, J. Fan, and W. Huda, "Mammographic feature enhancement by multiscale analysis," *IEEE Trans. Med. Imag.*, vol. 13, no. 4, pp. 725–740, Apr. 1994.
- [34] R. Strickland and H. I. Hahn, "Wavelet transforms for detecting microcalcifications in mammograms," *IEEE Trans. Med. Imag.*, vol. 15, no. 2, pp. 218–229, Apr. 1996.
- [35] J.-C. Olivo-Marin, "Extraction of spots in biological images using multiscale products," *Pattern Recognit.*, vol. 35, no. 9, pp. 1989–1996, 2002.
- [36] A. Genovesio, T. Liedl, V. Emiliani, W. Parak, M. Coppey-Moisand, and J.-C. Olivo-Marin, "Multiple particle tracking in 3-d+t microscopy: Method and application to the tracking of endocytosed quantum dots," *IEEE Trans. Image Process.*, vol. 15, no. 5, pp. 1062–1070, May 2006.
- [37] J. Kybic, T. Blu, and M. Unser, "Generalized sampling: A variational approach—Part I: Theory," *IEEE Trans. Signal Process.*, vol. 50, no. 8, pp. 1965–1976, Aug. 2002.
- [38] B. Pesquet-Popescu and J. L. Véhel, "Stochastic fractal models for image processing," *IEEE Signal Process. Mag.*, vol. 19, no. 5, pp. 48–62, 2002.
- [39] S. Tirosh, D. Van De Ville, and M. Unser, "Polyharmonic smoothing splines and the multidimensional Wiener filtering of fractal-like signals," *IEEE Trans. Image Process.*, vol. 15, no. 9, pp. 2616–2630, Sep. 2006.
- [40] W. B. Richardson Jr., "Applying wavelets to mammograms," *IEEE Eng. Med. Biol. Mag.*, vol. 14, no. 5, pp. 551–560, May 1995.
- [41] D. Sage, F. Neumann, F. Hediger, S. Gasser, and M. Unser, "Automatic tracking of individual fluorescence particles: Application to the study of chromosome dynamics," *IEEE Trans. Image Process.*, vol. 14, no. 9, pp. 1372–1383, Sep. 2005.
- [42] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multiscale transforms," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 587–607, Feb. 1992.



Michael Unser (M'89–SM'94–F'99) received the M.S. (summa cum laude) and Ph.D. degrees in electrical engineering in 1981 and 1984, respectively, from the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland.

From 1985 to 1997, he was a Scientist with the National Institutes of Health, Bethesda, MD. He is now a Professor and Director of the Biomedical Imaging Group at the EPFL. His main research area is biomedical image processing. He has a strong interest in sampling theories, multiresolution

algorithms, wavelets, and the use of splines for image processing. He has published over 150 journal papers on those topics, and is one of ISI's Highly Cited authors in Engineering (<http://isihighlycited.com>).

Dr. Unser has been actively involved with the IEEE TRANSACTIONS ON MEDICAL IMAGING, holding the positions of Associate Editor (1999, 2002, 2006–present), member of steering committee, and associate Editor-in-Chief (2003–2005). He has acted as Associate Editor or member of the editorial board for eight more international journals, including the IEEE Signal Processing Magazine, the IEEE TRANSACTIONS ON IMAGE PROCESSING (1992–1995), and the IEEE SIGNAL PROCESSING LETTERS (1994–1998). He organized the first IEEE International Symposium on Biomedical Imaging (ISBI2002). He was the founding chair of the technical committee of the IEEE-SP Society on Bio Imaging and Signal Processing (BISP). He is the recipient of three Best Paper Awards from the IEEE Signal Processing Society.



Dimitri Van De Ville (M'02) received the M.S. degree in engineering and computer sciences and the Ph.D. degree from Ghent University, Ghent, Belgium, in 1998 and 2002, respectively.

He obtained a grant as Research Assistant with the Fund for Scientific Research Flanders Belgium (FWO). In 2002, he joined Prof. M. Unser's Biomedical Imaging Group at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. In December 2005, he became responsible for the Signal Processing Antenna at the University Hospital of Geneva, Geneva, Switzerland, as part of the Centre d'Imagerie Biomédicale (CIBM). His research interests include wavelets, statistical analysis, multidimensional splines, and applications in biomedical imaging, such as functional magnetic resonance imaging, spectroscopy, electro-encephalography, and microscopy.

Dr. Van De Ville serves as an Associate Editor for the TRANSACTIONS ON IMAGE PROCESSING (since February 2006) and was previously an Associate Editor for IEEE SIGNAL PROCESSING LETTERS (2004–2006). Since 2003, he has also been Editor and Webmaster of *The Wavelet Digest*. He is co-chair of the Wavelets XII International Conference (August 2007), together with V. Goyal and M. Papadakis.