

A FILTER DESIGN TECHNIQUE FOR STEERABLE PYRAMID IMAGE TRANSFORMS

Anestis Karasaridis

Eero Simoncelli

GRASP Laboratory
University of Pennsylvania
Philadelphia, PA 19104-6228

ABSTRACT

We describe a novel recursive filter design technique for multi-scale “pyramid” transforms. The recursion in the design technique follows that of the pyramid construction, and allows us to solve a reduced design problem at each step. We demonstrate the use of this technique by designing filters of various orientation bandwidths for use in a “steerable pyramid” image transform.¹

1. INTRODUCTION

Recursive multi-scale transforms (e.g., *wavelets*) are now a standard tool in signal and image processing. The *steerable pyramid* is a particular variant of this type of transform which has been found useful in a variety of image processing applications [1, 2, 3, 4]. In this linear decomposition, an image is subdivided into a collection of subbands localized in scale (they have octave bandwidth) and orientation (they have orientation bandwidths of $2\pi/m$, m an integer). The transform is computed recursively using convolution and decimation operations, and it is “self-inverting”.² The advantages of this representation are that the subbands are translation- and rotation-invariant.

The filters used in constructing the steerable pyramid are highly constrained. The radial component of their Fourier transforms must obey a recursive system diagram, which we describe below. The angular (orientation) tuning of the filters is constrained by the property of *steerability* [1]. A set of filters form a steerable basis if (1) they are rotated copies of each other, and (2) a copy of the filter at any orientation may be computed via a *linear combination* of the basis filters. The simplest example of a steerable basis is a set of $n+1$ n th-order directional derivatives of a circular-symmetric function. For the purposes of this paper, we limit ourselves to steerable filters that are directional derivatives.

Given these constraints, proper filter design is important to produce a usable transform. In the following sections, we

¹ Source code and filter kernels for the steerable pyramid are available via anonymous ftp from `ftp.cis.upenn.edu` in directory `pub/eero/steerpyr.tar.Z`

² By this, we mean that the matrix corresponding to the inverse transformation is equal to the transpose of the forward transformation matrix. In the wavelet literature, such a transform is called a *tight frame* [5].




Figure 1. Idealized depiction of the single-stage first derivative (i.e., two orientation band) steerable pyramid transform.

describe in detail the design constraints, develop a design algorithm, and use it to construct three sets of filters with different orientation bandwidths.

2. FILTER CONSTRAINTS

The decomposition is best defined in the Fourier domain, where the subbands are (ideally) polar-separable. The frequency tiling of the single-stage subband decomposition is shown in figure 1, for the case of $n = 1$. $\{B_k(\vec{\omega}) | k = 0, 1\}$ are band-pass oriented filters, $H_0(\vec{\omega})$ is a non-oriented high-pass filter, and $L_1(\vec{\omega})$ is a narrowband low-pass filter. The system diagram for a single stage of the steerable pyramid is illustrated in figure 2. Recursion is implemented by inserting the portion of the diagram enclosed in the dashed box at the location of the filled circle.

Analysis of the system of figure 2 shows that the reconstructed image in the frequency domain is:

$$\hat{X}(\vec{\omega}) = \left\{ |H_0(\vec{\omega})|^2 + |L_0(\vec{\omega})|^2 (|L_1(\vec{\omega})|^2 + \sum_{k=0}^n |B_k(\vec{\omega})|^2) \right\} X(\vec{\omega}) + a.t. \quad (1)$$

where *a.t.* indicates the aliasing terms. To ensure elimination of the aliasing terms, the L_1 filter should be constrained to have a zero response for frequencies higher than $\pi/2$ in both ω_x and ω_y . Furthermore, to avoid amplitude distortion the transfer function of the system should be equal to one. We describe these constraints, and the steerability constraint in more detail in the following sections.

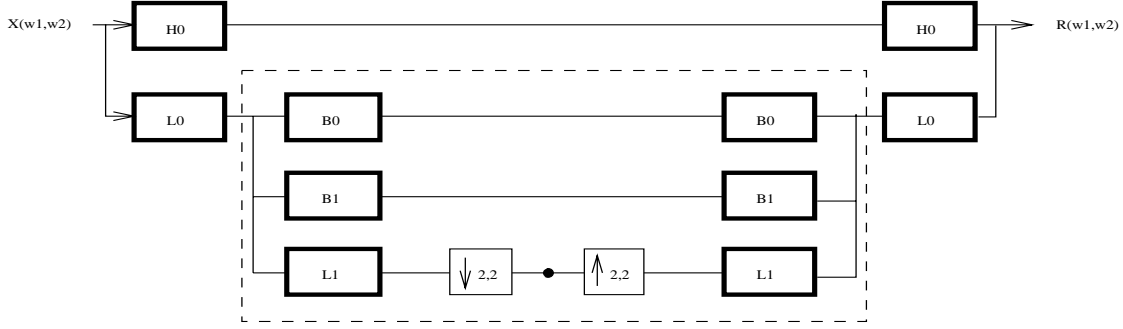


Figure 2. System diagram for a first derivative steerable pyramid. The image is initially divided into high and low-pass portions using filters $H_0(\vec{\omega})$ and $L_0(\vec{\omega})$. The low-pass branch is then further divided into low-pass and oriented band-pass portions using filters $L_1(\vec{\omega})$ and $B_k(\vec{\omega})$. A recursive pyramid is constructed by inserting the portion of the diagram enclosed in the dashed box at the location of the filled circle.

2.1. Perfect Reconstruction Constraints

From the above, the constraints for perfect reconstruction are:

1. Unity system response amplitude:

$$|L_0(\vec{\omega})|^2 \left[|L_1(\vec{\omega})|^2 + \sum_{k=0}^n |B_k(\vec{\omega})|^2 \right] + |H_0(\vec{\omega})|^2 = 1. \quad (2)$$

2. Recursion relationship. The low-pass branch of the diagram must be unaffected by insertion of the recursive portion of the system (see caption of figure 2):

$$|L_1(\vec{\omega}/2)|^2 \left[|L_1(\vec{\omega})|^2 + \sum_{k=0}^n |B_k(\vec{\omega})|^2 \right] = |L_1(\vec{\omega}/2)|^2. \quad (3)$$

3. Aliasing cancellation (for a circular symmetric filter):

$$L_1(\vec{\omega}) = 0, \quad \text{for } |\vec{\omega}| > \pi/2. \quad (4)$$

2.2. Angular Constraints

The angular constraint on the band-pass filters $B_k(\vec{\omega})$ is derived from the condition of steerability [1]. This condition for derivative steerable filters can be expressed as:

$$B_k(\vec{\omega}) = B(\vec{\omega}) [-j \cos(\theta - \theta_k)]^n, \quad (5)$$

where $\theta = \arg(\vec{\omega})$, $\theta_k = \pi k / (n + 1)$ for $k \in \{0, 1, \dots, n\}$, and

$$B(\vec{\omega}) = \sqrt{\sum_{k=0}^n |B_k(\vec{\omega})|^2}.$$

The constraint in equation (5) states that $B_k(\vec{\omega})$ is the n th-order directional derivative, in direction θ_k , of the function $B(\vec{\omega})/|\vec{\omega}|^n$.

3. FILTER DESIGN

In previous work, we have found that weighted least-squares frequency-domain designs are suitable for steerable filters [6]. This approach is not directly applicable in

the current context since the recursion constraint of equation (3) and the system response of equation (2) lead to an error function that is not quadratic in the filter taps. In our experience, direct attempts to iteratively minimize a weighted combination of the constraints given in equations (2) through (5) did not converge.

We thus introduce further assumptions to simplify the design process. We first note the similarity of equations (2) and (3). During recursion, $L_1(\vec{\omega}/2)$ plays the role of the initialization filter $L_0(\vec{\omega})$. We enforce this by setting $L_0(\vec{\omega}) = L_1(\vec{\omega}/2)$.

This restriction suggests a *recursive* design procedure, as illustrated in the flow graph of figure 3. We begin by fixing the number of orientation bands, n , and the size of all filter kernels, and choosing an initial $L_0(\vec{\omega})$. We then design filters $L_1(\vec{\omega})$ and $B_k(\vec{\omega})$ with error criterion a weighted sum of the maximum absolute errors of each of the constraints (3), (4), and (5). The resulting filter $L_1(\vec{\omega})$ is then decimated to produce a new $L_0(\vec{\omega})$ filter, and the process is repeated. At each recursive step, the filter $L_0(\vec{\omega}) = L_1(\vec{\omega}/2)$ is used as a frequency-weighting function in the design of $L_1(\vec{\omega})$ and $B_k(\vec{\omega})$, as indicated in the constraint of equation (3). Note that if the procedure converges, we may design $H_0(\vec{\omega})$ at the end to satisfy the power-complementary condition in equation (2). In practice, we have found that the recursive procedure converges quickly (typically after three iterations).

To initialize the design procedure, we solve for filter kernels that minimize the weighted sum of the mean square errors of the constraints given in equations (3), (4) and (5). Within the recursive design procedure we use a minimax error function (i.e., a weighted sum of the maximum absolute errors in the constraints). The recursion constraint (3) was given twice the weight of the other two constraints.

At each recursive step of the procedure illustrated in figure 3, the numerical optimization of the filters $L_1(\vec{\omega})$ and $B_k(\vec{\omega})$ is carried out using the BFGS method with cubic and quadratic line searches. This method has been found to be very efficient for spaces with high dimensionality. In the first derivative case, the dimensionality of the space for a 17×17 filter L_1 and 9×9 filters B_k is 60. We have assumed linear phase filters, and taken full advantage of symmetry

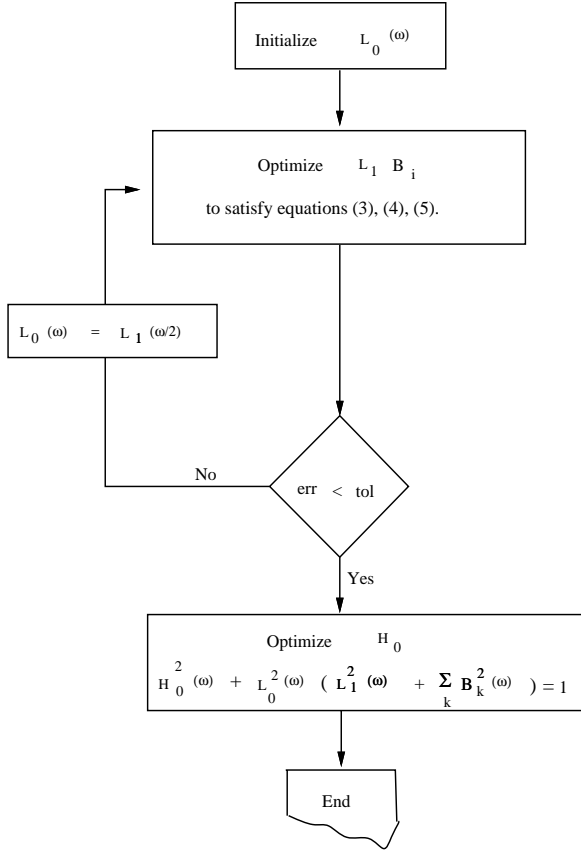


Figure 3. Flowgraph of the steerable pyramid filter design algorithm.

wherever possible. In particular, the two band-pass filters are exact rotated copies of each other, and thus we need only optimize over the independent coefficients of one of them.

3.1. Higher-order steerable pyramid filter design

In figures 1 and 2, we show the frequency tiling and the system diagram for the first derivative steerable pyramid. Higher derivative filters are necessary when we require finer orientation tuning of the filters (i.e., higher angular resolution). In particular, consider the third derivative case, in which the number of band-pass filters is 4. The axis of symmetry for filter B_k is oriented at $45k$ degrees. The combination of a rectangular sampling lattice and this set of filter orientations introduces a practical problem in the design: filter B_1 cannot be obtained from B_0 by simple rotation of the filter kernel by 45 degrees. Note, however, that B_2 and B_0 differ by a 90 degree rotation, as do B_3 and B_1 . Thus, we must include independent coefficients for two of the four band-pass filters in our optimization function.

Because of this significant increase in the dimensionality of our optimization space (as compared with the first derivative case), the first-derivative design algorithm would be impractical in terms of design time and the reliability of the solution. To remedy the above problems we fixed filter L_1 (and therefore filter L_0 which is subsampled version of




Figure 4. Fourier spectra of a set of zeroth derivative steerable pyramid filters plotted over the range $[-\pi, \pi] \times [-\pi, \pi]$. (a) Low-pass filter $L_0(\vec{\omega})$. (b) High-pass $H_0(\vec{\omega})$. (c) Low-pass $L_1(\vec{\omega})$. (d) Band-pass filter $L_0(\vec{\omega})B_0(\vec{\omega})$.




Figure 5. Fourier spectra of a set of first derivative steerable pyramid filters plotted over the range $[-\pi, \pi] \times [-\pi, \pi]$. (a) Low-pass filter $L_0(\vec{\omega})$. (b) High-pass $H_0(\vec{\omega})$. (c) Low-pass $L_1(\vec{\omega})$. (d) Band-pass filter $L_0(\vec{\omega})B_0(\vec{\omega})$. (e) Band-pass filter $L_0(\vec{\omega})B_1(\vec{\omega})$.

L_1) to be the one derived in the first derivative case. We then carried out the optimization for the B_0 and B_1 filters only.

4. EXAMPLE FILTERS

We have designed steerable pyramid filters using the technique described above for $n = 0$, $n = 1$, and $n = 3$. In figures 4, 5 and 6, we show two-dimensional images of the spectra of the steerable filter sets for these cases. The filter sizes are 9×9 for $B_k(\vec{\omega})$ and $L_0(\vec{\omega})$, and 17×17 for $L_1(\vec{\omega})$.

Given the above filters, a first derivative pyramid decomposition of a synthetic “zone plate” test image ($\cos(r^2)$) is shown in figure 7. In figure 8 we illustrate a third derivative pyramid decomposition of this image. Reconstruction errors in both cases are quite small (see figure captions).




Figure 6. Fourier spectra of a set of third derivative steerable pyramid filters, plotted over the range $[-\pi, \pi] \times [-\pi, \pi]$. (a) Initial low-pass $L_0(\vec{\omega})$. (b) High-pass $H_0(\vec{\omega})$. (c) Low-pass $L_1(\vec{\omega})$. (d)-(g) Band-pass filters $L_0(\vec{\omega})B_k(\vec{\omega})$.




Figure 7. Image decomposition of the “zone-plate” using the first-derivative steerable pyramid filters. Shown are the two orientation bands at three different scales and the final low-pass band. In the lower right is the original image. The high-pass band is not shown. Reconstruction error for this image is ≈ 47 dB.




Figure 8. Image decomposition of the “zone-plate” using the third-derivative steerable pyramid filters. Shown are the four orientation bands at three different scales and the final low-pass band. The high-pass band is not shown. Reconstruction error for this image is ≈ 40 dB.

5. CONCLUSION

We have presented a recursive filter design technique for the construction of steerable pyramid image transformations. For each cycle of the algorithm, the low-pass filter from the previous cycle is decimated and used as a frequency weighting function in designing the next set of low-pass/band-pass filters. We have used this technique to design zeroth-order (non-oriented), first-order (two orientation bands) and third-order (four orientation bands) steerable pyramid filters.

REFERENCES

- [1] W T Freeman and E H Adelson. The design and use of steerable filters. *IEEE Pat. Anal. Mach. Intell.*, 13(9):891–906, 1991.
- [2] Eero P Simoncelli, William T Freeman, Edward H Adelson, and David J Heeger. Shiftable multi-scale transforms. *IEEE Trans. Information Theory*, 38(2):587–607, March 1992. Special Issue on Wavelets.
- [3] H. Greenspan, S. Belongie, R. Goodman, P. Perona, S. Rakshit, and C. H. Anderson. Overcomplete steerable pyramid filters and rotation invariance. In *Proceedings, CVPR*, pages 222–228, 1994.
- [4] Eero P Simoncelli and William T Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *Second International Conference on Image Processing*, Washington, DC, November 1995.
- [5] I Daubechies. Ten lectures on wavelets. In *Regional Conference: Society for Industrial and Applied Mathematics*, Philadelphia, PA, 1992.
- [6] Eero P Simoncelli. Design of multi-dimensional derivative filters. In *First International Conference on Image Processing*, Austin, Texas, November 1994.