In Proceedings of the 5th International Conference on Computer Vision, 1995, Boston, MA.

Steerable Wedge Filters

E. P. Simoncelli and H. Farid
GRASP Laboratory
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6228

Abstract

Steerable filters, as developed by Freeman and Adelson, are a class of rotation-invariant linear operators that may be used to analyze local orientation patterns in imagery. The most common examples of such operators are directional derivatives of Gaussians and their 2-D Hilbert transforms. The inherent symmetry of these filters produces an orientation response that is periodic with period π, even when the underlying image structure does not have such symmetry. This problem may be alleviated by reconsidering the full class of steerable filters. In this paper, we develop a family of even- and odd- symmetric steerable filters that have a spatially asymmetric "wedge-like" shape and are optimally localized in their orientation response. Unlike the original steerable filters, these filters are not based on directional derivatives and the Hilbert transform relationship is imposed on their angular components. We demonstrate the ability of these filters to properly represent oriented structures.

1 Introduction

Oriented filters are used in many vision and image processing tasks, such as edge detection, texture analysis, motion analysis, etc. Several authors have addressed the problem of extracting orientation information using rotation-invariant operators [1, 3, 4, 2]. In [2], Freeman and Adelson propose a framework for synthesizing an arbitrary oriented filter from a linear combination of a fixed set of *steerable basis filters*. They designed a set of such steerable filters based on directional derivatives of Gaussians, and used these filters and their Hilbert transforms to compute "local orientation maps". Directional derivatives of Gaussians impose a periodicity (of period π) on the orientation map. For example, an orientation map computed at the end of a line segment will produce a *bimodal* response as a function of orientation. This problem may be alleviated by reconsidering the full class of steerable filters.

In this paper, we describe a new set of steerable filters that are *not* based on directional derivatives, but that are optimally localized in their oriented energy response. [1] We begin by reviewing the basic principles of steerability, and develop a design criterion for steerable filters with localized angular response. This design strategy is used to create a set of even- and odd-symmetric filters, which are then used to derive orientation maps on both synthetic and real images. These examples suggest that the filters are well-suited for applications in junction analysis, detection, and classification.

2 Steerability

This section reviews the principle of steerability, as developed by Freeman and Adelson [2]. The presentation is based on that given in [7, 6]. The simplest illustration of the steerability principle is the directional derivative of a two-dimensional Gaussian, $G(x, y)$. The directional derivatives of the Gaussian (in polar coordinates) in the horizontal and vertical directions are:

$$\begin{aligned} G_1^{(0)}(r,\theta) &= \cos(\theta)(-r\ e^{-r^2/2}) \\ G_1^{(\pi/2)}(r,\theta) &= \sin(\theta)(-r\ e^{-r^2/2}), \end{aligned}$$

where the subscript indicates the order of the derivative, and the parenthesized superscript indicates the angle of the derivative direction. It is straightforward to show (using basic trigonometric identities) that the function G_1 can be synthesized at an arbitrary orientation ϕ using the following equation:

$$\begin{aligned} G_1^{(\phi)}(r,\theta) &= \cos(\phi)G_1^{(0)}(r,\theta) \\ &+ \sin(\phi)G_1^{(\pi/2)}(r,\theta). \end{aligned} \quad (1)$$

[1] The filters are *wedge-like* in shape, an idea suggested in [5].

This equation embodies the steering property of the directional first derivative filters: the directional derivative G_1 can be generated at an arbitrary orientation ϕ as a linear combination of the *basis filters* $G_1^{(0)}$ and $G_1^{(\pi/2)}$, where the coefficients $\cos(\phi)$ and $\sin(\phi)$ are referred to as the *interpolation functions*. Freeman & Adelson also showed that the steering property extends naturally to directional derivatives of higher order: an $n-$th order derivative requires $n+1$ rotated basis filters in order to perform the interpolation properly.

Amongst several applications, Freeman and Adelson present a method of orientation analysis based on the outputs of quadrature pairs. The orientation strength, or "energy", is defined as the sum of squared outputs of the n-th order directional derivative of a Gaussian, $G_n^{(\phi)}$, and its Hilbert transform, $H_n^{(\phi)}$. Results of applying a quadrature pair of fourth-order filters (G_4/H_4) to several synthetic images are shown in Figure 1. The response of the filters to the vertical line and cross are as expected. However, the filters respond bi-modally to the half-line at $\phi = \pi/2$ and $\phi = 3\pi/2$ (instead of exhibiting a single peak at $\phi = \pi/2$). The response to the corner is bi-modal, but the peaks occur near $\phi = 3\pi/4$ and $\phi = 7\pi/4$ instead of $\phi = \pi/2$ and $\phi = \pi$.

These "incorrect" responses at asymmetric regions in the image are due to the inherent symmetry of directional derivatives: the filter response at any angle ϕ must be the same as the response at angle $\phi + \pi$. The cause of this symmetry is apparent when considering the properties of directional derivatives. An n-th order directional derivative has an angular component that contains a subset of the harmonics between the first and the n-th. In particular, for odd n, the directional derivative contains only the odd harmonics, and for even n, it contains only the even harmonics. Therefore, an even-order derivative kernel will always exhibit symmetry (when reflected through the origin) and an odd-order derivative kernel will exhibit *anti-symmetry*. In either case, the squared response of the operator will always be symmetric (that is, the angular component of the squared energy response will be periodic with period π).

Although the directional derivatives of a Gaussian provide a useful example, the principle of steerability is not limited to such functions. The conditions under which any function can be made steerable are outlined in [2], and essentially amount to a bandlimiting condition on the angular component of the filters. Given

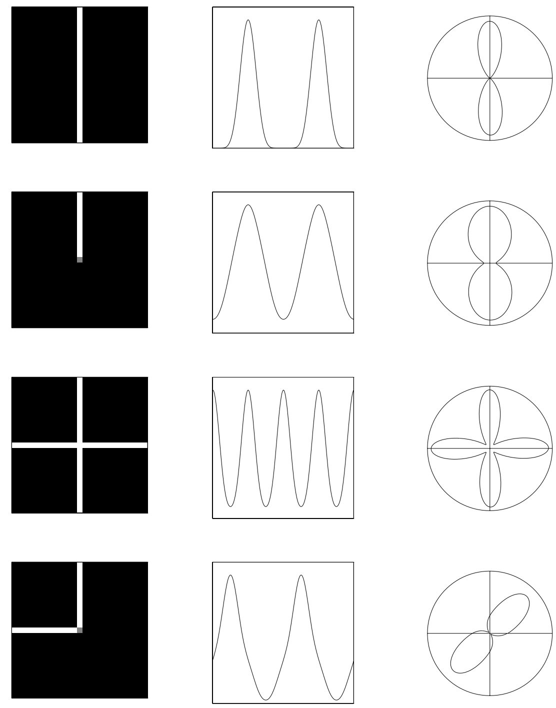

Figure 1: Middle column contains graphs of oriented energy as a function of angle at the center of the test images on the left. Far right column contains polar plots of the same functions. Oriented energy was measured using the set of eleven 13-tap G_4/H_4 quadrature pair filters presented in [2]. Note the forced symmetry in the oriented energy response even when the underlying image structure is asymmetric (i.e. the half-line and corner).

that directional derivatives always exhibit symmetry, we set out to design a more general steerable function that is *not* based on directional derivatives.

3 Design of Steerable Wedge Filters

This section discusses the design of a set of even- and odd-symmetric steerable wedge filters. We assume a polar-separable form for the filters, and describe the independent design of the angular and radial filter components.

3.1 Angular Function

The general steerability condition, in polar-separable form, is given by:

$$h(\alpha - \phi)g(r) = \sum_{n=1}^{\hat{N}} k_n(\alpha)h(\alpha_n - \phi)g(r),$$

where $h(\cdot)$ is the angular portion of the steerable kernel, $g(r)$ is the radial portion, $k_n(\cdot)$ are interpolation functions, and α_n are a fixed set of $\hat{N}$ orientations[2]. We now consider the derivation of a set of even- and odd-symmetric steerable wedge filters (g_e and g_o, respectively). To achieve steerability with $2N$ filters, the angular portion of each filter kernel must be constrained to be a weighted sum of the first N circular harmonics. Rather than construct 2-D Hilbert transform pairs, we design a pair of filters whose *angular* components are Hilbert transforms. The angular portions of the filters are then of the following form:

$$g_e(\phi) = \sum_{n=1}^{N} w_n \cos(n\phi) \tag{2}$$

$$g_o(\phi) = \sum_{n=1}^{N} w_n \sin(n\phi). \tag{3}$$

It is easily verified that these angular functions will be steerable with $\hat{N} = 2N$ basis filters. In order to eliminate symmetry in these filters, they are designed so as to maximally localize the oriented energy "impulse" response: $E(\phi) = g_e^2(\phi) + g_o^2(\phi)$. In order to minimize the width of this impulse response, we write a quadratic penalty function on the Fourier coefficients, $\{w_n\}$, as: $\mathcal{P}(\{w_n\}) = \int d\phi \, \lambda(\phi)^2 \left[g_e^2(\phi) + g_o^2(\phi)\right]$. The function $\lambda(\phi)$ is a weighting function, which we choose as $\lambda(\phi) = |\phi|$. To facilitate computer implementation, $\mathcal{P}(\{w_n\})$ is written in matrix form by sampling the variable ϕ at M locations:

$$\begin{aligned} P(\vec{w}) &= |\Lambda C \vec{w}|^2 + |\Lambda S \vec{w}|^2 \\ &= \vec{w}^t \left[C^t \Lambda^t \Lambda C + S^t \Lambda^t \Lambda S\right] \vec{w}, \end{aligned} \tag{4}$$

where $\vec{w}$ is a vector whose components are the weights w_n, Λ is a diagonal matrix with the values $|\phi|$ on the diagonal, and S and C are matrices whose columns contain the sampled sinusoidal and cosinusoidal basis functions, $S_{mn} = \sin(2\pi nm/M)$, and $C_{mn} = \cos(2\pi nm/M)$, respectively. From equation (4), the unit vector $\vec{w}$ that minimizes P is simply the minimal-eigenvalue eigenvector of the matrix $[C^t \Lambda^t \Lambda C + S^t \Lambda^t \Lambda S]$.

Examples of even- and odd-symmetric angular functions, for $M = 512$ and $N \in \{2, 4, 6, 8\}$, are given in Figure 2. The corresponding oriented energy responses are given in Figure 3. Note that as N increases (i.e. higher frequency components are added), the oriented energy response and the angular functions become narrower. However, this narrowing comes at a price: the size of the basis set increases by two with each additional frequency component.

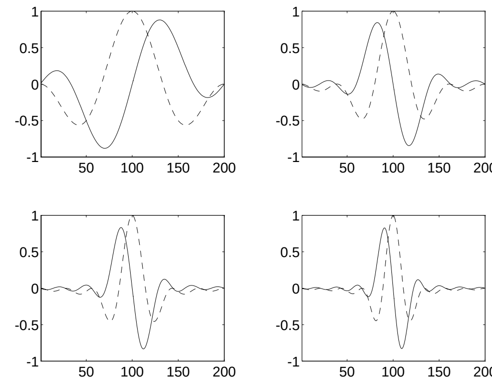

Figure 2: One-dimensional even- and odd-symmetric angular functions with $N = 2, 4, 6, 8$ (left to right/top to bottom). The width of the angular function decreases with increasing N, but the number of filters in the steerable basis set increases.

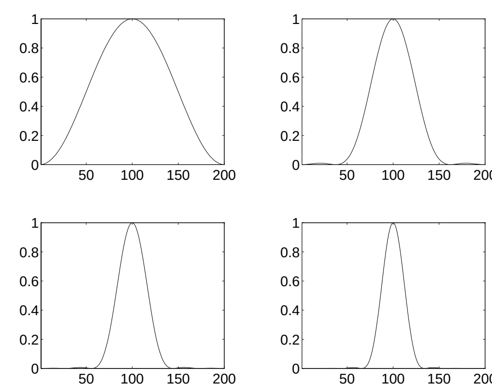

Figure 3: 1D oriented energy response with $N = 2, 4, 6, 8$ (left to right/top to bottom). The width of the energy response decreases with increasing N, but the number of filters in the steerable basis set increases.

3.2 Interpolation Functions

We now derive the interpolation functions $k_n(\alpha_n)$ for the even-symmetric filters; the derivation for the odd-symmetric filters is similar. From equation (2), the angular portion of the even-symmetric kernel is given by:

$$h(\phi - \alpha) = \sum_{n=1}^{N} w_n \cos(n(\phi - \alpha))$$

[2] The directional derivative of a Gaussian is the simplest such filter set, with $\hat{N} = 2$, $h(\phi) = \cos(\phi)$, $g(r) = r\gamma'(r)$, $\gamma(r)$ the unit normal function, and $\alpha_n \in \{0, \pi/2\}$.

$$
\begin{aligned}
&= \sum_{n=1}^{N} w_n[\cos(n\phi)\cos(n\alpha) + \sin(n\phi)\sin(n\alpha)] \\
&= \vec{k}(\alpha) \cdot \vec{f}(\phi), \qquad (5)
\end{aligned}
$$

where the vectors $\vec{k}(\alpha)$ and $\vec{f}(\phi)$ are defined as:

$$
\vec{k}(\alpha) = \begin{pmatrix} \cos(\alpha) \\ \sin(\alpha) \\ \cos(2\alpha) \\ \sin(2\alpha) \\ \dots \\ \cos(N\alpha) \\ \sin(N\alpha) \end{pmatrix}^t, \quad \vec{f}(\phi) = \begin{pmatrix} w_1 \cos(\phi) \\ w_1 \sin(\phi) \\ w_2 \cos(2\phi) \\ w_2 \sin(2\phi) \\ \dots \\ w_N \cos(N\phi) \\ w_N \sin(N\phi) \end{pmatrix}.
$$

Letting the steerable basis functions be defined at $\alpha_j = j\pi/N$, for $j \in \{1, 2, \dots, 2N\}$, we can write equation (5) for each basis filter, and combine the set of equations into a matrix equation:

$$
\vec{h}(\phi) = K\vec{f}(\phi), \qquad (6)
$$

where, $\vec{h}(\phi)$, contains the steerable filter set:

$$
\vec{h}(\phi) = (\, h(\phi - \alpha_1) \quad h(\phi - \alpha_2) \quad \dots \quad h(\phi - \alpha_{2N})\,)^t,
$$

and K is a *fixed* matrix that depends on the samples α_j and the number of frequencies N:

$$
K = \begin{pmatrix} \cos(\alpha_1) & \sin(\alpha_1) & \dots & \cos(N\alpha_1) & \sin(N\alpha_1) \\ \cos(\alpha_2) & \sin(\alpha_2) & & \cos(N\alpha_2) & \sin(N\alpha_2) \\ \vdots & & \ddots & & \vdots \\ \cos(\alpha_{2N}) & \sin(\alpha_{2N}) & \dots & \cos(N\alpha_{2N}) & \sin(N\alpha_{2N}) \end{pmatrix}.
$$

Equation (6) gives a linear relationship between the steerable basis set $\{h(\phi - \alpha_n); 0 \le n \le 2N\}$ and the Fourier basis $\{\cos(n\phi), \sin(n\phi); 0 \le n \le N\}$. Assuming that the matrix K is invertible (this is the condition required for steerability), we can solve for the Fourier basis as a function of the steerable basis:

$$
\vec{f}(\phi) = K^{-1}\vec{h}(\phi). \qquad (7)
$$

Combining equation (7) with equation (5) gives the steering relationship:

$$
h(\phi - \alpha) = \vec{k}(\alpha)K^{-1}\vec{h}(\phi).
$$

The interpolation functions are then simply the vector $\vec{k}(\alpha)K^{-1}$.

We have derived interpolation functions for the even-symmetric filters; a similar derivation would produce interpolation functions for the odd-symmetric functions. Computing oriented energy would then require $4N$ filters ($2N$ even and $2N$ odd). Such a basis is redundant: the entire basis of $4N$ functions spans only a $2N$-dimensional Fourier subspace. As such, we will derive interpolation functions using a *combination* of the even- and odd-symmetric filters, therefore requiring only $2N$ basis filters (N even, and N odd).

We denote the angular functions for the even- and odd- symmetric filters as $h_e(\phi)$ and $h_o(\phi)$, respectively, and then write an equation analogous to equation (5) for each of the two filters:

$$
\begin{aligned}
h_e(\phi - \alpha) &= \vec{k}_e(\alpha) \cdot \vec{f}(\phi) \qquad (8) \\
h_o(\phi - \alpha) &= \vec{k}_o(\alpha) \cdot \vec{f}(\phi), \qquad (9)
\end{aligned}
$$

where $\vec{f}(\phi)$ is the same weighted Fourier basis as above, and the interpolation vectors $\vec{k}_*$ are now given (from basic trigonometric identities) by:

$$
\vec{k}_e(\alpha) = \begin{pmatrix} \cos(\alpha) \\ \sin(\alpha) \\ \cos(2\alpha) \\ \sin(2\alpha) \\ \dots \\ \cos(N\alpha) \\ \sin(N\alpha) \end{pmatrix}^t, \quad \vec{k}_o(\alpha) = \begin{pmatrix} -\sin(\alpha) \\ \cos(\alpha) \\ -\sin(2\alpha) \\ \cos(2\alpha) \\ \dots \\ -\sin(N\alpha) \\ \cos(N\alpha) \end{pmatrix}^t.
$$

The collection of such equations for the filters at angles α_j can be written in matrix form (where sin and cos are abbreviated as S and C, respectively):

$$
\begin{pmatrix} h_e(\phi - \alpha_1) \\ h_o(\phi - \alpha_1) \\ \vdots \\ h_e(\phi - \alpha_{2N}) \\ h_o(\phi - \alpha_{2N}) \end{pmatrix} =
$$

$$
\begin{pmatrix} C(\alpha_1) & S(\alpha_1) & \dots & C(N\alpha_1) & S(N\alpha_1) \\ -S(\alpha_2) & C(\alpha_2) & & -S(N\alpha_1) & C(N\alpha_1) \\ \vdots & & \ddots & & \vdots \\ C(\alpha_{2N}) & S(\alpha_{2N}) & & C(N\alpha_{2N}) & -S(N\alpha_{2N}) \\ -S(\alpha_{2N}) & C(\alpha_{2N}) & \dots & -S(N\alpha_{2N}) & C(N\alpha_{2N}) \end{pmatrix} \cdot \vec{f}(\phi),
$$

which we denote as:

$$
\vec{h'}(\phi) = K'\vec{f}(\phi).
$$

Combining the inverse of this equation with equations (8) and (9) gives the desired linear expressions for h_e and h_o rotated to an arbitrary orientation α:

$$
\begin{aligned}
h_e(\phi - \alpha) &= \vec{k}_e(\alpha)(K')^{-1}\vec{h'}(\phi) \qquad (10) \\
h_o(\phi - \alpha) &= \vec{k}_o(\alpha)(K')^{-1}\vec{h'}(\phi). \qquad (11)
\end{aligned}
$$

where the basis set of *both even and odd* steerable filters is contained in the vector $\vec{h'}(\phi)$. Equations (10) and (11) give the interpolation functions for interpolating even and odd responses from a set of $2N$ even- and odd-symmetric filters.

3.3 Radial Function

Thus far, only the angular component of the steerable wedge filters have been specified. Since the filters are polar separable, the choice of the radial function does not affect the steerability of the angular function. For our purposes, we arbitrarily choose the following function:

$$g(r) = \begin{cases} 0 & \text{if } r > R \\ 1 & \text{if } r < R - \delta \\ \frac{1+\cos(\pi(r+\delta-R)\delta)}{2} & \text{otherwise} \end{cases} \qquad (12)$$

where, R is the radius of the filter, and δ is the transition width of the raised cosine function. Given this radial component, a set of 2D steerable wedge filter kernels are computed by sampling the polar-separable product $h(\phi - \alpha_j)g(r)$ (Figure 4).[3]

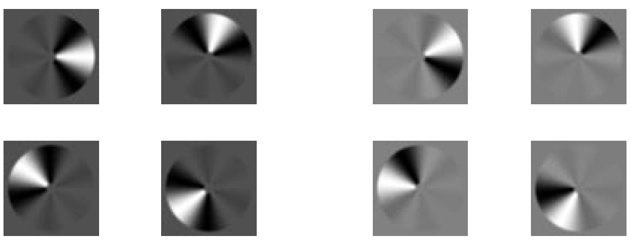

Figure 4: Four of ten (i.e., for $N = 5$) even-symmetric (left) and odd-symmetric (right) steerable wedge filters.

3.4 Orientation Analysis

The orientation strength or "energy" at an angle ϕ is defined as the sum of squared outputs of a pair of even- and odd-symmetric filters "steered" to the angle ϕ. Results of applying a pair of even- and odd-symmetric wedge filters (with $N = 5$: a basis set of 10 filters) to several synthetic images are shown in Figure 5. Consider the top two rows of Figures 5 and 1. The response to the vertical line in the two figures is similar. But, the original steerable filter (G_4/H_4) response to the half-line is *bi-modal*, while the wedge filters respond unimodally, consistent with the underlying image structure. The response of the wedge filters to the cross and corner are "incorrect" due to the rather broad tuning of the filters.

This problem is easily remedied by using more narrowly tuned filters. Results for the same set of images with a basis set corresponding to $N = 9$ (18 filters) are shown in Figure 6. With the larger basis set the over-smoothed responses to the cross and the corner are now eliminated. It is important to note that due to their inherent symmetry, more narrowly tuned G_n/H_n filters will not eliminate the bi-modal response to the half line or corner.

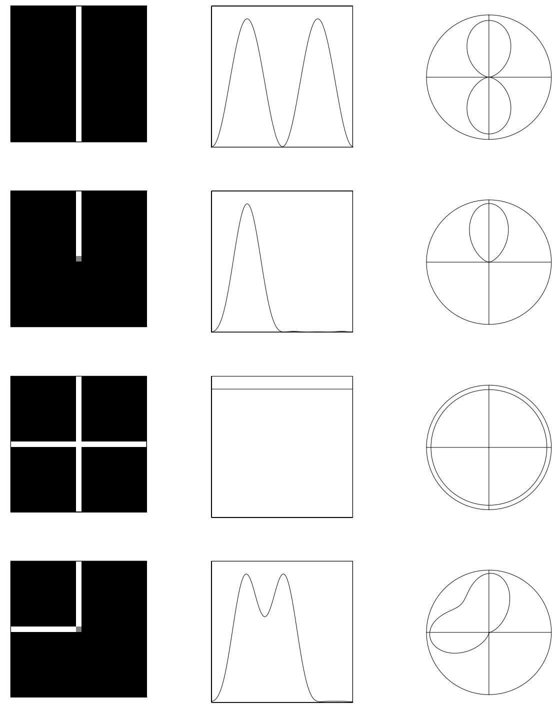

Figure 5: Middle column illustrates oriented energy as a function of angle at the center of the test images on the left. Far right column are polar plots of middle column. Oriented energy was measured using a set of 7-tap steerable wedge filters with $N = 5$ (i.e. 10 basis filters).

Figures 7 and 8 depict the results of applying 13-tap G_4/H_4 filters and 7-tap wedge filters ($N = 9$) to a "Y-junction" and "T-junction". The wedge filters respond with a tri-modal response to both the "T-" and "Y-junctions" (Figure 8), while the G_4/H_4 filters respond with a bi-modal and quadra-modal response (Figure 7). Results of applying 7-tap wedge filters ($N = 9$) to several regions in a real image are shown in Figure 10. For such tasks as junction analysis, detection, or classification, the response of the wedge filters is clearly favorable.

4 Discussion

The inherent symmetry of derivative-based steerable filters produces a symmetric orientation response (with period π) even when the underlying image structure does not have such symmetry. This type of re-

[3] Note that aliasing due to the sampling process will not affect the steerability property.

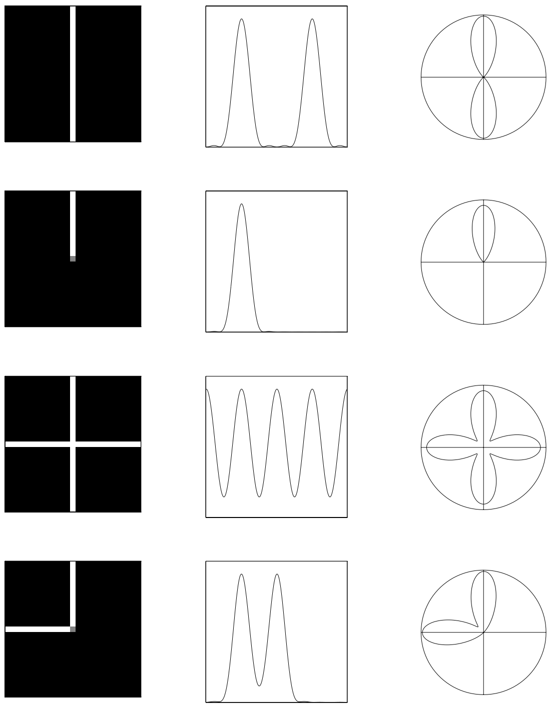

Figure 6: Middle column illustrates oriented energy as a function of angle at the center of the test images on the left. Far right column are polar plots of middle column. Oriented energy was measured using a set of 7×7 steerable wedge filters with $N = 9$ (i.e. 18 basis filters).

sponse is clearly not desirable for such tasks as junction analysis, detection or classification. The symmetry of Gaussian derivatives is due to the fact that even-order derivatives contain only even angular harmonics and odd-order derivatives contain only odd angular harmonics. Thus, the squared response of any derivative filter will be periodic (with period π). Since the steerable wedge filters presented in this paper contain both even and odd harmonics, this periodicity is not enforced. Unlike the original steerable filters, the steerable wedge filters are localized in their oriented energy response, therefore, the wedge filters respond "correctly" to asymmetric regions in the image (e.g. line endings, corners, "T-", "Y-junctions").

One drawback of the wedge filters (relative to the steerable derivative filters) is that they are not x-y separable. However, we found that 7-tap filters are sufficient to achieve satisfactory results.

We believe that the filters presented here will prove to be extremely useful in analyzing the local orientation structure of images. Future work will include development of an algorithm for rotation-invariant pattern classification, suitable for identifying junctions.

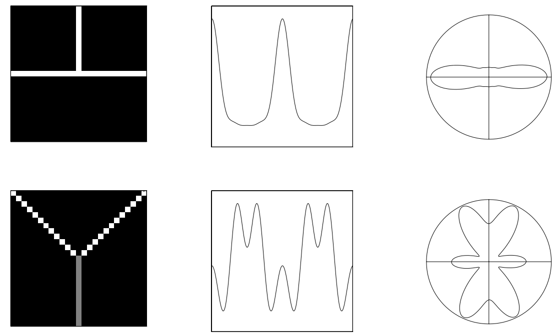

Figure 7: Middle column illustrates oriented energy as a function of angle at the center of the test images on the left. Far right column are polar plots of middle column. Oriented energy was measured using a set of 13-tap G_4/H_4 quadrature pair filters. Note the forced symmetry in the oriented energy response.

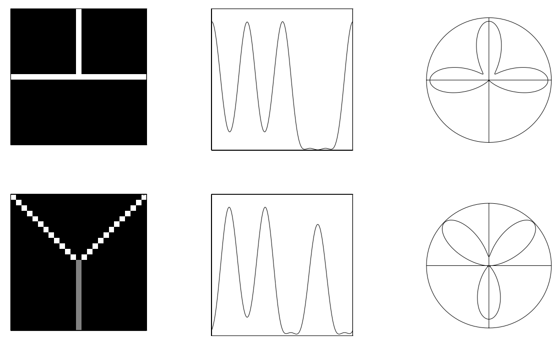

Figure 8: Middle column illustrates oriented energy as a function of angle at the center of the test images on the left. Far right column are polar plots of middle column. Oriented energy was measured using a set of 7-tap steerable wedge filters with $N = 9$.

The results are also readily extended to three dimensions, where the resulting filters may be used to analyze volumetric data or motion sequences. In particular, a three-dimensional set of such filters may prove useful for recognizing motion occlusion boundaries.

Acknowledgments

This work was supported by ARPA Grant N00014-92-J-1647; ARO Grant DAAL03-89-C-0031PRI; NSF Grant CISE/CDA-88-22719.

References

[1] P. Danielsson. Rotation-invariant linear operators with directional response. In 5^{th} *Int'l Conf. Patt. Rec.*, Miami, FL, 1980.

[2] W. Freeman and E. Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern*

Figure 9: Park bench image. See figure 10 for oriented energy responses.

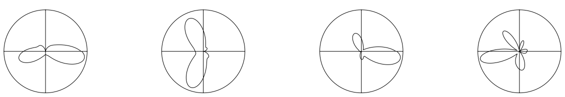

Figure 10: Results of applying 7-tap steerable wedge filters with $N = 9$ (i.e. 18 basis filters) to the parkbench image (figure 9). Polar plots of oriented energy are given for: (a) a horizontal edge; (b) a vertical edge; (c) a corner; and (d) a "T-junction".

Analysis and Machine Intelligence, 13(9):891–906, 1991.

[3] H. Knutsson and G. Granlund. Texture analysis using two-dimensional quadrature filters. In *IEEE Comput. Soc. Workshop Comp. Architecture Patt. Anal. Image Database Mgmt.*, pages 388–397, 1983.

[4] P. Perona and J. Malik. Detecting and localizing edges composed of steps, peaks, and roofs. In *International Conference on Computer Vision*, 1990.

[5] G. Provan, H. Farid, and E. Simoncelli. A novel radial intensity based edge operator. Technical Report MS-CIS-94-07, Department of Computer Science, University of Pennsylvania, 1994.

[6] E. Simoncelli. *Distributed Analysis and Representation of Visual Motion*. PhD thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, MA, January 1993. Also available as MIT Media Laboratory Vision and Modeling Technical Report #209.

[7] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger. Shiftable multiscale transforms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):587–607, 1992.