

Optimal Design of High-Performance Separable Wavelet Filter Banks for Image Coding

Michael D. Adams^{*,a}, Di Xu^b

^a*Department of Electrical and Computer Engineering, University of Victoria, Victoria, BC, Canada*

^b*Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC, Canada*

Abstract

An optimization-based method is proposed for the design of high-performance separable wavelet filter banks for image coding. This method yields linear-phase perfect-reconstruction systems with high coding gain, good frequency selectivity, and certain prescribed vanishing-moment properties. Several filter banks designed with the proposed method are presented and shown to work extremely well for image coding, outperforming the well-known 9/7 filter bank from JPEG 2000 in most cases. With the proposed design method, the coding gain can be maximized with respect to the separable or isotropic image model, or jointly with respect to both models. The joint case, which is shown to be equivalent to the isotropic case, is experimentally demonstrated to lead to filter banks with better average coding performance than the separable case.

During the development of the proposed design method, filter banks from a certain popular separable 2D wavelet class (to which our optimal designs belong) were observed to always have a higher coding gain with respect to the separable image model than with respect to the isotropic one. This behavior is examined in detail, leading to the conclusion that, for filter banks belonging to the above class, it is highly improbable (if not impossible) for the isotropic coding gain to exceed the separable coding gain.

Key words: optimal wavelet filter-bank design, image coding, coding gain, separable/isotropic image models

1. Introduction

Separable two-dimensional (2D) wavelet filter banks have proven to be an extremely valuable tool for image coding applications [1, 2, 3]. In order to be effective in such applications, however, a filter

***Corresponding author.** Address: Department of Electrical and Computer Engineering, University of Victoria, PO Box 3055 STN CSC, Victoria, BC, V8W 3P6, Canada; tel.: +1 250 721 6025; fax: +1 250 721 6052.

Email addresses: mdadams@ece.uvic.ca (Michael D. Adams), dixu@ece.ubc.ca (Di Xu)

bank must typically have a number of desirable characteristics such as perfect reconstruction (PR), linear phase, high coding gain [4], good frequency selectivity, and certain vanishing-moment properties. To date, a great many optimization-based design methods for wavelet filter banks have been proposed, some of which include [5, 6, 7, 8, 9, 10, 11, 12, 13, 14]. One distinguishing characteristic of a design method is the type of filter-bank parameterization that it employs. Lattice [15] and lifting [16] parameterizations have proven to be quite popular, with [6, 10] using the former and [9, 11, 13, 12, 14] using the latter. Another distinguishing trait of a design method is the set of filter-bank characteristics that it considers (e.g., PR, linear phase, and so on). Although many of the previously-proposed design methods consider a subset of the desirable characteristics for image coding mentioned above, few (if any) consider all of these characteristics simultaneously. For example, the method outlined in [7] considers coding gain, PR, linear phase, and imposes one dual vanishing moment, but does not explicitly consider frequency selectivity. The schemes proposed in [12, 13, 14] consider PR, linear phase, and vanishing-moment properties, but not coding gain, and are also restricted to a lifting parameterization with only two or three lifting filters.

To design filter banks having all of the desirable characteristics mentioned above is a challenging task. In this manuscript, we propose a new design method that yields high-performance separable wavelet filter banks with all of these characteristics. This method employs optimization and is based on ideas from [17]. During the course of our work, the filter banks obtained both at the intermediate and final stages of our design method were observed to always have higher coding gain with respect to the separable image model than with respect to the isotropic one. In this manuscript, we also study this phenomenon in detail and explain the reason for it. The work presented herein has been, in part, described in our conference papers [18, 19].

The remainder of the manuscript is structured as follows. Section 2 introduces some of the notational conventions used herein, and Section 3 provides some background information on filter banks and coding gain. Our proposed design method is presented in Section 4. In Section 5, our method is used to design several filter banks and these filter banks are shown to perform very well for image coding. Section 6 studies in detail the relationship between the coding gains for the separable and isotropic models and presents a number of interesting results in this regard. Finally, Section 7 concludes our work with a summary of our key results.

2. Notation and Terminology

Before proceeding further, we introduce some of the notation employed herein. The sets of integers and real numbers are denoted as \mathbb{Z} and \mathbb{R} , respectively. For $x \in \mathbb{R}$, the notation $\lfloor x \rfloor$ denotes the largest integer not greater than x (i.e., the floor function). For $x, y \in \mathbb{Z}$, we define the functions $\text{div}(x, y) = \lfloor x/y \rfloor$ and $\text{mod}(x, y) = x - y \lfloor x/y \rfloor$ (i.e., x divided by y yields the quotient $\text{div}(x, y)$ and remainder $\text{mod}(x, y)$). Matrices and vectors are typically denoted by uppercase and lowercase boldface letters, respectively. The transpose of the matrix/vector \mathbf{A} is denoted as \mathbf{A}^T . The symbols \mathbf{I} , $\mathbf{0}$, and $\mathbf{1}$ denote an identity matrix, a vector of all zeros, and a vector of all ones, respectively, the size of which should be clear from the context. The square root of a positive semi-definite matrix \mathbf{A} (e.g., as defined in [20]) is denoted as $\mathbf{A}^{1/2}$. For matrix multiplication, we define the product notation as $\prod_{k=M}^N \mathbf{A}_k \triangleq \mathbf{A}_N \mathbf{A}_{N-1} \dots \mathbf{A}_{M+1} \mathbf{A}_M$, where $N \geq M$. The element of the 2D sequence f with index $\mathbf{n} = (n_0, n_1) \in \mathbb{Z}^2$ is denoted as either $f[n_0, n_1]$ or $f[\mathbf{n}]$, whichever is more convenient. A similar notational convention is also employed for 2D functions. The Fourier transform of a sequence/function f is denoted as \hat{f} . The symbols $*$ and \star denote convolution and correlation, respectively. The p -norm of the vector $\mathbf{x} = [x_1 \ x_2 \ \dots \ x_d]^T$, denoted $\|\mathbf{x}\|_p$, is defined as $\|\mathbf{x}\|_p = (\sum_{i=1}^d |x_i|^p)^{1/p}$ for $p \in \{1, 2\}$. The symbol ∇ is used to denote the gradient operator, which is defined to always produce a column vector. In a context where it may be unclear with respect to what quantity a gradient is taken, a subscript on ∇ is used to indicate this quantity (e.g., $\nabla_{\mathbf{x}}$ denotes the gradient with respect to \mathbf{x}). For a one-dimensional (1D) sequence f satisfying $f[n] = sf[2c - n]$ for all $n \in \mathbb{Z}$, where $c \in \frac{1}{2}\mathbb{Z}$ and $s \in \{-1, 1\}$, f is said to be symmetric if $s = 1$ and antisymmetric if $s = -1$, and f is also said to have symmetry about c . The polynomial $F(z)$ is said to be symmetric or antisymmetric about (the term) z^c if the sequence $\mathcal{Z}^{-1}F$ is symmetric or antisymmetric about $-c$, respectively, where \mathcal{Z} denotes the z transform (e.g., $1 + 2z^{-1} + 1z^{-2}$ is symmetric about z^{-1} and $3 - 3z^{-1}$ is antisymmetric about $z^{-1/2}$). A 2D function/sequence f is said to have quadrantal symmetry if $f(t_0, t_1) = f(-t_0, t_1) = f(t_0, -t_1) = f(-t_0, -t_1)$ for all t_0, t_1 in the domain of f . For a filter H , we denote its transfer function and impulse response as H and h , respectively. The sans-serif letters “h” and “v” are used to denote the horizontal and vertical components of separable functions/sequences/operators. For example, the horizontal and vertical components of the the separable 2D sequence f would be denoted as f_h and f_v , respectively (i.e., $f[n_0, n_1] = f_h[n_0]f_v[n_1]$). For a signal with P bits per sample, the peak-signal-to-noise ratio (PSNR) is defined as $\text{PSNR} = 20 \log_{10} \left(\frac{2^P - 1}{\sqrt{\text{MSE}}} \right)$, where MSE denotes the mean-squared error. The relative difference d of two quantities is simply defined as $d = (n - r)/r$, where n and r denote the non-reference and reference quantities, respectively.




Figure 1: The canonical form of a 1D two-channel filter bank. (a) Analysis and (b) synthesis sides.




Figure 2: The lifting realization of a 1D two-channel filter bank. (a) Analysis and (b) synthesis sides.

3. Background

A 1D two-channel filter bank has the canonical form shown in Fig. 1. Such a filter bank consists of analysis filters $\{H_0, H_1\}$, and synthesis filters $\{F_0, F_1\}$ (where H_0 and F_0 are lowpass), as well as downsamplers and upsamplers. A filter bank can also be represented in terms of a lifting realization [16]. The lifting realization of a 1D two-channel filter bank is shown in Fig. 2, and consists of 2λ lifting filters $\{P_k\}_{k=0}^{2\lambda-1}$. Without loss of generality, we assume that only $P_0(z)$ and/or $P_{2\lambda-1}(z)$ may be identically zero. The analysis and synthesis filters $\{H_0, H_1\}$ and $\{F_0, F_1\}$ of the canonical form can be readily determined from the filters $\{P_k\}_{k=0}^{2\lambda-1}$ of the lifting parameterization using the following relationships:

$$H_0(z) = H_{0,0}(z^2) + zH_{0,1}(z^2), \quad H_1(z) = H_{1,0}(z^2) + zH_{1,1}(z^2), \quad \text{where} \quad (1)$$

$$\mathbf{H}_p(z) = \begin{bmatrix} H_{0,0}(z) & H_{0,1}(z) \\ H_{1,0}(z) & H_{1,1}(z) \end{bmatrix} = \prod_{k=0}^{\lambda-1} \left(\begin{bmatrix} 1 & P_{2k+1}(z) \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ P_{2k}(z) & 1 \end{bmatrix} \right), \quad \text{and}$$

$$F_0(z) = -z^{-1}H_1(-z), \quad \text{and} \quad F_1(z) = z^{-1}H_0(-z). \quad (2)$$

Since images are 2D signals, their processing requires multidimensional systems. To construct a 2D filter bank from a 1D two-channel filter bank, we simply apply the 1D filter bank in each of the two dimensions of the signal in succession. This results in a separable four-channel 2D filter bank. Furthermore, in practice, we usually apply the 2D filter bank in an L -level tree structure, decomposing the lowest-frequency subband signal at each level in the tree. The resulting L -level tree-structured filter bank can be equivalently expressed in the form of an M -channel nonuniform filter bank, where $M = 3L + 1$. This equivalent M -channel filter

bank has the general form shown in Fig. 3, consisting of analysis filters $\{H'_k\}_{k=0}^{M-1}$, synthesis filters $\{F'_k\}_{k=0}^{M-1}$, and downsamplers/upsamplers, where the k th downsampler/up-sampler (associated with sampling matrix $\mathbf{\Lambda}_k$) has the horizontal and vertical sampling factors of $M_{k,h}$ and $M_{k,v}$, respectively. Due to the separable nature of the filters, the subbands have four possible orientations: horizontally-and-vertically lowpass (LL), horizontally highpass and vertically lowpass (LH), horizontally lowpass and vertically highpass (HL), and horizontally-and-vertically highpass (HH).

As a matter of notation, the level in the analysis filter-bank tree associated with channel k is denoted as $\text{level}(k)$, where $\text{level}(k) \in \{0, 1, \dots, L-1\}$ with the value of zero corresponding to the tree root; and the orientation of channel k is denoted as $\text{orient}(k)$, where $\text{orient}(k) \in \{0, 1, 2, 3\}$ with the values of 0, 1, 2, and 3 corresponding to the LL, LH, HL, and HH orientations, respectively. For the l th level in the tree, LL_l , LH_l , HL_l , and HH_l denote the subbands with LL, LH, HL, and HH orientations, respectively. The mapping between the channel index k and subband level/orientation is given by

$$\text{level}(k) = \begin{cases} L-1 & \text{for } k = 3L \\ \text{div}(k, 3) & \text{otherwise} \end{cases} \quad \text{and} \quad \text{orient}(k) = \begin{cases} 0 & \text{for } k = 3L \\ 3 - \text{mod}(k, 3) & \text{otherwise.} \end{cases} \quad (3)$$

For convenience, we have numbered the channels such that larger values of k correspond to a deeper descent into the analysis filter-bank tree. Note that this numbering convention is backwards from what is typically used. That is, the lowest frequency (LL) band is always associated with the largest channel index $3L$ (instead of the smallest one 0). With our convention, decompositions with L and L' levels have an identical numbering scheme for their common filters/subbands. This allows for a much more concise presentation of some of our later results. The correspondence between channels and frequency bands is further illustrated in Fig. 4, with the analysis filter associated with each subband being shown in parentheses. In order to avoid an overly complicated diagram, only the first quadrant of the frequency plane is shown, with the remainder following from (quadrantal) symmetry.

Given the 1D analysis filters $\{H_0, H_1\}$, the analysis filters $\{H'_k\}_{k=0}^{M-1}$ of the 2D nonuniform filter bank can be computed as

$$H'_k(z_0, z_1) = H'_{k,h}(z_0)H'_{k,v}(z_1), \quad (4)$$

where $H'_{k,h}(z) = H_{\text{mod}(\text{orient}(k), 2)}(z^{2^{\text{level}(k)}})H_{\Pi}(z)$, $H'_{k,v}(z) = H_{\text{div}(\text{orient}(k), 2)}(z^{2^{\text{level}(k)}})H_{\Pi}(z)$, and $H_{\Pi}(z) = 1$ if $\text{level}(k) = 0$, and $H_{\Pi}(z) = \prod_{\ell=0}^{\text{level}(k)-1} H_0(z^{2^\ell})$ otherwise. The synthesis filters $\{F'_k\}_{k=0}^{M-1}$ of the 2D nonuniform filter bank are determined in a similar way from the 1D synthesis filters $\{F_0, F_1\}$ (i.e., simply replace

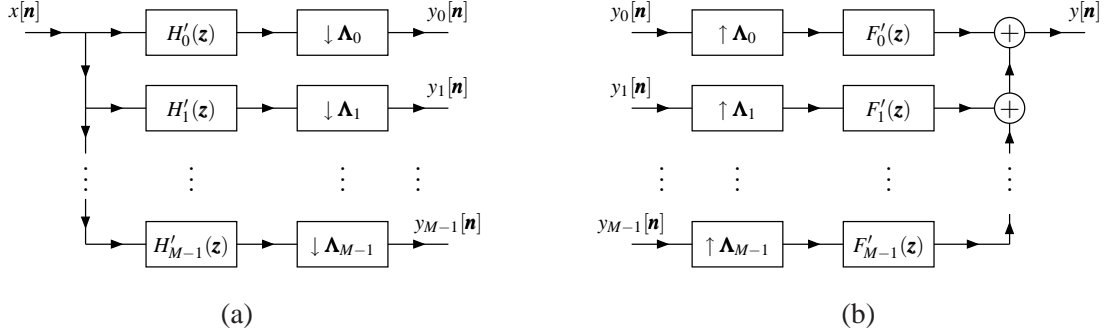


Figure 3: The equivalent M -channel nonuniform filter bank associated with the L -level tree-structured filter bank (where $M = 3L + 1$).




Figure 4: Subband tiling of the first quadrant of the frequency plane.

the H s by F s in the preceding equation).

In the context of image coding, we are often interested in the energy-compacting ability of a filter bank, which is typically quantified using a measure known as the coding gain. For a filter bank of the form shown in Fig. 3, the coding gain G is given by [4]

$$G = \prod_{k=0}^{M-1} \left(\frac{\alpha_k}{A_k B_k} \right)^{\alpha_k}, \quad \text{where} \quad (5)$$

$$A_k = \sum_{\mathbf{l} \in \mathbb{Z}^2} \sum_{\mathbf{p} \in \mathbb{Z}^2} h'_k[\mathbf{l}] h'_k[\mathbf{p}] r[\mathbf{p} - \mathbf{l}] = \sum_{m \in \mathbb{Z}} h'_{k,h}[m] \sum_{n \in \mathbb{Z}} h'_{k,v}[n] \sum_{p \in \mathbb{Z}} h'_{k,h}[p] \sum_{q \in \mathbb{Z}} h'_{k,v}[q] r[m - p, n - q],$$

$$B_k = \alpha_k \sum_{\mathbf{l} \in \mathbb{Z}^2} f_k'^2[\mathbf{l}] = \alpha_k \sum_{m \in \mathbb{Z}} f_{k,h}'^2[m] \sum_{n \in \mathbb{Z}} f_{k,v}'^2[n], \quad \alpha_k = (M_{k,h} M_{k,v})^{-1},$$

and r is the normalized autocorrelation sequence of the source image model. The two most common choices

for r are given by the separable and isotropic models, which are, respectively,

$$r_{\text{sep}}[\mathbf{n}] = \rho^{\|\mathbf{n}\|_1} \quad \text{and} \quad r_{\text{iso}}[\mathbf{n}] = \rho^{\|\mathbf{n}\|_2}, \quad (6)$$

where ρ is a correlation coefficient satisfying $\rho \in [0, 1]$. Since numerous quantities used herein depend on the image model (e.g., $G, \{A_k\}_{k=0}^{M-1}$), we use the qualifiers ‘‘sep’’ and ‘‘iso’’ to denote these quantities in the separable and isotropic cases, respectively. For example, $A_{k,\text{sep}}$ denotes the quantity A_k given by (5) with $r = r_{\text{sep}}$ and G_{iso} denotes the coding gain G given by (5) with $r = r_{\text{iso}}$.

4. Design Method

In our design method, rather than representing a filter bank in its canonical form as shown in Fig. 1, we instead use the lifting framework as depicted in Fig. 2. The use of the lifting framework has a number of advantages over the canonical form. The key benefit, however, is that the PR condition is automatically satisfied. Additionally, the linear phase requirement can be easily met by choosing the lifting filters $\{P_k\}_{k=0}^{2\lambda-1}$ to have certain symmetry properties, as we shall see shortly. Since the PR and linear-phase conditions can be imposed via the lifting framework, there is no need for explicit optimization constraints to ensure that these conditions are satisfied. This greatly reduces the complexity of the subsequent optimization problem. The lifting framework is also advantageous as it trivially allows for the construction of reversible integer-to-integer mappings [21]. Such mappings are extremely useful for image coding applications, especially in situations where lossless coding may be desired. In fact, the image coders used to obtain the coding results presented later in this manuscript all employ reversible integer-to-integer wavelet transforms.

As suggested above, the linear-phase condition can be easily imposed through a clever choice of the lifting filters $\{P_k\}_{k=0}^{2\lambda-1}$. In what follows, let L_k denote the length of the lifting filter P_k . It can be shown [22] that if the $\{P_k\}_{k=0}^{2\lambda-1}$ are chosen to be of either of the following two forms, then the resulting filter bank will have linear phase:

$$P_k(z) = \begin{cases} \sum_{i=0}^{(L_k-2)/2} p_{k,i}(z^{-i} + z^{i+1}) & \text{for even } k \\ \sum_{i=0}^{(L_k-2)/2} p_{k,i}(z^{-i-1} + z^i) & \text{for odd } k, \quad \text{or} \end{cases} \quad (7a)$$

$$P_k(z) = \begin{cases} -1 & \text{for } k = 0 \\ \frac{1}{2} + \tilde{P}_1(z) & \text{for } k = 1 \\ \tilde{P}_k(z) & \text{for } k \geq 2 \end{cases} \quad \text{and} \quad \tilde{P}_k(z) = \sum_{i=1}^{(L_k-1)/2} \tilde{p}_{k,i}(z^{-i} + z^i), \quad (7b)$$

where, in (7a), L_k is even for $k \in \{0, 1, \dots, 2\lambda - 1\}$ (i.e., $P_k(z)$ is symmetric about $z^{1/2}$ and $z^{-1/2}$ for even and odd k , respectively) and, in (7b), L_k is odd for $k \in \{1, 2, \dots, 2\lambda - 1\}$ (i.e., $\tilde{P}_k(z)$ is antisymmetric about z^0). In the case of parameterization (7a), h_0 is symmetric about 0 and h_1 is symmetric about -1 , while in the case of parameterization (7b), h_0 is symmetric about $-\frac{1}{2}$ and h_1 is antisymmetric about $-\frac{1}{2}$.

In the case of parameterization (7a), the symmetry properties of the analysis filter impulse responses h_0 and h_1 can be deduced as follows. To begin, one can show by induction that the analysis polyphase matrix $\mathbf{H}_p(z)$ given by (1) is such that: 1) $H_{0,0}(z)$ and $H_{1,1}(z)$ are symmetric about z^0 ; 2) $H_{0,1}(z)$ is symmetric about $z^{-1/2}$; and 3) $H_{1,0}(z)$ is symmetric about $z^{1/2}$. From properties 1 and 2, $H_{0,0}(z^2)$ and $zH_{0,1}(z^2)$ are both symmetric about z^0 , and consequently, their sum $H_0(z)$ (as given by (1)) is also symmetric about z^0 . From properties 1 and 3, $H_{1,0}(z^2)$ and $zH_{1,1}(z^2)$ are both symmetric about z^1 , and consequently, their sum $H_1(z)$ (as given by (1)) is symmetric about z^1 . Thus, the stated symmetry properties hold for h_0 and h_1 . In the case of parameterization (7b), a proof of the stated symmetry properties for h_0 and h_1 can be found in [22, Section VII and Appendices A and B] and in more verbose form in [23, pp. 74–77].

It is worth noting that (7a) completely parametrizes (up to a trivial normalization) all PR linear-phase FIR filter banks with odd-length analysis/synthesis filters, while (7b) parametrizes (up to a normalization) only a subset of all PR linear-phase FIR filter banks with even-length analysis/synthesis filters. For this reason, one might suspect parameterization (7a) to have greater potential to yield good filter banks than parameterization (7b). In fact, this suspicion turns out to be correct, as later confirmed by our experimental results in Section 5.

With the lifting framework, the synthesis filters are completely determined by the analysis filters as given by (2). Therefore, we focus primarily on the design of the analysis side of the filter bank in what follows. Since we have elected to use a lifting parameterization in our design method, we need to relate the various filter-bank properties of interest (i.e., the analysis-filter frequency responses, vanishing-moment properties, and coding gain) to the lifting-filter coefficients. In the case of the moment properties of the primal and dual wavelet coefficient sequences and the frequency responses, these relationships can be derived in a straightforward manner using (1), (2), and (7). The analysis-filter frequency responses and the expressions for the moments of the primal/dual wavelet coefficient sequences are polynomials in the lifting-filter coefficients, where the polynomial order depends on which of the two parameterizations (7a) and (7b) is employed as well as the number of lifting filters. For example, for filter banks from parameterization (7a) with four lifting filters, the lowpass analysis-filter frequency response, highpass analysis-filter frequency

response, primal-moment expression, and dual-moment expression are polynomials in the lifting-filter coefficients having orders of 4, 3, 4, and 3, respectively. The coding gain can be expressed in terms of the lifting-filter coefficients by combining (5), (4), (1), (2), and (7). The resulting expression is highly nonlinear, due mainly to the form of (5).

4.1. Abstract Optimization Problem

As indicated earlier, we seek to design filter banks having numerous desirable characteristics, namely, PR, linear phase, high coding gain, good frequency selectivity, and certain prescribed vanishing-moment properties. Since the PR and linear-phase properties are structurally imposed via the lifting framework, we need not consider them further. Thus, the design problem at hand reduces to one explicitly involving only coding gain, frequency selectivity, and vanishing-moment properties.

Now, let us consider the formulation of the design problem as an optimization to be performed with respect to the lifting-filter coefficients (i.e., the $\{p_{k,i}\}$ from (7a) or the $\{\tilde{p}_{k,i}\}$ from (7b)). Let \mathbf{x} denote the vector of (independent) lifting-filter coefficients, where the coefficients are lexicographically ordered by their first and then second index (e.g., in the case of (7a), we have the ordering $p_{0,0}, p_{0,1}, \dots, p_{0,(L_0-2)/2}, p_{1,0}, p_{1,1}, \dots, p_{1,(L_1-2)/2}, \dots, p_{2\lambda-1,0}, p_{2\lambda-1,1}, \dots, p_{2\lambda-1,(L_{2\lambda-1}-2)/2}$). We choose G , a measure related to coding gain, as the function to maximize. Let G_{sep} and G_{iso} denote the coding gain (in dB) obtained from (5) using the separable and isotropic models, respectively. In our work, we consider three possible choices for G as given by

$$G(\mathbf{x}) = \begin{cases} G_{\text{sep}}(\mathbf{x}) & \text{separable only} \\ G_{\text{iso}}(\mathbf{x}) & \text{isotropic only} \\ \min\{G_{\text{sep}}(\mathbf{x}), G_{\text{iso}}(\mathbf{x})\} & \text{joint.} \end{cases} \quad (8)$$

That is, we consider the maximization of each of the separable and isotropic coding gains individually as well as the joint maximization of both coding gains. The joint case in (8) is motivated by the observation that many images are nonstationary, exhibiting both separable and isotropic behaviors in different regions. Thus, we might suspect that there is an advantage to having both coding gains high.

The remaining filter bank properties are handled as constraints. To quantify the frequency selectivity of the analysis filters, we employ a stopband-energy measure. In particular, we define the stopband energy of the analysis filter H_k as

$$b_k(\mathbf{x}) \triangleq \int_{S_k} |\hat{h}_k(\omega, \mathbf{x})|^2 d\omega, \quad k \in \{0, 1\}, \quad (9)$$

where $S_0 = [\pi - \omega_b, \pi]$, $S_1 = [0, \omega_b]$, and ω_b denotes the stopband width of the analysis filters. The reason for using only a stopband constraint is twofold. First, limiting the stopband-energy leakage alone can be quite effective in avoiding aliasing. Second, for filter banks with relatively short filters, the number of degrees of freedom in the design process is limited. Consequently, if a passband constraint were also employed, the feasible region for the optimization may be overly restricted, leading to poorer designs.

To facilitate the introduction of moment constraints, we define the moment-constraint functions

$$c_k(\mathbf{x}) \triangleq \|\mathbf{m}_k(\mathbf{x})\|_2, \quad k \in \{1, 2, \dots, \eta\}, \quad (10)$$

where \mathbf{m}_k is a v_k -dimensional vector function with its elements corresponding to the moments of interest (i.e., moments that are to be constrained). Each \mathbf{m}_k may contain only one moment (i.e., $v_k = 1$) or a group of moments (i.e., $v_k \geq 2$). In this way, moments can be controlled either individually or jointly.

Combining (8), (9), and (10), we obtain the following abstract optimization problem to be solved:

$$\text{maximize } G(\mathbf{x}) \quad (11a)$$

$$\text{subject to: } b_k(\mathbf{x}) \leq \varepsilon_k, \quad k \in \{0, 1\} \quad \text{and} \quad (11b)$$

$$c_k(\mathbf{x}) \leq \gamma_k, \quad k \in \{1, 2, \dots, \eta\}, \quad (11c)$$

where the $\{\varepsilon_0, \varepsilon_1\}$ and $\{\gamma_k\}_{k=1}^{\eta}$ are strictly positive tolerances for the stopband-energy and moment constraints, respectively. Since the $\{\gamma_k\}_{k=1}^{\eta}$ are chosen to be strictly positive, we do not attempt to satisfy the vanishing-moment conditions exactly. Instead, we only ensure that the moments of interest are very nearly vanishing (e.g., typically on the order of 10^{-5} or less in some of our later design examples). In a practical sense, there is no significant disadvantage to allowing the moments to deviate slightly from zero, as exact vanishing moments are usually lost during implementation anyhow, due to finite-precision effects. In fact, this relaxed form of moment constraint is actually quite beneficial, as it allows increased design flexibility, which in most cases leads to better designs. In passing, we note that parameterization (7b) structurally imposes vanishing zeroth primal and dual moments [22]. So, when this parameterization is employed, the vanishing moment conditions for the zeroth moments will always be satisfied exactly.

In the abstract optimization problem (11), there are three cases for the objective function (8). The joint case is associated with a max-min problem, which is somewhat difficult to solve directly. For this reason,

we convert the joint case into the following equivalent problem, which can be more easily solved:

$$\text{maximize } t \tag{12a}$$

$$\text{subject to: } G_{\text{sep}}(\mathbf{x}) \geq t, \tag{12b}$$

$$G_{\text{iso}}(\mathbf{x}) \geq t, \tag{12c}$$

$$b_k(\mathbf{x}) \leq \varepsilon_k, \quad k \in \{0, 1\}, \quad \text{and} \tag{12d}$$

$$c_k(\mathbf{x}) \leq \gamma_k, \quad k \in \{1, 2, \dots, \eta\}, \tag{12e}$$

where t is an auxiliary variable.

4.2. Solution of the Abstract Optimization Problem

As introduced above, the abstract optimization problem for our filter-bank design scheme is given in (11), with the joint case reformulated as in (12). Unfortunately, these problems are highly nonlinear and somewhat difficult to solve directly. For this reason, we adopt a strategy based on the iterative solution of reduced-order problems. The algorithm for this approach has the following general form:

STEP 1 (INITIALIZATION). Set the iteration number i to zero. Choose the stopping tolerances τ_1 and τ_2 . Select an initial operating point \mathbf{x}_0 somewhere in the feasible region. Select a nominal maximum step size β_0 for use in step 3.

STEP 2 (ORDER REDUCTION). Represent each of the functions $G_{\text{sep}}(\mathbf{x})$ and/or $G_{\text{iso}}(\mathbf{x})$, $\{b_0, b_1\}$, and $\{c_k\}_{k=1}^{\eta}$ with a Taylor-series approximation about the current operating point \mathbf{x}_i . In particular, each function f is represented using a linear approximation as given by $f(\mathbf{x}_i + \boldsymbol{\delta}_i) \approx f(\mathbf{x}_i) + \nabla^T f(\mathbf{x}_i)\boldsymbol{\delta}_i$.

STEP 3 (OPTIMIZATION). (a) Let $\beta = \beta_0$ (i.e., set the maximum step size β to the nominal maximum step size β_0). (b) Solve the reduced-order optimization problem in the variable $\boldsymbol{\delta}_i$. Let $\boldsymbol{\delta}_i^*$ denote the corresponding optimal solution. Since the Taylor-series approximations obtained in step 2 are accurate only when $\boldsymbol{\delta}_i$ is small, the additional constraint $\|\boldsymbol{\delta}_i\|_2 \leq \beta$ is imposed in the reduced-order optimization problem to ensure a solution in the vicinity of the operating point \mathbf{x}_i . (c) The point $\boldsymbol{\delta}_i^*$ will always be in the feasible region of the reduced-order problem, but $\mathbf{x}_i + \boldsymbol{\delta}_i^*$ may not be in the feasible region of the original (i.e., non-reduced order) problem if the reduced-order approximation is not sufficiently accurate. Therefore, if $\mathbf{x}_i + \boldsymbol{\delta}_i^*$ is not in the feasible region of the original problem, set $\beta = \beta/1.6$ and go to step 3(b) (i.e., restart the reduced-order optimization with a smaller step size β in order to improve the reduced-order approximation accuracy). Otherwise, continue to step 4.

STEP 4 (OPERATING-POINT UPDATE). Set the new operating point \mathbf{x}_{i+1} to $\mathbf{x}_i + \boldsymbol{\delta}_i^*$.

STEP 5 (STOPPING-CRITERIA CHECK). If $|G(\mathbf{x}_{i+1}) - G(\mathbf{x}_i)| \leq \tau_1$ or $\|\boldsymbol{\delta}_i^*\|_2 \leq \tau_2$, then output the solution $\mathbf{x}^* = \mathbf{x}_{i+1}$ and stop. Otherwise, increment i by one and go to step 2.

Essentially, the above algorithm finds a reduced-order approximation of the original problem about the current operating point, solves the reduced-order problem, adjusts the operating point accordingly, and iterates. In what follows, we now explain in more detail how the preceding algorithm can be used to solve our design problem.

In step 2, since it is extremely difficult to derive closed-form expressions for the derivatives of $G_{\text{sep}}(\mathbf{x})$ and/or $G_{\text{iso}}(\mathbf{x})$, these quantities are computed numerically. In the case of the less nonlinear functions $\{b_0, b_1\}$ and $\{c_k\}_{k=1}^{\eta}$, closed-form expressions are used for the derivatives. Due to space constraints, we do not provide these expressions here (as they are quite long and messy), but they can be derived in a straightforward manner.

In step 3, we considered the use of several different optimization methods, namely linear programming, convex quadratic programming, sequential quadratic programming, and second-order cone programming (SOCP) [24, 25]. Each of these methods requires Taylor series of particular orders to be used in step 2 in order to approximate the various functions of interest. In the interest of brevity, we will only present a SOCP-based approach in what follows. The interested reader, however, can find details regarding the other approaches in [26, pp. 32-37]. As an aside, we note that numerous software packages exist for the solution of SOCP problems (e.g., SeDuMi [27], CVX [28], and YALMIP [29]), with our work having used SeDuMi.

To formulate step 3 as a SOCP problem, we choose (in step 2) to represent each of the functions G_{sep} and/or G_{iso} , $\{b_0, b_1\}$, and $\{c_k\}_{k=1}^{\eta}$ using a linear approximation. For the separable and isotropic cases of (11), it can then be shown that the optimization in step 3 can be expressed in terms of the following SOCP problem (with optimization being performed with respect to $\boldsymbol{\delta}$):

$$\text{maximize } \nabla^T G(\mathbf{x}_i) \boldsymbol{\delta} \quad (13a)$$

$$\text{subject to: } \left\| \mathbf{Q}_k^{1/2}(\mathbf{x}_i) \boldsymbol{\delta} + \mathbf{q}_k(\mathbf{x}_i) \right\|_2 \leq \varepsilon_k - b_k(\mathbf{x}_i) + \mathbf{q}_k^T(\mathbf{x}_i) \mathbf{q}_k(\mathbf{x}_i), \quad k \in \{0, 1\}, \quad (13b)$$

$$\left\| \nabla^T \mathbf{m}_k(\mathbf{x}_i) \boldsymbol{\delta} + \mathbf{m}_k(\mathbf{x}_i) \right\|_2 \leq \gamma_k, \quad k \in \{1, 2, \dots, \eta\}, \quad \text{and} \quad (13c)$$

$$\|\boldsymbol{\delta}\|_2 \leq \beta, \quad (13d)$$

where $\mathbf{Q}_k(\mathbf{x}) = \int_{S_k} [\nabla_{\mathbf{x}} \hat{h}_k(\boldsymbol{\omega}, \mathbf{x})] \nabla_{\mathbf{x}}^T \hat{h}_k(\boldsymbol{\omega}, \mathbf{x}) d\boldsymbol{\omega}$, $\mathbf{q}_k(\mathbf{x}) = \mathbf{Q}_k^{-1/2}(\mathbf{x}) \int_{S_k} \hat{h}_k(\boldsymbol{\omega}, \mathbf{x}) \nabla_{\mathbf{x}}^T \hat{h}_k(\boldsymbol{\omega}, \mathbf{x}) d\boldsymbol{\omega}$, S_k is as defined in (9), and $\boldsymbol{\delta}$ is a perturbation from the operating point \mathbf{x}_i . For the joint case of (11) as reformulated in (12),

the optimization in step 3 can be expressed in terms of the following SOCP problem, with the optimization being performed with respect to the augmented vector variable $\tilde{\boldsymbol{\delta}} \triangleq [t \quad \boldsymbol{\delta}]^T$:

$$\text{maximize} \quad [1 \quad \mathbf{0}] \tilde{\boldsymbol{\delta}} \quad (14a)$$

$$\text{subject to:} \quad [0 \quad \nabla^T G_{\text{iso}}(\mathbf{x}_i)] \tilde{\boldsymbol{\delta}} + G_{\text{iso}}(\mathbf{x}_i) - [1 \quad \mathbf{0}] \tilde{\boldsymbol{\delta}} \geq 0, \quad (14b)$$

$$[0 \quad \nabla^T G_{\text{sep}}(\mathbf{x}_i)] \tilde{\boldsymbol{\delta}} + G_{\text{sep}}(\mathbf{x}_i) - [1 \quad \mathbf{0}] \tilde{\boldsymbol{\delta}} \geq 0, \quad (14c)$$

$$\left\| \tilde{\mathbf{Q}}_k(\mathbf{x}_i) \tilde{\boldsymbol{\delta}} + \tilde{\mathbf{q}}_k(\mathbf{x}_i) \right\|_2 \leq \varepsilon_k - b_k(\mathbf{x}_i) + \mathbf{q}_k^T(\mathbf{x}_i) \mathbf{q}_k(\mathbf{x}_i), \quad k \in \{0, 1\}, \quad (14d)$$

$$\left\| \nabla^T \tilde{\mathbf{m}}_k(\mathbf{x}_i) \tilde{\boldsymbol{\delta}} + \tilde{\mathbf{m}}_k(\mathbf{x}_i) \right\|_2 \leq \gamma_k, \quad k \in \{1, 2, \dots, \eta\}, \quad \text{and} \quad (14e)$$

$$\left\| [\mathbf{0} \quad \mathbf{I}] \tilde{\boldsymbol{\delta}} \right\|_2 \leq \beta, \quad (14f)$$

where $\tilde{\mathbf{Q}}_k(\mathbf{x}) = [\mathbf{0} \quad \mathbf{q}_k^{1/2}(\mathbf{x})]$, $\tilde{\mathbf{q}}_k(\mathbf{x}) = [0 \quad \mathbf{q}_k(\mathbf{x})]^T$, $\tilde{\mathbf{m}}_k(\mathbf{x}) = [0 \quad \mathbf{m}_k(\mathbf{x})]^T$, $\mathbf{Q}_k(\mathbf{x}) = \int_{S_k} [\nabla_{\mathbf{x}} \hat{h}_k(\omega, \mathbf{x})] \nabla_{\mathbf{x}}^T \hat{h}_k(\omega, \mathbf{x}) d\omega$, $\mathbf{q}_k(\mathbf{x}) = \mathbf{Q}_k^{-1/2}(\mathbf{x}) \int_{S_k} \hat{h}_k(\omega, \mathbf{x}) \nabla_{\mathbf{x}}^T \hat{h}_k(\omega, \mathbf{x}) d\omega$, and S_k is as defined in (9). Note that, as far as step 3 is concerned, the output of the preceding optimization process is simply the $\boldsymbol{\delta}$ part of the augmented vector $\tilde{\boldsymbol{\delta}}$ (namely, $\boldsymbol{\delta} = [\mathbf{0} \quad \mathbf{I}] \tilde{\boldsymbol{\delta}}$).

In step 4 of the above algorithm, we update our operating point. As the solution in step 3 is limited to the neighborhood of the current operating point, we repeat the process in steps 2 to 5 until one of the stopping criteria is satisfied. This permits solutions farther away from the current operating point to be found.

4.3. Design-Parameter Selection

Having introduced our design method, we now briefly comment on the selection of numerous design parameters. For the frequency-selectivity constraints given by (13b) and (14d), an appropriate choice of tolerances $\{\varepsilon_0, \varepsilon_1\}$ is critical to achieving good designs. Based on our experiments, for a stopband width of $\omega_b = \frac{3\pi}{8}$, a choice of $\varepsilon_k \in [0.02, 0.14]$ is typically quite effective. The stopping tolerances τ_1 and τ_2 (used in step 5) might reasonably be chosen to be on the orders of 10^{-4} to 10^{-5} and 10^{-5} to 10^{-6} , respectively. Also, we must select L (i.e., the number of decomposition levels) which is used in the coding gain formula. In this regard, we found that choosing $L \in \{3, 4, 5\}$ makes a good tradeoff between computational complexity (which increases with L) and design quality. In step 1, we must choose the nominal maximum step size β_0 where $\beta_0 > \tau_2$. In practice, the choice of $\beta_0 = 2 \times 10^{-3}$ was found to work quite well.

In (13c) and (14e), the moment constraints $\{\mathbf{m}_k\}_{k=1}^{\eta}$ must be chosen along with appropriate tolerances $\{\gamma_k\}_{k=1}^{\eta}$. This choice is also key to obtaining good designs. Let μ_k denote the k th moment associated with the primal or dual wavelet coefficient sequence of the filter bank. In the case of filter banks from

parameterization (7a), the highpass analysis/synthesis filter is symmetric. Consequently, for even k , if $\mu_0 = \mu_1 = \dots = \mu_{k-1} = 0$, then $\mu_k = 0$ implies that $\mu_{k+1} = 0$ (i.e., the odd-indexed moments automatically vanish). For a proof, see Theorem 1 (symmetric case) in Appendix A. Thus, we need only consider constraints on the $\{\mu_k\}$ for even k . Similarly, in the case of filter banks from parameterization (7b), the highpass analysis/synthesis filter is antisymmetric. Consequently, $\mu_0 = 0$, and for odd k , if $\mu_0 = \mu_1 = \dots = \mu_{k-1} = 0$, then $\mu_k = 0$ implies that $\mu_{k+1} = 0$ (i.e., the even-indexed moments automatically vanish). Again, see Theorem 1 (antisymmetric case) in Appendix A for a proof. Thus, we need only consider constraints on the $\{\mu_k\}$ for odd k . Based on our experiments, good designs from the parameterization (7a) can be obtained by constraining the zeroth moments of the primal and dual wavelet coefficient sequences. More specifically, we found it to be quite effective to place these two moments in a single moment constraint function m_1 and choose a corresponding tolerance of $\gamma_1 = 2 \times 10^{-5}$.

Due to the highly nonlinear nature of the abstract optimization problem (11), the solution found by our method will most likely not be globally optimal. The particular solution obtained depends on the choice of the initial point \mathbf{x}_0 . Therefore, the quality of the design can be improved by finding multiple (locally optimal) solutions and then selecting the best one. As a practical matter, we found that an effective strategy in this regard is to consider many initial points with lifting-filter coefficients of magnitude 2 or less, as the best designs typically have coefficients in this range.

5. Design Results

Having introduced our design method, we are now ready to present some examples of filter banks generated with our method as well as some coding results obtained with these designs. Before proceeding further, however, we first introduce some important details about the methodology employed in our experiments. For all of our filter-bank designs, we selected the various design parameters as described in Section 4.3. The correlation coefficient ρ in (5) was chosen as 0.95, and five decomposition levels were used in the coding gain computation (i.e., $L = 5$). In our experiments, various image coding results were collected. For test data, we employed the twenty-six reasonably-sized grayscale images from the JPEG-2000 test set [30]. Often, we focus our attention on the results associated with three of these images, namely the `gold`, `sar2`, and `target` images. These three images were deliberately chosen, due to their significantly differing statistical properties. In particular, the autocorrelation sequence of the `gold` and `sar2` images most closely follow the separable and isotropic models, respectively, while the autocorrelation sequence of the `target`

image follows neither the separable nor isotropic model. For coding images, our implementations of the EZW [1], SPIHT [2], and MIC [31] coders were employed, all of which utilize reversible integer-to-integer transforms. Since similar results were obtained with all three coders, we present only results for the MIC coder. In all of our coding experiments, a six-level wavelet decomposition was employed.

For convenience, in what follows, we use the notation l_0/l_1 to indicate that a filter bank has lowpass and highpass analysis filters of lengths l_0 and l_1 , respectively. Also, for a lifting realization, we refer to the number and lengths of its lifting filters as its lifting configuration. As a matter of notation, the lifting configuration of a filter bank is denoted as $\{L_0, L_1, \dots\}$, where L_k is the length of the k th lifting filter. Recall that the $\{L_k\}$ are all odd for filter banks from parameterization (7a) and all even for filter banks from parameterization (7b). So, the parameterization being used can be trivially determined from the lifting configuration.

5.1. Choice of Objective Function

As indicated earlier, our design method allows for three possible objective functions as given in (8). We also suggested that, of these possibilities, the joint case might be the most desirable. Now, we study the impact of the choice of objective function in more detail.

To begin, for each of several different lifting configurations, we used our design method to construct three filter banks, one for each of the three choices of objective function in (8). In so doing, we were able to make an interesting observation. Namely, in all of our tests, optimizing with respect to each of the isotropic and joint coding gains always led to the same optimal designs. This is due to the fact that, for filter banks with reasonable analysis-filter frequency responses (i.e., at least one dual vanishing moment), the condition $G_{\text{sep}} > G_{\text{iso}}$ always seems to be satisfied. Since the reason that this condition holds is not at all obvious, we provide a detailed analysis later in Section 6 that explains why $G_{\text{sep}} > G_{\text{iso}}$. Due to the preceding behavior, maximizing the joint coding gain is equivalent, in a practical sense, to maximizing the isotropic coding gain alone. With the above observation in mind, we combine the isotropic and joint cases of (8) in the remainder of this discussion. After having designed each set of optimal filter banks as described above, we then compared the coding performance of the filter banks within each set.

Although several sets of optimal filter banks were considered in our work, we present results for only a representative subset herein. In particular, we consider two sets of optimal designs, one with analysis-filter lengths of 9/7 corresponding to the lifting configuration $\{2, 2, 2, 2\}$, and one with analysis-filter lengths of 6/14 corresponding to the lifting configuration $\{1, 3, 5\}$. The characteristics of these filter banks are shown

Table 1: Characteristics of the filter banks designed using different objective functions.

Transform	$G_{\text{sep}}^{\ddagger}$	$G_{\text{iso}}^{\ddagger}$	b_0^{\dagger}	b_1^{\dagger}	Dual VM*	Primal VM*
9/7-sep	14.9735	12.1781	0.0628	0.0347	2, 5.79E-5	2, 8.31E-5
9/7-iso/jnt	14.9326	12.1809	0.0570	0.0351	2, 0.0041	2, 0.1250
6/14-sep	15.0912	11.9285	0.0252	0.0131	1, 0.0483	1, 0.6209
6/14-iso/jnt	14.9766	12.0738	0.0212	0.0229	1, 0.0643	1, 0.1471

*index and magnitude of the first moment of dual/primal wavelet coefficient sequence with magnitude greater than 2×10^{-5} ; \ddagger coding gain (in dB) for a six-level decomposition; \dagger stopband energy as defined by (9)

in Table 1, where the transform-name suffixes “sep”, “iso”, and “jnt” designate the optimal designs obtained using the separable, isotropic, and joint objective functions of (8), respectively.

Having produced several sets of filter banks as described above, we then proceeded to compare the coding performance of the filter banks within each set. For each set, we chose to measure the performance of the optimal designs within the set relative to another previously-proposed filter bank having the same lifting configuration and also known to be effective for image coding. In particular, the reference filter banks used in the 9/7 and 6/14 cases are the 9/7 filter bank from JPEG 2000 [3] and the 6/14 filter bank from [32]. For each of the filter banks in each set, we compressed all twenty-six test images in a lossy manner at several bit rates, and in each case, we measured the relative difference in the distortions (in PSNR) obtained with our design and the corresponding reference filter bank. The results are summarized in statistical form in Table 2(a). In particular, we provide the mean and median relative differences in PSNR distortion (with positive values corresponding to our designs outperforming the reference filter bank). As well, we indicate the percentage of cases in which our filter bank outperforms the reference filter bank. From Table 2(a), we can see that, in both the 9/7 and 6/14 cases, designs based on the joint objective function (designated by the suffix “jnt”) have better coding performance than those based on the separable objective function. For example, in the 9/7 case, the joint design (i.e., 9/7-jnt) is able to outperform the reference filter bank about 87% of the time, while the separable design (i.e., 9/7-sep) can only beat the reference filter bank approximately 46% of the time. Similarly, in the 6/14 case, the joint design (i.e., 6/14-jnt) is able to outperform the reference in about 61% of the cases, while the separable design is only able to beat the reference in approximately 20% of the cases. In Table 2(b), we provide the actual distortions obtained for three representative images, with the best result in each case being highlighted. Here, we can see that the filter banks with the jointly-highest coding gains (i.e., the “jnt” case) perform better overall for all three images, in spite of the images having significantly different statistical properties. The above results clearly

Table 2: Lossy compression results for the filter banks designed using different objective functions. (a) Summary statistical results over all twenty-six test images and five bit rates. (b) Specific results for three images.

Transform	Mean (%)	Median (%)	Outperform (%)
9/7-sep	-0.0049	-0.0001	46.15
9/7-iso/jnt	0.1488	0.1070	87.69
6/14-sep	-0.5848	-0.5112	20.77
6/14-iso/jnt	0.0331	0.0279	61.54

Image		gold					target					sar2				
		8	16	32	64	128	8	16	32	64	128	8	16	32	64	128
PSNR (dB)	9/7-sep	36.76	33.77	31.26	29.15	27.33	41.48	33.55	27.08	22.70	19.13	30.33	26.61	24.69	23.55	22.73
	9/7-iso/jnt	36.88	33.84	31.27	29.15	27.32	41.59	33.55	27.19	22.84	19.14	30.35	26.62	24.70	23.55	22.73
	6/14-sep	36.77	33.50	31.07	28.82	27.06	40.94	32.79	27.37	22.39	18.49	30.24	26.49	24.64	23.48	22.59
	6/14-iso/jnt	36.96	33.78	31.15	28.83	27.26	41.68	32.91	27.00	22.43	18.26	30.39	26.49	24.78	23.46	22.68

demonstrate that there is a benefit to taking the isotropic coding gain into consideration during the design process (as is done in the joint and isotropic cases above).

5.2. Design Examples and Coding Results

To demonstrate the effectiveness of our design method, we now present some examples of filter banks generated by our method and evaluate their performance for image coding. In particular, five filter banks constructed with our method are considered. For all of these optimal designs, the joint objective function in (8) was employed, as this was shown earlier to be the most effective choice (i.e., better than the separable case). For comparison purposes, we also consider the well-known 9/7 filter bank from JPEG 2000 [3], which we refer to in this section by the name 9/7-J in order to distinguish it from another filter bank having the same analysis-filter lengths. Several characteristics of our optimal designs as well as the 9/7-J filter bank are shown in Table 3. Due to space constraints, the lifting-filter coefficients for our optimal designs are not presented here, but this information can be obtained from [26, p. 47].

For the reasons discussed earlier in Section 4.1, in the case of our optimal designs, the moments of interest only nearly (but not exactly) vanish. To be more precise, for the purposes of this discussion, we deem any moment with magnitude less than 2×10^{-5} to be nearly vanishing. Although the 9/7-J filter bank has four primal and four dual (exactly) vanishing moments (as indicated in Table 3), this assumes an implementation in exact arithmetic without quantization of the lifting-filter coefficients. In practice,

Table 3: Characteristics of the various filter banks

Transform	$\{L_k\}$	G_{sep}^\ddagger	G_{iso}^\ddagger	b_0^\ddagger	b_1^\ddagger	Dual VM*	Primal VM*	$\hat{h}_0(0)$	$\hat{h}_0(\pi)$	$\hat{h}_1(0)$	$\hat{h}_1(\pi)$
9/7	{2,2,2,2}	14.933	12.181	0.057	0.035	2, 0.004	2, 0.125	1.25	2.1E-6	-9.1E-6	1.60
9/11	{4,2,2}	14.928	12.112	0.111	0.043	2, 0.276	2, 0.858	1.24	-5.6E-6	-9.4E-6	1.61
13/11	{4,2,2,2}	15.041	12.206	0.030	0.027	2, 0.068	2, 0.743	1.20	2.2E-5	-1.9E-5	1.67
17/11	{2,2,4,4}	15.117	12.218	0.031	0.028	2, 0.337	2, 0.572	1.19	9.9E-6	4.9E-6	1.69
13/15	{6,2,2}	14.641	12.074	0.094	0.035	2, 0.169	2, 0.566	1.35	-5.4E-6	-9.4E-6	1.49
9/7-J	{2,2,2,2}	14.973	12.178	0.063	0.035	4, 9.560	4, 16.47	1.23	2.3E-9	3.9E-9	1.63

*index and magnitude of the first moment of dual/primal wavelet coefficient sequence with magnitude greater than 2×10^{-5} ; \ddagger coding gain (in dB) for a six-level decomposition; \ddagger stopband energy as defined by (9)




Figure 5: Synthesis scaling and wavelet functions. (a) Scaling and (b) wavelet functions for the 9/7 design; and (c) Scaling and (d) wavelet functions for the 13/11 design.

however, arithmetic is not exact and these coefficients must also be quantized (since they are irrational). Consequently, in any practical implementation of the 9/7-J filter bank, the moments that are supposed to exactly vanish, only nearly vanish. To provide the reader with more insight into these quantization effects, instead of giving the theoretical values for $\hat{h}_0(\pi)$ and $\hat{h}_1(0)$ for the 9/7-J filter bank in Table 3 (which are exactly zero), we give the actual ones in our implementation. Also, one might wonder to what extent the smoothness of the basis functions associated with the filter bank are impacted by allowing moments to only nearly vanish (instead of exactly vanish). In short, for all practical purposes, the smoothness is unaffected. To demonstrate this, plots of the underlying synthesis scaling and wavelet functions for two of our optimal designs are provided in Fig. 5. Observe that the functions in these plots are visibly quite smooth. That is to say, for all practical purposes, the filter banks behave as if their moments exactly vanished.

Now, we consider the lossy and lossless coding performance of the five filter banks constructed using our method. To evaluate lossy coding performance, each filter bank was used to compress all twenty-six test images at several bit rates. Then, we measured the relative difference between the distortions (in PSNR) obtained with each of our optimal designs and the reference 9/7-J filter bank. The results are summarized

Table 4: Lossy compression results for the various filter banks. (a) Summary statistical results over all twenty-six test images and five bit rates. (b) Specific results for three images.

(a)

Transform	Mean (%)	Median (%)	Outperform (%)
9/7	0.1488	0.1070	87.69
9/11	0.5371	0.0554	59.23
13/11	0.1863	0.0872	74.62
17/11	0.5804	0.2422	77.69
13/15	0.5936	0.1633	68.46

(b)

Image		gold					target					sar2				
Comp. Ratio		8	16	32	64	128	8	16	32	64	128	8	16	32	64	128
PSNR (dB)	9/7	36.88	33.84	31.27	29.15	27.32	41.59	33.55	27.19	22.84	19.14	30.35	26.62	24.70	23.55	22.73
	9/11	37.34	34.00	31.35	29.24	27.39	42.92	33.47	26.65	22.35	18.86	30.30	26.59	24.65	23.52	22.70
	13/11	36.85	33.76	31.24	29.17	27.37	42.12	33.83	27.84	23.11	19.35	30.33	26.60	24.69	23.54	22.74
	17/11	37.17	33.91	31.32	29.17	27.35	42.81	34.00	27.88	23.19	19.46	30.33	26.62	24.69	23.53	22.74
	13/15	37.39	33.95	31.27	29.25	27.34	43.11	33.45	26.84	22.42	18.94	30.37	26.57	24.70	23.54	22.67
	9/7-J	36.75	33.75	31.23	29.16	27.32	41.46	33.54	27.07	22.70	19.16	30.32	26.61	24.69	23.55	22.73

in statistical form in Table 4(a). In particular, we provide the mean and median relative differences in distortion (with positive values corresponding to our designs outperforming the 9/7-J filter bank). We also indicate the percentage of cases in which our optimal design outperforms the 9/7-J filter bank. From these results, it is clear that all of our optimal designs outperform the 9/7-J filter bank in the majority of cases. For example, our 9/7 optimal design outperforms the 9/7-J filter bank 87.69% of the time. Our four other designs outperform the 9/7-J filter bank by margins ranging from about 59 to 78%. The above results are extremely encouraging, given that the 9/7-J filter bank is well known for its exceptional lossy coding performance. In Table 4(b), we provide the actual PSNR results obtained for a representative subset of the test images, where the best result for each bit rate is highlighted. From this table, we can see that, even for images with different statistical properties (such as the three images considered here), our optimal designs outperform the 9/7-J filter bank, sometimes by as much as 1.65 dB. Lastly, we would like to note that our optimal designs also lead to good subjective image quality, comparable to that of the 9/7-J filter bank. In Figure 6, we provide an example of the lossy image reconstructions obtained with the various filter banks. From this figure, we can see that the quality of the image reconstructions produced by our optimal designs is comparable to that obtained with the 9/7-J filter bank.




Figure 6: Parts of the lossy reconstructions obtained after coding the `gold` image at a compression ratio of 32:1 using the (a) 9/7, (b) 9/11, (c) 13/11, (d) 17/11, (e) 13/15, and (f) 9/7-J filter banks.

As can be seen from Table 3, our 9/7 design and the 9/7-J filter bank have the same lifting configuration. While the 9/7-J filter bank has four dual and four primal (exactly) vanishing moments, our 9/7 design has only two dual and two primal nearly-vanishing moments and slightly higher isotropic coding gain as well. With our design approach, by reducing the number of constrained moments and relaxing the requirement that moments be exactly zero, we are able to gain additional freedom, which ultimately allows a higher-performance filter bank to be constructed.

It is also interesting to note that, amongst all of the filter banks that we designed from parameterization (7b), we were not able to find any that outperforms the 9/7-J filter bank for lossy coding. This is likely due to the fact that (7b) is an incomplete parameterization of all PR linear-phase FIR 1D two-channel filter banks with even-length analysis/synthesis filters, while (7a) is a complete parameterization of all PR linear-phase FIR 1D two-channel filter banks with odd-length analysis/synthesis filters (up to a normalization).

In our work, we also evaluated the lossless coding performance of the various filter banks. Each of the filter banks was used to losslessly compress all twenty-six test images. The results are shown in Table 5. In particular, we provide the normalized bit rate (i.e., the reciprocal of compression ratio) for three images as well as the mean taken over all twenty-six test images. Evidently, all of our filter banks perform better overall than the 9/7-J filter bank for lossless coding, with the 9/11 design yielding the best results.

Table 5: Lossless compression results for the various filter banks

Image	Normalized Bit Rate					
	9/7	9/11	13/11	17/11	13/15	9/7-J
gold	0.5666	0.5630	0.5652	0.5657	0.5644	0.5673
target	0.3086	0.2949	0.2964	0.3130	0.2975	0.3173
sar2	0.6351	0.6349	0.6353	0.6353	0.6352	0.6350
Mean [†]	0.4874	0.4771	0.4787	0.4870	0.4828	0.4880

[†]mean taken over all twenty-six test images

5.3. Robustness of the Design Method

Since our proposed design method involves a rather complex optimization, a few comments are worthwhile regarding the robustness of the method. In our work, we have run hundreds of design test cases, using many combinations of lifting configurations (i.e., number of lifting filters and lifting-filter lengths), design parameters (i.e., frequency-response and vanishing-moment constraint parameters) and initial points for optimization, and never once has our method been observed to fail to converge to a solution. Furthermore, the solution obtained has always been both feasible and optimal. Additional experimentation has shown that small perturbations in the initial point either have no effect on the final (locally) optimal solution or serve only to displace it to a nearby (locally) optimum point, with the first of these two scenarios being far more likely. Therefore, from a practical point of view, our method has been demonstrated to be quite robust/stable. This said, there is some theoretical justification for this good behavior. In each reduced-order problem, we use a linear Taylor-series approximation of numerous functions about the current operating point. Since these functions are sufficiently smooth that their first-order partial derivatives always exist, the Taylor approximations are always well defined. To ensure the validity of the solution obtained to each reduced-order problem, the solution is restricted to a small region in which the Taylor approximations are accurate. This restriction is imposed by the β_0 parameter of our design method. As long as β_0 is chosen sufficiently small (e.g., using the value we suggested earlier), the Taylor approximations should always be sufficiently accurate to ensure the validity of the intermediate/final results produced by our method as well as guarantee its convergence.

6. Analysis of the Coding Gain

As was noted earlier in Section 5.1, during the development of our proposed design method, we observed that the filter banks obtained both at the intermediate and final stages of our method always seem to have

higher coding gain with respect to the separable model than with respect to the isotropic one, regardless of which of the three objective functions in (8) is employed. In what follows, we study this behavior in detail and explain the reason for it. Since the results of our subsequent analysis have application beyond the class of filter banks associated with our design method, we consider a superset of this class. In particular, we consider the class of wavelet filter banks constructed from 1D two-channel PR real-coefficient filter banks with at least one dual vanishing moment. This class of filter banks includes most (if not all) separable filter banks that are commonly used in practical image-coding systems.

To begin, let us examine the coding gain formula given by (5). First, we make a few observations regarding the quantities $\{A_k\}_{k=0}^{M-1}$ and $\{B_k\}_{k=0}^{M-1}$ in this formula. Since each of the $\{A_k\}_{k=0}^{M-1}$ is the ratio between two variances (which are nonnegative real numbers) and $A_k = 0$ implies that h_k is the zero sequence (which would preclude PR), we have that $A_k > 0$. For similar reasons, each of the $\{B_k\}_{k=0}^{M-1}$ satisfies $B_k > 0$. Next, we make a comment concerning the correlation coefficient ρ appearing in (6). In the remainder of our analysis, we exclude the possibility that $\rho \in \{0, 1\}$, since these cases are of no practical value and would only serve to complicate the subsequent analysis. Since in practice ρ is typically chosen as $\rho \in [0.90, 0.95]$, we emphasize this range of values in our analysis. This said, however, our results are not strictly limited to ρ in this range. So long as ρ is not too far outside this range, all of our results should still hold. Let us now consider how the coding gain is affected by the choice of image model. From (5), we observe that the only dependence that the coding gain G has on the image model embodied by r is in the value of the $\{A_k\}_{k=0}^{M-1}$. Since $\alpha_k > 0$, G increases as the elements of $\{A_k\}_{k=0}^{M-1}$ decrease. With this in mind, we would like to more carefully consider how the $\{A_k\}_{k=0}^{M-1}$ are affected by the choice of image model.

Before proceeding further, we need to define several new quantities that will be used throughout the remainder of our analysis. The k th factor in the product G from (5) is denoted as G_k (i.e., $G_k = (\frac{\alpha_k}{A_k B_k})^{\alpha_k}$). Hence, $G = \prod_{k=0}^{M-1} G_k$. Let $G_{k,\text{sep}}$ and $G_{k,\text{iso}}$ denote the quantity G_k in the separable and isotropic cases, respectively. Define the quantities: $\Delta A_k = A_{k,\text{iso}} - A_{k,\text{sep}}$, $\Delta r = r_{\text{iso}} - r_{\text{sep}}$, and $\tilde{G}_k = G_{k,\text{iso}}/G_{k,\text{sep}}$.

Consider the expression for $\{A_k\}_{k=0}^{M-1}$ given by (5). Through a change of variable, we can rewrite this expression in a more convenient form as $A_k = \sum_{\mathbf{p} \in \mathbb{Z}^2} r[\mathbf{p}] c_k[\mathbf{p}]$, where $c_k = h'_k \star h'_k$ (i.e., the autocorrelation of h'_k). Using the fact that the Fourier transform preserves inner products and $\hat{c}_k = \left| \widehat{h'_k} \right|^2$ (since h'_k is real), the preceding equation can be rewritten as

$$A_k = \frac{1}{4\pi^2} \int_{[-\pi, \pi]^2} \hat{r}(\boldsymbol{\omega}) \hat{c}_k(\boldsymbol{\omega}) d\boldsymbol{\omega} = \frac{1}{4\pi^2} \int_{[-\pi, \pi]^2} \hat{r}(\boldsymbol{\omega}) \left| \widehat{h'_k}(\boldsymbol{\omega}) \right|^2 d\boldsymbol{\omega}. \quad (15)$$

Now, we make some observations regarding \hat{r} and $\hat{c}_k = \left| \widehat{h'_k} \right|^2$. One can easily verify that the sequence r

has 8-fold symmetry in both the separable (i.e., $r = r_{\text{sep}}$) and isotropic (i.e., $r = r_{\text{iso}}$) cases. That is, r has quadrantal symmetry while additionally satisfying $r[n_0, n_1] = r[n_1, n_0]$. Thus, it follows that \hat{r} has quadrantal symmetry in both the separable and isotropic cases. Moreover, since all of the filters $\{h'_k\}_{k=0}^{M-1}$ are separable, their autocorrelation sequences $\{c_k\}_{k=0}^{M-1}$ are also separable. Due to the separability of h'_k , the sequence c_k has quadrantal symmetry, which in turn implies that \hat{c}_k has quadrantal symmetry. Due to the quadrantal symmetry of \hat{r} and $\hat{c}_k = \left| \widehat{h'_k} \right|^2$, we can rewrite (15) as

$$A_k = \frac{1}{\pi^2} \int_{[0, \pi]^2} \hat{r}(\boldsymbol{\omega}) \left| \widehat{h'_k}(\boldsymbol{\omega}) \right|^2 d\boldsymbol{\omega}. \quad (16)$$

Thus, we have that

$$\Delta A_k = \frac{1}{\pi^2} \int_{[0, \pi]^2} \widehat{\Delta r}(\boldsymbol{\omega}) \left| \widehat{h'_k}(\boldsymbol{\omega}) \right|^2 d\boldsymbol{\omega}, \quad (17)$$

where $\widehat{\Delta r} = \hat{r}_{\text{iso}} - \hat{r}_{\text{sep}}$.

Now, we seek to determine a formula for the quantity \hat{r} in (16) and (17) for the separable and isotropic cases. The 1D sequence from which r_{sep} is composed is $r_{\text{sep}, 1\text{d}}[n] = \rho^{|n|}$, which has the Fourier transform $\hat{r}_{\text{sep}, 1\text{d}}(\omega) = \frac{1-\rho^2}{1-2\rho\cos\omega+\rho^2}$. Thus, the Fourier transform of r_{sep} is simply

$$\hat{r}_{\text{sep}}(\omega_0, \omega_1) = \frac{(1-\rho^2)^2}{(1-2\rho\cos\omega_0+\rho^2)(1-2\rho\cos\omega_1+\rho^2)}. \quad (18)$$

Now, we consider the Fourier transform of r_{iso} . Using the 8-fold symmetry of r_{iso} , we can show that

$$\begin{aligned} \hat{r}_{\text{iso}}(\omega_0, \omega_1) &= 1 + \theta(\omega_0, \rho) + \theta(\omega_1, \rho) + \theta(\omega_0 + \omega_1, \rho^{\sqrt{2}}) + \theta(\omega_1 - \omega_0, \rho^{\sqrt{2}}) \\ &+ 4 \sum_{i=2}^{\infty} \sum_{k=1}^{i-1} \rho^{\sqrt{i^2+k^2}} [\cos(i\omega_0) \cos(k\omega_1) + \cos(k\omega_0) \cos(i\omega_1)], \end{aligned} \quad (19)$$

where $\theta(\omega, \alpha) = \frac{2\alpha\cos\omega-2\alpha^2}{1-2\alpha\cos\omega+\alpha^2}$.

In passing, we would like to briefly make note of an alternative scheme for computing \hat{r}_{iso} . By observing that the Fourier transform of the function $f(\mathbf{t}) = \rho^{||\mathbf{t}||_2}$ is $\hat{f}(\boldsymbol{\omega}) = -2\pi(\ln^2 \rho + ||\boldsymbol{\omega}||_2^2)^{-3/2} \ln \rho$ [33, Eqn. (3)] and r_{iso} is a sampled version of f , we can conclude $\hat{r}_{\text{iso}}(\boldsymbol{\omega}) = -2\pi(\ln \rho) \sum_{\mathbf{k} \in \mathbb{Z}^2} [\ln^2 \rho + ||\boldsymbol{\omega} - 2\pi\mathbf{k}||_2^2]^{-3/2}$. Experimentally, however, this formula for \hat{r}_{iso} has been observed to converge more slowly than (19), and is therefore less useful for computational purposes.

For future reference, \hat{r}_{sep} and \hat{r}_{iso} are plotted in Figs. 7 and 8 for $\rho = 0.95$. In each case, only the first quadrant is shown as the remainder of the plot can be trivially deduced from quadrantal symmetry. For each of \hat{r}_{sep} and \hat{r}_{iso} , the general shape of the plot remains the same for all $\rho \in [0.90, 0.95]$, but the decay rate

increases with increasing ρ . Also, $\widehat{\Delta r}$ is plotted for $\rho = 0.95$ in Fig. 9. The gray-shaded region in the contour plot corresponds to where the plotted function is negative. As ρ increases, the decay rate of $\widehat{\Delta r}$ increases, and there is a larger region where $\widehat{\Delta r}$ is positive. In all of the above plots, the scale of the horizontal and vertical axes have been normalized so that a value of one corresponds to the Nyquist frequency.

Now, we will use the expressions for A_k and ΔA_k in (16) and (17) in order to gain some additional insight into the behavior of the coding gain. In what follows, we employ the definitions introduced in Section 3 (e.g., L , $M = 3L + 1$, level, orient, etc.) and k is used to denote the channel index, where $k \in \{0, 1, 2, \dots, 3L\}$.

For three filter banks, we have computed the various coding-gain-related quantities for $L \in \{1, 2, 3, 4\}$ with $\rho = 0.95$, the results of which can be found in Table 6. (Note that the coding gains in this table are stated as unitless quantities, and not in dB.) The first filter bank is an orthonormal system with ideal filters. This corresponds to a system with 1D prototype filters each having a gain of $\sqrt{2}$ and 0 in their passbands and stopbands, respectively. (Thus, we have that h'_k has passband and stopband gains of $2^{\text{level}(k)+1}$ and 0, respectively.) The other two filter banks are the well-known Haar and 9/7 [3] systems, and were chosen for comparison purposes as examples of systems with nonideal filters. The nonideal filter banks are normalized such that their 1D prototype filters have DC/Nyquist gains of $\sqrt{2}$, as this facilitates more direct comparisons with the ideal filter bank under consideration. Note that there is no loss of generality in considering only this particular normalization, as the coding gain for PR filter banks is invariant to scaling and translation of the analysis/synthesis filter impulse responses. We will refer to the results in the above table in some of the discussion that follows.

Now, we more carefully examine the formula for the coding gain G in (5). As k increases (which corresponds to $\text{level}(k)$ increasing), α_k decays exponentially to zero (i.e., $\alpha_k = 4^{-\text{level}(k)-1}$). This implies that G_k rapidly approaches one as k increases. Thus (since $G = \prod_{k=0}^{M-1} G_k$), the most significant contributions to the coding gain G come from G_k associated with small k , especially those k associated with the 0th level (i.e., $k \in \{0, 1, 2\}$ or $k \in \{0, 1, 2, 3\}$ for $L \geq 2$ and $L = 1$, respectively). As k increases, the numerator and denominator of $\tilde{G}_k = \frac{G_{k,\text{iso}}}{G_{k,\text{sep}}}$ each rapidly approach one. Consequently, \tilde{G}_k also rapidly approaches one. Thus (since $G_{\text{iso}}/G_{\text{sep}} = \prod_{k=0}^{M-1} \tilde{G}_k$), any difference in the separable and isotropic coding gains is most strongly influenced by \tilde{G}_k for small k . Furthermore, simple algebraic manipulation shows that

$$\tilde{G}_k = \left(\frac{A_{k,\text{sep}}}{A_{k,\text{sep}} + \Delta A_k} \right)^{\alpha_k} = \left(1 + \frac{\Delta A_k}{A_{k,\text{sep}}} \right)^{-\alpha_k}. \quad (20)$$

Consequently, $G_{k,\text{sep}}$ and $G_{k,\text{iso}}$ differ most when $|\Delta A_k|$ is large relative to $A_{k,\text{sep}}$ (i.e., $\frac{|\Delta A_k|}{A_{k,\text{sep}}}$ is large), with $G_{k,\text{sep}} > G_{k,\text{iso}}$ (i.e., G_{sep} favored) if $\Delta A_k > 0$ and $G_{k,\text{sep}} < G_{k,\text{iso}}$ (i.e., G_{iso} favored) if $\Delta A_k < 0$. All of the

above observations can be seen to be consistent with the data in Table 6.

In what follows, it is instructive to consider the ideal filter bank introduced above. In this case, from (16) and (17), A_k and ΔA_k are simply (up to scale) the integrals of \hat{r} and $\widehat{\Delta r}$ over the first-quadrant portion of the passband of the ideal filter h'_k . In particular, we have $A_k = \frac{\beta_k^2}{\pi^2} \int_{P_k} \hat{r}(\boldsymbol{\omega}) d\boldsymbol{\omega}$ and $\Delta A_k = \frac{\beta_k^2}{\pi^2} \int_{P_k} \widehat{\Delta r}(\boldsymbol{\omega}) d\boldsymbol{\omega}$, where P_k is the first-quadrant portion of the passband of h'_k , and β_k is the passband gain of h'_k . To envision what the result of such integrations will be, it is helpful to imagine the frequency-plane tiling of Fig. 4 superimposed on the plots of \hat{r}_{sep} , \hat{r}_{iso} , and $\widehat{\Delta r}$ in Figs. 7, 8, and 9.

Consider $k = 0$ (i.e., the HH₀ band) in the case of the ideal filter bank. For $\boldsymbol{\omega}$ in the (HH₀) passband, $\hat{r}_{\text{sep}}(\boldsymbol{\omega})$ and $\hat{r}_{\text{iso}}(\boldsymbol{\omega})$ are both very small but their difference $\widehat{\Delta r}(\boldsymbol{\omega})$ is relatively large (compared to $\hat{r}_{\text{sep}}(\boldsymbol{\omega})$ and $\hat{r}_{\text{iso}}(\boldsymbol{\omega})$). So, in (20), $A_{k,\text{sep}}$ is small and $\frac{\Delta A_k}{A_{k,\text{sep}}}$ is large in magnitude and positive. Consequently, from (20), \tilde{G}_k is significantly less than one, meaning that G_{sep} is very strongly favored over G_{iso} . A similar argument also applies to the other HH bands, but the influence on the coding gain G is less significant in these cases since \tilde{G}_k rapidly approaches one as k increases. Now consider what happens in the case of nonideal filters. Even in this case, it is very difficult for the above qualitative behavior to change. To obtain significantly different behavior, \hat{h}'_k would have to be very large along the axes where $\widehat{\Delta r}$ is most negative (and therefore most favorable to higher G_{iso}). Due to the presence of at least one dual vanishing moment, however, \hat{h}'_k must be zero along both axes. So, even in case of nonideal filters, G_{sep} is still likely to be strongly favored by the HH bands.

In the case of the LH and HL bands, due to the presence of at least one dual vanishing moment, $\hat{h}'_k(\omega_0, \omega_1)$ must be zero along exactly one of $\omega_0 = 0$ or $\omega_1 = 0$. We observe that $\widehat{\Delta r}$ (shown in Fig. 9) is largest in magnitude along the axes, with large positive values near the origin and large negative values elsewhere on the axes. Thus, in the calculation of ΔA_k , the large positive values near the origin in $\widehat{\Delta r}$ are effectively cancelled since \hat{h}'_k is zero along one axis, and we are left with large negative values along the other axis. This leads to $\Delta A_k < 0$ so that $\tilde{G}_k > 1$ and the LH and HL bands favor G_{iso} over G_{sep} . Lastly, we note that the LH and HL subbands do not have as much impact on the coding gain G as the HH bands. This is due to that fact $A_{k,\text{sep}}$ is larger for the LH and HL subbands, making it more difficult for $|\Delta A_k|$ to be large relative to $A_{k,\text{sep}}$.

An examination of the data for the filter banks in Table 6 shows that all of observations made above are consistent with this data. For example, the coding gain is most strongly influenced by the HH bands, with HH₀ figuring most prominently. The LL and HH bands favor G_{sep} (i.e., $\tilde{G}_k < 1$), while the LH and HL bands




Figure 7: Contour plot of \hat{r}_{sep} for $\rho = 0.95$.




Figure 8: Contour plot of \hat{r}_{iso} for $\rho = 0.95$.




Figure 9: Contour plot of $\widehat{\Delta r}$ for $\rho = 0.95$.

favor G_{iso} (i.e., $\tilde{G}_k > 1$).

6.1. Additional Commentary

In addition to the three filter banks for which results are presented (in Table 6), we have also considered quite a number of other filter banks, some of which were considered for adoption in the JPEG 2000 Part-1 standard [3], while others were produced at various stages of our optimal design method. Similar trends (to those described above) were also found in the case of these other filter banks. Moreover, we were not able to find any filter bank (belonging to the class of filter banks under consideration here) for which $G_{\text{iso}} > G_{\text{sep}}$. Based on our analysis, there is good reason to believe that such filter banks probably do not exist.

As an aside, we note that, for the class of filter banks being considered here, if the constraint of having at least one dual vanishing moment is dropped, it is possible for $G_{\text{iso}} = G_{\text{sep}}$. For example, this result is trivially obtained by the PR system with the 1D prototype analysis filters $H_0(z) = 1$ and $H_1(z) = z$. So, obviously, for filter banks outside the class being considered here, the condition $G_{\text{sep}} > G_{\text{iso}}$ may be violated.

7. Conclusions

In this manuscript, we have proposed a novel optimization-based method for the design of wavelet filter banks for image coding. Our method yields linear-phase PR systems with high coding gain, good frequency selectivity, and certain prescribed vanishing-moment properties. Several examples of filter banks constructed using our method were presented and shown to be highly effective for image coding. In particular, our optimal designs outperformed the well-known 9/7 filter bank from the JPEG-2000 standard for both lossy and lossless compression, an impressive feat given that this filter bank is known for its exceptional

Table 6: Intermediate coding gain quantities with $p = 0.95$ for the (a) ideal (b) Haar, and (c) 9/7 filter banks; and the final coding gains for the (d) ideal, (e) Haar, and (f) 9/7 filter banks.

(a)										(b)									
k	Band	$A_{k,sep}$	$A_{k,iso}$	ΔA_k	$\frac{\Delta A_k}{A_{k,sep}}$	$\frac{1}{\alpha_k}$	$G_{k,sep}$	$G_{k,iso}$	\tilde{G}_k	k	Band	$A_{k,sep}$	$A_{k,iso}$	ΔA_k	$\frac{\Delta A_k}{A_{k,sep}}$	$\frac{1}{\alpha_k}$	$G_{k,sep}$	$G_{k,iso}$	\tilde{G}_k
0	HH ₀	0.001	0.025	0.024	22.85	4	5.53	2.50	0.45	0	HH ₀	0.002	0.030	0.027	11.01	4	4.47	2.40	0.53
1	HL ₀	0.064	0.039	-0.024	-0.378	4	1.98	2.23	1.12	1	HL ₀	0.097	0.069	-0.027	-0.282	4	1.78	1.94	1.08
2	LH ₀	0.064	0.039	-0.024	-0.378	4	1.98	2.23	1.12	2	LH ₀	0.097	0.069	-0.027	-0.282	4	1.78	1.94	1.08
3	LL ₀	3.870	3.894	0.024	0.006	4	0.71	0.71	0.99	3	LL ₀	3.802	3.830	0.027	0.007	4	0.71	0.71	0.99
3	HH ₁	0.004	0.050	0.046	10.95	16	1.40	1.20	0.85	3	HH ₁	0.020	0.100	0.079	3.862	16	1.27	1.15	0.90
4	HL ₁	0.250	0.132	-0.117	-0.470	16	1.09	1.13	1.04	4	HL ₁	0.540	0.383	-0.156	-0.289	16	1.03	1.06	1.02
5	LH ₁	0.250	0.132	-0.117	-0.470	16	1.09	1.13	1.04	5	LH ₁	0.540	0.383	-0.156	-0.289	16	1.03	1.06	1.02
6	LL ₁	14.76	15.07	0.307	0.020	16	0.84	0.84	0.99	6	LL ₁	14.10	14.45	0.343	0.024	16	0.84	0.84	0.99
6	HH ₂	0.052	0.306	0.254	4.866	64	1.04	1.01	0.97	6	HH ₂	0.237	0.622	0.384	1.618	64	1.02	1.00	0.98
7	HL ₂	1.680	0.958	-0.721	-0.429	64	0.99	1.00	1.00	7	HL ₂	3.424	2.579	-0.844	-0.246	64	0.98	0.98	1.00
8	LH ₂	1.680	0.958	-0.721	-0.429	64	0.99	1.00	1.00	8	LH ₂	3.424	2.579	-0.844	-0.246	64	0.98	0.98	1.00
9	LL ₂	53.98	56.54	2.557	0.047	64	0.93	0.93	0.99	9	LL ₂	49.34	52.02	2.679	0.054	64	0.94	0.94	0.99
9	HH ₃	0.752	2.333	1.581	2.103	256	1.00	0.99	0.99	9	HH ₃	2.702	4.508	1.806	0.668	256	0.99	0.99	0.99
10	HL ₃	11.62	7.325	-4.300	-0.369	256	0.99	0.99	1.00	10	HL ₃	20.39	16.57	-3.814	-0.187	256	0.98	0.98	1.00
11	LH ₃	11.62	7.325	-4.300	-0.369	256	0.99	0.99	1.00	11	LH ₃	20.39	16.57	-3.814	-0.187	256	0.98	0.98	1.00
12	LL ₃	179.7	197.4	17.71	0.098	256	0.97	0.97	0.99	12	LL ₃	153.9	170.4	16.53	0.107	256	0.98	0.98	0.99

(c)									
k	Band	$A_{k,sep}$	$A_{k,iso}$	ΔA_k	$\frac{\Delta A_k}{A_{k,sep}}$	$\frac{1}{\alpha_k}$	$G_{k,sep}$	$G_{k,iso}$	\tilde{G}_k
0	HH ₀	0.001	0.024	0.023	21.81	4	5.40	2.47	0.45
1	HL ₀	0.064	0.040	-0.023	-0.360	4	1.97	2.20	1.11
2	LH ₀	0.064	0.040	-0.023	-0.360	4	1.97	2.20	1.11
3	LL ₀	3.811	3.834	0.023	0.006	4	0.72	0.72	0.99
3	HH ₁	0.008	0.090	0.081	9.265	16	1.35	1.16	0.86
4	HL ₁	0.352	0.199	-0.153	-0.434	16	1.06	1.10	1.03
5	LH ₁	0.352	0.199	-0.153	-0.434	16	1.06	1.10	1.03
6	LL ₁	14.12	14.44	0.314	0.022	16	0.84	0.84	0.99
6	HH ₂	0.098	0.511	0.412	4.176	64	1.03	1.00	0.97
7	HL ₂	2.213	1.343	-0.869	-0.392	64	0.98	0.99	1.00
8	LH ₂	2.213	1.343	-0.869	-0.392	64	0.98	0.99	1.00
9	LL ₂	49.61	52.10	2.496	0.050	64	0.93	0.93	0.99
9	HH ₃	1.325	3.652	2.327	1.755	256	0.99	0.99	0.99
10	HL ₃	14.44	9.829	-4.615	-0.319	256	0.98	0.99	1.00
11	LH ₃	14.44	9.829	-4.615	-0.319	256	0.98	0.99	1.00
12	LL ₃	157.4	173.2	15.84	0.100	256	0.97	0.97	0.99

(d)			(e)			(f)		
L	G_{sep}	G_{iso}	L	G_{sep}	G_{iso}	L	G_{sep}	G_{iso}
1	21.84	12.54	1	10.25	6.49	1	15.25	8.71
2	36.54	19.45	2	16.70	10.00	2	27.44	14.56
3	37.64	19.83	3	18.25	10.86	3	30.76	16.16
4	36.23	19.07	4	18.50	11.01	4	31.34	16.46

lossy coding performance. Our design method supports the maximization of the coding gain with respect to the separable or isotropic image model, or jointly with respect to both models. Through experimental results, we have demonstrated that the joint maximization of both coding gains leads to filter banks with improved coding performance on average, relative to simply maximizing the separable coding gain as is typically done.

Also, in this manuscript, we have studied the coding gain of separable 2D wavelet filter banks derived from 1D two-channel real-coefficient PR filter banks with at least one dual vanishing moment. We have explained why, for such filter banks, it is extremely difficult (if not impossible) for the isotropic coding gain to exceed the separable coding gain. The new insight provided by our analysis may prove helpful in the design of improved filter banks for image coding. For example, if one is trying to simultaneously maximize the minimum of G_{sep} and G_{iso} as in the case of our work, our analysis shows that this is practically equivalent to optimizing G_{iso} alone.

8. Acknowledgment

This work was supported by the Natural Sciences and Engineering Research Council of Canada.

A. Proofs

Theorem 1. Consider a symmetric/antisymmetric sequence h defined on \mathbb{Z} of the form $h[n] = sh[2c - n]$, where $c \in \frac{1}{2}\mathbb{Z}$ is the center of symmetry and $s \in \{-1, 1\}$ (i.e., $s = 1$ and $s = -1$ correspond the symmetric and antisymmetric cases, respectively). Let m_k denote the k th moment of h . For $s = 1$: if $m_k = 0$ for $k \in \{0, 1, \dots, 2N\}$, then $m_{2N+1} = 0$. For $s = -1$: $m_0 = 0$ and if $m_k = 0$ for $k \in \{1, 2, \dots, 2N + 1\}$, then $m_{2N+2} = 0$.

Proof. Suppose that $s = 1$. Due to the symmetry properties of h , we have that $\hat{h}(\omega) = \sum_{n \in \mathbb{Z}} h[n] e^{-j\omega n} \cos([n - c]\omega)$. Taking the ℓ th derivative of \hat{h} (via the Leibniz rule) and using the fact that $m_\ell = j^\ell \hat{h}^{(\ell)}(0)$, we obtain

$$m_\ell = \sum_{k \in \{0, 1, \dots, \lfloor \ell/2 \rfloor\}} \alpha_{\ell, k}(c) f_k(c), \quad (21)$$

where $\alpha_{\ell, k}(c) = (-1)^\ell \binom{\ell}{2k} (-c)^{\ell-2k}$ and $f_k(c) = \sum_{n \in \mathbb{Z}} h[n] (n - c)^{2k}$. Thus, from the preceding equation we have $m_0 = \alpha_{0,0}(c) f_0(c)$, $m_1 = \alpha_{1,0}(c) f_0(c)$, $m_2 = \alpha_{2,0}(c) f_0(c) + \alpha_{2,1}(c) f_1(c)$, $m_3 = \alpha_{3,0}(c) f_0(c) + \alpha_{3,1}(c) f_1(c)$, \dots , and $m_{2L} = \sum_{k=0}^L \alpha_{2L, k}(c) f_k(c)$. From this, we can deduce (by induction) that

$$m_\ell = 0 \text{ for } \ell \in \{0, 1, \dots, 2L\} \text{ implies } f_\ell(c) = 0 \text{ for } \ell \in \{0, 1, \dots, L\}. \quad (22)$$

Suppose now that $m_\ell = 0$ for $\ell \in \{0, 1, \dots, 2L\}$ so that (22) applies. Since (22) applies, we have from (21) that $m_{2L+1} = \sum_{k \in \{0, 1, \dots, \lfloor (2L+1)/2 \rfloor\}} \alpha_{2L+1, k}(c) f_k(c) = \sum_{k=0}^L \alpha_{2L+1, k}(c) f_k(c) = 0$. Thus, we have proven the desired result for $s = 1$.

Suppose that $s = -1$. Trivially, $m_0 = 0$, as $\hat{h}(0) = 0$ (due to the antisymmetry of h). The remainder of the proof is obtained in a manner similar to the $s = 1$ case above. \square

References

- [1] J. M. Shapiro, Embedded image coding using zerotrees of wavelet coefficients, *IEEE Trans. on Signal Processing* 41 (12) (1993) 3445–3462.
- [2] A. Said, W. A. Pearlman, A new fast and efficient image codec based on set partitioning in hierarchical trees, *IEEE Trans. on Circuits and Systems for Video Technology* 6 (3) (1996) 243–250.
- [3] ISO/IEC 15444-1: Information technology—JPEG 2000 image coding system—Part 1: Core coding system (2000).
- [4] J. Katto, Y. Yasuda, Performance evaluation of subband coding and optimization of its filter coefficients, in: *Proc. of SPIE Visual Communications and Image Processing*, Vol. 1605, 1991, pp. 95–106.
- [5] M. J. T. Smith, T. P. Barnwell, Exact reconstruction techniques for tree-structured subband coders, *IEEE Trans. on Acoustics, Speech, and Signal Processing* 34 (3) (1986) 434–441.
- [6] P. P. Vaidyanathan, P.-Q. Hoang, Lattice structures for optimal design and robust implementation of two-channel perfect-reconstruction QMF banks, *IEEE Trans. on Acoustics, Speech, and Signal Processing* 36 (1) (1988) 81–94.
- [7] J. Katto, K. Komatsu, Y. Yasuda, Short-tap and linear-phase PR filter banks for subband coding of images, in: *Proc. of SPIE Visual Communications and Image Processing*, Vol. 1818, 1992, pp. 735–746.
- [8] O. Rioul, P. Duhamel, A Remez exchange algorithm for orthonormal wavelets, *IEEE Trans. on Circuits and Systems—II: Analog and Digital Signal Processing* 41 (8) (1994) 550–560.
- [9] S.-M. Phoong, C. W. Kim, P. P. Vaidyanathan, A new class of two-channel biorthogonal filter banks and wavelet bases, *IEEE Trans. on Signal Processing* 43 (3) (1995) 649–665.
- [10] C.-K. Goh, Y. C. Lim, An efficient algorithm to design weighted minimax perfect reconstruction quadrature mirror filter banks, *IEEE Trans. on Signal Processing* 47 (12) (1999) 3303–3314.
- [11] R. Ansari, C. W. Kim, M. Dedovic, Structure and design of two-channel filter banks derived from a triplet of halfband filters, *IEEE Trans. on Circuits and Systems—II: Analog and Digital Signal Processing* 46 (12) (1999) 1487–1496.
- [12] S. C. Chan, C. K. S. Pun, K. L. Ho, The design of a class of perfect reconstruction two-channel FIR linear-phase filterbanks and wavelets bases using semidefinite programming, *IEEE Signal Processing Letters* 11 (2) (2004) 297–300.
- [13] S. C. Chan, C. K. S. Pun, K. L. Ho, New design and realization techniques for a class of perfect reconstruction two-channel FIR filterbanks and wavelet bases, *IEEE Trans. on Signal Processing* 52 (7) (2004) 2135–2141.
- [14] S. C. Chan, K. S. Yeung, On the design and multiplierless realization of perfect reconstruction triplet-based FIR filter banks and wavelet bases, *IEEE Trans. on Circuits and Systems—I: Regular Papers* 51 (8) (2004) 1476–1491.
- [15] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*, Prentice-Hall, Englewood Cliffs, NJ, USA, 1993.

- [16] W. Sweldens, The lifting scheme: A custom-design construction of biorthogonal wavelets, *Applied and Computational Harmonic Analysis* 3 (2) (1996) 186–200.
- [17] Y. Chen, M. D. Adams, W.-S. Lu, Design of optimal quincunx filter banks for image coding, in: *Proc. of IEEE International Symposium on Circuits and Systems*, 2006, pp. 2041–2044.
- [18] D. Xu, M. D. Adams, Design of high-performance filter banks for image coding, in: *Proc. of IEEE International Symposium on Signal Processing and Information Technology*, Vancouver, BC, Canada, 2006, pp. 868–873.
- [19] M. D. Adams, On the coding gain of separable 2D wavelet filter banks, in: *Proc. of IEEE International Conference on Image Processing*, San Diego, CA, USA, 2008, pp. 1204–1207.
- [20] P. Lancaster, M. Tismenetsky, *The Theory of Matrices: with Applications*, 2nd Edition, Academic Press, San Diego, CA, USA, 1985.
- [21] A. R. Calderbank, I. Daubechies, W. Sweldens, B.-L. Yeo, Wavelet transforms that map integers to integers, *Applied and Computational Harmonic Analysis* 5 (3) (1998) 332–369.
- [22] M. D. Adams, R. K. Ward, Symmetric-extension-compatible reversible integer-to-integer wavelet transforms, *IEEE Trans. on Signal Processing* 51 (10) (2003) 2624–2636.
- [23] M. D. Adams, Reversible integer-to-integer wavelet transforms for image coding, Ph.D. thesis, University of British Columbia, Vancouver, BC, Canada, URL <http://www.ece.uvic.ca/~mdadams> (Sep. 2002).
- [24] M. S. Lobo, L. Vandenberghe, S. Boyd, H. Le Bret, Applications of second-order cone programming, *Linear Algebra and its Applications* 248 (1998) 193–228.
- [25] A. Antoniou, W.-S. Lu, *Practical Optimization: Algorithms and Engineering Applications*, 1st Edition, Springer, 2007.
- [26] D. Xu, Improved subband-based and normal-mesh-based image coding, M.A.Sc. thesis, University of Victoria, Victoria, BC, Canada, URL <http://hdl.handle.net/1828/284> (Aug. 2007).
- [27] J. F. Sturm, Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones, *Optimization Methods and Software* 11–12 (1999) 625–653.
- [28] M. Grant, S. Boyd, Y. Ye, *cvx Users' Guide for cvx version 1.1 (build 586)*, URL http://www.stanford.edu/~boyd/cvx/cvx_usrguide.pdf (Jan. 2008).
- [29] J. Lofberg, YALMIP: A toolbox for modeling and optimization in MATLAB, in: *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.
- [30] ISO/IEC JTC 1/SC 29/WG 1 N 545, JPEG-2000 test images (Jul. 1997).
- [31] M. D. Adams, ELEC 545 project: A wavelet-based lossy/lossless image compression system (Apr. 1999).
- [32] M. D. Adams, F. Kossentini, Reversible integer-to-integer wavelet transforms for image compression: Performance evaluation and analysis, *IEEE Trans. on Image Processing* 9 (6) (2000) 1010–1024.
- [33] J. B. O'Neal, T. R. Natarajan, Coding isotropic images, *IEEE Trans. on Information Theory* 23 (6) (1977) 697–707.