An Orthogonal Family of Quincunx Wavelets with Continuously-Adjustable Order

Manuela Feilner, Dimitri Van De Ville *Member*, Michael Unser *Fellow*

The authors are with the Biomedical Imaging Group, STI/IOA, Swiss Federal Institute of Technology Lausanne (EPFL), CH-1015 Lausanne, Switzerland

Corresponding author: Michael Unser (Michael.Unser@epfl.ch)

Abstract

We present a new family of 2D and 3D orthogonal wavelets which uses quincunx sampling. The orthogonal refinement filters have a simple analytical expression in the Fourier domain as a function of the order λ, which may be non-integer. We can also prove that they yield wavelet bases of $L_2(\mathbb{R}^2)$ for any $\lambda > 0$. The wavelets are fractional in the sense that the approximation error at a given scale a decays like $O(a^\lambda)$; they also essentially behave like fractional derivative operators. To make our construction practical, we propose an FFT-based implementation that turns out to be surprisingly fast. In fact, our method is almost as efficient as the standard Mallat algorithm for separable wavelets.

Index Terms

Wavelet transform, Quincunx sampling, Non-separable filter design, McClellan transform.

I. INTRODUCTION

THE great majority of wavelet bases that are currently used for image processing are separable. There are two primary reasons for this. The first is convenience because wavelet theory is most developed in 1D and that these results are directly transposable to higher dimensions through the use of tensor product basis functions. The second is efficiency because a separable transform can be implemented by successive 1D processing of the rows and columns of the image. The downside, however, is that separable transforms tend to privilege the vertical and horizontal directions. They also produce a so-called "diagonal" wavelet component, which does not have a straightforward directional interpretation.

Non-separable wavelets, by contrast, offer more freedom and can be better tuned to the characteristics of images [1], [2]. Their less attractive side is that they require more computations. The quincunx wavelets are especially interesting because they can be designed to be nearly isotropic [3]. In contrast with the separable case, there is a single wavelet and the scale reduction is more progressive: a factor $\sqrt{2}$ instead of 2. The preferred technique for designing quincunx wavelets with good isotropy properties is to use the McClellan transform [4] to map 1D biorthogonal designs to the multidimensional case. Since this approach requires the filters to be symmetric, it has only been applied to the biorthogonal case because of the strong incentive to produce filters that are compactly supported [5]–[8]. One noteworthy exception is the work of Nicolier et al. who used the McClellan transform to produce a quincunx version of the Battle-Lemarié wavelet filters [9]. However, we believe that their filters were truncated because they used a representation in terms of Tchebycheff polynomials.

In this paper, we construct a new family of quincunx wavelets that are orthogonal and have a fractional order of approximation. The idea of fractional orders was introduced recently in the context of spline wavelets for extending the family to non-integer degrees [10]. The main advantage of having a continuously-varying order parameter—not just integer steps as in the traditional wavelet families—is flexibility. It allows for a continuous adjustment of the key parameters of the transform; e.g., regularity and localization of the basis functions. The price that we are paying for these new features—orthogonality with symmetry as well as fractional orders—is that the filters can no longer be compactly supported. We will make up for this handicap by proposing a fast FFT-based implementation which is almost as efficient as Mallat's algorithm for separable wavelets [11].

II. QUINCUNX SAMPLING AND FILTERBANKS

First, we recall some basic results on quincunx sampling and perfect reconstruction filterbanks [12]. The quincunx sampling lattice is shown in Fig. 1. Let $x[\vec{k}]$ with $\vec{k} = (k_1, k_2) \in \mathbb{Z}^2$ denote the discrete signal on the initial grid. The 2D Z-transform of $x[\vec{k}]$ is denoted by $X(\vec{z}) = \sum_{\vec{k}\in\mathbb{Z}^2} x[\vec{k}]\vec{z}^{-\vec{k}}$, where $\vec{z}^{\vec{k}} = z_1^{k_1} z_2^{k_2}$. The continuous 2D Fourier transform is then given by $X(e^{j\vec{\omega}}) = \sum_{\vec{k}\in\mathbb{Z}^2} x[\vec{k}]e^{-j\langle\vec{\omega},\vec{k}\rangle}$ with $\vec{\omega} = (\omega_1, \omega_2)$, and finally the discrete 2D Fourier transform for $x[\vec{k}]$ given on an $N \times N$ grid ($k_1, k_2 = 0, 1, \dots, N-1$) by $X[\vec{n}] = \sum_{\vec{k}\in\mathbb{Z}^2} x[\vec{k}]e^{\frac{-j2\pi\langle\vec{n},\vec{k}\rangle}{N}}$, with $n_1, n_2 = 0, 1, \dots, N-1$.

Now we write the quincunx sampled version of $x[\vec{k}]$ as

$$[x]_{\downarrow\mathbf{D}}[\vec{k}] = x[\mathbf{D}\vec{k}], \quad \text{where } \mathbf{D} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \tag{1}$$

Our down-sampling matrix $\mathbf{D}$ is such that $\mathbf{D}^2 = 2\mathbf{I}$. The Fourier-domain version of (1) is

$$[x]_{\downarrow\mathbf{D}}[\vec{k}] \longleftrightarrow \frac{1}{2}\left[X\left(e^{j\mathbf{D}^{-T}\vec{\omega}}\right) + X\left(e^{j(\mathbf{D}^{-T}\vec{\omega}+\vec{\pi})}\right)\right], \tag{2}$$

where $\vec{\pi} = (\pi, \pi)$.

The upsampling is defined by

$$[x]_{\uparrow\mathbf{D}}[\vec{k}] = \begin{cases} x[\mathbf{D^{-1}}\vec{k}] & \text{when } k_1 + k_2 \text{ is even,} \\ 0 & \text{elsewhere,} \end{cases} \tag{3}$$

and its effect in the Fourier domain is as follows:

$$[x]_{\uparrow\mathbf{D}}[\vec{k}] \longleftrightarrow X\left(e^{j\mathbf{D}^T\vec{\omega}}\right). \tag{4}$$

If we now chain the down-sampling and up-sampling operators, we get

$$[x]_{\downarrow\mathbf{D}\uparrow\mathbf{D}}[\vec{k}] = \begin{cases} x[\vec{k}] & \text{when } k_1 + k_2 \text{ is even,} \\ 0 & \text{elsewhere,} \end{cases} \tag{5}$$

$$\updownarrow$$

$$\frac{1}{2}\left[X\left(e^{j\vec{\omega}}\right) + X\left(e^{j(\vec{\omega}+\vec{\pi})}\right)\right]. \tag{6}$$

Since quincunx sampling reduces the number of image samples by a factor of two, the corresponding reconstruction filterbank has two channels (cf. Fig. 2). The low-pass filter $\widetilde{H}$ reduces the resolution by a factor of $\sqrt{2}$; the wavelet coefficients correspond to the output of the high-pass filter $\widetilde{G}$.

Applying the relation (6) to the block diagram in Fig. 2, it is easy to derive the conditions for a perfect reconstruction:

$$\begin{cases} \widetilde{H}\left(\vec{z}\right) H\left(\vec{z}\right) + \widetilde{G}\left(\vec{z}\right) G\left(\vec{z}\right) = 2 \\ \widetilde{H}\left(-\vec{z}\right) H\left(\vec{z}\right) + \widetilde{G}\left(-\vec{z}\right) G\left(\vec{z}\right) = 0, \end{cases} \tag{7}$$

where H and G (respectively $\widetilde{H}$ and $\widetilde{G}$) are the transfer functions of the synthesis (respectively analysis) filters. In the orthogonal case, the analysis and synthesis filters are identical up to a central symmetry; the wavelet filter G is simply a modulated version of the low-pass filter H.

III. FRACTIONAL QUINCUNX FILTERS

To generate quincunx filters, we will use the standard approach which is to apply the diamond McClellan transform to map a one dimensional design onto the quincunx structure.

A. A new 1D wavelet family

As starting point for our construction, we introduce a new one-dimensional family of orthogonal filters:

$$\begin{aligned} H_\lambda(z) &= \frac{\sqrt{2}(z+2+z^{-1})^{\frac{\lambda}{2}}}{\sqrt{(z+2+z^{-1})^\lambda + (-z+2-z^{-1})^\lambda}} \\ &= \frac{\sqrt{2}(2+2\cos\omega)^{\frac{\lambda}{2}}}{\sqrt{(2+2\cos\omega)^\lambda + (2-2\cos\omega)^\lambda}}, \end{aligned} \tag{8}$$

which is indexed by the continuously-varying order parameter λ.

These filters are symmetric and are designed to have zeros of order λ at $z = -1$; the numerator is a fractional power of $(z + 2 + z^{-1})$ (the simplest symmetric refinement filter of order 2) and the denominator is the appropriate l_2-orthonormalization factor. By varying λ, we can adjust the frequency response as shown in Fig. 3. As λ increases, $H_\lambda(z)$ converges to the ideal half-band low-pass filter. Also note that these filters are maximally flat at the origin; they essentially behave like $H_\lambda\left(\omega\right)/\sqrt{2} = 1 + O(\omega^\lambda)$ as $\omega \to 0$. Their frequency response is similar to the Daubechies' filters with two important differences: (1) the filters are symmetric, and (2) the order is not restricted to integer values.

We can prove mathematically that these filters will generate valid 1D fractional wavelet bases of L_2 similar to the fractional splines presented in [10]. The order property (here fractional) is essential because it determines the rate of decay of the approximation error as a function of the scale. It also conditions the behavior of the corresponding wavelet ψ which will act like a fractional derivative of order λ. In other words, it will kill all polynomials of degree $n \leq \lceil \lambda - 1 \rceil$; i.e.,

$$\int x^n \psi_\lambda(x) dx = 0. \tag{9}$$

B. Corresponding 2D wavelet family

Applying the diamond McClellan transform to the filter above is straightforward; it amounts to replacing $\cos\omega$ by $\frac{1}{2}(\cos\omega_1 \ + \cos\omega_2)$ in (8). Thus, our quincunx refinement filter is given by

$$H_\lambda(e^{j\vec{\omega}}) = \frac{\sqrt{2}(2 + \cos\omega_1 + \cos\omega_2)^{\frac{\lambda}{2}}}{\sqrt{(2 + \cos\omega_1 + \cos\omega_2)^\lambda + (2 - \cos\omega_1 - \cos\omega_2)^\lambda}}. \tag{10}$$

This filter is guaranteed to be orthogonal because the McClellan transform has the property of preserving biorthogonality. Also, by construction, the λth order zero at $\omega = \pi$ gets mapped into a corresponding zero at $\vec{\omega} = (\pi, \pi)$; this is precisely the condition that is required to get a two dimensional wavelet transform of order λ. Also note the isotropic behavior and the flatness of $H_\lambda(e^{j\vec{\omega}})$ around the origin; i.e., $H_\lambda(e^{j\vec{\omega}})/\sqrt{2} = 1 + O(||\vec{\omega}||^\lambda)$ for $\vec{\omega} \to 0$. Figure 4 shows contour plots of the scaling filter for several choices of the order λ.

The orthogonal wavelet filter is obtained by modulation

$$G_\lambda(\vec{z}) = z_1 H_\lambda(-\vec{z}^{-1}). \tag{11}$$

The corresponding orthogonal scaling function $\varphi_\lambda(\vec{x})$ is defined implicitly as the solution of the quincunx two-scale relation:

$$\varphi_\lambda(\vec{x}) = \sqrt{2} \sum_{\vec{k} \in \mathbb{Z}^2} h_\lambda[\vec{k}] \varphi_\lambda\left(\mathbf{D}\vec{x} - \vec{k}\right). \tag{12}$$

Since the refinement filter is orthogonal with respect to the quincunx lattice, it follows that $\varphi_\lambda(\vec{x}) \in L_2(\mathbb{R}^2)$ and that it is orthogonal to its integer translates. Moreover, for $\lambda > 0$, it will satisfy the partition of unity condition, which comes as a direct consequence of the vanishing of the filter at $(\omega_1, \omega_2) = (\pi, \pi)$. Thus, we have the guarantee that our scheme will yield orthogonal wavelet bases of $L_2(\mathbb{R}^2)$. The underlying orthogonal quincunx wavelet is simply

$$\psi_\lambda(\vec{x}) = \sqrt{2} \sum_{\vec{k} \in \mathbb{Z}^2} g_\lambda[\vec{k}] \varphi_\lambda\left(\mathbf{D}\vec{x} - \vec{k}\right). \tag{13}$$

Given the behavior of $H_\lambda(e^{j\vec{\omega}})$ at $\vec{\omega} = 0$, we also have $\hat{\psi}_\lambda(\vec{\omega}) \propto ||\vec{\omega}||^\lambda$, and as such the wavelet behaves as the λth order differentiator for low frequencies [13]. The vanishing moment property in the 2D case becomes

$$\int x_1^{n_1} x_2^{n_2} \psi_\lambda(\vec{x}) d\vec{x}, \quad \text{for } n_1 + n_2 \leq \lceil \lambda - 1 \rceil. \tag{14}$$

Figure 5 shows the wavelet $\psi_\lambda(\vec{x})$ for various choices of the order λ. Note that the wavelet is centered around $(\frac{1}{2}, \frac{1}{2})$. As illustrated by these plots, the wavelets clearly gets smoother as λ increases. However, a mathematical rigorous estimation of their regularity is beyond the scope of this paper.

IV. IMPLEMENTATION IN FOURIER DOMAIN

The major objection that can be made to our construction is that the filters are not FIR and that it may be difficult and costly to implement the transform in practice. We will see here that we can turn the situation around and obtain a very simple and efficient algorithm that is based on the FFT, following the idea of [14]. Working in the frequency domain is also very convenient because of the way in which we have specified our filters (see Eqs. (10) and (11)). Implementations of the wavelet transform for the quincunx subsampling matrix using FFTs has been proposed before [9], [15], but our algorithm is another variation, which in particular minimizes the number and size of FFTs and seems to be faster.

First let us assume that the image size is $N \times N$. Now, we will describe the decomposition part of our algorithm which corresponds to the block diagram presented in Fig. 6, where we have pooled together two levels of the decomposition. The initialization step is to evaluate the FFT of the initial input image $x[\vec{k}]$ and to precompute the corresponding sampled frequency responses of the analysis filters $\widetilde{H}[\vec{n}]$ and $\widetilde{G}[\vec{n}]$ using (10) and (11). We also precompute the rotated version of the filters, denoted as $\widetilde{H}_p[\vec{n}]$ and $\widetilde{G}_p[\vec{n}]$, that can be obtained as

$$\begin{aligned} \widetilde{H}_p[\vec{n}] &= \widetilde{H}[\mathbf{D}\vec{n} \bmod (N,N)], & (15)\\ \widetilde{G}_p[\vec{n}] &= \widetilde{G}[\mathbf{D}\vec{n} \bmod (N,N)]. & (16) \end{aligned}$$

Let us now consider the 2D FFT of the input, given by

$$X_i[\vec{n}] = \sum_{\vec{k}} x_i[\vec{k}] e^{-j\frac{2\pi\langle\vec{k},\vec{n}\rangle}{N}}, \quad \text{for } n_1, n_2 = 0, 1, \ldots, N-1. \qquad (17)$$

Globally, at the end of the process, the output variables are the quincunx wavelet coefficients $y_1[\vec{k}]$, $y_2[\vec{k}], \ldots, y_J[\vec{k}]$ and $x_J[\vec{k}]$; e.g., as shown in Fig. 7 (a). Their Fourier transforms for the odd iterations are derived from the auxiliary $N \times N$ signals (see also Fig. 6):

$$\begin{aligned} X'_{i+1}[\vec{n}] &= \sum_{\vec{k}} x'_{i+1}[\vec{k}] e^{-j\frac{2\pi\langle\vec{k},\vec{n}\rangle}{N}}, & (18)\\ Y'_{i+1}[\vec{n}] &= \sum_{\vec{k}} y'_{i+1}[\vec{k}] e^{-j\frac{2\pi\langle\vec{k},\vec{n}\rangle}{N}}. & (19) \end{aligned}$$

Down- and upsampling with $\mathbf{D}$ in the first iteration step introduces zeros in the space domain while it preserves the size of $Y'_{i+1}[\vec{n}]$. However, it implies some symmetry/redundancy in frequency domain.

Therefore, only half of the coefficients needs to be computed which saves operations. The reduced signal $Y'_{i+1}[\vec{k}]$ and its corresponding low-pass signal are obtained by

$$\begin{aligned} Y'_{i+1}\left[\vec{n}'\right] &= \frac{1}{2}\left(\widetilde{G}\left[\vec{n}'\right]X_i\left[\vec{n}'\right] + \widetilde{G}\left[\vec{n}' + \left(\tfrac{N}{2},\tfrac{N}{2}\right)\right]X_i\left[\vec{n}' + \left(\tfrac{N}{2},\tfrac{N}{2}\right)\right]\right), & (20)\\ X'_{i+1}\left[\vec{n}'\right] &= \frac{1}{2}\left(\widetilde{H}\left[\vec{n}'\right]X_i\left[\vec{n}'\right] + \widetilde{H}\left[\vec{n}' + \left(\tfrac{N}{2},\tfrac{N}{2}\right)\right]X_i\left[\vec{n}' + \left(\tfrac{N}{2},\tfrac{N}{2}\right)\right]\right), & (21)\end{aligned}$$

where $\vec{n}' \in [0, \frac{N}{2}-1] \times [0, N-1]$.

To generate the signal $y_{i+1}[\vec{k}]$ of (19) in the way that is depicted in Fig. 7 (a) with every second row shifted by one pixel, we separate the image in even ($y_{i+1,\text{even}}$) and odd ($y_{i+1,\text{odd}}$) rows already in the Fourier domain, using the auxiliary variable $Z[\vec{m}]$:

$$\begin{aligned} Z[\vec{m}] &= Y'_{i+1}[\vec{m}] + Y'_{i+1}\left[\vec{m} + \left(0,\tfrac{N}{2}\right)\right] & \\ &\quad + j\left(Y'_{i+1}[\vec{m}] - Y'_{i+1}\left[\vec{m} + \left(0,\tfrac{N}{2}\right)\right]\right)e^{j\frac{2\pi(m_1+m_2)}{N}} & (22)\\ &\qquad \updownarrow & \\ z_{i+1}[\vec{k}] &= y_{i+1,\text{even}}[\vec{k}] + j y_{i+1,\text{odd}}[\vec{k}], & (23)\end{aligned}$$

with $\vec{m} \in [0, \frac{N}{2}-1]^2$. The sum in the real part ($Y'_{i+1}[\vec{m}] + Y'_{i+1}[\vec{m} + (0,\frac{N}{2})]$) represents downsampling by two in the vertical direction, keeping all the even rows, whereas the sum in the imaginary part represents the odd rows. In the space domain, we alternate the rows $y_{i+1}\left[k_1, 2k_2+1\right] = \text{Re}\left\{z[\vec{k}]\right\}$ and $y_{i+1}\left[k_1, 2k_2\right] = \text{Im}\left\{z[\vec{k}]\right\}$. Since $z[\vec{k}]$ is four times smaller than $y'_{i+1}[\vec{k}]$, we save computations with the reduced-size IFFT.

Instead of rotating the frequency variables after each iteration, we use the precomputed rotated version of the filters (i.e., $\widetilde{H}_p$ and $\widetilde{G}_p$), which we apply at all even iterations. In this way, we also save two rotations per iteration in the frequency domain.

The Fourier transforms of the output for the even iterations are:

$$Y_{i+2}\left[\vec{m}\right] = \sum_{\vec{k}} y_{i+2}[\vec{k}]e^{-j\frac{2\pi\langle\vec{k},\vec{m}\rangle}{\frac{N}{2}}} \quad \text{for } m_1, m_2 = 0, 1, \ldots, \frac{N}{2}-1. \tag{24}$$

They are computed by

$$\begin{aligned} X_{i+2}\left[\vec{m}\right] &= \frac{1}{2}\left(\widetilde{H}_p\left[\vec{m}\right]X'_{i+1}\left[\vec{m}\right] + \widetilde{H}_p\left[\vec{m} + \left(0,\tfrac{N}{2}\right)\right]X'_{i+1}\left[\vec{m} + \left(0,\tfrac{N}{2}\right)\right]\right) & (25)\\ Y_{i+2}\left[\vec{m}\right] &= \frac{1}{2}\left(\widetilde{G}_p\left[\vec{m}\right]X'_{i+1}\left[\vec{m}\right] + \widetilde{G}_p\left[\vec{m} + \left(0,\tfrac{N}{2}\right)\right]X'_{i+1}\left[\vec{m} + \left(0,\tfrac{N}{2}\right)\right]\right). & (26)\end{aligned}$$

The process is then iterated until one reaches the final resolution. When the last iteration is even, we lower the computation costs with the FFT by utilizing its imaginary part:

$$z[\vec{k}] = \sum_{\vec{m}} \left(X_{i+2}[\vec{m}] + jY_{i+2}[\vec{m}]\right)e^{j\frac{2\pi\langle\vec{m},\vec{k}\rangle}{\frac{N}{2}}}, \tag{27}$$

where $x_{i+2}[\vec{k}] = \mathrm{Re}\left\{z[\vec{k}]\right\}$ and $y_{i+2}[\vec{k}] = \mathrm{Im}\left\{z[\vec{k}]\right\}$.

Obviously, as the resolution gets coarser after each iteration, the Fourier transforms of the filters need not be recalculated; they are simply obtained by down-sampling the previous arrays.

The synthesis algorithm operates according to the same principles and corresponds to the flow graph transpose of the decomposition algorithm using up-sampling instead. For instance, the synthesis counterpart of (25) is:

$$\begin{aligned} X'_{i+2}\left[m_1, m_2 + \left(\tfrac{N}{2}\right)\right] &= X_{i+2}\left[m_1, m_2\right], \\ Y'_{i+2}\left[m_1, m_2 + \left(\tfrac{N}{2}\right)\right] &= Y_{i+2}\left[m_1, m_2\right], \\ X_{i+1}\left[m_1, n_2\right] &= X'_{i+2}[m_1, n_2]H_p[m_1, n_2] + Y'_{i+2}[m_1, n_2]G_p[m_1, n_2]. \end{aligned}$$

V. EXPERIMENTS

A. Benchmark and testing

We have implemented two versions of the algorithm, based on Java and Matlab. For the Matlab version, we report computation times below 0.8s for 16 quincunx iterations of a 256×256 image on an Apple G4 700MHz desktop; the decomposition is essentially perfect with a reconstruction error below 10^{-12} RMS. The method is generic and works for any set of filters that can be specified in the frequency domain, independent of their spatial support (or infinite spatial support such as in our case). As a comparison, the Matlab implementation available in the latest Wavelet Toolbox [16] for the Daubechies 9/7 filters (used in JPEG 2000) applied to the same image and for an equivalent of 8 separable iterations, takes about 1.7s. For N datapoints, the complexity of our approach boils down to $\mathcal{O}(N \log N)$ for the FFT-based implementation, versus $\mathcal{O}(NB)$ for the spatial-domain implementation, where B is related to the filter support. The exact trade-off will depend on the image size and the filter size. However, taking into account the benchmark measures and its flexibility, we believe that the FFT-based implementation deserves consideration for a broad class of applications.

We provide also an applet written in Java, which makes it possible to run the algorithm over the Internet, at the site `http://bigwww.epfl.ch/demo/jquincunx/`. A screen-shot of this applet is presented in Fig. 8.

Two examples of fractional quincunx wavelet decompositions with $\lambda = \sqrt{2}$ and $\lambda = \pi$ are shown in Fig. 9. Note how the residual image details are more visible for the lower value of λ. The larger λ reduces the energy of the wavelet coefficient, but this also comes at the expense of some ringing. Thus, it is convenient to have an adjustable parameter to search for the best tradeoff.

An advantage of the present approach is that the filters for small λ are nearly isotropic; this is the reason why the wavelet details in Fig. 9 do not present any preferential orientation. The degree of isotropy of the various lowpass filters can be seen from Fig. 4. The shape of the contour-plots of the low-pass filter $H_\lambda(e^{j\vec{\omega}})$ confirms that the degree of isotropy is the best for small values of λ. At the other extreme, when $\lambda \to \infty$, $H_\lambda(e^{j\vec{\omega}})$ tends to the diamond-shaped ideal filter.

Another nice feature of the algorithm is that the computational cost remains the same irrespective of the value of λ.

B. Dependence of the order parameter

The usefulness of a tunable order parameter is demonstrated in the following experiment: we apply the quincunx transform to the test image “cameraman” (see Fig. 10 (a)) and reconstruct using only 15% of the largest coefficients. Then the SNR is measured depending on the order parameter. The plot in Fig. 11 shows how the SNR changes according to the order λ; the optimum, indicated by the circle, is achieved for $\lambda = 2.5$. Figure 10 (b) and (c) show the reconstructions for the optimal order and an order too high. The last one gets penalized by the introduction of ringing artefacts around the edges. We also plot the SNR curves for 20% and 25% of the coefficients. The same type of qualitative behavior holds for other images.

C. Approximation properties

The main differences between the quincunx and the conventional separable algorithm is the finer scale progression and the non-separability. To test the impact that this may have on compression capability, we compared the approximation qualities of both approaches. Since the wavelet transform is orthogonal, the approximation error (distortion) is equal to $D^2 = \|x - \hat{x}\|^2 = \|y - \hat{y}\|^2$, where y are the wavelet coefficients of the input image x; $\hat{x}$ is the reconstructed image obtained from the quantized—or truncated—wavelet coefficients $\hat{y}$. Also D^2 in the space domain is equivalent to the sum of squares of discarded wavelet coefficients [17].

1) Linear approximation: In classical rate-distortion theory, the coefficients are grouped into channels and coded independently. In the orthogonal case, D^2 is equivalent to the difference between the signal's energy and the energy of the reconstructed signal: $\|x - \hat{x}\|^2 = \|x\|^2 - \|\hat{x}\|^2 = \|x\|^2 - \left(\|\hat{x}_J\|^2 + \sum_{j=1}^{J} \|\hat{y}_j\|^2\right)$. The distortion across N channels with variance σ_i^2 is

$$D = N \cdot C \cdot 2^{-2\bar{R}} \cdot \rho^2, \tag{28}$$

where C is a constant, $\bar{R}$ is the mean rate and ρ is the geometric mean of the subband variances:

$$\rho = \left(\prod_{i=1}^{N} \sigma_i^2\right)^{\frac{1}{N}}. \tag{29}$$

When ρ is small, the distortion is small as well. What this means qualitatively is that the wavelet transform which has the larger spread in the variances will achieve the better coding gain [12]. The linear approximation subband coding gain for sample-by-sample quantization (PCM) is described by:

$$G_{PCM} = \frac{D_{PCM}}{D_{SBC}} = \frac{\frac{1}{N}\sum_{i=1}^{N}\sigma_i^2}{\left(\prod_{i=1}^{N}\sigma_i^2\right)^{\frac{1}{N}}}. \tag{30}$$

To better illustrate this issue, we have decomposed the test image “cameraman” for the maximal number of iterations, both for the quincunx and the separable case as shown in Fig. 12. The order was fixed (i.e., $\lambda = 4$) for our method and for the orthogonal separable approach (corresponding to the commonly used degree parameter $\alpha = 3$ for the underlying B-splines). In Fig. 13 (a) we compare the energy packing properties of both decompositions for linear approximation. “Energy packing” refers to the property that the more the first coefficients contain energy, the better the DWT yields compression. We start to sum up the energy of the subbands with the lowest resolution. Each step of the stairs represents a subband[1]. The first subbands of the quincunx decomposition report higher energy packing than the separable case, but the overall coding gain is slightly better for the separable case than the quincunx case (47.69 versus 45.23). Figure 13 (c) shows similar results for the “Lena” test image.

Since the branches are orthogonal, the transformation that provides the maximum energy compaction in the low-pass channel is also the one that results in the minimum approximation error [17]. Since most images have a power spectrum that is roughly rotationally-invariant and decreases with higher frequencies, separable systems are usually not best suited for isolating a low-pass channel containing most energy and having high-pass channels with low energy. In contrast, a quincunx low-pass filter will retain more of the original signal’s energy [12].

Consequently, the type of images that benefit the most from the quincunx scheme have a more isotropic spectrum. For example, for the well-known zoneplate test image of Fig. 14 (a), the coding gain of quincunx scheme is about 20% better than the one obtained by the separable scheme (4.30 versus 3.64). Also the quincunx scheme gives better energy compaction for textures of highly isotropic nature (and as such a higher coding gain). Two such examples of the Brodatz textures are shown in Fig. 14 (b) and (c), corresponding to a coding gain of 13.67 versus 12.45 and 12.04 versus 9.62, respectively. On the other hand, a separable treatment leads to a better energy compaction for the texture shown in Fig. 15 (8.78 versus 15.48). Other authors have also found that texture analysis using the quincunx scheme improves the results as compared to the separable scheme [18].

---

[1] A quincunx wavelet decomposition with J iterations generates $N = J + 1$ channels, while a separable wavelet decomposition with J iterations results into $N = 3J + 1$ channels.

2) Non-linear approximation: A more recent trend in wavelet theory is the study of *non-linear approximation*. In this case we do not take the "n-first" coefficients, but the "n-largest" coefficients to approximate a signal with n coefficients. This yields better energy packing, since in the wavelet domain the "n-first" coefficients are not necessarily the largest one, especially along the position-indices [19]. The distortion is described by [20]:

$$D^2 = \|y - y_{N(T)}\|^2 = \sum_{|y|<T} |y[n]|^2. \tag{31}$$

Moreover, it can be shown that

$$D \propto C \cdot (N\,(T))^{-\frac{\gamma}{2}}, \tag{32}$$

when the smoothness of y is measured by its inclusion in some critical Besov space $B_q^\gamma(L^q(I))$ with $\frac{1}{q} = \frac{\gamma}{2} + \frac{1}{2}$, roughly speaking when y is a function with γ derivatives in $L^q(I)$ [20], [21].

For the non-linear approximation, the quincunx scheme also yields a better approximation than the separable one for a small n in many cases. Figure 13 (b) represents the energy depending on the n largest coefficients (in log).

VI. EXTENSION TO 3D

The extension of quincunx sampling to 3D is rather straightforward. First, the filters are obtained by replacing $\cos\omega$ by $\frac{1}{3}(\cos\omega_1 + \cos\omega_2 + \cos\omega_3)$ in (8). Next, the quincunx sampling lattice for 3D is shown in Fig. 16 (a). Let $x[\vec{k}]$ denote the discrete signal on the initial grid. Then, its quincunx sampled version, following [6], is

$$[x]_{\downarrow\mathbf{D}}[\vec{k}] = x[\mathbf{D}\vec{k}] \quad , \quad \text{where } \mathbf{D} = \begin{pmatrix} 1 & 0 & 1 \\ -1 & -1 & 1 \\ 0 & -1 & 0 \end{pmatrix}. \tag{33}$$

Our down-sampling matrix $\mathbf{D}$ is such that $\mathbf{D}^3 = 2\mathbf{I}$ and $|\det \mathbf{D}| = 2$. The Fourier-domain version of this formula is similar to the 2D case:

$$[x]_{\downarrow\mathbf{D}}[\vec{k}] \longleftrightarrow \frac{1}{2}\left[X\left(e^{j\mathbf{D}^{-T}\vec{\omega}}\right) + X\left(e^{j(\mathbf{D}^{-T}\vec{\omega}+\vec{\pi})}\right)\right], \tag{34}$$

where $\vec{\pi} = (\pi, \pi, \pi)$.

The implementation for the 3D case goes as follows. The output variables are the discrete Fourier transforms of the wavelet coefficients

$$\begin{aligned}
Y_{i+1}[\vec{n}] &= \sum_{\vec{k}} y_{i+1}[\vec{k}] e^{-j\frac{2\pi\langle\vec{k},\vec{n}\rangle}{N}}, \quad \text{for } n_1, n_2, n_3 = 0, 1, \ldots, N-1, & (35)\\
Y_{i+2}[\vec{n}] &= \sum_{\vec{k}} y_{i+2}[\vec{k}] e^{-j\frac{2\pi\langle\vec{k},\vec{n}\rangle}{N}}, \quad \text{for } n_1, n_2, n_3 = 0, 1, \ldots, N-1, & (36)\\
Y_{i+3}[\vec{m}] &= \sum_{\vec{k}} y_{i+3}[\vec{k}] e^{-j\frac{2\pi\langle\vec{k},\vec{m}\rangle}{\frac{N}{2}}}, \quad \text{for } m_1, m_2, m_3 = 0, 1, \ldots, \frac{N}{2}-1. & (37)
\end{aligned}$$

The coefficients themselves are recovered by inverse FFT. The Fourier transforms after the first level of filtering are given by

$$\begin{aligned}
X_{i+1}[\vec{n}] &= \frac{1}{2}\left(\widetilde{H}[\vec{n}]\, X_i[\vec{n}] + \widetilde{H}\left[\vec{n} + \left(\tfrac{N}{2}, \tfrac{N}{2}, \tfrac{N}{2}\right)\right] X_i\left[\vec{n} + \left(\tfrac{N}{2}, \tfrac{N}{2}, \tfrac{N}{2}\right)\right]\right), & (38)\\
Y_{i+1}[\vec{n}] &= \frac{1}{2}\left(\widetilde{G}[\vec{n}]\, X_i[\vec{n}] + \widetilde{G}\left[\vec{n} + \left(\tfrac{N}{2}, \tfrac{N}{2}, \tfrac{N}{2}\right)\right] X_i\left[\vec{n} + \left(\tfrac{N}{2}, \tfrac{N}{2}, \tfrac{N}{2}\right)\right]\right). & (39)
\end{aligned}$$

After the second level of filtering, we have:

$$\begin{aligned}
X_{i+2}[\vec{n}] &= \frac{1}{2}\left(\widetilde{H}_p[\vec{n}]\, X_{i+1}[\vec{n}] + \widetilde{H}_p\left[\vec{n} + \left(0, \tfrac{N}{2}, 0\right)\right] X_{i+1}\left[\vec{n} + \left(0, \tfrac{N}{2}, 0\right)\right]\right), & (40)\\
Y_{i+2}[\vec{n}] &= \frac{1}{2}\left(\widetilde{G}_p[\vec{n}]\, X_{i+1}[\vec{n}] + \widetilde{G}_p\left[\vec{n} + \left(0, \tfrac{N}{2}, 0\right)\right] X_{i+1}\left[\vec{n} + \left(0, \tfrac{N}{2}, 0\right)\right]\right). & (41)
\end{aligned}$$

Note that these are computed at the resolution of the input. The size reduction only takes place during the third step:

$$\begin{aligned}
X_{i+3}[\vec{m}] &= \frac{1}{2}\left(\widetilde{H}_{pp}[\vec{m}]\, X_{i+2}[\vec{m}] + \widetilde{H}_{pp}\left[\vec{m} + \left(\tfrac{N}{2}, 0, 0\right)\right] X_{i+2}\left[\vec{m} + \left(\tfrac{N}{2}, 0, 0\right)\right]\right), & (42)\\
Y_{i+3}[\vec{m}] &= \frac{1}{2}\left(\widetilde{G}_{pp}[\vec{m}]\, X_{i+2}[\vec{m}] + \widetilde{G}_{pp}\left[\vec{m} + \left(\tfrac{N}{2}, 0, 0\right)\right] X_{i+2}\left[\vec{m} + \left(\tfrac{N}{2}, 0, 0\right)\right]\right), & (43)
\end{aligned}$$

where $\widetilde{H}_p[\vec{m}] = \widetilde{H}[\mathbf{D}\vec{m} \bmod (N,N,N)]$ and $\widetilde{H}_{pp}[\vec{m}] = \widetilde{H}\left[\mathbf{D}^2\vec{m} \bmod (N,N,N)\right]$. Analogously, we have that: $\widetilde{G}_p[\vec{m}] = \widetilde{G}[\mathbf{D}\vec{m} \bmod (N,N,N)]$ and $\widetilde{G}_{pp}[\vec{m}] = \widetilde{G}\left[\mathbf{D}^2\vec{m} \bmod (N,N,N)\right]$.

Figure 16 (b) shows how the coefficients can be arranged in a non-redundant way inside the cube. Note that it is still possible to optimize the 3D implementation taking into account this arrangement already in the Fourier domain, as such reducing the calculation of FFTs. Again, the rotated filters H_p, H_{pp}, G_p, and G_{pp} are precomputed.

A. Approximation properties in 3D

We compared the compression capability for the quincunx and the separable scheme applied to 3D data, similar to the type of experiments that are described for 2D in Sect. V-C. Figure 17 shows the results for a spiral CT dataset of part of a human spine. The linear approximation quality is shown in (b). The separable scheme takes much advantage of the availibity of many small (i.e., seven for each iteration) bandpass subbands, as compared to the quincunx scheme. To illustrate this point, we have grouped the bandpass subbands for the separable case together in one single bandpass in (c). For non-linear approximation, both schemes perform similarly with a small advantage for the separable one, as shown in (d). If the dataset contains more (isotropic) high-frequency components, the breakpoint between the quincunx and the separable case shifts to the right.

The main advantage of the 3D quincunx scheme is in applications that can benefit from the (much) slower scale progression. One example is the statistical analysis of brain activity using functional magnetic resonance imaging (fMRI). Here we show an example using the classical wavelet-based approach for detecting activity, using the linear model analysis and the t-test in the wavelet domain for a 3D dataset ($64 \times 64 \times 64$) with an auditory stimulus [22]; we refer to [23] for more details. We compared the use of the 3D dyadic separable wavelet decomposition based on orthogonal linear B-spline wavelets versus our 3D quincunx wavelets (same order). The parameter maps where obtained using the same threshold after reconstruction (5% of the maximal parameter value). The number of detected voxels, and as such the sensitivity of the approach, is almost 10% higher (578 versus 536) when we use the 3D quincunx DWT, which confirms that the slower scale progression improves the quality of the results. Figure 18 shows the detected activation patterns around the auditory cortex (slice 33).

Other potential applications might include image analysis and 3D feature detection.

VII. CONCLUSION

We have introduced a new family of orthogonal wavelet transforms for quincunx lattices. A key feature is the continuously-varying order parameter λ which can be used to adjust the bandpass characteristics as well as the localization of the basis functions.

We have also demonstrated that these wavelet transforms could be computed quite efficiently in 2D and 3D using FFTs. This should help dispel the commonly held belief that non-separable wavelet decompositions are computationally much more demanding than the separable ones.

Because of their nice properties and their ease of implementation, these wavelets present an alternative to the separable ones that are being used in a variety of image processing applications (image analysis, image enhancement, filtering and denoising, feature detection, texture analysis, and so on).

REFERENCES

[1] J. Kovačević and M. Vetterli, “Nonseparable two- and three-dimensional wavelets,” *IEEE Transactions on Signal Processing*, vol. 43, no. 5, pp. 1269–1273, 1995.

[2] A. Mojsilović, M. Popović, S. Marković, and M. Krstić, “Characterization of visually similar diffuse diseases from B-scan liver images using nonseparable wavelet transform,” *IEEE Transactions on Medical Imaging*, vol. 17, no. 4, pp. 541–549, 1998.

[3] J. C. Feauveau, “Analyse multirésolution avec un facteur de résolution $\sqrt{2}$,” *Journal de Traitement du Signal*, vol. 7, no. 2, pp. 117–128, 1990.

[4] J. H. McClellan, “The design of two-dimensional digital filters by transformations,” in *Proc.* 7^{th} *Annual Princeton Conf. Inf. Sciences and Systems*, Princeton (USA), 1973, pp. 247–251.

[5] A. Cohen and I. Daubechies, “Nonseparable bidimensional wavelet bases,” *Rev. Mat. Iberoamericana*, vol. 9, pp. 51–137, 1993.

[6] J. Kovačević and M. Vetterli, “Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for $\mathbb{R}^n$,” *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 533–555, 1992.

[7] J. Shapiro, “Adaptive McClellan transformations for quincunx filter banks,” *IEEE Trans. on Signal Processing*, vol. 42, no. 3, pp. 642–648, 1994.

[8] D. B. H. Tay and N. G. Kingsbury, “Flexible design of multidimensional perfect reconstruction FIR 2-band filters using transformations of variables,” *IEEE Trans. on Image Processing*, vol. 2, no. 4, pp. 466–480, 1993.

[9] F. Nicolier, O. Laligant, and F. Truchetet, “B-spline quincunx wavelet transform and implementation in Fourier domain,” in *Proc. SPIE*, vol. 3522, Boston, Massachussets, USA, november 1998, pp. 223–234.

[10] M. Unser and T. Blu, “Fractional splines and wavelets,” *SIAM Review*, vol. 42, pp. 43–67, 2000.

[11] S. Mallat, “A theory for multiresolution signal decomposition: the wavelet representation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, pp. 674–693, 1989.

[12] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*. Prentice-Hall, 1995.

[13] M. Unser and T. Blu, “Fractional wavelets, derivatives, and Besov spaces,” in *Proceedings of the SPIE Conference on Mathematical Imaging: Wavelet Applications in Signal and Image Processing X*, vol. 5207. San Diego, CA: SPIE, Aug. 2003, pp. 147–152.

[14] O. Rioul and P. Duhamel, “Fast algorithms for discrete and continuous wavelet transforms,” *IEEE Trans. on Information Theory*, vol. 38, no. 2, pp. 569–586, 1992.

[15] F. Nicolier, O. Laligant, and F. Truchetet, “Discrete wavelet transform implementation in Fourier domain for multidimensional signal,” *Journal of Electronic Imaging*, vol. 11, no. 3, pp. 338–346, July 2002.

[16] Mathworks Inc., “Wavelet toolbox 2.2,” 2000.

[17] M. Unser, “On the optimality of ideal filters for pyramid and wavelet signal approximation,” *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3591–3596, Dec. 1993.

[18] F. Nicolier, A. Legrand, O. Laligant, S. Kohler, and F. Truchetet, “Ovocyte texture analysis through almost shift-invariant decimated wavelet transform,” in *Fourth International Conference on Knowledge-Based Intelligent Engineering Systems & Allied Technologies*, Brighton (UK), Aug. 2000, pp. 385–388.

[19] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies, “Data compression and harmonic analysis,” *IEEE Trans. Inform. Th.*, vol. 44, no. 6, pp. 2435–2475, 1998.

[20] R. A. DeVore, B. Jawerth, and B. J. Lucier, “Image compression through wavelet transform coding,” *IEEE Trans. Inform. Th.*, vol. 38, no. 2, pp. 719–746, 1992.

[21] R. A. DeVore, “Nonlinear approximation,” *Acta Numerica*, vol. 7, pp. 51–150, 1998.

[22] G. Rees and K. Friston, “Single subject epoch (block) auditory fMRI activation data,” http://www.fil.ion.ucl.ac.uk/spm/data/, 1999.

[23] D. Van De Ville, T. Blu, and M. Unser, “Wavelets versus resels in the context of fMRI: establishing the link with

SPM," in *SPIE's Symposium on Optical Science and Technology: Wavelets X*, vol. 5207. SPIE, Aug. 2003, pp. 417–425.

LIST OF FIGURES

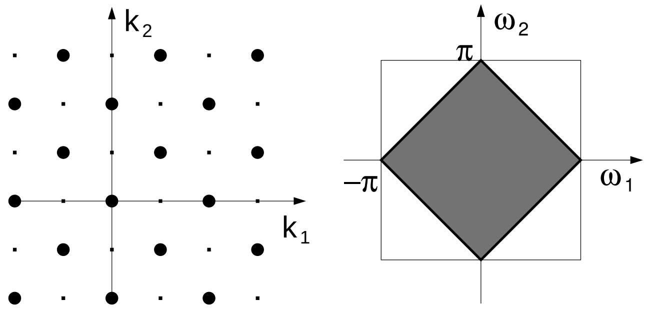

(a) (b)

Fig. 1. (a) The quincunx lattice and (b) the corresponding Nyquist area in the frequency domain.

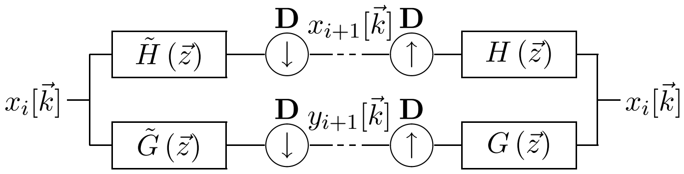

Fig. 2. Perfect reconstruction filterbank on a quincunx lattice.

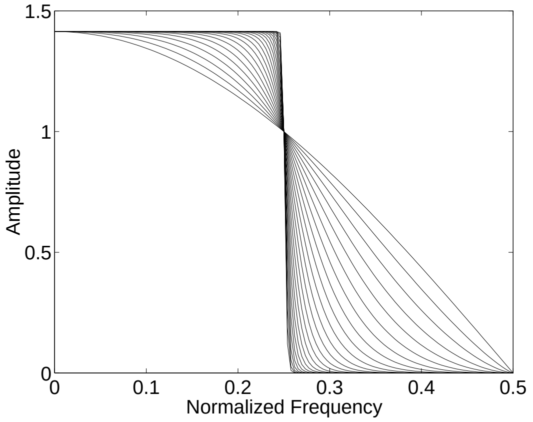

Fig. 3. Frequency responses of the orthogonal refinement filters for $\lambda = 1, \ldots, 100$.

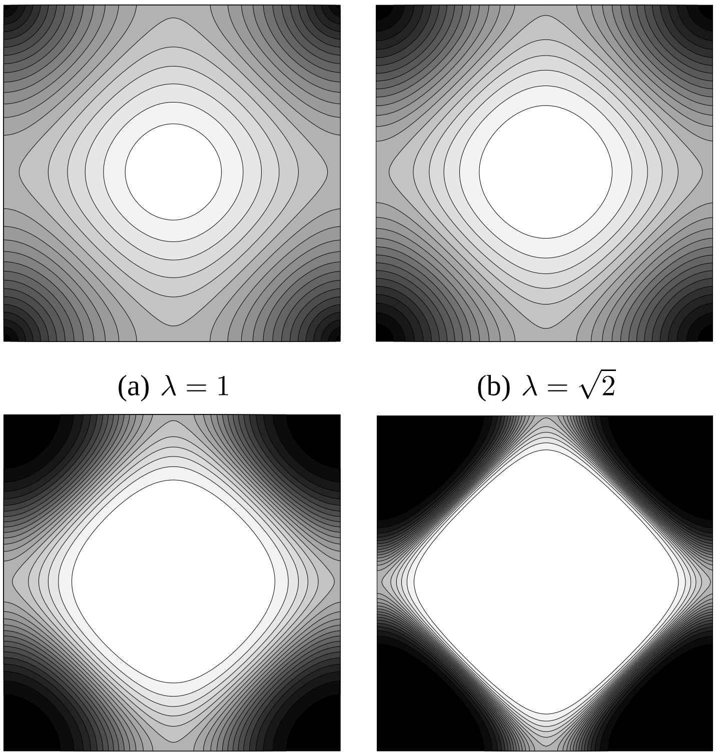

(a) $\lambda = 1$ (b) $\lambda = \sqrt{2}$

(c) $\lambda = \pi$ (d) $\lambda = 10$

Fig. 4. Contour plots of the low-pass filters $H_\lambda(e^{j\vec{\omega}})$ for various values of the order parameter λ.

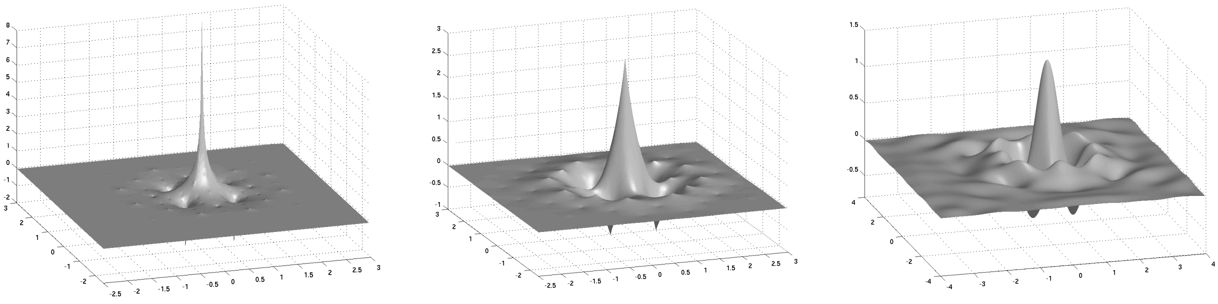

(a) $\lambda = \sqrt{2}$ (b) $\lambda = \pi$ (c) $\lambda = 10$

Fig. 5. Surface plots of the wavelets ψ_λ for various values of the order parameter.

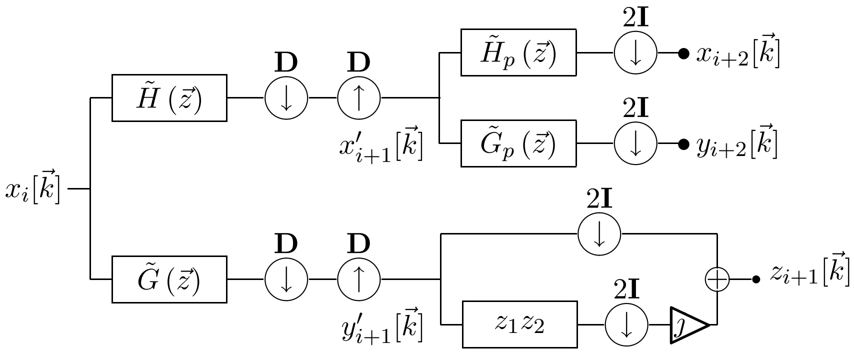

Fig. 6. Analysis part of the 2D quincunx wavelet transform for two iterations.

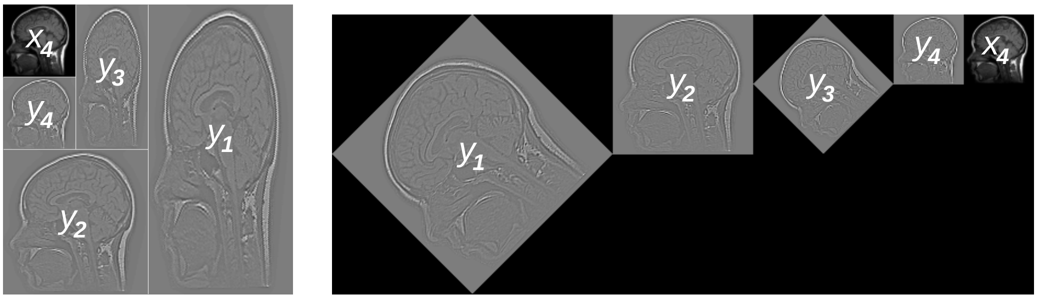

(a) (b)

Fig. 7. The wavelet coefficients for the quincunx subsampling scheme can be arranged in two ways. An example for $J = 4$ iterations. (a) Compact representation. (b) Classic representation.

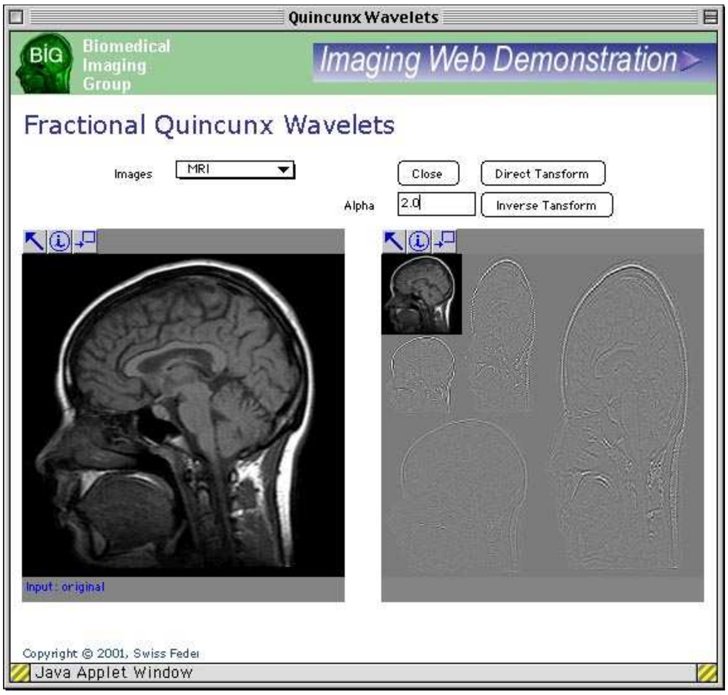

Images MRI

Alpha 2.0

Copyright © 2001, Swiss Feder

Fig. 8. Applet of the Fourier-based implementation of the quincunx wavelet transform, available on the site: `http://bigwww.epfl.ch/demo/jquincunx/`.

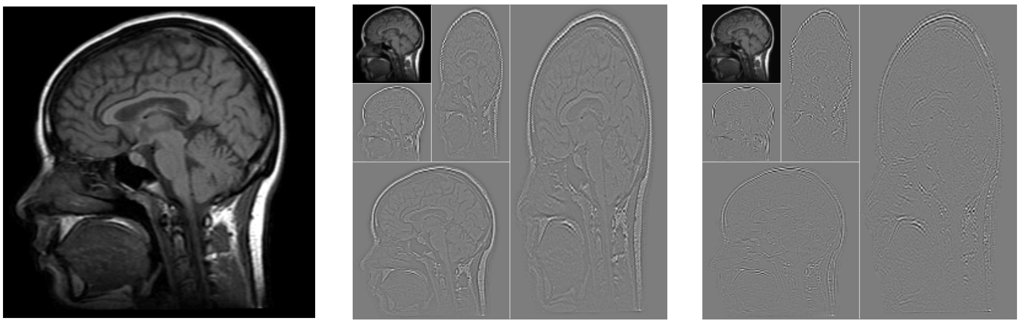

(a) input image (b) $\lambda = \sqrt{2}$ (c) $\lambda = \pi$

Fig. 9. Quincunx wavelet transforms with four iterations: (a) original test image, (b) $\lambda = \sqrt{2}$, (c) $\lambda = \pi$.

(a) original image (b) $\lambda = 2.5$ (c) $\lambda = 14$

Fig. 10. (a) Original test image “cameraman”. (b) & (c) Reconstruction of “cameraman” using 15% of the largest coefficients with $\lambda = 2.5$ (optimal) and $\lambda = 14$.

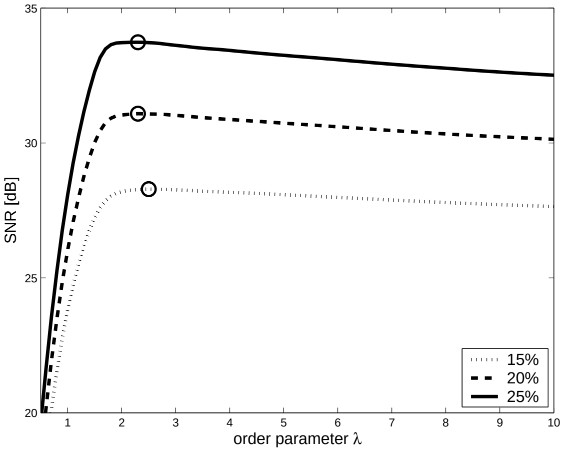

Fig. 11. Relation between the order parameter λ and the SNR of the reconstructed image (test image “cameraman”) using only the largest coefficients. The full line, dashed line, and dotted line correspond respectively to 25%, 20%, and 15% of the largest coefficients.

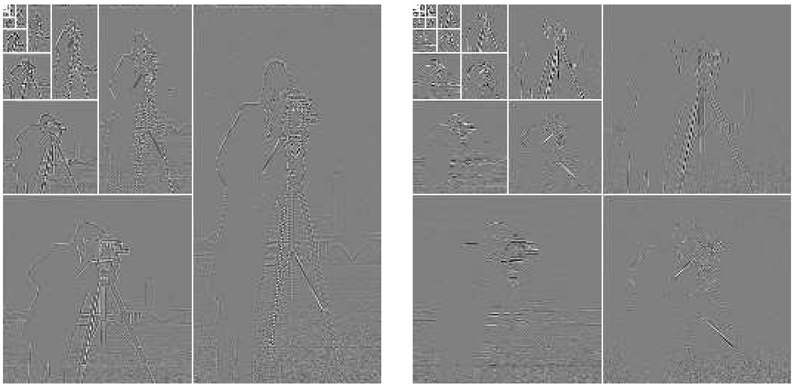

(a) (b)

Fig. 12. Decomposition of the test image “cameraman” for the maximal possible number of iterations. (a) Quincunx case. (b) Separable case. The contrast of each subband has been enhanced.

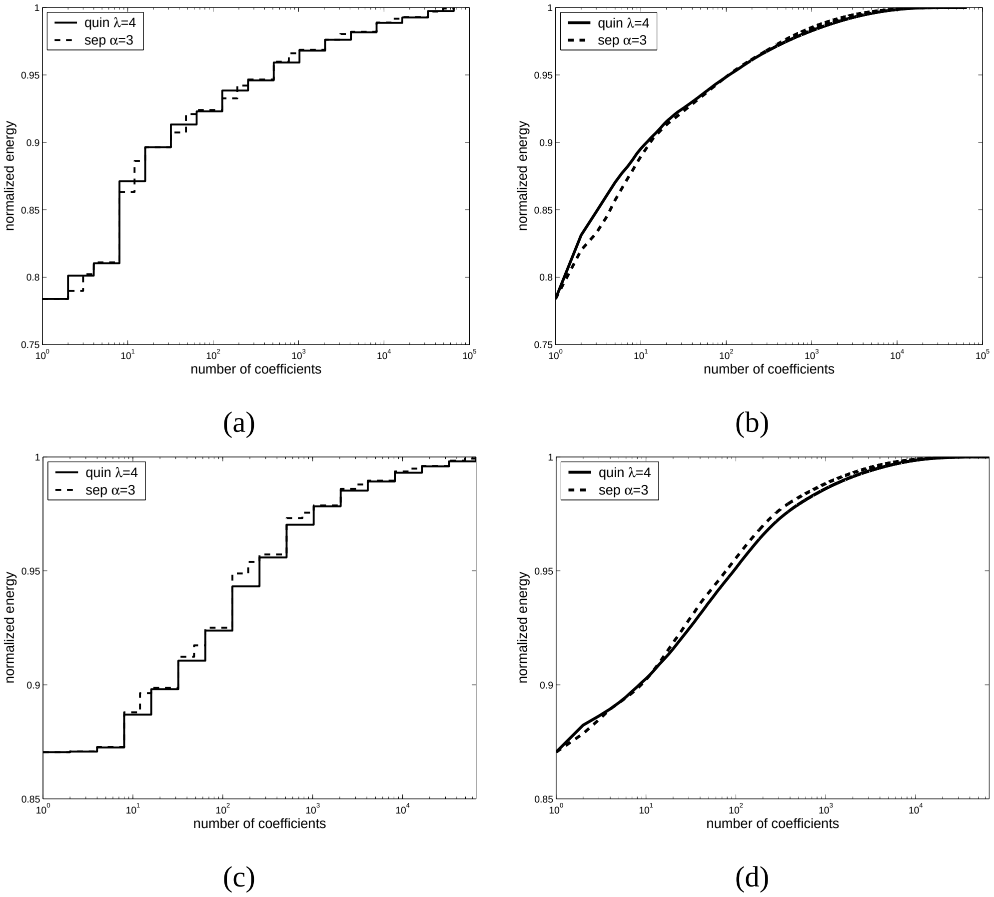

(a) (b)

(c) (d)

Fig. 13. Comparison of energy-compaction property between the quincunx and the separable case of image decomposition (as shown in Fig. 12). (a) & (c) Linear approximation depending on number of coefficients (in log, grouped per subband), respectively for “cameraman” and “Lena”. (b) & (d) Non-linear approximation depending of the n largest coefficients (in log), respectively for “cameraman” and “Lena”. The quincunx scheme yields better results for a low number of coefficients. In the case of “lena”, the separable scheme performs better than the quincunx one over most of the range.

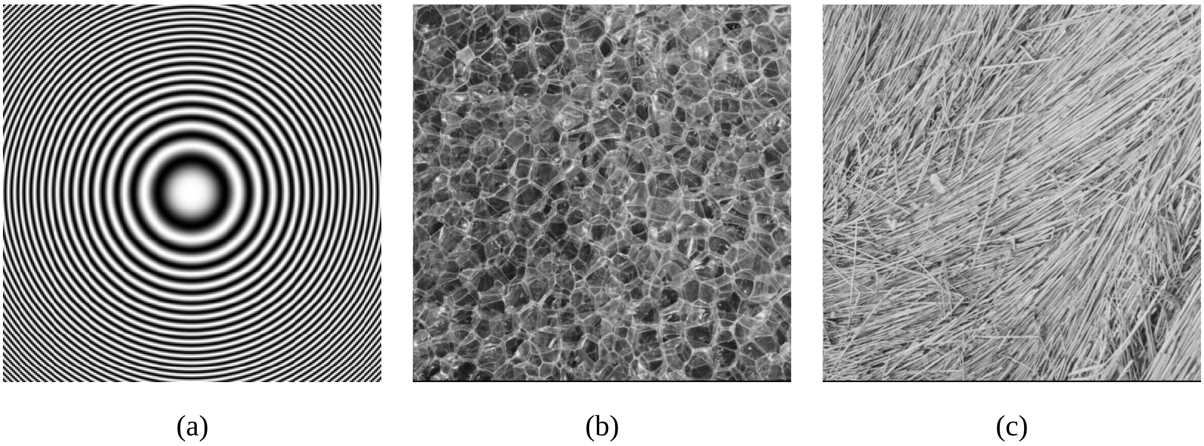

(a) (b) (c)

Fig. 14. Some examples of typical images where the quincunx scheme outperforms the separable case in term of coding gain. (a) Zoneplate. (b) Brodatz texture D112. (c) Brodatz texture D15. The Brodatz textures have 512×512 pixels and are obtained from the USC-SIPI Image Database.

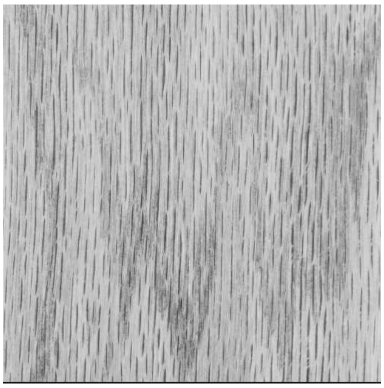

Fig. 15. An example of a texture (Brodatz D68) that is better suited for a separable treatment.

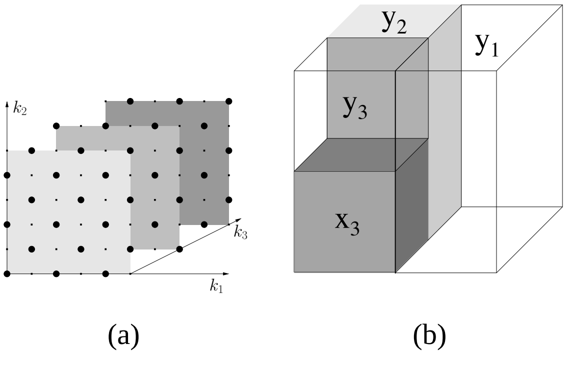

Fig. 16. (a) 3D-FCO (face-centered orthorhombic) lattice, corresponding to the sampling matrix of Eq. (33). (b) Compact representation of the wavelet coefficients for the 3D case.

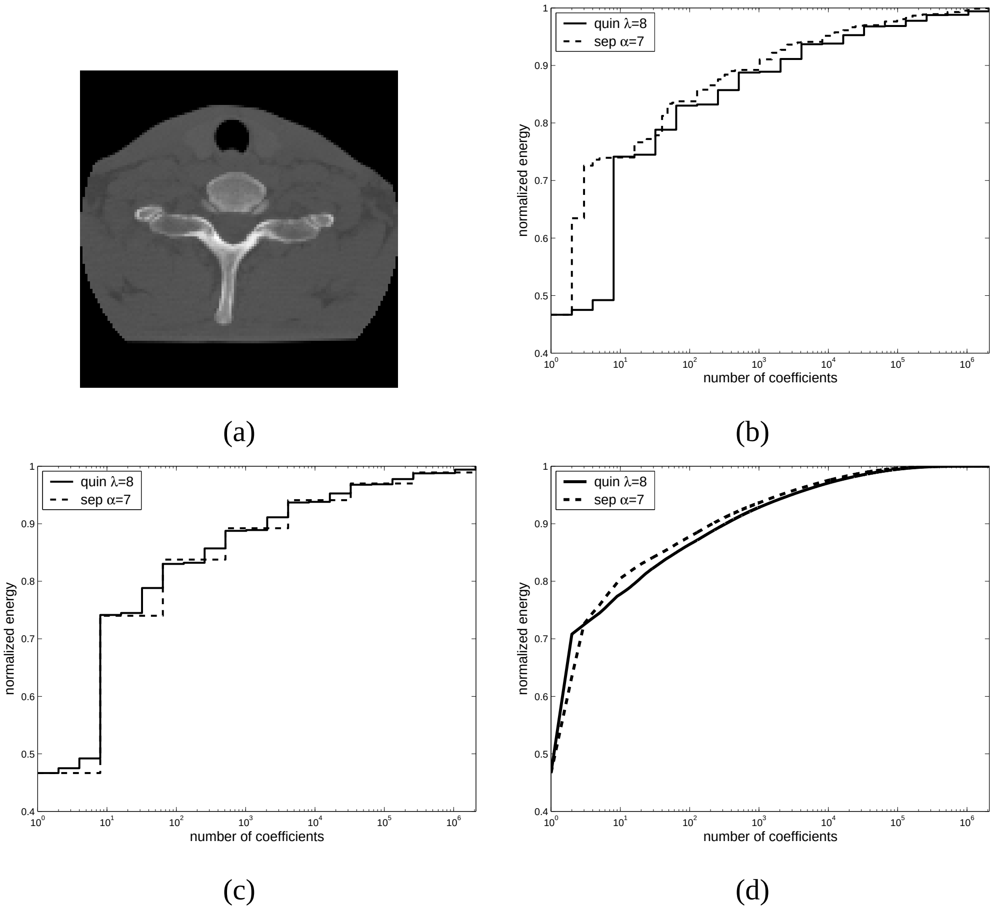

(a) (b)

(c) (d)

Fig. 17. (a) Slice of an spiral CT dataset of part of a human spine (courtesy and copyright of Ramani Pichumani, Stanford University School of Medicine). (b) Linear approximation, for the separable case each bandpass subband is considered independently. (c) Linear approximation, for the separable case the bandpass subbands are grouped together into one single subband. (d) Non-linear approximation.

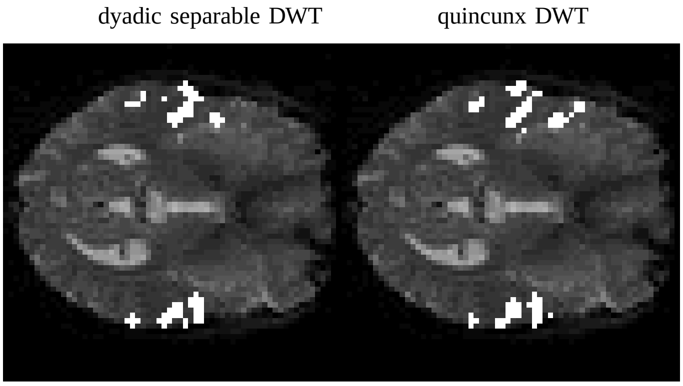

dyadic separable DWT quincunx DWT

Fig. 18. FMRI brain activation detected using the classical wavelet-based approach. The activated voxels in the slice are left white, superposed on a background of the T2*-scan.