ORTHOGONAL QUINCUNX WAVELETS WITH FRACTIONAL ORDERS

Manuela Feilner, Mathews Jacob and Michael Unser

Swiss Federal Institute of Technology Lausanne
Biomedical Imaging Group, DMT/IOA
CH-1015 Lausanne EPFL
Switzerland
e-mail: manuela.feilner@epfl.ch

ABSTRACT

We present a new family of 2D 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 α, which may be non-integer. The wavelets have good isotropy properties. We can also prove that they yield wavelet bases of $L_2(\mathbb{R}^2)$ for any $\alpha > 0$. The wavelets are fractional in the sense that the approximation error at a given scale a decays like $O(a^\alpha)$; 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.

1. 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 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 to map 1D biorthogonal designs to the multidimensional case [4]. Since this approach requires the filters to be symmetric, it has been applied mainly to the biorthogonal case because of the strong incentive to produce filters that are compactly supported [5, 6, 7, 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].

2. QUINCUNX SAMPLING AND FILTERBANKS

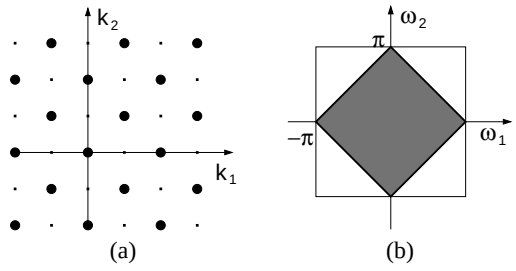

Fig. 1. The quincunx lattice (a) and its corresponding bandwidth (b).

First, we recall some basic results on quincunx sampling and perfect reconstruction filterbanks [12]. The quincunx sampling lattice is shown in Fig. 1. Let $x[\vec{k}]$ denote the discrete signal on the initial grid. Then, its quincunx sampled version is

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

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

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

This work was supported in part by the Swiss National Science Foundation under grant 2100-053540.

where $\vec{\pi} = (\pi, \pi)$.
The upsampling is defined by

$$[x]_{\uparrow \mathbf{D}}[\vec{k}] = \begin{cases} x[\mathbf{D^{-1}}\vec{k}] & \text{if } k_1 + k_2 \text{ even} \\ 0 & \text{else where} \end{cases}$$

and its effect in the transform domain is as follows:

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

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

$$[x]_{\downarrow \mathbf{D} \uparrow \mathbf{D}}[\vec{k}] = \begin{cases} x[\vec{k}] & \text{if } k_1 + k_2 \text{ even} \\ 0 & \text{else where} \end{cases}$$
$$\updownarrow$$
$$\frac{1}{2}\left[X\left(\vec{\omega}\right) + X\left(\vec{\omega} + \vec{\pi}\right)\right] \tag{1}$$

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 lowpass filter $\widetilde{H}$ reduces the resolution by a factor of $\sqrt{2}$; the wavelet coefficients correspond to the output of the highpass filter $\widetilde{G}$.

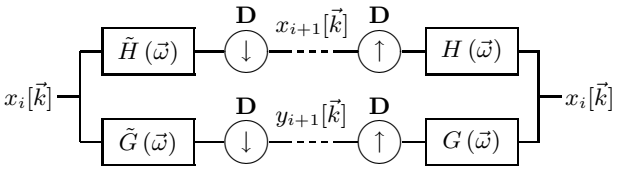

Fig. 2. Perfect reconstruction filterbank on a quincunx lattice.

Applying the relation (1) to the blockdiagram in Fig. 2, it is easy to derive the conditions for a perfect reconstruction:

$$\begin{cases} \widetilde{H}\left(\vec{\omega}\right) H\left(\vec{\omega}\right) + \widetilde{G}\left(\vec{\omega}\right) G\left(\vec{\omega}\right) = 2 \\ \widetilde{H}\left(\vec{\omega} + \vec{\pi}\right) H\left(\vec{\omega}\right) + \widetilde{G}\left(\vec{\omega} + \vec{\pi}\right) G\left(\vec{\omega}\right) = 0 \end{cases} \tag{2}$$

where H and G (resp., $\widetilde{H}$ and $\widetilde{G}$) are the transfer functions of the synthesis (resp., 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 lowpass filter H.

3. 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.

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

$$\begin{aligned} H_\alpha(z) &= \frac{\sqrt{2}(z + 2 + z^{-1})^{\frac{\alpha}{2}}}{\sqrt{(z + 2 + z^{-1})^\alpha + (-z + 2 - z^{-1})^\alpha}} \\ &= \frac{\sqrt{2}(2 + 2\cos\omega)^{\frac{\alpha}{2}}}{\sqrt{(2 + 2\cos\omega)^\alpha + (2 - 2\cos\omega)^\alpha}} \end{aligned} \tag{3}$$

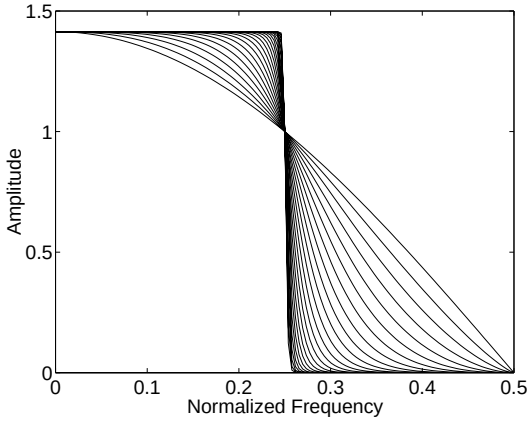

Fig. 3. Frequency responses of the orthogonal refinement filters for $\alpha = 1, \ldots, 100$.

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

These filters are symmetric and are designed to have zeros of order α at $z = -1$; the numerator is a fractional power of $(z + 2 + z^{-1})$ (the simplest symmetric refinement filter) and the denominator is the appropriate orthonormalization factor. By varying α, we can adjust the frequency response as shown in Fig. 3. As α increases, $H_\alpha(z)$ converges to the ideal half-band lowpass filter. Also note that these filters are maximally flat at the origin; they essentially behave like $H_\alpha\left(\omega\right) = 1 + O(\omega^\alpha)$ 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 which will act like a fractional derivative of order α; in other words, it will kill all polynomials of degree $n \leq \lceil \alpha - 1 \rceil$.

Applying the diamond McClellan transform to the filter above is straightforward; it amounts to replacing $\cos\omega$ by $\frac{1}{2}\left(\cos\omega_1 + \cos\omega_2\right)$ in (3). Thus, our quincunx refinement filter is given by

$$H_\alpha(\omega_1, \omega_2) = \frac{\sqrt{2}(2 + \cos\omega_1 + \cos\omega_2)^{\frac{\alpha}{2}}}{\sqrt{(2 + \cos\omega_1 + \cos\omega_2)^\alpha + (2 - \cos\omega_1 - \cos\omega_2)^\alpha}} \tag{4}$$

This filter is guaranteed to be orthogonal because the McClellan transform has the property of preserving biorthogonality. Also, by construction, the αth order zero at $\omega = \pi$ gets mapped into a corresponding zero at $(\omega_1, \omega_2) = (\pi, \pi)$; this is precisely the condition that is required to get a two dimensional wavelet transform of order α.

The orthogonal wavelet filter is obtained by modulation

$$G_\alpha(\omega_1, \omega_2) = e^{j\omega_1} H_\alpha(-\omega_1 - \pi, -\omega_2 - \pi) \tag{5}$$

The corresponding orthogonal scaling function $\varphi_\alpha\left(\vec{x}\right)$ is de-

fined implicitly as the solution of the quincunx two-scale relation:

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

Since the refinement filter is orthogonal with respect to the quincunx lattice, it follows that $\varphi_\alpha(\vec{x}) \in L_2(\mathbb{R}^2)$ and that it is orthogonal to its integer translates. Moreover, for $\alpha > 0$, 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

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

4. IMPLEMENTATION IN FOURIER DOMAIN

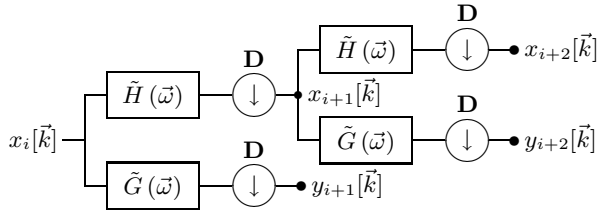

Fig. 4. Analysis part of the 2D QWT for two iterations.

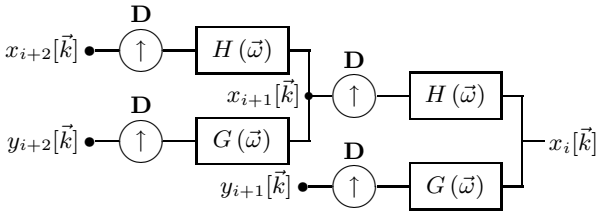

Fig. 5. Synthesis part of the 2D quincunx wavelet transform (IQWT).

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. Following the proposal of Nicolier et al. [9], we will see here that one can turn the situation around and obtain a very simple and efficient algorithm that is based on the FFT. Indeed, Rioul et al. [13] have suggested early on to use the FFT as an efficient tool for computing 1D wavelet transforms when the filters are long. Working in the frequency domain is also very convenient for us because of the way in which we have specified our filters (cf. Eqs. (4) and (5)).

Here, we will only describe the decomposition part of the algorithm which corresponds to the block diagram in Fig. 4 where we have pooled together two levels of 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 (4) and (5). Conceptually, the proposed method is similar to the one described in [9]. However, we take advantage of symmetries and redundancies in the Fourier domain which accelerates the evaluation of $X_{i+2}[\vec{m}]$ and $Y_{i+2}[\vec{m}]$ by at least a factor of two.

Assuming that the current image size is $N \times N$, the input variable is

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

The output variables are the discrete Fourier transforms of the wavelet coefficients

$$Y_{i+1}[\vec{n}] = \textstyle\sum_{\vec{k}} y_{i+1}[\vec{k}] e^{-j\frac{2\pi\langle\vec{k},\vec{n}\rangle}{N}} \quad \text{for } n_1, n_2 = 0 \ldots N-1$$
$$Y_{i+2}[\vec{m}] = \textstyle\sum_{\vec{k}} y_{i+2}[\vec{k}] e^{-j\frac{2\pi\langle\vec{k},\vec{m}\rangle}{\frac{N}{2}}} \quad \text{for } m_1, m_2 = 0 \ldots \frac{N}{2} - 1$$

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

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

Note that these are computed at the resolution of the input. The size reduction only takes place during the second step (even iteration):

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

where $\widetilde{H}_p[\vec{m}] = \widetilde{H}[\mathbf{D}\vec{m} \bmod (N, N)]$ and $\widetilde{G}_p[\vec{m}] = \widetilde{G}[\mathbf{D}\vec{m} \bmod (N, N)]$. The process is then iterated until one reaches the final resolution. Obviously, as one gets coarser, 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 using up-sampling instead. It corresponds to the block diagram in Fig. 5.

We have implemented the algorithm in Matlab and report computation times below $1.2\,sec$ for a 256×256 image on an aging Sun Ultra 30 workstation; the decomposition is essentially perfect with an (RMS) reconstruction error below 10^{-12}. The method is generic and works for any set of filters that can be specified in the frequency domain. Based on the timings reported in [9], we estimate our implementation to be at least five times faster than the one of Nicolier which is also coded in Matlab.

Two examples of fractional quincunx wavelet decompositions with $\alpha = \sqrt{2}$ and $\alpha = \pi$ are shown in Fig. 6. Note how the residual image details are more visible for the lower value of α. The larger α reduces the energy of the wavelet coefficients, 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's are nearly isotropic, especially for small values of α; this is the reason

why the wavelet details in Fig. 6 do not present any preferential orientation. Another nice feature of the algorithm is the computational cost remains the same irrespective of the value of α.

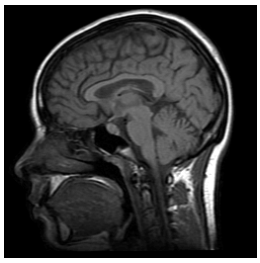

input image (256×256)

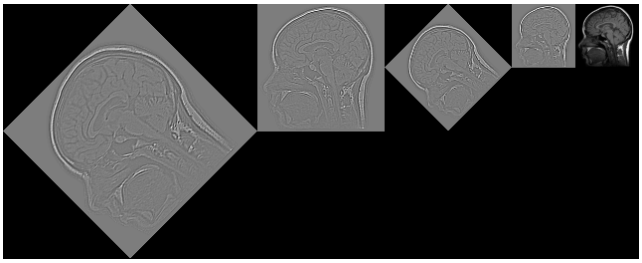

(a)

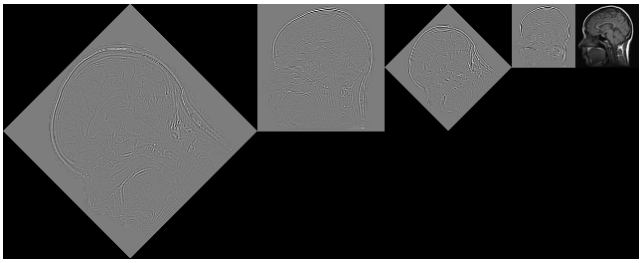

(b)

Fig. 6. Quincunx wavelet transforms with 4 iterations: (a) $\alpha = \sqrt{2}$, (b) $\alpha = \pi$. The bandpass coefficients are on the left (first four images) and the lowpass coefficients on the right.

5. CONCLUSION

We have introduced a new family of orthogonal wavelet transforms for quincunx lattices with good isotropy properties. A key feature is the continuously-varying order parameter α 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 using FFTs. This should help dispell 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 should present a good alternative to the separable ones that are being used in a variety of image processing applications (data compression, filtering, texture analysis etc.).

6. 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 Transaction 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 fransformations 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*, Boston, Massachusets, USA, November 1998, vol. 3522, pp. 223–234.

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

[11] S. G. Mallat, "A theory for multiresolution signal decomposition: the wavelet representaion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 7, pp. 674–693, 1989.

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

[13] 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.