Image Compression With Adaptive Local Cosines: A Comparative Study

François G. Meyer, *Member, IEEE*

***Abstract*—The goal of this work is twofold. First, we demonstrate that an advantage can be gained by using local cosine bases over wavelets to encode images that contain periodic textures. We designed a coder that outperforms one of the best wavelet coders [1] on a large number of images. The coder finds the optimal segmentation of the image in terms of local cosine bases. The coefficients are encoded using a scalar quantizer optimized for Laplacian distributions. This new coder constitutes the first concrete contribution of the paper. Second, we used our coder to perform an extensive comparison of several optimized bells [2]–[4] in terms of rate-distortion and visual quality for a large collection of images. This study provides for the first time a rigorous evaluation in realistic conditions of these bells. Our experiments show that bells that are designed to reproduce exactly polynomials of degree 1 resulted in the worst performance in terms of PSNR. However, a visual inspection of the compressed images indicates that these bells often provide reconstructed images with very few visual artifacts, even at low bit rates. The bell with the most narrow Fourier transform gave the best results in terms of PSNR on most images. This bell tends however to create annoying visual artifacts in very smooth regions at low bit rate.**

***Index Terms*—DCT, image compression, lapped transform, local cosine transform, transform coding.**

I. Introduction

MANY classes of images have very diffuse representations in a standard wavelet basis. Images with oscillatory patterns are examples of non-wavelet-friendly signals. Rapid variations of the intensity can only be described by many small scale wavelet coefficients. Unfortunately, these small scale coefficients carry very little energy, and are often quantized to zero, even at high bit rates. In principle, the description of oriented periodic patterns can be best performed with local windowed Fourier or cosine bases. A lot of work has recently been expended in the design of orthonormal and biorthogonal bases that exploit a combination of overlapping bells and a Fourier-like transform [2]–[9]. While several variations exist, one focuses in this paper on the following family of local cosine basis functions

$$\psi_{n,k}(x) = b_n(x)\cos\left(k+\tfrac{1}{2}\right)\pi x$$

where the function b_n is a smooth window, or *bell*. If one accepts to use Riesz biorthogonal bases instead of orthonormal bases, then one has much freedom to choose the bell b_n. Recently, several authors have designed "optimized bells" [2]–[4]. The optimality of these bells has been claimed from a strictly theoretical point of view. The goal of this work is twofold. First, we wish to demonstrate that an advantage can be gained by using local cosine bases over wavelets to encode images that contain periodic textures. The most concrete product of this research is a compression algorithm that outperforms one of the best wavelet coder [1] on a large number of images. This new coder constitutes the first contribution of the paper. Second, we used our coder to perform an extensive comparison of several bells [2]–[4] in terms of rate-distortion and visual quality for a large collection of images.

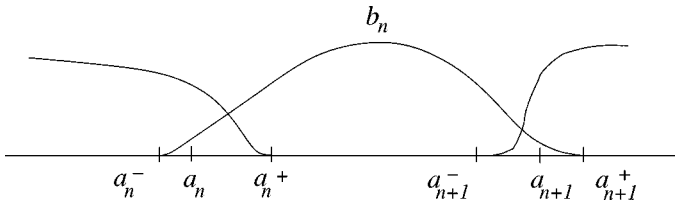

Fig. 1. Bell b_n lives over the interval $[a_n^-,\ a_{n+1}^+]$.

This paper is organized as follows. In Section II we review the construction of smooth localized cosine transforms. Several choices of optimized bells are presented in Section III. A theoretical comparison of the bells is performed in Section IV. In Section V we describe the new compression algorithm. Results of experiments are presented in Section VI.

II. Adaptive Smooth Local Cosine Transforms

We review here the construction of smooth localized cosine transforms (LCT) in the one-dimensional (1-D) case (references on the topic include for instance [6], [9], [10]). In two dimensions, we use tensor products of 1-D bases. We consider the more general setting where we have two biorthogonal bases [4], [6], [11]. Let $\bigcup_{n=-\infty}^{n=+\infty}[a_n, a_{n+1})$ be a cover of $\mathbb{R}$. We define a neighborhood around each point a_n: $[a_n^-, a_n^+]$, where a_n^-, a_n, a_n^+ are such that (see Fig. 1)

$$a_n^- < a_n < a_n^+ \leq a_{n+1}^-. \tag{1}$$

Let b_n be a sequence of bells such that

$$\forall x \in [a_n^-, a_n^+],$$
$$b_n(x)b_{n-1}(2a_n - x) + b_n(2a_n - x)b_{n-1}(x) \neq 0 \tag{2}$$

and

$$\forall x \in [a_n^+, a_{n+1}^-], \quad b_n(x) \neq 0. \tag{3}$$

Manuscript received June 12, 2001; revised February 4, 2002. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Nasir Memon.

The author is with the Department of Electrical Engineering, University of Colorado at Boulder, Boulder, CO 80309-0425 USA (e-mail: francois.meyer@colorado.edu).

Publisher Item Identifier S 1057-7149(02)04789-9.

We only consider the two overlapping case: the bell b_n only talks to the bells b_{n-1} and b_{n+1} (see Fig. 1). Let

$$\theta_n(x) = \frac{1}{b_n(x)b_{n-1}(2a_n - x) + b_n(2a_n - x)b_{n-1}(x)} \quad (4)$$

then the dual bell $\tilde{b}_n$ is defined as follows:

$$\tilde{b}_n(x) = \begin{cases} \theta_n(x)b_{n-1}(2a_n - x), & \text{if } a_n^- \le x \le a_n^+ \\ \dfrac{1}{b_n(x)}, & \text{if } a_n^+ \le x \le a_{n+1}^- \\ \theta_{n+1}(x)b_{n+1}(2a_{n+1} - x), & \text{if } a_{n+1}^- \le x \le a_{n+1}^+ \\ 0, & \text{otherwise} \end{cases}. \quad (5)$$

Let $C_{n,k}(x)$ be the family of basis functions of the DCT-IV

$$C_{n,k}(x) = \sqrt{\frac{2}{a_{n+1} - a_n}} \cos\left[(k + 1/2)\pi \frac{x - a_n}{a_{n+1} - a_n}\right]. \quad (6)$$

We define the local cosine basis functions as

$$\psi_{n,k}(x) = b_n(x)C_{n,k}(x) \quad (7)$$

and the dual basis functions are defined as

$$\tilde{\psi}_{n,k}(x) = \tilde{b}_n(x)C_{n,k}(x). \quad (8)$$

One can use either one of the two bases to perform the analysis, and compute the coefficients, of a function f. One can then use the other basis for the synthesis, or reconstruction, of the function.

Lemma 1: $\psi_{n,k}$ and $\tilde{\psi}_{n,k}$ are Riesz biorthogonal bases

$$\int \psi_{n,k}(x)\tilde{\psi}_{m,j}(x)dx = \delta_{j,k}\,\delta_{n,m} \quad (9)$$

$\forall f \in L^2(\mathbb{R})$

$$f(x) = \sum_{n,k} \tilde{f}_{n,k}\psi_{n,k}(x) \quad \tilde{f}_{n,k} = \int f(x)\tilde{\psi}_{n,k}(x)\,dx \quad (10)$$

$$f(x) = \sum_{n,k} f_{n,k}\tilde{\psi}_{n,k}(x) \quad f_{n,k} = \int f(x)\psi_{n,k}(x)\,dx. \quad (11)$$

Furthermore, $\exists B > A > 0$ such that

$$A\sum_{n,k} |f_{n,k}|^2 \le \|\sum_{n,k} f_{n,k}\psi_{n,k}\|^2 \le B\sum_{n,k} |f_{n,k}|^2. \quad (12)$$

The constants A and B are called the Riesz bounds. If $\psi_{n,k}$ is an orthonormal sequence, then $A = B = 1$. If f has unit norm, then

$$\frac{1}{B} \le \sum_{n,k} |f_{n,k}|^2 \le \frac{1}{A}. \quad (13)$$

If A is much smaller than 1, then the coefficients $f_{n,k}$ in (13) can be very large. Conversely, if B is very large, the coefficients $f_{n,k}$ can become extremely small. In order to obtain decompositions that are numerically stable, and coefficients that neither explode nor vanish, one would like to have Riesz bounds close to one.

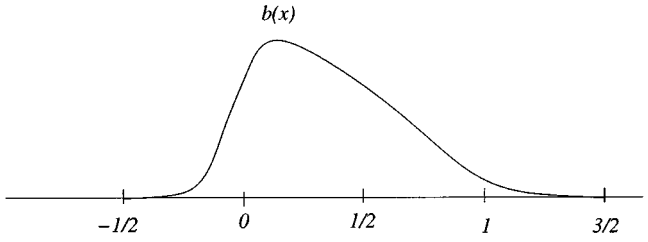

Fig. 2. Bell b lives over the interval $[-1/2, 3/2)$.

A. Implementation by Folding

In practice, in order to expand a function f into the basis $\{\psi_{n,k}\}$ we do not calculate the correlation between f and the basis functions $\psi_{n,k}$. Instead we transform f restricted to $[a_n^-, a_{n+1}^+]$ into a smooth function onto $[a_n, a_{n+1}]$ with the adequate polarity (even around a_n, and odd around a_{n+1}), and we expand it into the basis $\{C_{n,k}\}$. To do this we fold the overlapping parts of the bells b_n, b_{n-1} and b_{n+1} back into the interval $[a_n, a_{n+1}]$, across the end points of the interval, with some folding and unfolding operators [9], and then use a fast DCT to calculate the expansion into the basis $\{C_{n,k}\}$. The folding and unfolding with nonsymmetric bells [6] is a straightforward generalization of the folding defined in [9]. The folding operator U_{a_n} that folds a function f around a_n is given by

$$\begin{cases} b_{n-1}(x)f(x) \\ \quad -b_{n-1}(2a_n - x)f(2a_n - x), & \text{if } a_n^- \le x \le a_n \\ b_n(x)f(x) \\ \quad +b_n(2a_n - x)f(2a_n - x), & \text{if } a_n \le x \le a_n^+. \end{cases} \quad (14)$$

The unfolding operator $U_{a_n}^*$, that unfolds a function f around a_n, is obtained in a similar way:

$$\begin{cases} \tilde{b}_{n-1}(x)f(x) + \tilde{b}_n(x)f(2a_n - x), & \text{if } a_n^- \le x \le a_n \\ \tilde{b}_n(x)f(x) - \tilde{b}_{n-1}(x)f(2a_n - x), & \text{if } a_n \le x \le a_n^+. \end{cases} \quad (15)$$

III. Choice of the Bell Function

In this section we describe several "optimized" bells that were constructed recently [2]–[4]. These bells require that $a_n^+ = a_{n-1}^-$. In the case where the bell is symmetric this condition implies that $a_n^+ = (a_n + a_{n+1})/2$, and that all intervals need to have the same size. For the rest of the paper, one will assume that

$$\forall n, \quad a_{n+1} - a_n = l, \quad a_{n+1}^- = a_n^+ = (a_n + a_{n+1})/2. \quad (16)$$

All bells b_n are obtained from a prototype bell b by translation by a_n, and by dilation by l

$$b_n(x) = b\left(\frac{x - a_n}{l}\right). \quad (17)$$

The prototype bell b is defined on $[-1/2, 3/2]$ (see Fig. 2). In this work, our application is image compression, and a good model for images is provided by piece-wise polynomial functions (where edges appear at the junction of two polynomials). Our aim should be to obtain good approximation of polynomials in the basis $\psi_{n,k}$. One can first consider the simpler

problem of reproducing the constant function $p^0(x) \equiv 1$ with exactly one coefficient $p^0_{k_0}$

$$\forall x \in \mathbb{R}, \quad p^0(x) = \sum_n p^0_{n,k_0} \tilde{\psi}_{n,k_0}(x) \tag{18}$$

with

$$p^0_{n,k} = \int p^0(x) \psi_{n,k}(x)\, dx = p^0_{k_0}\, \delta_{k,k_0}. \tag{19}$$

It is natural (but not necessary) to choose $k_0 = 0$. Then, the bell

$$b(x) = \sin \frac{\pi}{2}(x + 1/2) \tag{20}$$

satisfies the condition (19) and results in an orthonormal basis. Unfortunately, this bell is not differentiable at $x = -1/2$ and 3/2, and therefore the Fourier transform, $\hat{b}(\xi)$, of b will decay very slowly. Because the decay of the bell in the frequency domain is directly related to the decay of the coefficients of the expansion of sine functions, it is critical to have the fastest decay of $\hat{b}(\xi)$. Incidentally, $\sin \pi/2\,(x + 1/2)$ is the only symmetric (with respect to 1/2) bell that satisfies (19). Wickerhauser [9] made $b(x)$ in (20) smoother by "flattening" it at both end points. For $x \in [-1/2, 1/2]$, and $s \in \mathbb{N}$, he defines

$$b^s(x) = \sin \frac{\pi}{2}(x_s + 1/2), \quad x_s = \frac{1}{2}\sin(\pi x_{s-1}), \quad x_0 = x. \tag{21}$$

This bell is symmetric: if $x \in [1/2, 3/2]$ we define $b^s(x) = b^s(1 - x)$ (see Fig. 3). One can show by induction that b^s has $2^s - 1$ vanishing derivatives at $-1/2$ and 3/2, and thus $b^s \in C^{2^s-1}$. This bell gives rise to an orthonormal basis. The magnitude of the Fourier transform of the basis function $\psi_{n,16}$ (for an interval of $N = 512$ samples) is shown in Fig. 3. As s increases the main lobe becomes slightly wider, but the side lobes become much smaller. Unfortunately, as soon as $s > 0$, one can no longer reproduce constant functions with b^s. Two avenues can be pursued to try to fix this problem. Both avenues consist in constructing biorthogonal bases which leave us with more freedom to design the bell b than the orthogonal setting. The first approach consists in replacing the condition (19) with an optimization problem: one should try to minimize the number of coefficients $p^0_{n,k}$ on each interval $[a_n, a_{n+1})$. This approach was followed by Matviyenko [4], and is described in the next section. The other approach consists in relaxing the symmetry condition with respect to 1/2, and consider asymmetric bells. In this case, Bittner showed that one can find bells that satisfy (19), and that generate biorthogonal bases. In [2], Bittner also constructed bells that could reproduce linear functions with exactly three coefficients. These results are reviewed in Section III-B.

A. Optimized Bell of Matviyenko

We consider the approximation of p^0 over the interval $[a_n, a_{n+1})$ generated from the first K coefficients $p^0_{n,k}$, $k = 0, \ldots, K-1$

$$\sum_{k=0}^{K-1} p^0_{n,k} \tilde{\psi}_{n,k} \quad \text{with } p^0_{n,k} = \int p^0(x) \psi_{n,k}(x)\, dx. \tag{22}$$

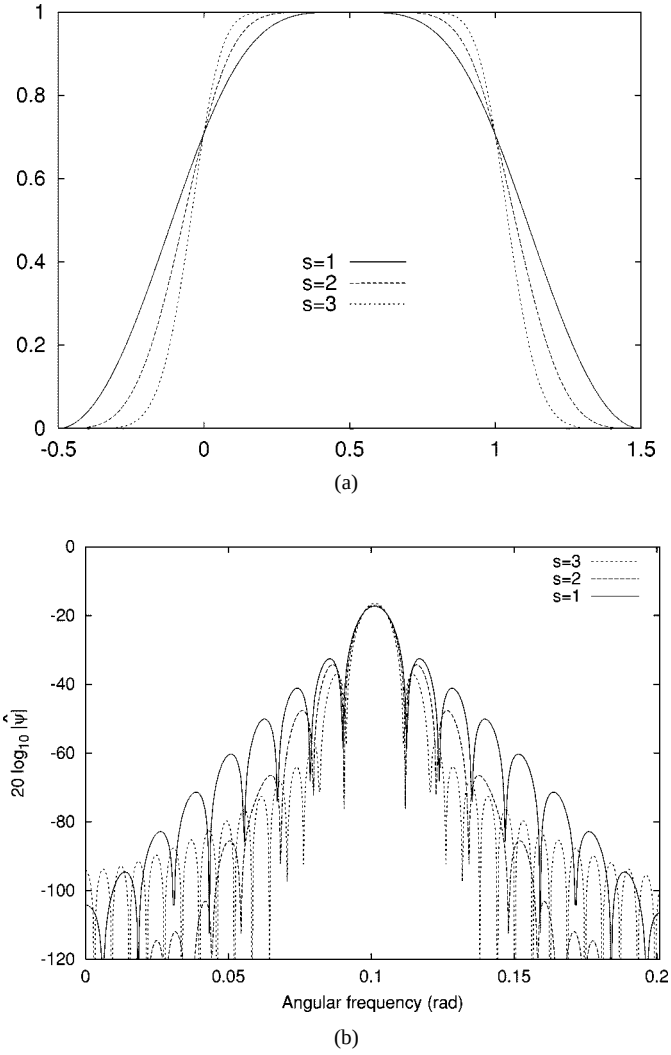

Fig. 3. (a) Orthonormal bells, $s = 1, 2, 3$. (b) Magnitude of the Fourier transform of $\psi_{n,16}$, (512 samples).

The norm of the residual error is

$$\left\| \sum_{k=K}^{\infty} p^0_{n,k} \tilde{\psi}_{n,k} \right\|_2^2. \tag{23}$$

Instead of reproducing exactly p^0 with one coefficient, Gregory Matviyenko designed a family of bells that minimize the residual error (23). Matviyenko shows that minimizing the sum

$$\sum_{k=K}^{\infty} |p^0_{n,k}|^2 \tag{24}$$

is related to minimizing the residual error (23). He then finds the bell b that minimizes (24) under the constraint

$$b(x) + b(-x) = 1 \qquad \text{for all } x \in [0, 1/2]. \tag{25}$$

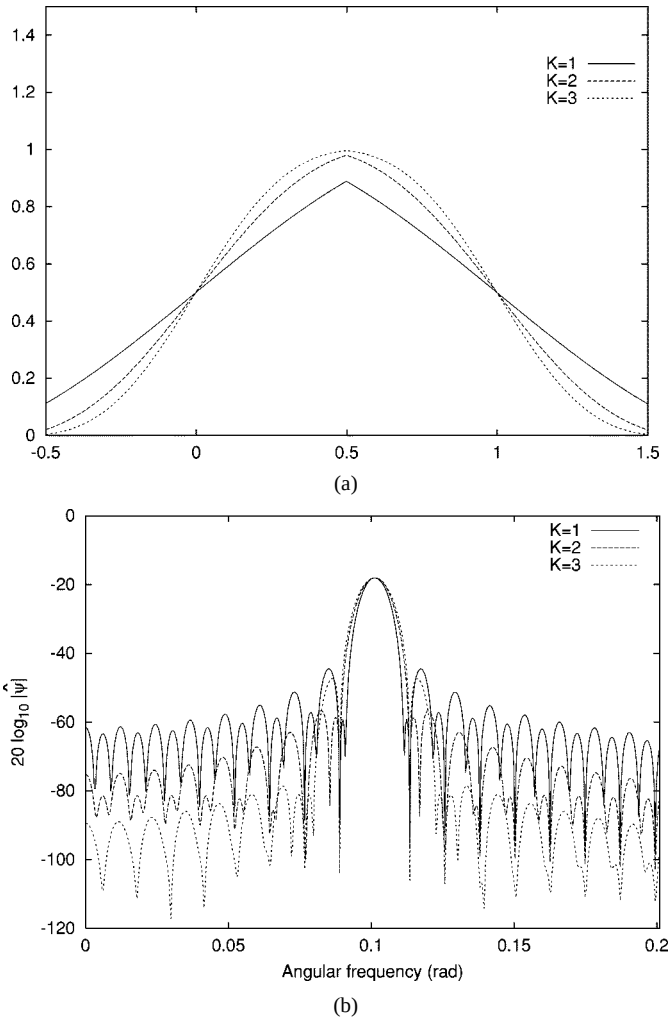

Fig. 4. (a) Matviyenko's optimized bell b^K and (b) magnitude of the Fourier transform of $\psi_{n,16}$.

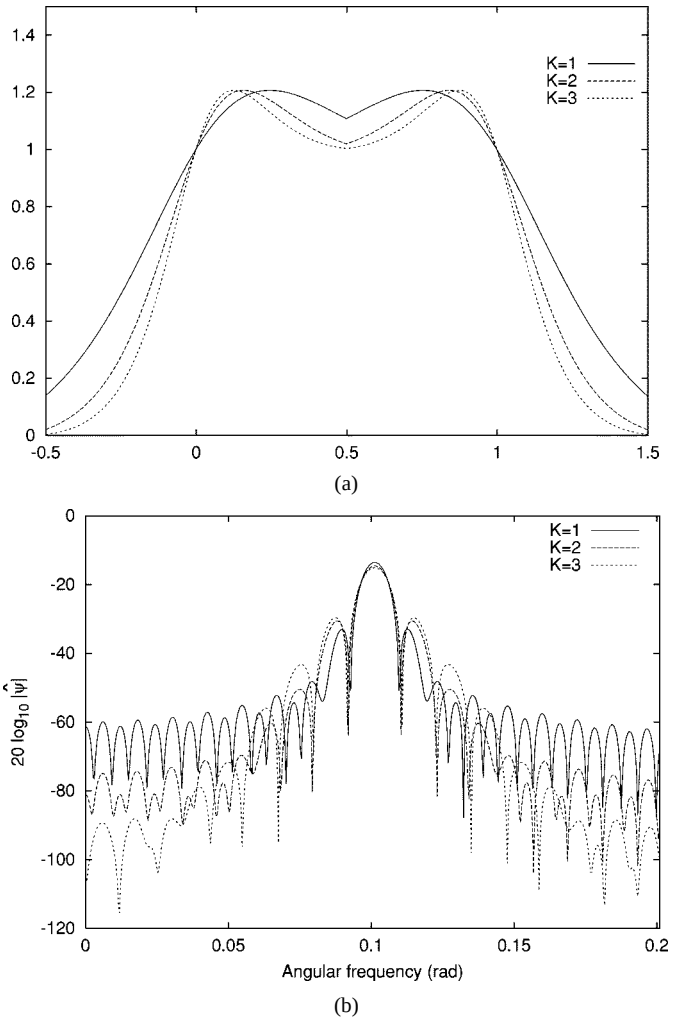

Fig. 5. (a) Dual Matviyenko's bell $\tilde{b}^K$ and (b) magnitude of the Fourier transform of $\tilde{\psi}_{n,16}$.

The solution of the optimization problem is a bell $b^K(x)$ given by

$$\begin{cases} \frac{1}{2}\left(1+\sum_{k=0}^{K-1} g_k \sin\left(k+\frac{1}{2}\right)\pi x\right), & \text{if } x\in\left[-\frac{1}{2},\,\frac{1}{2}\right) \\ \frac{1}{2}\left(1+\sum_{k=0}^{K-1} (-1)^k g_k \cos\left(k+\frac{1}{2}\right)\pi x\right), & \text{if } x\in\left[\frac{1}{2},\,\frac{3}{2}\right) \\ 0, & \text{otherwise.} \end{cases} \tag{26}$$

The g_k are calculated numerically in [4]. K influences the steepness of the bell. All bells b^K are bounded by one, and the dual bells $\tilde{b}^K$ are bounded by $(\sqrt{2}+1)/2$. These bounds guarantee that the Riesz bounds will be $A=1$ and $B=2$ for all K. Matviyenko argues in [4] that this family of bells should yield a sparse representation of oscillatory signals of the form $c(x)=\cos(\omega x+\varphi)$. Fig. 4 shows the bell b^K and the magnitude of the Fourier transform of the basis function $\psi_{n,16}$ (for an interval of $N=512$ samples). Fig. 5 shows the dual bell $\tilde{b}^K$ and the magnitude of the Fourier transform of the basis function $\tilde{\psi}_{n,16}$ (for an interval of $N=512$ samples). As K increases the side lobes become much smaller. This observation, and equation (24) seems to indicate that a large K should provide a smaller error, and a better frequency resolution. In practice, as shown in the experiments, small values of K often provide better compression performances.

As stated in lemma 1, one can either choose to use $\psi_{n,k}$ (and therefore b^K) or $\tilde{\psi}_{n,k}$ (and $\tilde{b}^K$) to compute the coefficients. Because b^K is optimized to minimize the residual error, one should use b^K for the analysis and compute the coefficients with

$$f_{n,k}=\int \psi_{n,k}(x)f(x)\,dx \tag{27}$$

and then use $\tilde{b}^K$ for the synthesis and reconstruct the image using

$$\sum_{n,k} f_{n,k}\tilde{\psi}_{n,k}. \tag{28}$$

Our experiments confirmed that this choice resulted in the largest PSNR. One potential drawback is that $\tilde{\psi}_{n,k}$ may have a very irregular shape. In such a case, reconstructing the compressed image from a small number of $\tilde{\psi}_{n,k}$ would result in unpleasant visual artifacts. This issue was never encountered with any of the bells described in this work.

B. Optimized Bells of Bittner

1) Reproduction of Constant Functions: Kai Bittner keeps the problem (19), but considers nonsymmetric bells. Starting with the bell (20), he adds two different terms on $[-1/2, 1/2]$ and on $[1/2, 3/2]$ in order to obtain vanishing derivatives at $-1/2$ and $3/2$. In [2], Bittner constructs the following bell $b(x)$:

$$\sin\frac{\pi}{2}\left(x+\tfrac{1}{2}\right)+\begin{cases}\frac{1}{4}\sin 2\pi x, & \text{if } -\dfrac{1}{2}<x<\frac{1}{2}\\ -\frac{1}{2}\cos\pi(x-1), & \text{if } \frac{1}{2}<x<\frac{3}{2}.\end{cases} \tag{29}$$

This bell satisfies (19). The dual bell $\tilde{b}(x)$ can be computed from (5). One can easily check that the bells b, and $\tilde{b}$ are now in $C^2(\mathbb{R})$. Unfortunately, the Riesz bounds for the basis $\psi_{n,k}$ are mediocre: $A = 0.704$ and $B = 11.65$.

2) Reproduction of Linear Functions: It is natural to ask ourselves if we could improve upon the previous result and try to reproduce exactly polynomials of degree $d \geq 1$ with a fixed number, K, of coefficients on each interval $[a_n, a_{n+1})$. As is shown by Bittner in [2], it turns out that this is only possible for polynomials of degree $d = 1$. We describe in the following the optimal bell that allows to reproduce exactly linear functions. Let $p^1(x) \equiv x$, and let $K \geq 1$. Our goal is to find a bell b such that the function p^1 can be reproduced with at most K coefficients

$$p^1_{n,k} = \langle \psi_{n,k}, p^1 \rangle = 0 \qquad \text{for } k \geq K. \tag{30}$$

In [2], Bittner shows that there exists one unique solution to (30) and (19) that is in C^{K-1}, with $\operatorname{supp}(b) \subset [-1/2, 3/2]$.

Furthermore, the family $\{\psi_{n,k}\}$ generated by the bell b forms a Riesz basis if and only if K is odd. Unfortunately, the Riesz bound B increases very rapidly with K. This makes the bell of little practical use for compression as soon as $K \geq 5$. In our experiments we chose $K = 3$, and the bell $b(x)$ is given by

$$\begin{cases}\dfrac{1}{\sqrt{2}}\cos\dfrac{\pi}{2}x \\ \quad +\displaystyle\sum_{k=0}^{2} a_k \frac{\cos(k+1/2)\pi x}{2x}, & \text{if } x\in\left[-\dfrac{1}{2},\dfrac{1}{2}\right),\\ \dfrac{1}{\sqrt{2}}\dfrac{x-2}{x-1}\cos\dfrac{\pi}{2}x \\ \quad +\displaystyle\sum_{k=0}^{2} a_k \frac{\cos(k+1/2)\pi x}{2(x-1)}, & \text{if } x\in\left[\dfrac{1}{2},\dfrac{3}{2}\right),\\ 0, & \text{otherwise}\end{cases} \tag{31}$$

with

$$a_0 = \frac{1}{2\sqrt{2}}, \quad a_1 = -\frac{2+\pi}{4\sqrt{2}\pi}, \quad a_2 = \frac{2-\pi}{4\sqrt{2}\pi}. \tag{32}$$

The dual bell $\tilde{b}$ is computed using equation (5). The Riesz bounds can be computed numerically and are equal to $A = 0.742$ and $B = 3.067$. In the remainder of the paper we will call *Bittner's bells* the bells b and $\tilde{b}$ that are optimized for linear functions, and are defined in (31). Fig. 6 (top) shows the bell b and the dual bell $\tilde{b}$. The magnitude of the Fourier transform of $\psi_{n,16}$ and $\tilde{\psi}_{n,16}$ (for an interval of $N = 512$ samples) is shown in Fig. 6 (bottom). As explained before one should use b and therefore $\psi_{n,k}$ to compute the coefficients $f_{n,k}$, and reconstruct the image from the compressed coefficients using $\tilde{\psi}_{n,k}$.

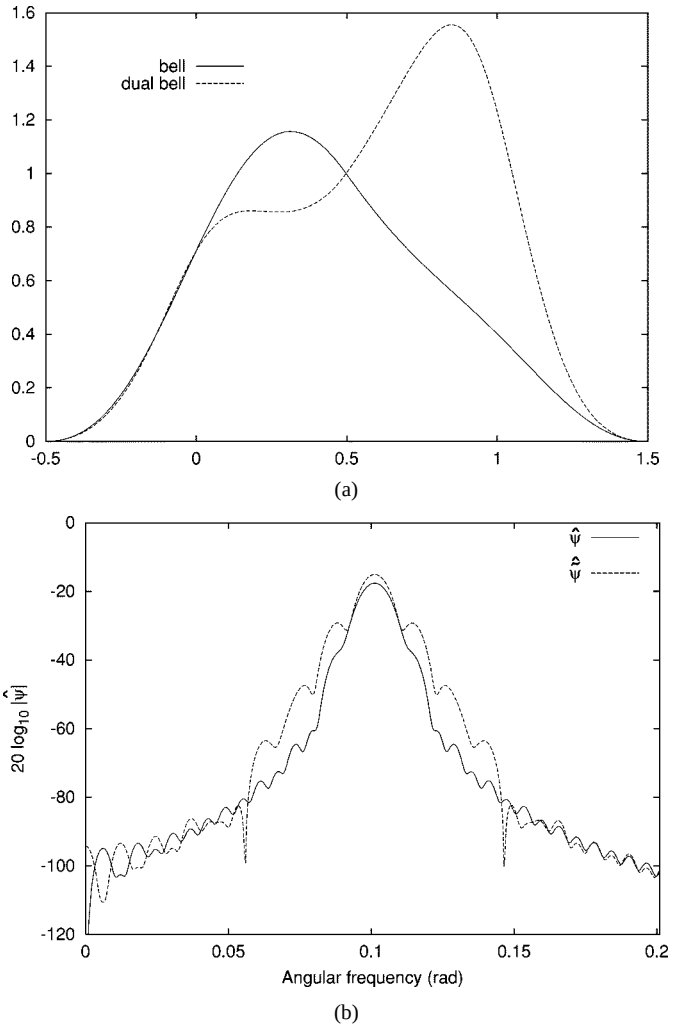

Fig. 6. Bittner's bell optimized for linear functions: (a) graphs of b and $\tilde{b}$ (b) magnitude of the Fourier transform of $\psi_{n,16}$ and $\tilde{\psi}_{n,16}$.

C. Modulated Lapped Biorthogonal Transform (MLBT)

A simple way to smooth $\sin\pi/2\,(x+1/2)$ at both end-points is to take the square of the bell

$$\left[\sin\frac{\pi}{2}(x+1/2)\right]^2 = \frac{1-\cos\pi(x+1/2)}{2}. \tag{33}$$

As a result the following bell is in $C^1(\mathbb{R})$:

$$b(x) = \begin{cases}\dfrac{1-\cos\pi(x+1/2)}{2}, & \text{if } x\in\left[-\dfrac{1}{2},\dfrac{1}{2}\right]\\ b(x) = b(1-x), & \text{if } x\in\left[\dfrac{1}{2},\dfrac{3}{2}\right].\end{cases} \tag{34}$$

Henrique Malvar [3] proposed the bell $b(x)$ defined by

$$b(x) = \begin{cases}\dfrac{1-\cos[\pi(x+1/2)^\alpha]+\beta}{2+\beta}, & \text{if } x\in\left[-\dfrac{1}{2},\dfrac{1}{2}\right]\\ b(x) = b(1-x), & \text{if } x\in\left[\dfrac{1}{2},\dfrac{3}{2}\right].\end{cases} \tag{35}$$

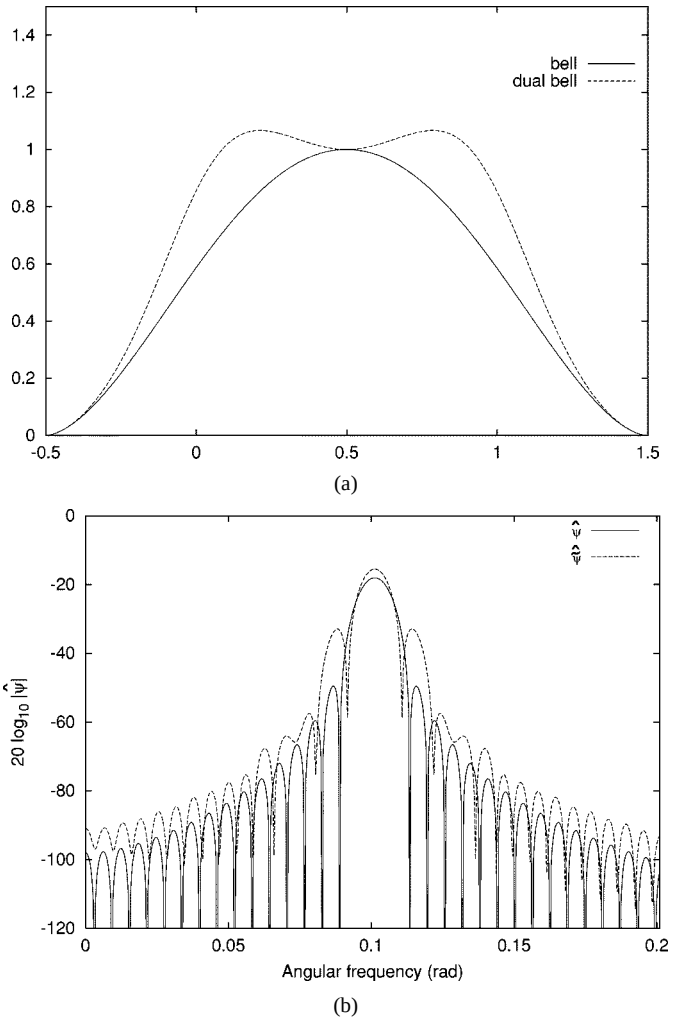

Fig. 7. (a) MLBT's bell b and the dual bell $\tilde{b}$ (b) magnitude of the Fourier transform of $\psi_{n,16}$ and $\tilde{\psi}_{n,16}$.

For $\alpha = 1$ and $\beta = 0$ we clearly find the square of $\sin(\pi/2)(x + 1/2)$. The bell is in $C^1(\mathbb{R})$ if $\alpha \geq 1$. Because the bell is smoother, it will have a faster decay in the Fourier domain, and a better frequency selectivity of the associated basis functions [3]. As usual, the dual bell is defined by (5). In this paper, we use the following values of the parameters: $\alpha = 0.85$, and $\beta = 0$. Fig. 7 (top) shows the graphs of b and $\tilde{b}$. The two bells are similar to the optimized bells of Matviyenko. The Riesz bounds are $A = 1$ and $B = 1.458$. The magnitude of the Fourier transform of the analysis basis function $\psi_{n,16}$ and the synthesis basis function $\tilde{\psi}_{n,16}$ are shown in Fig. 7 (bottom) (for an interval of $N = 512$ samples). The Fourier transform of $\psi_{n,16}$ has a wider main lobe than the Fourier transform of $\tilde{\psi}_{n,16}$ but a better stop-band attenuation (smaller side lobes).

In [3], Malvar uses b as the synthesis bell, and $\tilde{b}$ as the analysis bell. This choice is motivated by the fact that the main application in [3] is audio compression, and, therefore, one needs to be able to reconstruct a signal using basis functions that have the best frequency localization. In our experiments we noticed that we obtained better compression is b was used as the analysis bell, and $\tilde{b}$ was used as the synthesis bell.

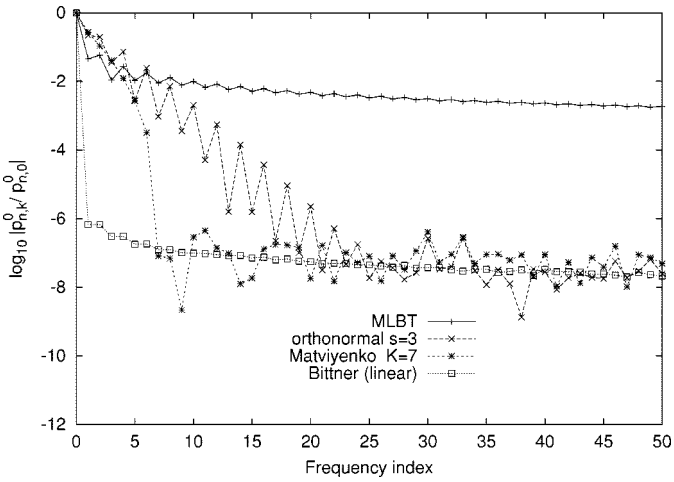

Fig. 8. Coefficients of $p^0(x) \equiv 1$, as a function of k.

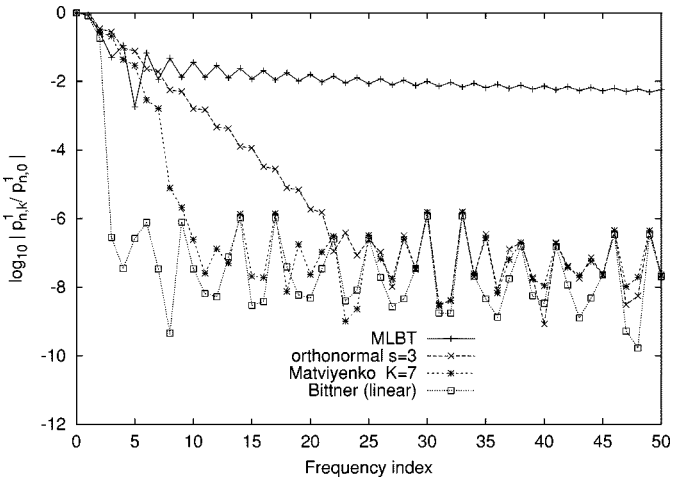

Fig. 9. Coefficients of $p^1(x) \equiv x$, as a function of k.

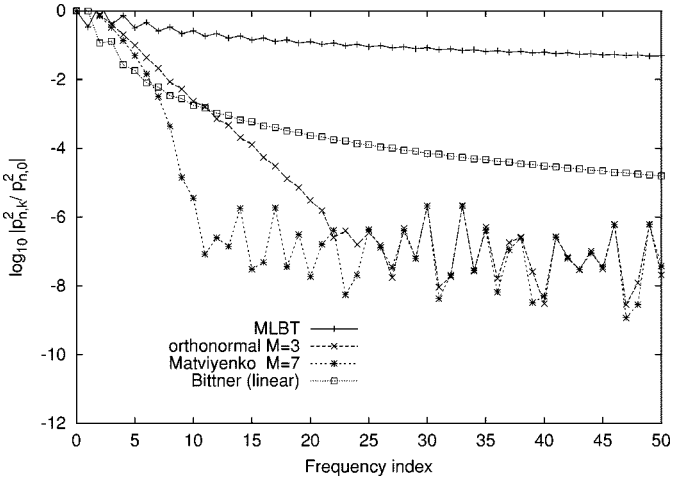

Fig. 10. Coefficients of $p^2(x) \equiv x^2$, as a function of k.

IV. Comparison of the Optimized Bells

A. Reproduction of Low Degree Polynomials

We first compare the rate of decay of the coefficients $p^0_{n,k}$ of the expansion of the constant function $p^0 \equiv 1$, over a fixed

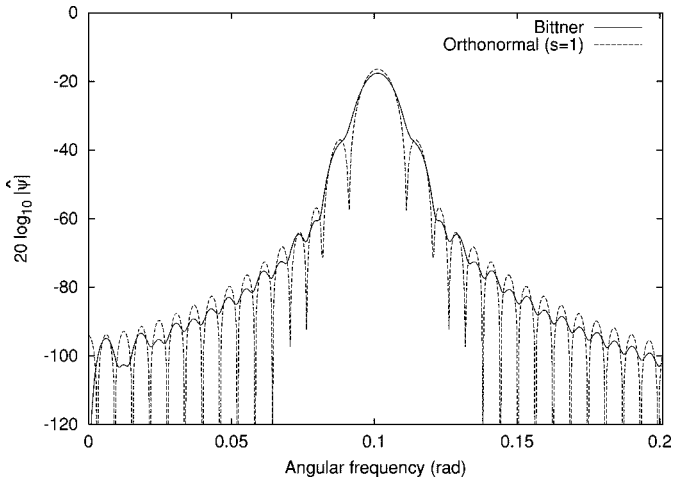

Fig. 11. Magnitude of the Fourier transform of $\psi_{n,16}$ for Bittner's bell, and the orthonormal bell (512 samples).

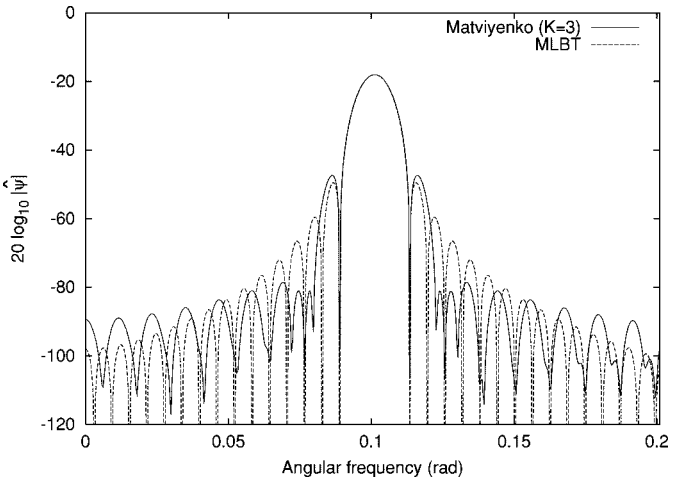

Fig. 12. Magnitude of the Fourier transform of $\psi_{n,16}$ for Matviyenko's bell ($K = 3$), and the MLBT bell (512 samples).

interval $[a_n, a_{n+1}]$. We expect to have only one nonzero coefficient for the optimized bell of Bittner. We also expect a fast decay of the $p^0_{n,k}$ for the bell of Matviyenko. Fig. 8 displays the logarithm of the magnitude of the coefficients, normalized by the DC coefficient

$$\log_{10}\left|\frac{p^0_{n,k}}{p^0_{n,0}}\right|, \qquad k = 0, 1, \ldots, 49 \tag{36}$$

as a function of the frequency k for the first 50 (out of 1024) coefficients. Our experimental findings confirm the theoretical results. We also note that the MLBT performed very poorly. We then compare the rate of decay of the coefficients $p^1_{n,k}$ of the expansion of the function $p^1(x) \equiv x$. As announced three terms only are necessary to reproduce x if one uses the optimized bell of Bittner (see Fig. 9). The bell of Matviyenko performed reasonably well (exponential decay), and the MLBT bell performed very poorly. Finally, we compared the decay of the coefficients $p^2_{n,k}$ of the expansion of the function $p^2(x) \equiv x^2$. While none of the bells were optimized for a quadratic function, Bittner's bell still provided a reasonably good decay (see Fig. 10).

B. Frequency Response

We then compared the Fourier transforms of the basis function $\psi_{n,16}$, and $\tilde{\psi}_{n,16}$ with the four afore-mentioned bells. The Fourier transform tells us what the rate of decay of the coefficients of a sine function should be. Fig. 11 shows the magnitude of the Fourier transform of $\psi_{n,16}$ for Bittner's bell, and the orthonormal bell. Is clear that Bittner's bell results in a Fourier transform $\hat{\psi}_{n,16}(\xi)$ with a very wide main lobe. Fig. 12 shows the magnitude of the Fourier transform of $\psi_{n,16}$ for Matviyenko's bell ($K = 3$), and the MLBT. Both bells yield remarkably similar main lobe. While the first side lobe of the MLBT appears to be slightly smaller, the other side lobes are larger than the ones generated with Matviyenko's bell. In general, the Fourier transform of the dual basis function $\tilde{\psi}_{n,16}$ is much broader than the Fourier transform of $\psi_{n,16}$ (see Figs. 13 and 14). As shown in Fig. 14, the dual MLBT's bell has a much better stop band attenuation than the dual bell of Matviyenko.

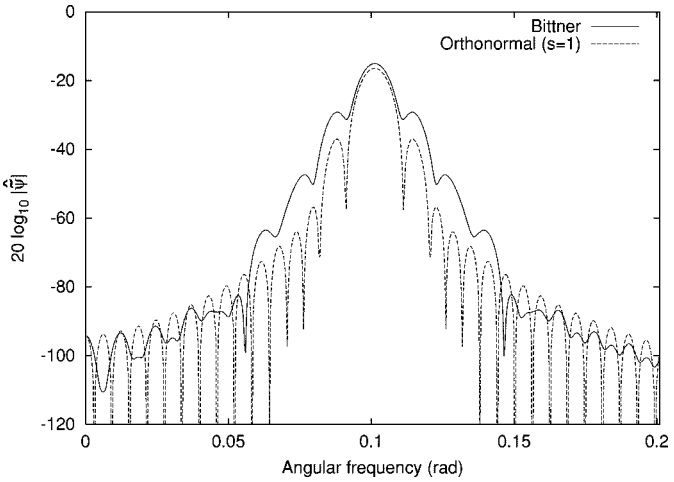

Fig. 13. Magnitude of the Fourier transform of $\tilde{\psi}_{n,16}$ for Bittner's bell, and the orthonormal bell (512 samples).

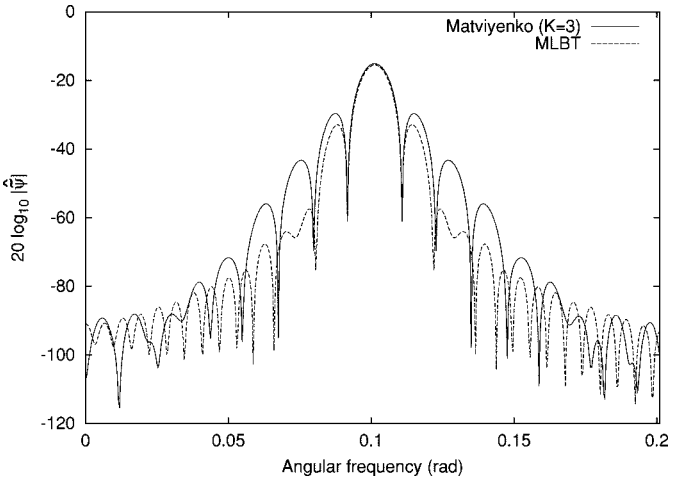

Fig. 14. Magnitude of the Fourier transform of $\tilde{\psi}_{n,16}$ for Matviyenko's bell, and the MLBT bell (512 samples).

V. Compression Algorithm

The compression algorithm is divided into three parts. In the first part, we select that local cosine basis which is best adapted

to encode the image. The actual optimality of the basis depends to a large extent on the design of the cost function [12]. Our cost function computes an estimate of the actual cost of coding the coefficients for a given quantization level. The cost is therefore closely related to the actual rate R, given a distortion D. During the second part, the coefficients are quantized; we assume a Laplacian distribution and we use an efficient near optimal scalar quantizer [13]. Finally, the significance map is entropy coded using a higher order arithmetic coder [14] that relies on a context consisting of pixels in a causal neighborhood.

A. Adaptive Segmentation

In two dimensions, we consider the separable tensor products of the bases $\psi_{n,k}$, and $\psi_{m,j}$. We adaptively select the size and location of the windows with the best basis algorithm [15]: any cover of the image with local bases associated with a collection of sub-squares selected at several, possibly different, levels result in a valid basis. One could replace the squares with rectangles to obtain an anisotropic segmentation [16]. A fast dynamic programming algorithm, with order $N\log(N)$ (N is the number of pixels in the image), searches for that *best basis* [15] which is optimal according to a given cost function $\mathcal{M}$.

B. Choice of a Cost Function $\mathcal{M}$

The metric defined by the cost function $\mathcal{M}$ defines the optimality criterion. In this work, one basis is better than another if it provides a better reconstruction quality for the same number of bits spent in coding the coefficients.

Ramchandran and Vetterli [17] wedded the bit allocation algorithm of Shoham and Gersho [18] to the best basis algorithm [15]. Unfortunately their approach is extremely computationally intensive: the problem in [17] involves three layers of nonlinear approximations, only one of which lends itself to a fast algorithm). Instead of using the rate distortion framework, we designed a cost function that returns an estimate of the actual rate achieved by each node. The cost function mimics the actual scalar quantization, and entropy coding, which are presented in Section V-D. However, the cost function is much faster to compute. It is composed of two complementary terms:

- $c_1(\mathbf{x})$ is the cost of coding the sign and the magnitude of the nonzero output levels of the scalar quantizer;
- $c_2(\mathbf{x})$ is the cost of coding the locations of the nonzero output levels (significance map).

Let $\mathbf{x} = \{x_k\}$. A first order approximation of the cost of coding the magnitude of the output levels $\{|Q(x_k)|\}$ is given by the number of bits needed to represent the set $\{|Q(x_k)|,\, k/Q(x_k) \neq 0\}$

$$c_1(\mathbf{x}) = \sum_{k/Q(x_k)\neq 0} \max(\log_2 |Q(x_k)|,\, 0). \tag{37}$$

The second term, $c_2(\mathbf{x})$, is calculated using the first order entropy of a Bernoulli process: each coefficient x_k is significant with a probability p, and we assume that the significance of the coefficients are independent events. We get

$$c_2(\mathbf{x}) = -N\left(p\log_2(p) + (1-p)\log_2(1-p)\right). \tag{38}$$

The computation of the cost function requires to quantize the coefficients. A first estimate of the quantization step is required to compute the cost function. A second pass uses the actual quantization step obtained after quantization.

C. Boundary Extensions and Finite Image Size

In order to be able to apply the folding operator defined by (14) around the boundaries of the image, one needs to extend the image across the boundaries. A periodic extension gives rise to large coefficients (due to discontinuities) and results in poor compression performances. We describe the 1-D extension that is performed along the rows and columns.

We consider a finite interval $[a_0,\, a_N]$. Folding around the border point a_0 requires the knowledge of the image intensity $I(x)$ inside the interval $[a_0^-,\, a_0+]$. One needs to extend $I(x)$ in the intervals $[a_0^-,\, a_0)$, and $(a_N,\, a_N^+]$. In a similar way, the inverse of the folding operator, $U_{a_0}^*$, requires the knowledge of the values of $U_{a_0}I(x)$, for $x \in [a_0^-,\, a_0)$. One could include these values with the rest of the folded samples from the original image, and compress them. As the size of the bell increases this becomes a major overhead. Instead, we prefer to construct the extension of I on $[a_0^-,\, a_0)$ in such a way that $U_{a_0}^*$ can be directly calculated from the knowledge of $U_{a_0}I(x)$ on $[a_0,\, a_0^+)$ only—without knowing $U_{a_0}I(x)$ on $(a_0^-,\, a_0]$. This can be achieved with an even extension of $I(x)$ around a_0. One defines

$$I(x) = I(2a_0 - x) \qquad \text{if } x \in [a_0^-,\, a_0). \tag{39}$$

One can easily verify that with this extension one has:

$$U_{a_0}I(x) = (b(x-a_0) + b(a_0-x))I(x), \qquad \text{if } a_0 \leq x < a_0^+ \tag{40}$$

and therefore the knowledge of $UI(x)$ in $[a_0,\, a_0^+)$ is actually sufficient to recover $I(x)$ in that interval. Similarly, one performs an odd extension of $I(x)$ around a_N and defines

$$I(x) = -I(2a_N - x) \quad \text{if } x \in (a_N, a_N^+]. \tag{41}$$

Again, one clearly has

$$U_{a_N}I(x) = (b(x-a_N) + b(a_N-x))I(x), \quad x \in [a_N^-,\, a_N) \tag{42}$$

and one does not need $UI(x)$ on $(a_N,\, a_N^+]$ to recover $I(x)$.

D. Laplacian Based Scalar Quantization

The distribution of the cosine coefficients is approximated with a Laplacian distribution. The Laplacian distribution yields tractable computations of the optimal entropy constrained scalar quantizers [13]. We use a particularly efficient near optimal scalar quantizer, with a symmetric dead-zone, and a reconstruction offset [13].

E. Ordering of the Coefficients

In order to exploit large scale correlations that may exist between adjacent blocks, we gather together coefficients with similar two-dimensional (2-D) frequencies. One first divides each LCT block into a fixed number of frequency subsets: in each

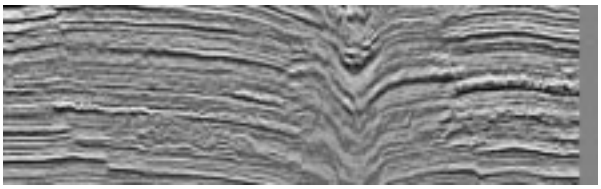

(a)

(b)

(c)

Fig. 15. (a) Seismic image. (b) From left: wood grain, straw, roofs. (c) From left: fingerprint, Barbara, and Clown.

subset the coefficients have similar 2-D frequencies. We then gather from all the LCT blocks all the coefficients that are in the same subset. We start with the subset with the smallest frequency, continue until we scan the subset with the highest frequency [we use a zigzag order to navigate in the frequency plane from the (0, 0) frequency until the largest frequency]. The signs of the output levels are not entropy coded, but are simply packed. The magnitude of the output levels are variable length encoded, using an arithmetic coder to encode the length. The best basis geometry is described by a quadtree. We code the quadtree, with an adaptive arithmetic coder.

VI. Experiments

A. The Images

The following images have been used for the experiments (see Fig. 15).

1) **Seismic image**, 32 bpp, 192 × 640. This image contains rapid oscillations, as well as discontinuities.
2) **Wood grain**, 8 bpp, 512 × 512. This image is part of Brodatz's book.
3) **Straw**, 8 bpp, 512 × 512. This image is also part of Brodatz's book. It is extremely difficult to code: the intensity, as well as the orientation of the patterns are varying very rapidly.
4) **Roofs**, 8 bpp, 512 × 512. This image is part of the MIT VisTex database. It is composed of a mixture of periodic texture (roofs), as well as smooth regions (façades and sky).
5) **Fingerprint**, 8 bpp, 512 × 512. This image contains oscillatory patterns in several different orientations.

TABLE I
Seismic Data: PSNR (in Decibels) for Various Bit Rates

Rate (bpp)	0.40	0.50	0.67	1.00	2.00
Orthonormal (s=1)	**19.09**	**19.90**	**21.02**	**22.94**	**27.72**
Matviyenko (K=12)	18.93	19.76	20.90	22.79	27.54
MLBT	19.06	19.87	21.00	22.91	27.68
Bittner	18.37	19.15	20.19	21.93	26.39
No bell	18.05	18.88	20.05	22.05	26.97

6) **Barbara**, 8 bpp, 512 × 512. It is the standard image of the lady with the stripes and checker tablecloth.
7) **Clown**, 8 bpp, 512 × 512. It is the standard image of the Clown.

B. Comparison With Fixed Window Size

We first report the results of experiments conducted with a fixed window size. Because different bells can result in different best bases, one needs to separate the effect of the bell from the effect of the basis selection. For the orthonormal bell and Matviyenko's bell we only report the results obtained with the optimal value of s and K. In addition to comparing the three bells discussed in the previous section, we also included the results obtained with the abrupt cutoff bell: the characteristic function of the interval $[a_n, a_{n+1}]$; the other parts of the algorithm being equal. Finally, we also compared our coder to one of the best wavelet coder that was available to us [19]: the SPIHT wavelet coder of Said and Pearlman [1].

Seismic: A fixed window size of 32 × 32 was used for all bells. Except for Bittner's bell, all bells performed equally well (see Table I), and resulted in a 1 dB improvement over an abrupt cutoff bell (no bell). Interestingly, the optimized bell of Matviyenko did not result in any significant improvement over the orthonormal bell, even though the image contains many oscillatory patterns. Each pixel in the seismic image is encoded as a 4 byte floating number. The version of SPIHT that was available to us [19] could not process images encoded as floating numbers. Because the range of the intensity is quite large, we could not map the floating number to one-byte integers.

Wood Grain: A fixed window size of 32 × 32 was used for all bells. Almost all bells performed equally well (see Table II), and resulted in a significant improvement of 1.5 to 2 dB over an abrupt cutoff bell. The LCT significantly outperformed SPIHT by 1.5 to 2.5 dB. As all coders, our coder introduces artifacts at low bit-rate. The wood grain image, compressed at 64 : 1 clearly illustrates the type of artifacts associated with our compression method: the main texture is very well preserved, but artifacts appear at block boundaries. These artifacts can be explained by observing that the expansion of the wood grain image creates two types of coefficients:

TABLE II
WOOD GRAIN: PSNR (IN DECIBELS) FOR VARIOUS BIT RATES

Rate (bpp)	0.125	0.25	0.50	0.75	1.00
Orthonormal (s=1)	18.52	21.07	**23.94**	**25.78**	**27.24**
Matviyenko (K=1)	**18.55**	21.05	23.87	25.69	27.12
MLBT	18.54	**21.08**	23.93	25.73	27.18
Bittner	17.99	20.35	23.08	24.79	26.14
No bell	17.07	19.30	21.94	23.72	25.17
SPIHT	16.49	18.57	21.49	23.71	25.78

TABLE III
ROOFS: PSNR (IN DECIBELS) FOR VARIOUS BIT RATES

Rate (bpp)	0.125	0.25	0.50	0.75	1.00
Orthonormal (s=1)	23.86	27.77	32.19	35.33	37.80
Matviyenko (K=1)	24.54	**27.93**	**32.32**	35.43	37.80
MLBT	**24.55**	27.84	32.32	**35.45**	**37.83**
Bittner	24.12	27.38	31.64	34.69	37.12
No bell	19.30	21.51	24.67	27.30	29.12
SPIHT	23.77	27.02	31.18	34.18	36.68

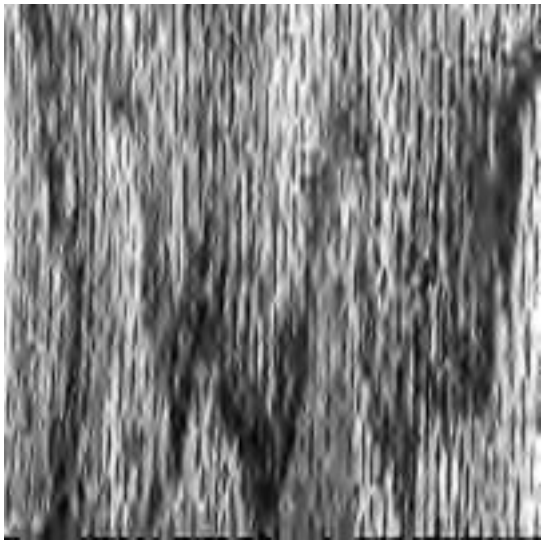

(a)

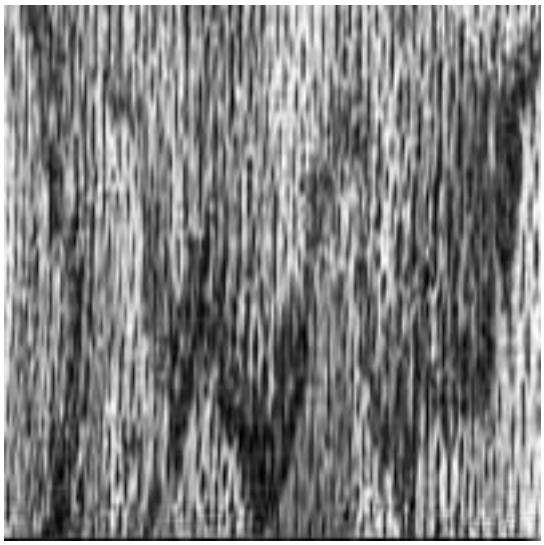

(b)

Fig. 16. Wood grain at 0.125 bpp. (a) SPIHT, PSNR = 16.49 dB. (b) Matviyenko's bell, PSNR = 18.55 dB.

Fig. 17. Roof at 0.125 bpp, MLBT bell, PSNR = 24.55 dB.

1) a small set of very large coefficients that describe the wood texture in the direction of the grain;
2) a set of small coefficients that encode the variations of the intensity created by the folding.

At low bit-rates, the small coefficients are quantized to zero, and the reconstructed image contains artifacts in the region where the folding occurs. Even at low bit-rates the large coefficients are preserved, and the main texture of the wood is very well rendered by our algorithm. As shown in Fig. 16, the wavelet coder could not preserve the texture and smeared the patterns of the wood. In fact, even without a smooth bell, our coder performs on this image as well as SPIHT.

Roof: A fixed window size of 16 $\times$ 16 was used for all bells. Because this image contains large regions with periodic texture such as the tiles on the roofs of the buildings, we expect the MLBT and Matviyenko's bell to perform well. In fact, these two bells only marginally outperformed the other bells (see Table III). As expected for this type of image, Bittner's bell and SPIHT performed poorly in terms of PSNR (see Table III). However, a visual inspection of Figs. 17 and 18 indicates that

(a)

(b)

Fig. 18. Roof at 0.125 bpp. (a) SPIHT, PSNR = 23.77 dB. (b) Bittner's bell, PSNR = 24.12 dB.

Bittner's bell is able to reconstruct the sky and the facades with very few visual artifacts. The MLBT bell (the best in terms of PSNR) leaves some very annoying artifacts (created by the folding) in the sky. Moreover, Fig. 18 shows that, as opposed to the wavelet coder, the LCT equipped with Bittner's bell could preserve the tiles on the roof in the background. A striking consequence of using an abrupt cutoff bell was a dramatic drop in the PSNR of 3.6 to 5.5 dB.

Straw: The straw image is quite difficult to compress. Neither the local cosine bases nor the wavelet basis are well adapted to the features present in this image: long lines.

A fixed window size of 16×16 was used for all bells. As expected, the PSNR degrades very rapidly as the rate decreases. Except for Bittner's bell, all bells performed in a similar way (see Table IV). Even for this image, using an abrupt cutoff bell has the effect of a drop of 1 dB in the PSNR. Our coder again outperformed SPIHT in terms of PSNR. As shown in Fig. 19 the wavelet coder was unable to represent the long edges created

TABLE IV
STRAW: PSNR (IN DECIBELS) FOR VARIOUS BIT RATES

Rate (bpp)	0.125	0.25	0.50	0.75	1.00
Orthonormal (s=1)	**14.29**	**16.25**	**18.94**	**20.92**	**22.57**
Matviyenko (K=1)	14.28	16.22	18.88	20.84	22.45
MLBT	14.25	16.20	18.89	20.87	22.51
Bittner	13.46	15.21	17.77	19.71	21.28
No bell	13.96	15.73	18.08	19.83	21.29
SPIHT	14.16	16.01	18.34	20.33	21.87

(a)

(b)

Fig. 19. Straw at 0.125 bpp. (a) SPIHT, PSNR = 14.16 dB. (b) Orthonormal bell, PSNR = 14.29 dB.

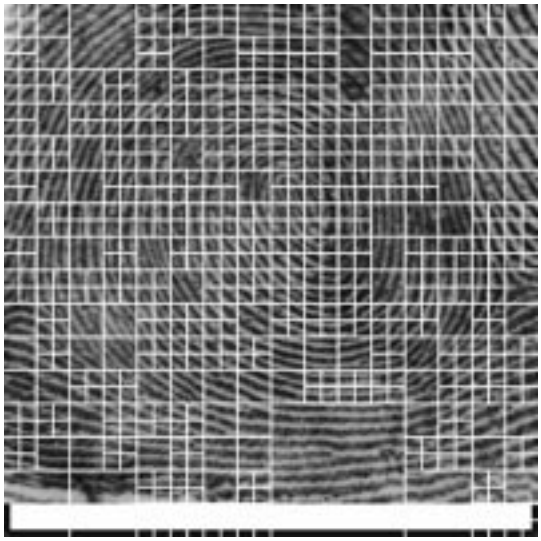

Fig. 20. Fingerprints at 0.125 bpp, MLBT bell, best basis.

Fig. 21. Fingerprints at 0.125 bpp, MLBT bell, PSNR = 24.28 dB.

by the straw. The LCT (see Fig. 19) provides a much crisper rendering of these edges.

C. Comparison With Adaptive Window Selection

The second set of experiments was designed to include the effect of the best basis selection. The best basis was selected using the cost function described in Section V-B. The smallest window size for the optimal segmentation was 16×16, while the largest window size was 512×512.

Fingerprints: As shown in Fig. 20, the best basis was able to select windows according to the curvature of the patterns: small boxes are needed to approximate circles with small radius, and large boxes could be selected in regions where the lines are almost parallel (and thus can be well approximated by a planar wave). The MLBT bell marginally outperformed Matviyenko's bell and the orthonormal bell (see Fig. 21 and Table V). Again, the wavelet coder did not perform well.

Barbara: We report here the results obtained with the image Barbara (see Table VI). All the bells performed equally well,

TABLE V
FINGERPRINTS: PSNR (IN DECIBELS) FOR VARIOUS BIT RATES

Rate (bpp)	0.125	0.25	0.50	0.75	1.00
Orthonormal (s=1)	24.28	27.64	31.78	34.71	37.12
Matviyenko (K=2)	24.11	27.56	31.68	34.60	37.01
MLBT	**24.28**	**27.66**	**31.82**	**34.75**	**37.17**
Bittner	23.86	27.16	31.23	33.96	36.33
No bell	23.98	27.21	31.09	34.17	36.41
SPIHT	23.98	27.12	31.27	33.88	36.01

TABLE VI
BARBARA: PSNR (IN DECIBELS) FOR VARIOUS BIT RATES

Rate (bpp)	0.125	0.25	0.50	0.75	1.00
Orthonormal (s=1)	25.75	28.71	32.43	35.00	37.04
Matviyenko (K=2)	25.76	28.82	**32.61**	35.04	37.03
MLBT	**25.87**	**28.87**	32.58	**35.10**	**37.11**
Bittner	25.31	28.18	31.71	34.18	36.16
No bell	24.78	27.01	29.71	32.19	33.89
SPIHT	24.86	27.58	31.39	34.25	36.41
JPEG [20]		26.7	30.6	33.6	35.9
STQ [21]		26.5	30.4	32.4	34.9
EZDCT [22]	24.51	27.26	31.10	33.93	36.21

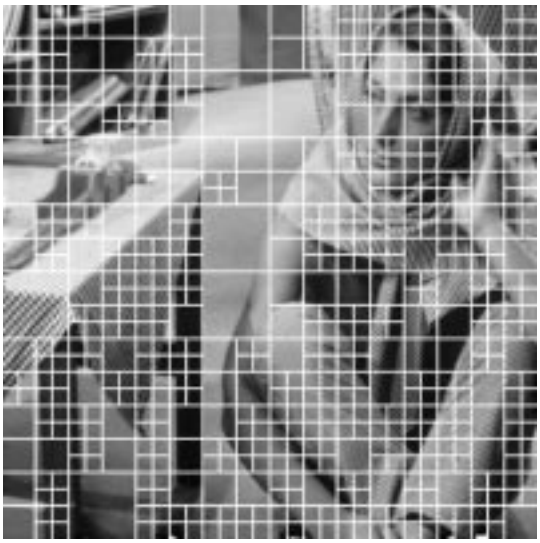

Fig. 22. Barbara at 0.25 bpp, MLBT bell, best basis.

Fig. 23. Barbara at 0.25 bpp, MLBT bell, PSNR = 28.87 dB.

TABLE VII
PSNR (IN DECIBELS) FOR VARIOUS BIT RATES

Rate (bpp)	0.125	0.25	0.50	0.75	1.00
Orthonormal (s=1)	27.16	30.51	34.15	36.56	38.44
Matviyenko (K=2)	27.35	30.95	34.99	37.45	39.23
MLBT	27.43	31.00	35.03	37.52	39.33
Bittner	27.22	30.65	34.54	36.99	38.65
No bell	26.84	29.55	32.81	35.07	36.89
SPIHT	**28.23**	**31.95**	**35.93**	**38.27**	**40.05**

and resulted in a gain of 1 to 3 dB over an abrupt cutoff bell. Our coder outperformed SPIHT by 0.7 to 1 dB. The best basis is shown in Fig. 22; the decoded Barbara at 0.125 is shown in Fig. 23. To further illustrate the gain obtained from using local cosine bases instead of the DCT, we compare in Table VI several DCT-based coders [20]–[22] that use very elaborate coding strategies. Despite its simple design our coder outperforms all these more complicated coders.

A "Wavelet Friendly" Image—Clown: Images that are composed of regions within which the intensity is smoothly varying can be very compactly represented with a wavelet transform. While the goal of this work is the coding of images that contain periodic textures, we wanted to benchmark our coder against a "wavelet-friendly" image. We report in Table VII the results obtained with the image "Clown." SPIHT only marginally outperformed our coder in terms of PSNR (improvement of 0.7 to 0.9 dB). The MLBT bell was the optimal bell, and resulted in a gain of 0.5 to 2.5 dB over an abrupt cutoff bell. Fig. 24 shows the decoded Clown at 0.125 bbp using SPIHT and the MLBT bell. The decoded SPIHT image is more blurred than the LCT image. In particular, the reflection

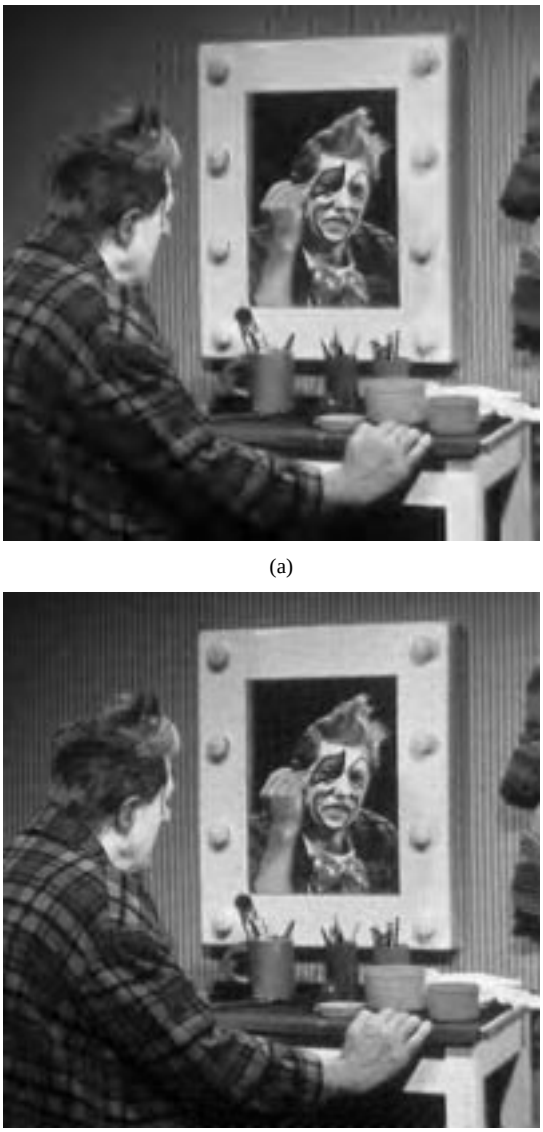

(a)

(b)

Fig. 24. Clown at 0.125 bpp. (a) SPIHT, PSNR = 28.23 dB. (b) MLBT bell, PSNR = 27.43 dB.

of the Clown in the mirror, as well as the right hand of the Clown have been smeared by the wavelet coder. SPIHT was also unable to preserve the texture of the wallpaper or the texture of the Clown's shirt. The pens and brushes in the cups below the mirror have been almost erased by the wavelet coder, and are much crisper with our coder. While some ringing artifacts are visible on the top of the mirror (where the intensity abruptly changes) in the LCT encoded image, similar artifacts are visible in the wavelet coded image. In summary, this experiment demonstrates that—while the wavelet based coder (SPIHT) outperforms our coder in terms of PSNR—a careful choice of the window allows us to obtain results that are better in terms of visual quality than a wavelet coder.

VII. CONCLUSION

The first contribution of this work is a new compression algorithm that is based on an adaptive local cosine transform. The algorithm was compared to one of the best wavelet coder

[1]. Our numerous experiments demonstrated that our algorithm could more efficiently encode textured images than wavelets. A critical component of the local cosine transform is the bell that is used to isolate blocks of pixels. We investigated the performances of several optimized bells [2]–[4]. These bells have all been designed in order to provide sparse representations of smooth or periodic signals. We first compared the bells from a theoretical point of view. We then performed a systematic comparison using our coder and a large number of images. Our experiments show that Bittner's bells, that are optimized for polynomials of degree 1, resulted in the worst performance in terms of PSNR. A visual inspection of the compressed images indicates however that Bittner's bells often provide reconstructed images with very few visual artifacts, even at low bit rates. The bell with the most narrow Fourier transform, the MLBT bell, gave the best results in terms of PSNR on most images. This bell tends however to create annoying visual artifacts in very smooth regions at low bit rate. We certainly expected these bells to perform very well with images that contain large patches of periodic textures, which can locally be modeled by a planar wave. Our coder used tensor products of 1-D cosine functions. In principle, coding a periodic texture (that behaves like a planar wave) with a tensor product of complex exponentials would result in four times fewer coefficients. We are currently investigating if this claim can be translated into an experimental reality.

REFERENCES

[1] A. Said and W. A. Pearlman, "A new fast and efficient image codec based on set partioning in hierarchical trees," *IEEE Trans. Circuits, Syst., Video Technol.*, vol. 6, pp. 243–250, June 1996.

[2] K. Bittner, "Error estimates and reproduction of polynomials for biorthogonal local trigonometric bases," *Appl. Comput. Harmon. Anal.*, vol. 6, pp. 75–102, 1999.

[3] H. S. Malvar, "Biorthogonal and nonuniform lapped transforms for transform coding," *IEEE Trans. Signal Processing*, vol. 46, pp. 1043–1053, Apr. 1998.

[4] G. Matviyenko, "Optimized local trigonometric bases," *Appl. Comput. Harmon. Anal.*, vol. 3, pp. 301–323, 1996.

[5] G. Aharoni, A. Averbuch, R. Coifman, and M. Israeli, "Local cosine transform—A method for the reduction of the blocking effect," *J. Math. Imag. Vis.*, vol. 3, pp. 7–38, 1993.

[6] C. K. Chui and X. Shi, "Characterization of biorthogonal cosine wavelets," *J. Fourier Anal. Applicat.*, vol. 3, no. 5, pp. 560–575, 1997.

[7] R. R. Coifman and Y. Meyer, "Remarques sur l'analyze de Fourier à fenêtre," *C. R. Acad. Sci. Paris I*, pp. 259–261, 1991.

[8] B. Jawerth and W. Sweldens, "Biorthogonal local trigonometric bases," *J. Fourier Anal. Applicat.*, vol. 2, no. 2, pp. 109–133, 1995.

[9] M. V. Wickerhauser, *Adapted Wavelet Analysis From Theory to Software*. New York: Peters, 1995.

[10] P. Auscher, G. Weiss, and M. V. Wickerhauser, "Local sine and cosine bases of Coifman and Meyer," in *Wavelets—A Tutorial*. New York: Academic, 1992, pp. 237–256.

[11] F. G. Meyer and R. R. Coifman, "Brushlets: A tool for directional image analysis and image compression," *Appl. Comput. Harmon. Anal.*, pp. 147–187, 1997.

[12] H. Šikić and M. V. Wickerhauser, "Information cost functions," *Appl. Comput. Harmon. Anal.*, vol. 11, pp. 147–166, 2001.

[13] G. J. Sullivan, "Efficient scalar quantization of exponential and Laplacian random variables," *IEEE Trans. Inform. Theory*, vol. 42, pp. 1365–1374, Sept. 1996.

[14] M. Nelson and J.-L. Gailly, *The Data Compression Book*. New York: M & T, 1996.

[15] R. R. Coifman and M. V. Wickerhauser, "Entropy-based algorithms for best basis selection," *IEEE Trans. Inform. Theory*, vol. 38, pp. 713–718, Mar. 1992.

[16] N. N. Bennett, "Fast algorithm for best anisotropic Walsh bases and relatives," *Appl. Comput. Harmon. Anal.*, vol. 8, no. 1, pp. 86–103, 2000.

[17] K. Ramchandran and M. Vetterli, "Best wavelet packet bases in a rate-distortion sense," *IEEE Trans. Image Processing*, pp. 160–175, Apr. 1993.

[18] Y. Shoham and A. Gersho, "Efficient bit allocation for an arbitrary set of quantizers," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 36, pp. 1445–1453, Sept. 1988.

[19] A. Said and W. A. Pearlman. Binaries for the codec based on set partioning in hierarchical trees. [Online]. Available: http://www.cipr.rpi.edu/research/SPIHT/.

[20] M. Crouse and K. Ramchandran, "Joint thresholding and quantizer selection for transform image coding," *IEEE Trans. Image Processing*, vol. 6, pp. 285–297, Feb. 1997.

[21] G. Davis and S. Chawla, "Image coding using optimized significance tree," in *IEEE Data Compression Conference—DCC'97*, 1997, pp. 387–396.

[22] Z. Xiong, K. Ramchandran, M. T. Orchard, and Y. Zhang, "A comparative study of DCT- and wavelet-based image coding," *IEEE Trans. Circuits, Syst., Video Technol.*, pp. 692–695, Aug. 1999.

François G. Meyer (M'94) received the M.S. degree (with honors) in computer science and applied mathematics from Ecole Nationale Superieure d'Informatique et de Mathématiques Appliquées, Grenoble, France, in 1987. He received the Ph.D. degree in electrical engineering from INRIA, France, in 1993.

From 1988 to 1989, he was with the Department of Radiotherapy, Institut Gustave Roussy, Villejuif, France. From 1989 to 1990, he was with Alcatel, Paris, France. From 1993 to 1995, he was a Postdoctoral Associate in the Departments of Diagnostic Radiology and Mathematics, Yale University, New Haven, CT. In 1996, he was an Associate Research Scientist with the Departments of Diagnostic Radiology and Mathematics, Yale University. From 1997 to 1999, he was an Assistant Professor with the Departments of Radiology and Computer Science, Yale University. He is currently an Assistant Professor with the Department of Electrical Engineering, University of Colorado, Boulder. He is also an Assistant Professor with the Department of Radiology, School of Medicine, University of Colorado Health Sciences Center. His research interests include image processing, biomedical signal and image analysis, image and video compression, and diverse applications of wavelets to signal and image processing. He is the editor (with A. Petrosian) of *Wavelets in Signal and Image Analysis, From Theory to Practice* (Norwell, MA: Kluwer, 2001).

Dr. Meyer is a member of Sigma Xi.